



613-000629 Rev.A 061006

fMAP
User Guide
Release 8.0
Issue 2



Introduction to the fMAP Series

Congratulations on your purchase of a Factory Automation Multiservice Access Platform (fMAP) product.

Who Should Read This Guide?

This document is for those who perform all tasks for the fMAP Series products once it has been successfully installed. Users of this Guide should understand Ethernet technology and have experience configuring Ethernet devices.

About this Guide

This guide includes all Administration, Feature, and Maintenance aspects of the fMAP Series products.

- Section 1 provides an overview of how the document is organized so that all types of user tasks for the fMAP Series products can be identified and accessed quickly in this document or related documents.
- Section 2 provides an overview of the fMAP Series products and includes a description of the physical and functional components.
- Section 3 shows how to set up the management interfaces that allow the fMAP Series products to be accessed and controlled from local and remote locations.
- Sections 4 - 8 show how to provision the Service, Network and Control Modules. Included is a description of how auto-provisioning allows these components to be installed and quickly put into service.
- Section 9 shows how to configure Performance Management.
- Sections 10 and 11 highlight software load control and how to configure the system for easy downloads and upgrades, as well as the procedures for performing software upgrades.
- Section 12 highlights existing and new features for 802.3 Ethernet, VLAN, and UFO.
- Section 13 highlights existing and new features for Topology Configuration.
- Section 14 highlights existing and new features for IGMP.
- Section 15 highlights existing and new features for Traffic Management.
- Section 16 highlights existing and new features for the fMAP.
- Section 17 highlights existing and new features for Routine Administration.
- Section 18 lists the trouble indicators for the product and procedures to resolve a fault.
- Section 19 contains technical specifications, including MIB objects.



Copyright Notices

All company names, logos, and product designs that are trademarks or registered trademarks are the property of their respective owners.

Copyright (c) 1990, 1993

The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by

Van Jacobson.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE



FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

1. About This fMAP Document

1.1 Overview

In general, this document follows the life cycle of the Factory Automation Multiservice Access Platform (fMAP) product, starting with when the fMAP product has just been installed and turned up (refer to the various **Installation Guides**) to adding new software, components, and features. Each section uses knowledge that has been gained from a previous section; although reading a section will give you, the user, enough information to perform the included tasks, a better understanding of the product is gained by reviewing related information from previous sections first.

1.2 Document Sections

1.2.1 Section 1 - Document and fMAP Product Overview

This section gives an outline of the User Guide, provides a roadmap of the features supported by **all** fMAPs, and lists the components for the product and which features they are compatible with.

Note: Certain sections are applicable to all fMAP products since they share the same features and functions. These sections will highlight that they are common for all fMAP products. Other sections may only apply to specific products.

1.2.2 Section 2 - Overview of the fMAP Series Product

Section 2. provides an overview of the following:

- How the fMAP product fits into the network.
- How the shelf and interfaces are configured.
- The main attributes of the configuration once the shelf is powered up.

1.2.3 Sections 3 through 18 - Functions and Features

After a section that is a general description of a set of features and functions, there is a section that highlights how the feature or function operates on the fMAP product.

Each section is divided into subsections that highlight the main components as well as the commands and parameters used. The subsections are usually in the following order:

- Overview - A description of a function or feature and how it affects the fMAP product. Whenever possible, a figure is provided to help in understanding how the feature or function works.
- Main tasks - The steps involved in any large task are highlighted so the user knows how the system will be affected. A table is then presented that shows the sequence of tasks and commands.
- Components - The main components are highlighted with headings, which in many cases will match a command or parameter.
- Examples - A relevant example with commonly used commands/parameters helps the user fully understand the feature and can be applied immediately (in many cases as scripts). Whenever possible, inputs and outputs from the user interfaces are supplied.
- Listing of commands - A table lists all the commands and parameters associated with what has been described. The table will be grouped by main objects, and within each object the actions (verbs) will be listed in the order they would most likely be performed. Each verb/action includes a description with the most common parameters used.

1.2.4 Section 17 and 18 - Overview of the fMAP Maintenance System

This section provides an overview of how the fMAP products are configured to provide alarm indicators at both local and remote interfaces so that indicators for problems or potential problems can be recognized.

*Note: Refer to the **fMAP Log / Troubleshooting Manual** for detailed indicators and procedures.*

1.2.5 Section 19 - Specifications

This section includes the specifications supported by the fMAP product line, such as the MIB objects supported.

Note: The tables that lists the ATN Enterprise traps has been deleted; this information is included in the fMAP Log Manual.

1.3 How Commands are Presented

1.3.1 Command Explanations in each Section

Commands are usually presented in the following ways:

- Tables for specific functions or features that include important parameters
- Specific commands that are part of examples
- A table at the end of each Section that includes all commands and descriptions of all parameters.

1.3.2 Command Presentation in Examples

When a command is presented in an example, it follows the exact syntax and parameter values that match the example configuration. If a command is very long, a (->) is used to note the command continues on the next line.

1.3.3 Command Syntax

The syntax rules for a Command and its parameters use the following conventions throughout this document:

- All upper case = Key Word
- | = Option (OR)
- [] = Optional
- { = Choice of until }

1.3.4 Security Levels

Section 3.8 gives an overview of the three levels of security controlling the types of commands that can be entered:

- User
- Manager
- Security Officer

Note: All of the commands used when explaining features in this document assume the Security Officer privilege.

1.3.5 Editing Functions, Keystrokes, and Abbreviations

The fMAP product supports line editing, line recall, and abbreviations, so that command line input and editing can be done very quickly once command syntax and the line editing commands are learned. These are described in 3.8.3.

Note: Throughout this document all syntax will use complete words, with verbs and parameters in upper case and the pairing of parameters and values with equal (=) signs.

Table of Contents

Overview 1

Document Sections 1

- Section 1 - Document and fMAP Product Overview 1
- Section 2 - Overview of the fMAP Series Product 1
- Sections 3 through 18 - Functions and Features 1
- Section 17 and 18 - Overview of the fMAP Maintenance System 2
- Section 19 - Specifications 2

How Commands are Presented 2

- Command Explanations in each Section 2
 - Command Presentation in Examples 2
 - Command Syntax 3
 - Security Levels 3
 - Editing Functions, Keystrokes, and Abbreviations 3
-

9000 Series (non-10G) Overview 1

- The fMAP (CFC12) Shelf 1

Initial Status of fMAP Products 3

- Initial Interface 3
 - Initial System Status 3
-

Overview 1

Identifying the fMAP Product 2

- Overview 2
- Command Summary for Identifying the fMAP Product 2

Configuring Physical Interfaces and Protocols 3

- Overview 3
- fMAP 4
- Configure the IP Interface - Example 5

Command Summary for the IP Interface 6

Setting Up the Management Interfaces - Overview 9

Configuring the SNMP Community 10

Overview (Product Support) 10
SNMP Components 10
Securing an SNMP Community 12
Setting Up an SNMP Community 12
Disabling an SNMP Community or its Traps 13
Setting Up an SNMP Community - Example 14
Command Summary for SNMP 17

Configuring Log Filtering and Output 18

Overview 18
Viewing Logs 20
Controlling Output of Logs 20
Example Log Configuration Setup 21
Capturing and Sending Logs to a Storage Device 22
Command Summary for Log Management 23

Setting up Simple Network Time Protocol (SNTP) 27

Overview 27
Command Summary for SNTP 28

Setting up the CLI 29

Command and Session Overview 29
Password Recovery 30
Editing Functions, Keystrokes, and Abbreviations 33
Control of CLI command confirmation 34
Command Alias 35
Provisioning the Login Banner 40
Customizing the CLI Prompt 41
Setting Up User Accounts, Profiles, and Sessions 43
Command Summary for User Administration 44
TACACS+ and RADIUS Authentication 47
Radius Details and Commands 48
TACACS+ Details and Commands 53
Using Help 57
Display and Clear ARP Table 59
Filtering on the MGMT interface 61
Telnet Client (telnet to another device from CLI) 62

Overview 1

- Provisioning in Release 8.0 for fMAP Products 1
- Provisioning Mode (Manual and Automatic) 2
- The AUTOPROV Profile 3
- User Created Profiles (Starting in Release 6.0) 3
- Administrative and Operational States 16
- Software Loads 16
- Provisioning Data at Startup 17
- Provisioning Mode (SHOW SYSTEM PROVMODE) 18

SM Category Attributes 18

- SM Card 18
- SM Interface 20

NM Category Attributes 23

- NM Card 23
- NM Interface 23

Provisioning Interfaces 25

- Overview 25
- Interface Provisioning 26
- Displaying ADSL Interface information 27
- Interface States 30
- Port Attributes no Longer Supported 30
- VC Configuration 30
- EPON/ONU Configuration 31

Control Module Provisioning Data 31

- Overview (Simplex versus Duplex) 31
- CFC Card Attributes and States (SHOW CARD ACTCFC) 32
- Changing the Administrative State of the Inactive CFC 35
- Manual versus Automatic Swaps 36

Provisioning Scenarios for Control Modules 37

- Overview 37
- Simplex to Duplex (AutoProv Mode) 37
- Duplex to Simplex (AutoProv Mode) 38
- Simplex to Duplex (Manual Mode) 38
- Duplex to Simplex (Manual Mode) 39

Provisioning the 9100 (GE Settings) 39

- Overview 39
- GE Speed Settings are Configurable 39

Command Summary for Provisioning 40

Overview 1

7000/9000 1

Fast Ethernet and Fiber-based Interfaces 2

Overview 2

FE10 Interface 3

FX10 Interface 4

9100 Ethernet Interfaces 5

Overview 1

ADSL Card 1

ADSL Interface Attributes 2

Output of SHOW INTERFACE Command 2

ADSL Mode Selection 3

Annex A versus Annex B 8

ADSL Cards Starting with Release 6.0 (ADSL24A) 8

Listing of ADSL Interface Attributes 8

Provisioning Scenarios for ADSL Cards 13

Command Summary 15

Overview 1

CES8 Configuration 1

Concepts and Terms 1

Common Terms 2

Provisioning CES8 2

CES8 Terms 2

CES Connectivity in Release 6.0 3

Provisioning Model 3

PSPAN 4

CES8 Card Attributes 7

CES8 Port 8

IP Interface (IP Address/VLAN) 9

DS1/E1 Connection (PSPAN<->DS1) 9

8KHz Timing References 9

Commands and Usage Notes 12

Example CES8 Configurations 13

Same Shelf, Card Timing (Internal Oscillator) 14

Across a Network, External Timing 16

Engineering the Packet Network for CES 18

Command Summary for CES8 Card/Interface 19

Common Command Summary for DS1/E1 24

Overview 1

EPON Configuration 1

Concepts and Terms 2

Common Terms 2

EPON Connectivity in Release 8.0 2

Provisioning Components 4

Provisioning Models 6

Video 6

Data, Voice 7

Provisioning Model - BRUUM 10

Traffic Management 11

Classifiers 11

QoS (Traffic Queues/Priorities) 12

Connection Admission Control (CAC) 13

Feature Interaction 13

Example Configuration 14

Example Figures 14

Pre-provisioning Tasks 15

Major Tasks 16

Existing Commands 22

Command Summary for EPON/ONU 23

Command Set 23

Overview 1

RMON (Ethernet-Based) Statistics 2

Overview 2

Example of Configuring RMON Statistics 4
Sample Management Logs for RMON (Ethernet-Based) Thresholds 6

PMON (ADSL Port) Statistics 7

Overview 7
Example of Configuring PMON Statistics (ADSL) 9
ADSL24 Egress Queue Counts are not supported 11
Sample Management Logs for PMON (ADSL-Based) Thresholds 13
Commands for Setting Up Performance Management 13

Monitoring Performance Management 18

Overview 18
RMON Collection 18
PMON Collection 20
Commands for Collecting Performance Management Data 23

Retrieving IP Statistics 24

IP Routing 25
TCP 26
UDP 26
ICMP 27
Access of MIB Statistics Using an SNMP Browser 27

Performance Monitoring for CES 28

IP Interface (RMON) 28
DS1/E1 Port Performance (PMON) 30
PSPAN Statistics (PMON) 32

Performance Monitoring for EPON 37

Overview 37
Example Outputs 38

Overview 1

File Management 2

File Names 2
File Storage 2
Compact FLASH 3
Commands for File Management 8

Software Load Management 11

Card Load Preferences 11
Load File Verification 12
Boot Server (Control Module Only) 12

Database Management 15

Overview 15

Database in upgrade mode 16

Text File Configuration 17

Overview 17

Creating a Text Configuration file 17

Restoring a Configuration Database Using a Text Configuration File 18

Stopping a Backup/Restore in Progress 19

Viewing the Progress of a BACKUP or RESTORE 19

Editing a Text Configuration File 20

Using Configuration Text Files During Upgrades/Downgrades 21

Summary of Commands 21

Software Compatibility 22

Overview 22

Supported Loads 23

Software Upgrade 23

Overview 23

Obtaining New Loads 23

Simplex Configuration 1

Overview 1

Simplex Upgrade Procedure 2

Upgrade 2

Downgrade 4

Layer 2 Switching (Learning) 1

Overview 1

Ingress Rules 1

Learning Process 2

Forwarding Process 3

Clearing of FDB (Selective or Global) 3

Egress Rules 4

Command Summary for Switch Learning 5

Virtual LAN (VLAN) 6

Overview 6

VLAN Tagging 6

Standard VLAN Configuration 7
MAC Address Limiting for an Interface 8
Command Summary for 802.1q VLAN 8

Upstream Forwarding Only (UFO) Mode 11

Overview 11
Configuration Rules for VLANs (UFO and Standard) in 6.0 12
9000 13

VLAN - Virtual Channel (VC) Mapping 13

Overview 13
Overview 13
Example VC provisioning 15
Usage Notes 17
Example configuration 17

VLAN Distribution 21

Overview 1

Spanning Tree Protocol (STP) 2

Overview 2
Protocol Concepts 3
Spanning Tree Parameters 5
An STP Network with Multiple VLANs 11
STP and LAG Interaction 12
Command Summary for STP 13

Link Aggregation (LAG) 14

Overview 14
Provisioning Rules 14
Destroying a LAG 15
Command Objects for Link Aggregation 16

Networking Topologies in Release 5.0 18

Overview 18
Linear Daisy Chain Network 19
Ring-based Topology with STP 22
Protection Schemes - EPSR 26
EPSR - Interoperability 37
EPSR Engineering Rules 37

Upstream Control Protocol (UCP) 41

Overview 41

UCP Protocol / Operation 42
UCP with EPSR Topology Feature in Release 6.0 46
UCP with STP Topology Feature in Release 6.0 49
Datafill for UFO VLANs when Upgrade to Release 6.0 50
Summary of UCP and Topology Engineering 51

EPSR and (R)STP Interaction 53

Overview 53
Summary of (R)STP 55
Configuring (R)STP for Interaction with EPSR 56
Port Costs 57
Example Configuration 57
Command Set for EPSR/(R)STP 57

MSTP 59

Overview 59
MSTP with a Primary and Secondary Upstream Port 62
MTSP Region 64
Provisioning Parameters 66
Network Engineering and Balancing VLAN/Port Configurations 69
Command Overview 69
Cisco-Compatible MSTP Mode 70
Command Summary 73

DHCP Relay 76

Overview 76
DHCP Relay Agent Functionality in 6.0 78
DHCP Relay Agent Functionality in 6.1 81
DHCP Relay Mode 81
DHCP Snooping Mode 84
Auto Ageing and IP Filter Removal 86
Default and Upgrade Configuration 86
Counters 87
DHCP Relay / Snooping on Network Interfaces (6.1.3) 87
DHCP Relay Commands 88
DHCP Relay Example 91

BPDU Cop 103

Link Layer Discovery Protocol (LLDP) 105

Overview 105
Concepts 107
Protocol Configuration Parameters 107
Functional Overview 112
Command Set for LLDP 116

Layer 1 Protection Groups 118

Overview 118

Usage Notes 119

State and Alarm Handling Rules for Layer 1 Protection Group 119

Internet Group Management Protocol (IGMP) Snooping 1

Overview of IGMP 1

IGMP Snooping 1

IGMP Snooping Disabled 2

IGMP Snooping Enabled 3

Interaction Between System and Ports/Interfaces 3

Reserved Multicast Range Behavior 3

IGMP Changes to Support FE/FX Upstream Interface 4

Overview 4

Router Interface 4

Customer Interface 4

Network Interface 5

Multicast Stream Counts 5

Key IGMP Snooping Syntax 5

IGMP Interactions 6

IGMP Group Limit (MCASTGROUPLIMIT) for Each ADSL24A Card 6

IGMP Multicast Handling for FLOODUNKNOWN (EPON2 Card) 7

MAC Limiting 7

ICCommand Summary for IGMP 8

Overview 1

Summary Table and Notes 1

Usage Notes: 4

EPON2 Traffic Management 5

Quality of Service (QoS) Model 6

Overview 6

Ingress Traffic 7

Egress Traffic 8

Traffic Management Throughout the Network 9

Classifier Management 9

Overview 9

Classifier Match Rules	10
Classifier Actions and COUNTs	13
Classifier Association with a Port or Interface (Precedence)	14
Default Classifier telesyn_default_video Behavior in Release 5.0	14
Classifier telesyn_default_video Behavior in Release 6.0	15
Derived Classifiers (D)	18
Set Match Rule Defaults (SETDEFAULTS)	18
System Monitoring for Errors (NORES, ERR, NOSPT)	20
Classifiers and Subinterface interactions	23

Access Control List 23

Overview	23
Usage Notes	25
Examples	26
ACL and Manual Classifier configuration	30

Traffic Management Alarms 32

Ingress Metering 33

Egress Port Rate Limiting 33

Priority Queueing 34

Overview	34
IGMP Multicast queue priority with telesyn_default_video	36
Changing Queue Mapping and Disabling/Enabling Interfaces	36

Product Support for Rate and Burst Size 36

Example Configurations 37

Example (IP Source)	37
Example (Class of Service)	39

Address Resolution Protocol (ARP) filtering 41

Dynamic IP Filtering Using DHCP Relay 45

Overview	45
----------	----

Treatment of DHCP Packets (MAC/VID only) 46

Classifier Provisioning and Topology Control 46

Summary of IP Filtering Options 46

Command Summary for Traffic Management 48

Overview 1

IGMP on the fMAP 1

Channel Usage for IGMP 1
Multicast stream availability 6

HVLAN (Port Based) 7

Overview for Release 6.0 7
VLAN Frame Flow (TAGALL and TPID Values) 7
Provisioning Rules 8
Sample Configuration and Commands 9
Port-based HVLAN for Release 7.0 (9100) 11

VLAN Translation 13

Overview 13
Provisioning Rules 13
Example Configuration 13

Port-Based HVLAN and Translation Feature Interactions 14

Command Summary for HVLAN 16

Traffic Management for the fMAP 17

Overview 17
Possible Conflict with Classifier Combinations - fMAP 19
Solution to Classifier Mismatch 21
QoS Classifier Capacity for FE10/FX10 Cards 22
Adding Classifiers to Service Module Ports/Interfaces 22
Classifiers and Feature Interaction for the fMAP Product 23

LAG 23

Overview 23
LAG and UFO Incompatibility 24
LAG and 9100 GE2RJ Interfaces 25
Changes to Feature Support from Release 6.0 29
Feature Change Details for 6.0 30
Software Upgrade 31
Traffic Provisioning 31
Maintenance/Alarms 31
Example Configuration 32
Convert Subtending Ring in D43 from GE3 to GE8 35
Upgrade the D42 System from the CFC24 to CFC56 36
Upgrade the D32 System from the CFC24 to CFC56 38
Convert the D32 Uplink to GE8 Interfaces (LAG) 39

Overview 1

Database Management 1

- Overview 1
- Database Back Up 1
- Database Purge 2
- Restore Database 2
- Database Transaction Failure 2

Delete Obsolete Users 4

DELETE Obsolete FILES 4

Scripting 4

Overview 1

Querying Alarm Status (SHOW ALARMS, CLEAR ALARMS) 2

- Overview 2
- Displaying alarms 2

Alarm System Features 7

- Overview of Alarm System 7
- Alarm types - (Interface, Card, System) and the Alarm System 8
- Interface/Port Outage Threshold Feature 10
- Configurable Alarm Severity (INTERFACE Alarm) 11

Overview of Troubleshooting 13

- Overview 13
- Card Diagnostics 13

DIAGNOSE command 13

- Diagnosing the CFC Control Module 13
- Logs 17

SYSTEM COOLING 18

Audits 18

System Recovery 19

TRACEROUTE 20

- Overview 20
- Using TRACEROUTE 20

Call Debugging 20

- Address Resolution Protocol (ARP) 20
 - Ping 21
 - Unsupported commands 21
 - Call State 21
-

VOICECALL Trace Logs 22

IGMP Trace 25

Overview 25

Command Examples 25

EPSR Trace 30

Overview 30

Command Examples 30

User Event Logging 32

Overview 32

Usage Notes 32

Overview of Commands 32

Event Logging commands 33

CES Troubleshooting 38

Overview 38

Card and fMAP Chassis Support 39

Concepts and Terms 39

Loopback 39

Port State Management 41

Fault Management 41

CSE Troubleshooting Examples 46

CES/CES8 IP Debugging 46

EPON2 Troubleshooting 46

Technical Support Scripts 47

Overview for 7.0 47

Overview 1

MIB Objects Supported 2

Physical Standards 16

Protocol / Software Standards 16

2. Overview of the fMAP

2.1 9000 Series (non-10G) Overview

The fMAP series multiservice access platform leverages widely accepted Ethernet switching technology to bring the service provider the ability to provide subscribers with Fiber To The Home (FTTH), Metro Ethernet, and ADSL services. It is a feature-rich platform that enables service providers to offer advanced, simultaneous “Triple Play” services, such as high quality voice, tiered IP/Ethernet data services, and broadcast quality IP video.

2.1.1 The fMAP (CFC12) Shelf

Figure 2-1 and Table 2-1 show the fMAP shelf and explain its key components. Refer to these when performing any task in this document. The 9100 is offered in these versions:

1. The 9102-A has a non-redundant AC power supply.
2. The 9103-A has a redundant AC power supply.

Note that the locations of Service Module cards for all versions are the same, and that the CFC12 is always in Slot 3.




FIGURE 2-1 9101 full configuration

TABLE 2-1 Listing of components with slot/position and configuration notes

No.	Module	Configuration Notes
A	CFC12 Slot (3)	Always provisioned with the CFC12 card
B	SM slots (0, 1, 2)	At least one is provisioned.
D	Console	Provides local connection to PC
E	OAM	Provides connection to MGMT (out -of-band) port

TABLE 2-1 Listing of components with slot/position and configuration notes

No.	Module	Configuration Notes
A	CFC12 Slot (3)	Always provisioned with the CFC12 card
F	GE2RJ	Provides Fast Ethernet to a close proximity, lower data rate (100M) Ethernet, as well as 1000BaseT (GE2RJ)
G	GE4	Provides fiber interface to WAN. Small Form Factor Pluggable (SFP) modules are required

2.2 Initial Status of fMAP Products

In the fMAP Installation Guide, the installation procedures end with the product powered up and the local interface connected.

Before any functions and features have been implemented on a fMAP product, the user can input a set of commands that show the initial state of the system.

2.2.1 Initial Interface

Refer to [Figure 2-2](#), which shows the local terminal connected to the CONSOLE port.

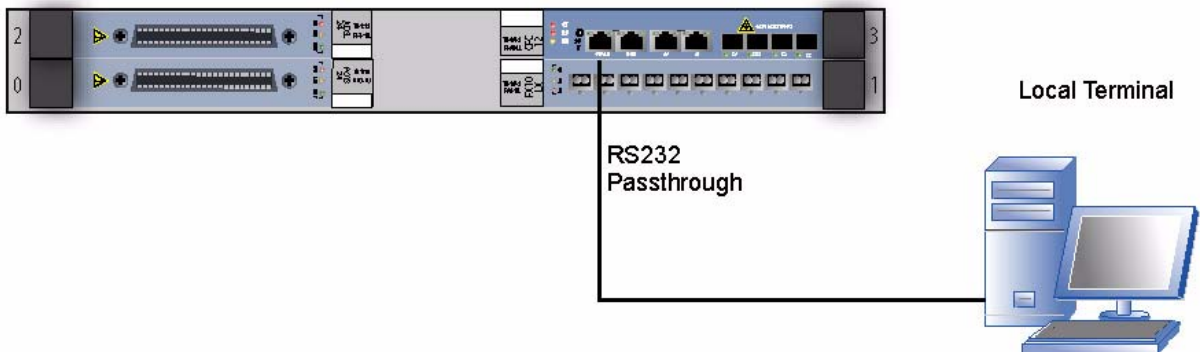


FIGURE 2-2 Initial Interface for Communicating with the fMAP product

2.2.2 Initial System Status

To see how the system is initially configured, input the **SHOW SYSTEM** command from the CLI. The system response below shows what is displayed for the **SHOW SYSTEM** command and its initial state for a fMAP 9101.

```
-User Access Veri fi cati on
```

Username: officer
Password:

10.52.202.110 SEC>> restart sys
Do you really want to restart the system? (Y/N)? y
Command has been submitted
10.52.202.110 SEC>>

Separator line of asterisks

Telesyn 12G Central Fabric Controller Boot Loader
Version 7.0.g.02
Created on Fri 09/02/2005 at 1:34p
Copyright Allied Telesis Holdings K.K. 2004

VxWorks Version 5.5.1 for IBM PowerPC 440GP Rev. 2.0
BSP version 1.2/3
Copyright Wind River Systems, Inc., 1984-2002

Separator line of asterisks

FGPA Version 12.16
Starting Telesyn Product Software Loading.
Attaching to Flash File System ... done. /tffs/ - Volume is OK

Press any key to stop automatic loading of software image...
Automatically loading software image...

Boot album is (current, attempt 1): '/tffs/load/cfc12_7.2.3.tar'
Checking Album's integrity... done
Loading vxWorks.bin.gz... (8206848 bytes)
Starting at 0x10000...

Attaching interface lo0... done

Adding 52254 symbols for standalone.

Separator line of asterisks

ASCII art logo for Allied Telesis K.K.

Allied Telesis K.K.

ATN 12G Central Fabric Controller
Version 7.2.3 (Customer-Release Build)
Created on Thu 12/22/2005 at 3:40p
Copyright Allied Telesis Holdings K.K. 2005

VxWorks Version 5.5.1 for IBM PowerPC 440GP Rev. 2.0
BSP version 1.2/3
Copyright Wind River Systems, Inc., 1984-2002

Memory Size: 255 MB

Separator line of asterisks

System Time is 2006-02-07 23:52:46.011
System initializing... /tffs/ - Volume is OK

Initialization completed successfully (7.2.3)

User Access Verification

Username: officer

Password:

```
10.52.202.110 SEC>> sh ca al
Error (010004): Parameter CARD, invalid value, "al" is not within the valid
range
```

```
10.52.202.110 SEC>> sh ca all
```

```
--- Card Information ---
```

Slot	Prov Card Type	State	Faults
0	SHDSL24	UP-DN-NotInstalled	-
1	NTE8	UP-DN-NotInstalled	-
2	SHDSL16	UP-DN-NotInstalled	-
3	CFC12	UP-UP-Degraded (Active)	Minor
4	GE2RJ	UP-UP-Online	-
5	GE4	UP-UP-Online	-
FAN	PEM71	UP-UP-Online	-

3. Configuring Management Interfaces

3.1 Overview

As described in Section 2, the initial configuration of the fMAP product has a PC/terminal connected to the CONSOLE port, allowing for a local Command Line Interface (CLI).

The next task is to identify the fMAP product to the network and set up a management interface that is integrated with the customer's network. Setting up these management interfaces involves these main areas:

1. Identifying the fMAP product to the network (IP address, system time, system configuration)
2. Setting up physical interfaces and protocols for the management interfaces
3. Setting up the management interfaces over the physical interfaces
4. Setting up the user interface (login IDs, passwords, levels of security, etc.) for the personnel who will query and control the system.
5. Setting up optional filtering on the management interface.

Table 3-1 lists these tasks and the commands that are used.

TABLE 3-1 Commands Used to Configure Management Interfaces

Interface	Main Tasks	Ref.
Set up Initial Interface	Identify fMAP product to the network	3.2
	Enable physical interfaces	3.3
Configure SNMP	Set up SNMP Community	3.5
	Routing of Traps	
Configure Management Logs	Filter logs according to certain criteria	3.6
	Set output for logs	
	Associate filtered logs with output	
Set up SNTP	Set Universal Time offset	3.7
	Add SNTP server	
	Enable SNTP, reset SNTP	

TABLE 3-1 Commands Used to Configure Management Interfaces

Interface	Main Tasks	Ref.
Set up Management Command Interface	Add TELNET users with passwords Set global command configuration	3.8
Setting up optional filtering on the management interface	Add filters to the MGMT interface.	3.8.15

3.2 Identifying the fMAP Product

3.2.1 Overview

Once the fMAP product is installed and powered up, it needs to be identified to the network and have a system time set.

3.2.2 Command Summary for Identifying the fMAP Product

Table 3-2 lists the commands used to set up and show the fMAP product attributes.

TABLE 3-2 Command Summary for System Objects (SYSTEM)

Object	Verb	Syntax	Description
SYSTEM	SHOW	SHOW SYSTEM	Displays the main attributes of the fMAP product configuration.
SYSTEM	SET	SET SYSTEM { CONTACT=contact LOCATION=location NAME=name HOSTNAME=name GATEWAY=ipaddress DOMAINNAME=name DNS=ipaddress- list }	<p>Sets the attributes that identify the fMAP product.</p> <p><i>Note: The SET SYSTEM DNS, SET SYSTEM DOMAINNAME, and SET SYSTEM GATEWAY commands attempt to change the specified attributes on a system-wide basis for ALL disabled interfaces. The recommended method of setting the DNS, DOMAINNAME, and, GATEWAY attributes is to use the SET IP INTERFACE command. See 3.3.4</i></p> <p>Once a DNS server has been provisioned in the system, to delete it use the command SET SYSTEM DNS "".</p>
SYSTEM TIME	SHOW	SHOW SYSTEM TIME	Displays the system time
SYSTEM TIME DATE	SET	SET SYSTEM [TIME=hh:mm:ss] [DATE=yyyy-mm-dd]	Sets the local time or date on the fMAP product.

3.3 Configuring Physical Interfaces and Protocols

3.3.1 Overview

One of two IP interfaces can be used:

- The **MGMT** Ethernet interface that transports only management data packets.
- An **inband** Ethernet interface that interleaves user data packets with management data packets on the uplink, using a VLAN interface. In using a VLAN interface the management data packets are always VLAN-tagged.

Over these two interfaces, the TELNET or SNMP agent can be configured.

Note: Only one interface can be enabled at a time; enabling an interface will disable an interface already enabled. If necessary, the ENABLE IP INTERFACE command will automatically disable the other IP Interface.

3.3.2 fMAP

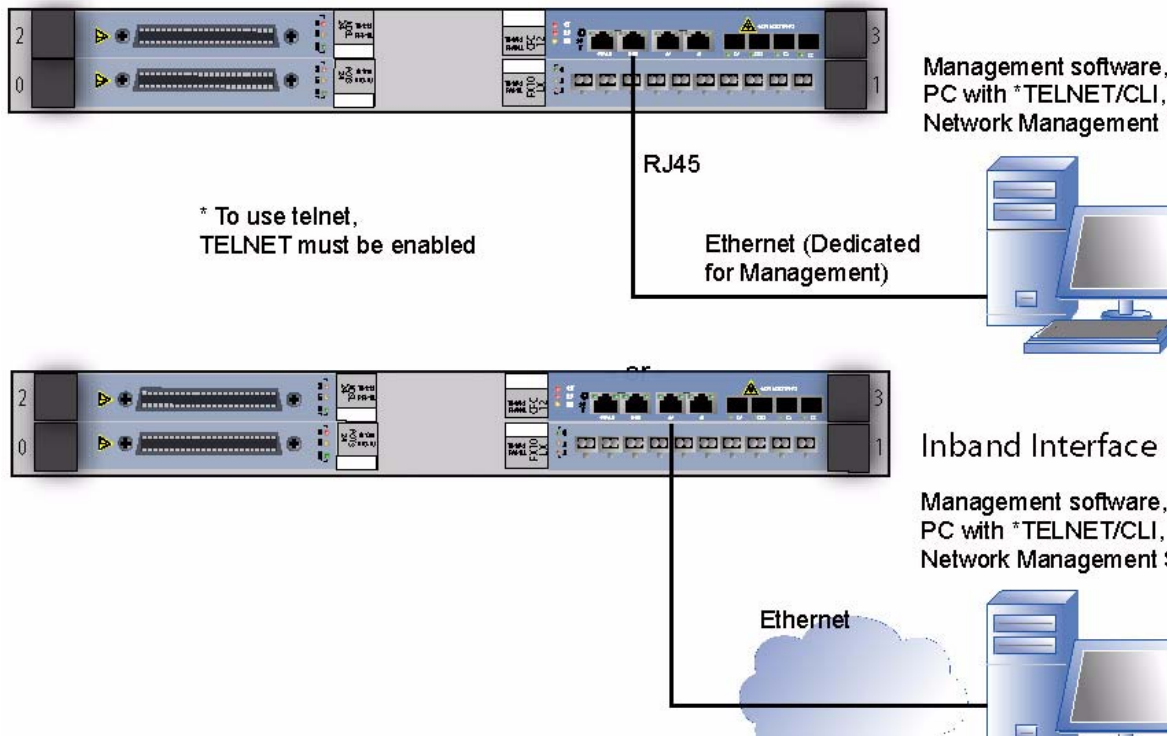


FIGURE 3-1 Connections for Management Interfaces for the fMAP

To enable TELNET access for the management ethernet interfaces, TELNET must be enabled. The user can then choose which interface to use and supply the IPADDRESS and SUBNETMASK for the fMAP product that will be used by the management device when a user logs in.

Note: These interfaces should be set up using the local RS232 interface. See the following Caution.

	<p><i>If the user disables or deletes an IP interface, and the user is currently using that interface to communicate with the fMAP product, the interface will be immediately disconnected.</i></p>
--	---

3.3.3 Configure the IP Interface - Example

The basic steps to configure the management interface connection are as follows:

ENABLE TELNET SERVER so an interface can be used.

6. Choose which interface to use

Note: For the interface that uses a VLAN, ensure that a VLAN has already been created. Moreover, a VLAN that is configured to be a management interface must be in standard rather than Upstream Forwarding Only (UFO) mode.

7. ADD the IPADDRESS and SUBNETAMASK for the interface.

8. ENABLE the interface so management-related data can be transmitted and received.

9. Refer to the system output below, which shows a dedicated MGMT interface being configured. Note that the **SHOW IP INTERFACE=MGMT** displays both status and alarm information about the IP interface.

```
offi cer SEC> ENABLE TELNET SERVER
offi cer SEC> ADD IP INTERFACE=MGMT IPADDRESS=172.16.5.11 SUBNETMASK=255.255.255.0
offi cer SEC> ENABLE IP INTERFACE=MGMT
offi cer SEC> SHOW IP INTERFACE=MGMT
```

```
-----
Fault..... No Faults

Interface..... MGMT
Status..... Enabled
IP Address..... 172.16.5.11
Subnet mask..... 255.255.255.0
VLAN..... N/A
```

3.3.3.1 Configuring the Inband Interface

1. For provisioning the inband interface, the user could add a VLAN interface. Refer to the system output below.

Caution: Enabling the interface would disable the MGMT interface.

Caution: The inband interface can be accessed from any port, GE, FE, ADSL, etc. Therefore, avoid provisioning subscriber ports on the inband VLAN.

```
offi cer SEC> ADD IP INTERFACE=VLAN: 3.0 IPADDRESS=172.16.6.11 SUBNETMASK=255.255.255.0
```

```
offi cer SEC> SHOW IP INTERFACE=ALL
```

```
-----
Interface..... MGMT
Status..... Enabled
IP Address..... 172.16.5.11
Subnet Mask..... 255.255.255.0

Interface..... mgtvlan (3)
Status..... Disabled
IP Address..... 172.16.6.11
Subnet Mask..... 255.255.255.0
```

```
offi cer SEC> ENABLE IP INTERFACE=3
```

```
offi cer SEC> SHOW IP INTERFACE=ALL
```

```
-----
Interface..... MGMT
Status..... Disabled
IP Address..... 172.16.5.11
Subnet Mask..... 255.255.255.0
```

```
Interface..... mgvl an (3)
Status..... Enabled
IP Address..... 172.16.6.11
Subnet Mask..... 255.255.255.0
```

officer SEC> PING 172.16.64.13

```
PING 172.16.64.13 (172.16.64.13)
64 bytes from 172.16.64.13 (172.16.64.13): icmp_seq=1
```

--- 172.16.64.13 ping statistics ---

1 packets transmitted, 1 packets received, 0% packet loss

3.3.4 Command Summary for the IP Interface

Table 3-3 lists the commands used at the PC/terminal to set up the management interfaces.

TABLE 3-3 Management Interface Commands

Object	Verb	Syntax	Description
IP INTER-FACE	SHOW	<pre>SHOW IP [INTERFACE [= { MGMT type:id-range ifname-list ALL }]] [FULL]</pre>	Shows whether the MGMT or VLAN IP interfaces is provisioned and its state.
IP INTER-FACE	ADD	<pre>ADD IP INTERFACE={ MGMT type:id } IPADDRESS=ipaddress SUBNETMASK=mask [CARD={ slot ACTCFC }] [IFNAME=ifname] [GATEWAY=ipaddress] [DOMAINNAME=name] [DNS=ipaddress-list]</pre>	<p>Configures the IP address, gateway address, and subnetmask of a MGMT or VLAN interface. The VLAN interface is specified by either the vlan name or vlan number (vid).</p> <p>When the IP interface is added, its state is by default disabled and must be enabled using the ENABLE IP INTERFACE command.</p>

TABLE 3-3 Management Interface Commands (Continued)

Object	Verb	Syntax	Description
IP INTER-FACE	SET	<pre> SET IP INTERFACE={ MGMT type:id-range ifname-list ALL } [IPADDRESS=ipaddress] [SUBNETMASK=mask] [IFNAME=ifname] [GATEWAY=ipaddress] [DOMAINNAME=name] [DNS=ipaddress-list] </pre>	Changes the existing setting for the MGMT or VLAN interface.
IP INTER-FACE	ENABLE	<pre> ENABLE IP INTERFACE={ MGMT type:id-range ifname-list } </pre>	Activates the VLAN or MGMT interface.
IP INTER-FACE	DISABLE	<pre> DISABLE IP INTERFACE={ MGMT type:id-range ifname-list ALL } </pre>	Deactivates the VLAN or MGMT interface, so that users can no longer log into the fMAP product using the IP address.
IP INTER-FACE	DELETE	<pre> DELETE IP INTERFACE={ MGMT type:id-range ifname-list ALL } </pre>	Deletes the MGMT or VLAN interface.
IPAD-DRESS INTER-FACE	PING	<pre> PING={ ipaddress hostname } [FROM { INTERFACE={ type:id id ifname } IPADDRESS=ipaddress }] [DELAY=1..900] [LENGTH=1..65535] [NUMBER={ 1..65535 CONTINUOUS }] [TIMEOUT=1..900] </pre>	Pings an interface or IP address from the fMAP product

TABLE 3-3 Management Interface Commands (Continued)

Object	Verb	Syntax	Description
PING	STOP	STOP PING	Stops a ping session that is in progress.
TELNET SERVER (see note)	ENABLE	ENABLE TELNET SERVER	Since the default is for TELNET to be disabled for security, the user must input this command before the TELNET interfaces can be used.
	DISABLE	DISABLE TELNET SERVER	Disables the TELNET interface. The fMAP product cannot communicate through a remote TELNET interface.
	SHOW	SHOW TELNET SERVER	Display TELNET information.

The following text illustrates the use of the PING command.

```
officer_SEC>> SHOW SYSTEM
SHOW SYSTEM
```

```
--- System Information -----
System Date..... 2004-07-07 09:43:46
System Uptime..... 0 days, 22 hours, 05 minutes, 05 seconds
Shelf Serial Number..... ATNLAB4030200080
Software Version..... 4.0.0.BETA.20040705
Software Options..... Lab-Only Build
Software Created..... Tue 07/06/2004 at 11:26 AM
SDRAM (free/total)..... 67405 KB / 131072 KB
Flash (free/total)..... 11540 KB / 63488 KB
Booted from..... preferred
Provisioning Mode..... AUTO
Contact..... <none>
Location..... <none>
Name..... <none>
Services..... Layer 2 - Datalink/Subnetwork
MGMT MAC address..... 00:0C:25:00:01:0D
Number of MACs on card..... 4
Description..... Allied Telesyn 7700 Multi-service Access Platform
Hostname..... <none>
MGMT IP Address..... 172.16.66.231
MGMT Subnet Mask..... 255.255.255.0
MGMT Gateway..... 172.16.66.1
MGMT Domainname..... telesyn.corp
MGMT DNS..... <none>
vlan: 402.0 IP Address..... 172.16.66.240
vlan: 402.0 Subnet Mask..... 255.255.255.0
vlan: 402.0 Gateway..... 172.16.66.1
vlan: 402.0 Domainname..... telesyn.corp
vlan: 402.0 DNS..... <none>
```

Slot	Prov Type	Phys Type	Model Num	Serial Num	Rev
0	not provisioned	not present	-	-	-
1	not provisioned	not present	-	-	-
2	not provisioned	not present	-	-	-
3	not provisioned	not present	-	-	-
4	not provisioned	not present	-	-	-
5	not provisioned	not present	-	-	-
6	not provisioned	not present	-	-	-
7	not provisioned	not present	-	-	-
8	CFC6	CFC6	TN-CFC-06	ATNLAB4030200031	X1
9	CFC6	CFC6	TN-CFC-06	ATNLAB4030200031	X1

10	GE1	not present	-	-	-
11	not provisioned	not present	-	-	-
12	CFC6	CFC6	TN-CFC-06	ATNLAB4030200033	X1
13	CFC6	CFC6	TN-CFC-06	ATNLAB4030200033	X1
14	not provisioned	not present	-	-	-
15	not provisioned	not present	-	-	-
16	not provisioned	not present	-	-	-
17	not provisioned	not present	-	-	-
18	not provisioned	not present	-	-	-
19	POTS24	POTS24	TN-113-A	ATNLAB4030200846 XX	X1
20	not provisioned	not present	-	-	-
21	not provisioned	not present	-	-	-
FAN	FAN8	not present	-	-	-

Layer 2 Base System

```
-----
Ageing time..... 300
Ageing time status..... Enabled
Learning status..... Enabled
```

```
officer SEC>> PING 172.16.66.1
```

```
PING 172.16.66.1
```

```
officer SEC>>
```

```
PING 172.16.66.1 (172.16.66.1)
```

```
64 bytes from 172.16.66.1 (172.16.66.1): icmp_seq=1
```

```
--- 172.16.66.1 ping statistics ---
```

```
1 packets transmitted, 1 packets received, 0% packet loss
```

```
officer SEC>> PING 172.16.66.1 FROM INTERFACE=vlan:402.0
```

```
PING 172.16.66.1 FROM INTERFACE=vlan:402.0
```

```
officer SEC>>
```

```
PING 172.16.66.1 (172.16.66.1)
```

```
64 bytes from 172.16.66.1 (172.16.66.1): icmp_seq=1
```

```
--- 172.16.66.1 ping statistics ---
```

```
1 packets transmitted, 1 packets received, 0% packet loss
```

```
officer SEC>> PING 172.16.66.1 FROM IPADDRESS=172.16.66.240
```

```
PING 172.16.66.1 FROM IPADDRESS=172.16.66.240
```

```
officer SEC>>
```

```
PING 172.16.66.1 (172.16.66.1)
```

```
64 bytes from 172.16.66.1 (172.16.66.1): icmp_seq=1
```

```
--- 172.16.66.1 ping statistics ---
```

```
1 packets transmitted, 1 packets received, 0% packet loss
```

Note: To enable the SNMP server, refer to [3.5.4.](#)

3.4 Setting Up the Management Interfaces - Overview

This next figure shows how the MGMT or VLAN interface is configured to provide the management interfaces. The following subsections explain these interfaces.

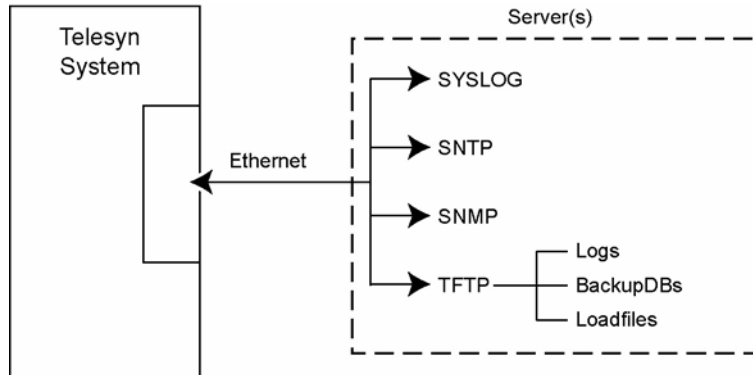


FIGURE 3-2 Configuration of the fMAP product Management Interfaces

3.5 Configuring the SNMP Community

3.5.1 Overview (Product Support)

As with other fMAP products, this fMAP product allows a network management system to query and control one or more managed devices using the Simple Network Management Protocol (SNMP). Each managed device has a **Managed Information Base (MIB)**, which includes the device's configuration as well as counters for system events and activities. The network management system can ask the device for information only or ask the device to change its configuration through the device's **SNMP agent**. The device can also send unsolicited messages called **traps** for critical events.

3.5.1.1 fMAP

The SNMP for the fMAP product uses the following standards:

- RFC 1213 (referred to as MIB-II), which defines the core set of objects.
- RFC 1157 (referred to as SNMP), which defines the protocol used between the management stations and devices.
- RFC 1902, which defines the structure of management information for SNMPv2.
- RFC 1903, which defines the textual conventions for SNMPv2.
- RFC 1904, which defines the conformance statements for SNMPv2.
- ATN Enterprise MIB, which defines the fMAP enterprise MIBs.

3.5.2 SNMP Components

Note: Refer to Section: **Specifications** for a listing of which objects are supported.

Figure 3-3 gives an overview of the main components of SNMP and how they could be configured for the fMAP product.




FIGURE 3-3 SNMP Configuration

The components of an SNMP community include the following:

- SNMP Protocol - The protocol consists of the following:
 - Transport protocol - The protocol supported by the fMAP product is UDP. Traps are sent out over UDP port 162; all other messages are sent out over UDP port 161.
 - SNMP Messages - SNMP messaging uses the commands **get**, **get-next**, **get-response**, **set**, and **trap**.
- Polling - The normal exchange of information between the SNMP agent on the fMAP product and the SNMP browser on an external device is as follows:
 - The SNMP browser sends a get (or get-next) request to the managed device
 - The device responds with a get-response message.

Note: In this release, the SNMP set message is not supported.

- Traps - The device can also send unsolicited messages called traps, for alarm-type messages.
- SNMP Communities - A community is the way the SNMP is defined for the fMAP product, and includes a name and the IP addresses for trap hosts (receive all traps or a subset of traps) and management stations.
- SNMP Authentication - This is how an SNMP message is declared to be authentic (from an SNMP application actually in the community to which the message claims to belong). The generation of Authentication fail-

ure traps is through the parameter **AUTHENTICATE_TRAP**. When enabled, the following conditions cause authentication failure traps:

- Invalid community used to access the fMAP product
- Invalid use of the community, such as attempting a set with a read-only community
- Attempted access from a **MANAGER** which is not a **trusted** host or is not associated with the community associated with the SNMP operation.

3.5.3 Securing an SNMP Community

There are three methods of providing security in an SNMP community so that unauthorized users cannot query or change the MIB variables:

1. A community for the device can be defined with an **ACCESS** of **READ** only, so even if the device is accessed values cannot be changed.
2. The community is associated with a set of trusted devices (by their IP address when creating or adding hosts to the community), and if an attempt is made to access the fMAP product from a device that does not have that IP address an authentication trap is produced. (This is controlled with **SNMP AUTHENTICATE_TRAP**.)
3. The community name, which acts as a trivial password.

Note: SNMP commands can only be input by a Security Officer.

Note: Up to 12 SNMP communities can be configured on one system.

It is important to understand these methods when enabling an SNMP community and enabling SNMP trap generation for the community.

3.5.4 Setting Up an SNMP Community

Setting up an SNMP community involves the following steps:

1. Enable the SNMP Server in the fMAP product using the **ENABLE SNMP** command.
2. Create the SNMP community using the **CREATE SNMP COMMUNITY=name** command, where “name” is the community name.
3. Enable the SNMP community using **ENABLE SNMP COMMUNITY=name** command. This command allows the fMAP product to be accessed by the community. If the community is defined as **READ**, the user can access any of the supported MIB variables. If the community is defined as **WRITE**, the user can perform set commands as well. (The default is **READ**.) Once the community is enabled, it is possible for a **MANAGER** device to access the fMAP product; however, traps are still not generated.



*When a community is defined with an **ACCESS** of **WRITE**, traps from the community include the community name (the trivial password), which could be seen by unauthorized users.*

Note: If the user plans to reset ATN Enterprise MIB counters or change the cache or SNMP trap filter settings for an SNMP community, ACCESS for the community must be set to WRITE. If ACCESS is set to READ, the user will not be able to make these changes or any other changes to the MIB variables for the community.

4. Enable the SNMP community traps using the **ENABLE SNMP COMMUNITY=name TRAP** command. The fMAP device will now send traps, which includes the community name.

Note: To change the cache settings, the SNMP community must be set to WRITE access.

3.5.5 Disabling an SNMP Community or its Traps

Once a community has been enabled and the sending of traps has been enabled, as explained in 3.5.3, the user can input the command:

```
DISABLE SNMP COMMUNITY=name TRAP
```

This will disable the ability of the snmp community to produce traps, but it does **not** disable the snmp community; a MANAGER device can still access the fMAP product.

The user can also disable the snmp community with the command:

```
DISABLE SNMP COMMUNITY=name
```

This will disable the ability of a MANAGER to access the fMAP product, but it will **not** disable the ability of the community to produce traps; this must be done with the DISABLE command above that includes TRAP.

3.5.6 Setting Up an SNMP Community - Example

The initial status of SNMP is *disabled*. The main steps for configuring an SNMP community are as follows. The community will be READ only and will follow the configuration shown in [Figure 3-3](#). The steps are as follows:

1. Input `SHOW SNMP` to see the initial SNMP configuration. Refer to the system output below:

```
officer SEC> SHOW SNMP
```

```
-----
SNMP Configuration:
-----
Status. . . . . DISABLED
Authentication Failure Traps. . . . . DISABLED
SNMP Counters:
-----
inPkts. . . . . 0          outPkts. . . . . 0
inBadVersions. . . . . 0      outTooBig. . . . . 0
inBadCommunityNames. . . . . 0  outNoSuchNames. . . . . 0
inBadCommunityUses. . . . . 0  outBadValues. . . . . 0
inASNParseErrs. . . . . 0      outGenErrs. . . . . 0
inTooBig. . . . . 0           outGetRequests. . . . . 0
inNoSuchNames. . . . . 0       outGetNexts. . . . . 0
inBadValues. . . . . 0         outSetRequests. . . . . 0
inReadOnly. . . . . 0          outGetResponses. . . . . 0
inGenErrs. . . . . 0           outTraps. . . . . 0
inTotalReqVars. . . . . 0
inTotalSetVars. . . . . 0
inGetRequests. . . . . 0
inGetNexts. . . . . 0
inSetRequests. . . . . 0
inGetResponses. . . . . 0
inTraps. . . . . 0
```

Note that the initial status of SNMP configuration:

- The Status is `DISABLED`
- The Authentication Failure Traps is `DISABLED`.
- The SNMP counters are all set to 0. Refer to [Table 3-4](#) for a description of these counter fields

TABLE 3-4 SNMP Counters (Displayed Using SHOW SNMP)

Parameter	Description
inPkts	The total number of SNMP packets received by the fMAP product.
inBadVersions	The number of SNMP packets with a bad version field received by the fMAP product.
inBadCommunityNames	The total number of SNMP PDUs delivered to the SNMP agent that used an unknown SNMP community name.
inBadCommunityUses	The total number of SNMP PDUs delivered to the SNMP agent that represented an SNMP operation not allowed by the SNMP community name in the PDU.
inASNParseErrs	The total number of ASN.1 parsing errors, either in encoding or syntax, encountered by the SNMP agent when decoding received SNMP PDUs.
inTooBig	The total number of valid SNMP PDUs delivered to the SNMP agent where the value of the errorStatus component was tooBig.
inNoSuchNames	The number of SNMP packets received with an error status of nosuchname.
inBadValues	The number of SNMP packets received with an error status of badvalue.

TABLE 3-4 SNMP Counters (Displayed Using SHOW SNMP) (Continued)

Parameter	Description
inReadOnlys	The number of SNMP packets received with an error status of readonly.
inGenErrs	The number of SNMP packets received with an error status of gener.
inTotalReqVars	The total number of SNMP MIB objects requested.
inTotalSetVars	The total number of SNMP MIB objects requested that were changed.
inGetRequests	The number of SNMP Get Request packets received by the fMAP product.
inGetNexts	The number of SNMP Get Next packets received by the fMAP product.
inSetRequests	The number of SNMP Set Request packets received by the fMAP product.
inGetResponses	The number of SNMP Get Response packets received by the fMAP product.
inTraps	The total number of SNMP trap message packets received by the fMAP product.
outPkts	The total number of SNMP packets transmitted by the fMAP product.
outTooBigs	The number of SNMP packets transmitted with a status of toobig.
outNoSuchNames	The number of SNMP packets received with an error status of nosuchname.
outBadValues	The number of SNMP packets transmitted with an error status of badvalue.
outGenErrs	The number of SNMP packets transmitted with an error status of genererror.
outGetRequests	The number of SNMP Get Request response packets transmitted by the fMAP product.
outGetNexts	The number of SNMP Get Next response packets transmitted by the fMAP product.
outSetRequests	The number of SNMP Set Request packets transmitted by the fMAP product.
outGetResponses	The number of SNMP Get Response packets transmitted by the fMAP product.
outTraps	The total number of SNMP trap message packets transmitted by the fMAP product.

2. Enable SNMP so it is possible to create an SNMP community. Note the status of SNMP is now ENABLED.

```
officer SEC> ENABLE SNMP
Info: Operation Successful
officer SEC> SHOW SNMP
```

```
-----
SNMP Configuration:
-----
Status..... ENABLED
```

(output omitted)

3. Enable SNMP authentication traps.

```
officer SEC> ENABLE SNMP AUTHENTICATE_TRAP
```

4. Create the snmp community, supplying a name and the list of trusted trap hosts and management stations, following what is shown in [Figure 3-3](#).

```
officer SEC> CREATE SNMP COMMUNITY=public V2CTRAHOST=172.16.34.2
MANAGER=172.16.32.3 TRAPHOST=172.16.22.8
Info: Operation Successful
officer SEC> SHOW SNMP COMMUNITY=public
```

```
--- SNMP Community Information -----
```

```
Name..... public
-----
Access..... READ-ONLY
Status..... DISABLED
Traps..... DISABLED
Open Access..... NO
Manager..... 172.16.32.3
TrapHost..... 172.16.22.8
v2cTrapHost..... 172.16.34.2
```

5. Enable the snmp community, and then enable traps for that community.

```
officer SEC> ENABLE SNMP COMMUNITY=public
Info: Operation Successful
officer SEC> ENABLE SNMP COMMUNITY=public TRAP
```

6. Do a final check of the SMNP community.

```
officer SEC> SHOW SNMP COMMUNITY=public
--- SNMP Community Information -----
Name..... public
Access..... READ-ONLY
Status..... ENABLED
Traps..... ENABLED
Open Access..... NO
Manager..... 172.16.32.3
TrapHost..... 172.16.22.8
v2cTrapHost..... 172.16.34.2
```

3.5.7 Command Summary for SNMP

TABLE 3-5 Summary of Commands for SNMP

Object	Verb	Syntax	Description
SNMP	SHOW	SHOW SNMP	Show current state and configuration of the SNMP agent
	DISABLE	DISABLE SNMP	Disable SNMP agent. This is the default.
	ENABLE	ENABLE SNMP	Enable SNMP agent. The user must do this before creating an SNMP community
SNMP COMMUNITY	SHOW	SHOW SNMP COMMUNITY [={ name-list ALL }]	Show the current SNMP community configuration
	CREATE	CREATE SNMP COMMUNITY=name [ACCESS={ READ WRITE }] [V2CTRAHOST=ipaddress-list] [TRAPHOST=ipaddress-list] [MANAGER=ipaddress-list]	Create the SNMP community's name, and optionally add the community's attributes. By default, the community is disabled and must be enabled to allow access to the fMAP product.
	ADD	ADD SNMP COMMUNITY=name [TRAPHOST=ipaddress-list] [V2CTRAHOST=ipaddress-list] [MANAGER=ipaddress-list]	Add a traphost or management station to the already defined SNMP community.

TABLE 3-5 Summary of Commands for SNMP (Continued)

Object	Verb	Syntax	Description
SNMP COMMUNITY	SET	SET SNMP COMMUNITY=name [ACCESS={ READ WRITE }] [OPEN={ ON OFF YES NO TRUE FALSE }]	Modify the attributes of the already defined SNMP community, as well as the name.
	DISABLE	DISABLE SNMP COMMUNITY=name [TRAP]	Disable a community that exists, or only the traps sent by the fMAP product.
	ENABLE	ENABLE SNMP COMMUNITY=name [TRAP]	Enable the SNMP community, or enable the sending of traps for the community. Refer to 3.5.3 .
	DELETE	DELETE SNMP COMMUNITY=name [TRAPHOST=ipaddress-list] [V2CTRAPHOST=ipaddress-list] [MANAGER=ipaddress-list]	Delete from an SNMP community a traphost or manager station.
	DESTROY	DESTROY SNMP COMMUNITY=name	Destroy an existing SNMP community.
SNMP AUTHENTICATE_ TRAP	ENABLE	ENABLE SNMP AUTHENTICATE_TRAP	Enable the generation of authentication traps (traps generated when a request message has not been authenticated). Disabled is the default.
	DISABLE	DISABLE SNMP AUTHENTICATE_TRAP	Disable the generation of authentication traps.

3.6 Configuring Log Filtering and Output

3.6.1 Overview

The fMAP product produces management logs that provide information about all changes that occur. [Figure 3-5](#) shows an example log, and [Table 3-6](#) describes the fields included with a management log.

A B _____ C _____ D E
 * CARD002 2003-07-13 23:11:24 5618 INFO
 Location: Slot:5
F | Description: Card Mismatch Cleared
 Reason Code: Card Mismatch
A = Severity D = Sequence Number
B = Category E = Log Type
C = Date and Time F = Message

FIGURE 3-4 Sample Log Produced by the fMAP product

TABLE 3-6 Field Definitions of Management Logs

Field	Value	Description
Category	ADSL	Change to ADSL configuration and ADSL statistics
	BDB	Configuration database has been backed up
	CARD	Change to a card
	CFCP	A change in CFC protection, such as duplex to simplex.
	CHAS	Chassis
	CLI	Command-line interface
	CUC	Cooling Unit Controller
	FAN	Fan Unit
	FILE	File Changes
	IGMP	Changes to IGMP configuration
	LOG	Log management
	PORT	Port change
	RDB	Configuration database has been restored
	RMON	Performance Monitoring of Ethernet-based statistics
	RSDB	Configuration Database has been reset (purged)
	SHLF	Changes in shelf
	SNTP	Changes in SNTP (time setting)
	STP	Spanning tree protocol
	SYS	Changes in overall system
	TRAP	A trap has been produced
USER	Changes in user configuration	

TABLE 3-6 Field Definitions of Management Logs (Continued)

Field	Value	Description
Log Type	INFO	Information only
	FAULT	Fault condition
	OTHER	All other logs
Severity	*C	CRITICAL: data service is affected and requires immediate attention.
	**	MAJOR: data service may be affected and must be investigated.
	*	MINOR: data service is not affected but could lead to a larger problem.
	<blank>	NONE: Information only
Date and Time	yyyy-mm-dd hh:mm:ss	Date and time the log was produced

3.6.2 Viewing Logs

Use the `SHOW LOG` command to filter logs immediately in the output, for example to show only logs that have a severity of CRITICAL. Refer to the `SHOW LOG` command in [Table 3-7](#).

3.6.3 Controlling Output of Logs

To control the output of logs the following are used:

- Log Filter - This is a **filterid** (usually a text string) that is associated with a Category and Severity.
- Log Output - This is an **outputid** (also usually a text string) that is associated with the destination for the logs. The destination can be a terminal or SYSLOG server. The outputid can also define the log format.

By combining the two, a filterid can be created and then associated with an outputid. [Figure 3-5](#) shows an example configuration.

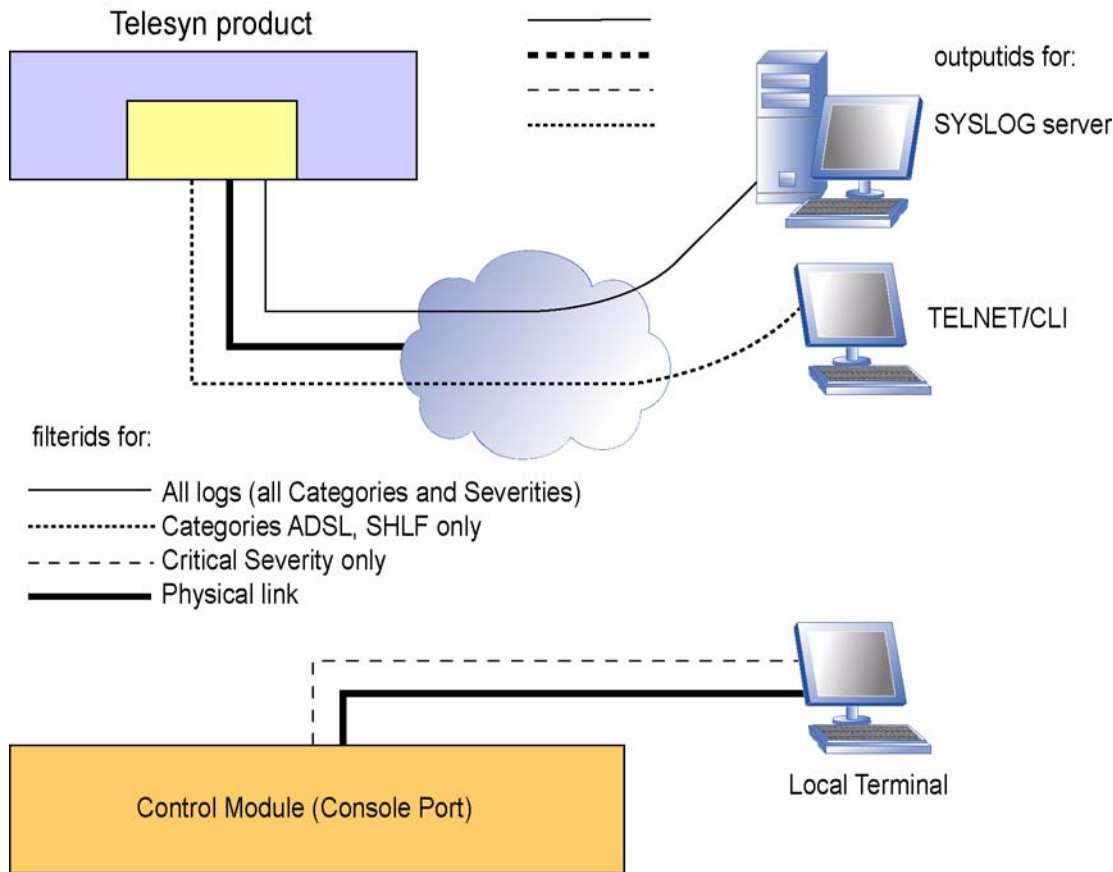


FIGURE 3-5 Example Log Configuration

3.6.4 Example Log Configuration Setup

An example sequence of setting up a log management system would be as follows:

- **Create a log filter** - Use the `CREATE LOG FILTER=<filter name>` command to create a name for a set of logs, called the filterid, and set up a criteria so logs that meet that criteria are collected together and associated with that filterid.
- **Create a log output** - Use the `CREATE LOG OUTPUT=<output name>` command to create a name for the destination for the logs (called an outputid), set up the attributes for that destination (such as an IP address), and specify the log format.
- **Associate the filterid with the outputid** - Use the `ADD FILTER OUTPUT=<output name>` command to associate the filterid with the outputid so that logs filtered in a certain way are sent to a certain destination.

The log format can be set to FULL, SUMMARY, or MSGONLY by using the FORMAT keyword with the **CREATE LOG OUTPUT** or the **SET LOG OUTPUT** command. The FULL format displays the entire log message. The SUMMARY format displays only the category, timestamp, and log type. The MSGONLY format displays only the log message. A comparison of the formats is shown in [Figure 3-6](#).

```

Full      → USER002 2003-07-15 10:57:23 9123 INFO
           User: user01 at IP: 192.16.18.103 has logged in

Summary  → USER002 2003-07-15 10:57:23 9123 INFO

Message  → User: user01 at IP: 192.16.18.103 has logged in
Only

```

FIGURE 3-6 Comparison of Log Formats

Following is an example of setting up the log configuration that matches what is shown in [Figure 3-5](#).

1. Show logs that match a criteria. For example, to view logs that have a severity level of CRITICAL, input the following:

```

officer SEC> SHOW LOG SEVERITY=CRITICAL
*C SYS010 2003-04-16 14:39:42 3538 FAULT
System: Cleared Port Outage Threshold

*C SYS009 2003-04-16 14:39:40 3519 FAULT
System: Raised Port Outage Threshold

*C SYS009 2003-04-16 14:37:01 3187 FAULT
System: Raised Port Outage Threshold

```

2. Create a log filter for critical severity only


```
CREATE LOG FILTER=CRITICAL SEVERITY=CRITICAL
```
3. Create a log output to associate with the log filterid CRITICAL.


```
CREATE LOG OUTPUT=terminal DESTINATION=CLI FORMAT=SUMMARY
```
4. Add the log filter created in step 2. to the log output created in step 3.


```
ADD LOG FILTER=CRITICAL OUTPUT=TERMINAL
```
5. Enable the output.


```
ENABLE LOG OUTPUT=TERMINAL
```

3.6.5 Capturing and Sending Logs to a Storage Device

Users can query the system for logs and send them to a storage device using the PUT FILE command. Note that the PUT FILE command can be used for not only log files, but any supported file type. Usually logs will be captured and sent to a network server for analysis. Logs can be captured from both the ACTCFC and the

INACTCFC. An example of using the PUT FILE command to capture logs and send them to a TFTP server follows. In this example, the user captured logs from the INACTCFC of a duplex system.

```
officer SEC> PUT LOG FILE=LOG_FILE TFTP SERVER=172.16.18.50 CARD=INACTCFC
```

Command has been submitted

```
officer SEC>
```

```
Info (010020): Successfully transferred file: LOG_FILE
```

The log file now exists on the TFTP server:




FIGURE 3-7 TFTP session with reception of log file

Any text editor can now be used to analyze the file.

3.6.6 Command Summary for Log Management

[Table 3-7](#) lists the commands used for management logs.

TABLE 3-7 Log Management Commands

Object	Verb	Syntax	Description
LOG	SHOW	SHOW LOG [CATEGORY=category] [DATE=[op] yyyy-mm-dd [-yyyy-mm-dd]] [FORMAT={ FULL MSGONLY SUMMARY }] [REVERSE] [SEQUENCE=0..9999 [-0..9999]] [SEVERITY=[op] { CRITICAL MAJOR MINOR NONE }] [TAIL [=count]] [TIME=[op] hh:mm:ss [-hh:mm:ss]]	Displays all the stored management logs. The default is all management logs are displayed in order from newest to oldest.
	PURGE	PURGE LOG	Removes all stored management logs from the system.

TABLE 3-7 Log Management Commands (Continued)

Object	Verb	Syntax	Description
LOG FILE	PUT	<pre> PUT LOG FILE={ destinationfile unit:destinationfile serverpath/destinationfile } [{ TFTP SERVER={ ipaddress hostname } ZMODEM FTP SERVER={ ipaddress hostname } USER=userid PASSWORD=password }] [TYPE={ MGMT ERROR TRACE CRASH }] [CARD={ ACTCFC INACTCFC }] </pre>	<p>Transfers management, error, trace or crash logs off the device. Currently, TFTP is the only supported transfer method.</p> <p>By default, logs with the TYPE of ERROR are placed in the file.</p>

TABLE 3-7 Log Management Commands (Continued)

Object	Verb	Syntax	Description
LOG FILTER	SHOW	SHOW LOG FILTER	Displays all the existing management log filters in the system. The log filter name, the log categories filtered, if any, and the severity values filtered are displayed.
	CREATE	CREATE LOG FILTER=filterid [CATEGORY=category] [SEVERITY=[op] { CRITICAL MAJOR MINOR NONE }]	Creates a management log filter.
	DESTROY	DESTROY LOG FILTER={ filterid-list ALL }	Removes management log filters from the system.
	SET	SET LOG FILTER=filterid [CATEGORY=category] [SEVERITY=[op] { CRITICAL MAJOR MINOR NONE }]	Changes the attributes of the filterid specified by the category and severity. Refer to the CREATE LOG FILTER command. By default, if no category, severity or format options are specified, the management log filter is set to match all logs.
LOG FILTER OUTPUT	ADD	ADD LOG FILTER={ filterid-list ALL } OUTPUT=outputid	Associates an existing management log filter with an existing management log output destination.
	DELETE	DELETE LOG FILTER={ filterid-list ALL } OUTPUT=outputid	Removes the association between a management log filter and a management log output destination.
LOG OUTPUT	DISABLE	DISABLE LOG OUTPUT={ outputid-list ALL }	Disables management log streaming for an existing management log output destination.

TABLE 3-7 Log Management Commands (Continued)

Object	Verb	Syntax	Description
LOG OUTPUT	ENABLE	ENABLE LOG OUTPUT={ outputid-list ALL }	Enables management log streaming for an existing management log output destination.
LOG OUTPUT	DESTROY	DESTROY LOG OUTPUT={ outputid-list LL }	Removes an existing management log output destination from the system.
LOG OUTPUT	SHOW	SHOW LOG OUTPUT	Shows information (outputid, server location, output format, and type of format), about the log outputs of the fMAP product.

3.7 Setting up Simple Network Time Protocol (SNTP)

3.7.1 Overview

When the fMAP product is first installed, local time can be set up using the command **SET SYSTEM TIME**, as explained in 3.7.

The fMAP product can also synchronize with a network time server using the SNTP protocol, which requires an SNTP server with a host name or IP address to be configured.

An example sequence of setting up the SNTP server would be as follows:

- **SET SNTP UTCOFFSET** - Sets the difference between local time and Coordinated Universal Time (UTC, also called Greenwich Mean Time).
- **ADD SNTP SERVER** - Inputs the hostname or IP address of the SNTP server that the fMAP product will use.
- **ENABLE SNTP** - Activates the SNTP so that the fMAP product will be able to synchronize its clock with the SNTP clock.
- **RESET SNTP** - Once the SNTP server is configured, this has the fMAP product send an sntp query to re-synchronize the fMAP product with the SNTP server. Note that the SNTP server must be enabled to do this.

To delete the SNTP server, the following sequence would be used:

- **DI SABLE SNTP** - Deactivates the SNTP so that the fMAP product will no longer synchronize its clock with the SNTP clock.
- **DELETE SNTP SERVER** - Deletes the hostname or IP address of the SNTP server that the fMAP product is using.

3.7.2 Command Summary for SNTP

TABLE 3-8 Commands for Network Timing

Object	Verb	Syntax	Description
SNTP	SHOW	SHOW SNTP	Shows the attributes of the Simple Network Time Protocol (SNTP) configuration, which includes the SNTP server hostname/address, and UTC offset)
	ENABLE	ENABLE SNTP	Activates the SNTP so that the fMAP product will be able to synchronize its clock with the SNTP clock once the SNTP server has been added.
	DISABLE	DISABLE SNTP	Deactivates the SNTP so that the fMAP product will no longer synchronize its clock with the SNTP clock. The SNTP server can now be deleted.
	RESET	RESET SNTP	Resets the timing counters and sends a query to the SNTP server to re-establish the time. Note that the SNTP server must be in use.
SNTP SERVER	ADD	ADD SNTP SERVER={ ipaddress hostname }	Inputs the hostname or IP address of the SNTP server that the fMAP product will use.
	DELETE	DELETE SNTP SERVER	Deletes the hostname or IP address of the SNTP server that the fMAP product is using.
SNTP UTCOffset	SET	SET SNTP UTCOffset={ + - } hh:mm	Set the difference between local time and Coordinated Universal Time (UTC, also called Greenwich Mean Time).

3.8 Setting up the CLI

3.8.1 Command and Session Overview

The fMAP product supports three levels of security: User, Manager, and Security Officer. Each level provides a specific degree of system access in a progressive fashion as shown in [Figure 3-8](#).




FIGURE 3-8 fMAP Product Security Levels

Each security level controls the commands that can be entered as follows:

- **User** - Users have the lowest level of access, which is equivalent to read-only privileges. They can change their password and use any of the `SHOW` commands to display information. When logged in, the User receives the command-line prompt:

```
username USR >
```

- **Manager** - Managers have a higher priority than Users and can perform all actions that a User can perform. In addition to User privileges, Managers can view and configure all areas of the fMAP product. When logged in, the Manager receives the command-line prompt:

```
username MGR >
```

- **Security Officer** - Security Officers have the highest priority and can access the full set of commands. When logged in, the Security Officer receives the command-line prompt:

```
username SEC >
```

For information on the specific commands that can be accessed at each security level, refer to [3.8.8](#).

For all security levels, a login name and password (case-sensitive) are required to access the system. There is a timer (default of 300 seconds or 5 minutes) that will log off the session if no commands are entered within the timeout period.

Note: The fMAP product can support up to 10 concurrent TELNET sessions.

3.8.2 Password Recovery

If, for some reason, all system user IDs and passwords have been deleted, destroyed, or corrupted, the user can recover the default user ID and password using the password recovery procedure.

3.8.2.1 Usage Notes

- A system power cycle must be performed in order to recover the default user ID and password.
- The system must be rebooted in order to recover the default user ID and password.
- Boot flags must be set in order to perform the password recovery procedure.

Note: The user should read through all the steps before beginning the procedure.

Since all system user IDs and passwords have been destroyed, there is no user access to the CLI command line prompt to initiate a reboot from software using the RESTART command. Therefore, the system must be power-cycled. Here are the steps:

1. Ensure that the management device (PC, lap top, etc.) is connected to the CONSOLE (serial) port of the active CFC card. Refer to [3.3](#) and [3.4](#).
2. Disconnect system power. Turn **OFF** circuit breaker A, then circuit B on the system power entry module.
3. Restore system power. Turn **ON** circuit breaker A, then circuit B on the system power entry module.
4. As the CFC recovers, watch the boot banner as it appears:

```

@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
24G Central Fabric Controller Boot Loader Version 3.0.g.1
  Created on Thu 03/11/2004 at 9:24a
  Copyright Allied Telesis Holdings K.K. 2003
VxWorks Version 5.5 for IBM PowerPC 440GP Rev. 2.0
  BSP version 1.2/3
  Copyright Wind River Systems, Inc., 1984-2002
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
FPGA Version 1.2
Starting Telesyn Product Software Loading.
Attaching to Flash File System ... done.
/tffs/ - Volume is OK
Press any key to stop automatic loading of software image...
```

5. Press any key on the management device. When the boot countdown timer is running, press any key. Depending on how the CFC NVRAM is configured, this countdown timer may be 7 seconds or it may be 1 second, so the best thing to do is to hit **ENTER** repeatedly in anticipation of this prompt. This stops the boot process and gives control of the session to the user.

6. The boot loader prompt appears.

[Allied Telesyn Boot Loader]:

7. Type “c” (for change parameter) and press **ENTER**. As each parameter appears, press **ENTER** to accept the current value, until the parameter “**BOOTFLAGS**” is reached. At this point, add the hexadecimal value “**0x10000**” to the existing value and press **ENTER**. This enables the password recovery mode. For example, if the existing value is “0x1000”, enter “0x101000” (0x1000 + 0x100000).

[Allied Telesyn Boot Loader]: c

'.' = clear field; '-' = go to previous field; ^D = quit

```

BOOTSERVER NAME      : l abserver1
BOOTSERVER IPADDR    : 172. 16. 5. 5
NETWORKLOAD          : cfc5-62/vxWorks
HOSTNAME              : cfc5-62
MGMT IPADDR          : 172. 16. 5. 62
GATEWAY IPADDR       : 172. 16. 5. 1
SUBNETMASK           : 255. 255. 255. 0
FTP USERNAME         : target
FTP PASSWORD         : tel esyn
BOOTFLAGS            : 0x1000 0x101000
    
```

8. Type “@” and press **ENTER** to reboot the system again.

[Allied Telesyn Boot Loader]: @

9. The boot sequence starts again. This time, let the countdown timer expire and the system reboot automatically. Note the message that appears indicating that password reset has been performed. All existing users and passwords have been removed from the system and the default user ID and password combination (**officer/officer**) has been restored.

```

@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
24G Central Fabric Controller Boot Loader Version 3.0.g.1
  Created on Thu 03/11/2004 at 9:24a
  Copyright Allied Telesis Holdings K.K. 2003
VxWorks Version 5.5 for IBM PowerPC 440GP Rev. 2.0
  BSP version 1.2/3
  Copyright Wind River Systems, Inc., 1984-2002
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
FPGA Version 1.2
    
```

Starting Telesyn Product Software Loading.
Attaching to Flash File System ... done.

/tffs/ - Volume is OK

Press any key to stop automatic loading of software image...

0

Automatically loading software image...

Boot album is (current, attempt 1): '/tffs/load/cfc6_4.0.0.tar'

Checking Album's integrity...done

Loading vxWorks.bin.gz... (3356160 bytes)

Starting at 0x10000...

Attached TCP/IP interface to emac unit 0

Attaching interface lo0...done

Adding 31077 symbols for standalone.

@@

```

, @@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
, @@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
, @@@@ @@@@ @@@@"" "" @@@@@@ @@@@@@ @@@@@@,
, @@" @@" , @@" "" @@@@@ @@@@@ @@@@@,
, @" , @@" , @@" "" @@@@@ @@@@@ @@@@@,
/" /" /" "" @@@@@ @@@@@ @@@@@
@@@@@@ /" @""@ @""@ @""@
/ @@ @@@@ @@ @@@@ @@@, @@ @ @@ @
@@ @@ @@ @@ @@ " @@ @ @@@ @ @@@@@ @@@@@ @@@@@
@@ @@@@ @@ @@@@ @@@ @@@ @ @@@ @@@@@. @@@@@. @@@@@
@@ @@ @@ @@ , @@ @@ @ @@ @@@@@@. @@@@@. @@@@@
@@ @@@@ @@@@@ @@@@@ " @@@ @@ @ @@ @@@@@ @@@@@ @@@@@

```

24G Central Fabric Controller Version 4.0.0 (Customer-Release Build)

Created on Mon 05/16/2004 at 11:12p

Copyright Allied Telesis Holdings K.K. 2003
VxWorks Version 5.5 for IBM PowerPC 405GP Rev E
BSP version 1.2/2
Copyright Wind River Systems, Inc., 1984-2002

Memory Size: 127 MB

@@

* WARNING: Password reset mode has been activated for this reboot.

System Time is 2004-05-26 10:39:04.070

System initializing...

User Access Verification

10. Enter the user ID **officer**.

Username: officer

11. Enter the user password **officer**.

Password: *(typed password "officer" hidden from view)*

12. The user is logged into the system. System user data can be re-configured and stored in the database. Other configuration data remains intact.

13. Note that there is a security risk while the default user ID and password are enabled. To minimize this risk, the default password should be modified **as soon as possible** or the default "officer" account should be replaced by a different security officer account.

officer SEC>> show user

--- User Authentication Database -----

Username:	officer	Description:	Security Officer	User
Privilege.....	SECURITY	Status.....	Enabled	Logins..... 1
	OFFICER			
Telnet User...	Yes	Last Login...	2004-05-26 10:40:01	Fails..... 0
				Lockouts... 0

officer SEC>>

3.8.3 Editing Functions, Keystrokes, and Abbreviations

The fMAP product supports line editing, line recall, and abbreviations, so that command line input and editing can be done very quickly once command syntax and the line editing commands are learned.

Table 3-9 lists the terminal editing and keystroke functions most commonly used.

TABLE 3-9 Terminal Editing Functions and Keystrokes

Action	Key Sequence
Move cursor within command line	left and right arrow
Delete character to left of cursor	[Delete] or [Backspace]
Clear command line	[Ctrl/U]
Recall previous command in command history	CTRL/P or up arrow
Recall next command in command history	CTRL/N or down arrow
Automatically complete a partially entered command keyword	[Tab] or [Ctrl/I]

The fMAP product also allows the following methods to make entering commands more efficient:

- Commands and parameters can be in lower case.
- Commands and parameters can be abbreviated (such as en for ENABLE, sh for SHOW, etc.)
- The equal sign (=) is not needed as long as a parameter is paired with a space and then a value.

Note: Throughout this document, all syntax will use complete words, with verbs and parameters in upper case and the pairing of parameters and values with equal signs.

3.8.4 Control of CLI command confirmation

3.8.4.1 Overview

CLI commands that may result in destructive actions will warn the user by responding to the input of such commands with a prompt asking the user to confirm the requested action with a “YES or Y” or “NO or N”. The user must respond with either a “YES or Y” or “NO or N”. The system will continue to prompt for this response until the user inputs a correct response. This provides the system a certain level of protection from unwanted destructive events.

Note: It is recommended that confirmation prompting be enabled.

3.8.4.2 Disable/Enable Confirmation

CLI Confirmation can be disabled if the user requires it. Disabling is especially useful when executing command scripts on the system. Scripts are discussed in detail in section 17.5.

The user disables confirmation by:

```
> DI SABLE CONFIRMATION
```

The user enables confirmation by:

```
> ENABLE CONFIRMATION
```

3.8.5 Command Alias

Command alias functionality allows fMAP users to define shortcuts to command strings to simplify the use of the CLI. It allows users to create shortcut strings which can be typed in place of commonly used (longer) commands. With alias command strings, the user is able to quickly perform operations without having to type the full command line.

3.8.5.1 Usage notes

- Command alias allows the user to create a shortcut to a command string; thus, requiring only the input of the alias instead of the extended command string. For example, suppose the user types the following, `CREATE ALIAS=reboot STRING="RESTART CARD ACTCFC COLD"`. From this point forward, the user would only be required to type “reboot” at the CLI to reboot the active CFC.
- When the user inputs an alias that happens to fail for some reason, the failure will be displayed exactly as if the user entered the full command string.
- The alias must encompass an entire command line. Alias strings cannot be used to substitute a portion of a command.
- The alias is not case-sensitive; similar to normal CLI commands.
- The system has an upper limit of 200, 255-character long alias entries.
- Alias commands persist between the ACTCFC and INACTCFC. Alias commands are available to all users and are protected by the user level settings.
- Question mark (?), the CLI help symbol, is not available for aliases. For example, if a user created an alias called “mkcard” that takes a parameter for card slot and card type, the following would not provide any useful information: `Manager>> mkcard ?`.
- If an alias was created using a command whose definition changed as the result of an upgrade or the alias is associated with a command that no longer exists after the upgrade, the alias will remain after the upgrade, but will no longer work. Auditing of the aliases over upgrades to ensure that they match any commands in the current command set is not supported.
- An existing alias cannot be overwritten. To reuse an existing alias name, the user must first destroy and then create the new definition of the alias using the `DESTROY` and `CREATE` commands.
- An alias command name that is all or part of a command verb is not allowed. For example, the alias command could not be “`CREATE`”, “`CREAT`”, “`CRE`”, “`CR`” or “`C`”, because it could potentially overwrite the command verb “`CREATE`”.
- All aliases are visible to all user privilege levels. Validation of the user privilege level (to execute a certain command), is done when attempting to use the alias. If the user does not have the privilege level required for the command, command execution will fail. This will be discussed in more detail later in this subsection.
- Alias names must be alphanumeric.
- Nesting of alias commands is not supported within commands with aliases.

- Alias command strings must substitute a CLI command string from the root of the command. For instance, the user may set a command string “ge1prof” to be equivalent to “SHOW PROFILE NAMES GE1” (the entire command string) by creating the alias:

```
CREATE ALIAS=ge1prof STRING="SHOW PROFILE NAMES GE1"
```

When the user inputs this alias, the user will type at the prompt:

```
officer SEC>> ge1prof
```

However, the alias may NOT be part of a command. For instance:

```
CREATE ALIAS=ge1prof STRING="PROFILE NAMES GE1"
```

is not acceptable. If the user attempts to input it as:

```
officer SEC>> SHOW ge1prof
```

the CLI responds with an error indicating that the command did not function correctly.

With this in mind, alias names must not exactly match any of the CLI command root keywords, such as **ADD**, **DELETE**, **CREATE**, **SET**, **SHOW**, etc. The **CREATE ALIAS** command validates this and rejects any attempts to do so. This prevents the user from creating an alias string that overrides an existing command.

- An Alias command string may be defined to take input values. For example,

```
officer SEC>> CREATE ALIAS=makecard STRING="CREATE CARD=$1 $2"
```

takes two input values. These values are typed in the order indicated by the $\$n$ part of the string, and separated by a space. If the user types:

```
officer SEC>> makecard 4 GE1
```

a GE1 card is created in slot 4, as if the user typed:

```
officer SEC>> CREATE CARD=4 GE1
```

- Alias command strings are made available to all users in the system. Since the alias command is mapped to actual CLI commands, the substituted CLI command is verified to be valid for the privilege level of the user using the alias. For instance, if there is an alias string “reboot” to substitute the command “RESTART CARD=ACTCFC CODE FORCE” (**CREATE ALIAS=reboot STRING=restart card=actcfc force**) and a user with **USER** privileges attempts to use it, the response is an error indicating that the command is not available for the current user privilege.
- Alias command strings persist over system reboots. Because of this, there are limits on the number of aliases stored and the maximum size of the alias name and the substitution string. As mentioned earlier, a maximum of 200 alias commands may be created and stored with the maximum name length being 40 characters and the maximum substitution string being 256 characters.
- If there are existing alias commands that conflict with a new verb/action, by overriding it, that was introduced as a result of system software upgrade, the alias will be automatically removed during the upgrade. For instance if there was an existing alias:

```
CREATE ALIAS="clear" STRING="PURGE USERS"
```

and the upgrade introduces a new command with the verb/action **CLEAR**, for example:

“CLEAR DATABASE”

then the “clear” alias would override the CLEAR verb/action, causing the command to be interpreted as “PURGE USERS DATABASE”, which is invalid. Therefore, during the upgrade, checks are made for conflicts and any alias commands that would result in a conflict are destroyed.

If an existing command that is referenced in an alias has changed syntax, thereby causing an invalid syntax, no correction is taken. Once a user attempts to use the alias, an “Unable to Parse” error is displayed.

3.8.5.2 Command Usage

CREATE ALIAS=*name* STRING=*substitution*

The CREATE ALIAS command allows the user to define an alias string and the command string that it is a short-cut for. The alias name string specifies the case-insensitive literal string which will be used in place of the CLI string provided in the STRING=*substitution string* parameter. The *substitution string* value may not be the same as an existing CLI command ROOT keyword, such as ADD, DELETE, SHOW, SET, etc. This command may be executed by users with a MANAGER privilege level or higher.

DESTROY ALIAS=*name*

The DESTROY ALIAS command allows the user to remove an existing alias from the persisted list of alias commands. The alias name string specifies the case-insensitive literal string which is to be destroyed. This command may be executed by users with a MANAGER privilege level or higher.

SHOW ALIAS[={*name*|ALL}]

The SHOW ALIAS command allows the user to view a list of all alias commands and their corresponding substitution strings. If the user enters a name value, that alias information is displayed. If a name value is not entered, the default is to display all alias commands. The list of alias commands is displayed in alphabetical order. This command may be executed by users with a USER privilege level or higher.

3.8.5.3 System Default Aliases

Certain aliases will be created by the fMAP upon system start-up. These default aliases are illustrated below (they were displayed using the SHOW ALIAS command):

```
officer SEC>> SHOW ALIAS
--- Alias Commands ---
Alias Name           Substitution String
-----
showdebug..... showoamp
                   $1; showint; showtraf; showimgp; showvlan; showlag; showswit
                   ch; showstp; showdhcp; showepsr; showrtip; showvc; showuser; s
                   howsys
showdhcp..... show dhcprelay
showepsr..... show epsr all;
showimgp..... show imgp; show igmpsnooping count interface all; show
                   igmpsnooping count messengeresponse; show igmpsnooping
                   card all full; show igmpsnooping interface all full
showint..... show interface all queuecount; show interface all
                   counter; show ip arp all; show ip connections; show ip
                   interface all; show ip interface all full; show ip
                   route all;
```

```

showlag..... show lag all
showoamp..... show card $1 software; show card $1 ports; show card
actcfc cpu; show card inactcfc cpu; show card actcfc
memory heap; show card memory quickheap; show card
actcfc memory messagebuffers; show card inactcfc
memory messagebuffers; show profile autoprov
adsl8s; show profile autoprov adsl16; show profile
autoprov adsl16b; show profile autoprov adsl24; show
profile autoprov pots24; show profile autoprov
adslport; show profile autoprov geport; show profile
autoprov potsport; show profile autoprov fe10; show
profile autoprov ge1; show profile autoprov cfc6; show
profile names; show feature all; show feature all keys;
showrtp..... show rtp interface all full;
showstp..... show stp; show stp counter; show stp interface all
showswitch..... show switch; show switch fdb; show switch counter
showsys..... show system; show system cooling; show fanmodule; show
sntp; show bootserver; show alarms all; show
contactalarm all severity all state all;
showtraf..... show classifier all full; show classifier all
interface all full; show traffi cdescriptor all; show
arpfilter; show qos; show accesslist all
showuser..... show user; show telnet server; show sessions; show
system userconfig; show log filter; show log
output; show radius; show tacplus; show snmp; show snmp
community all;
showvc..... show vc interface all full;
showvlan..... show vlan all full

```

The default aliases are created dynamically when the system reboots and no other aliases have been created by the user. If, for some reason, a default alias or aliases have been deleted, the user can recreate them all using the **SETDEFAULTS ALIAS** command.

3.8.5.4 Usage Examples

1. Create Alias commands:

An example of creating an alias would be to define a shortcut string for rebooting the active and inactive CFCs. The alias commands would be created as follows:

```

officer SEC>> CREATE ALIAS=rebootact STRING="reboot card actcfc force"
officer SEC>> CREATE ALIAS=rebootinact STRING="reboot card inactcfc force"

```

To use these commands, the user would simply type:

```

officer SEC>> rebootact
officer SEC>> rebootinact

```

2. Destroying Alias commands:

To destroy existing alias commands, the user would type the following:

```

officer SEC> DESTROY ALIAS=rebootact

```

The alias would no longer be available for use nor persisted.

3. Displaying all alias commands available:

To view the list of alias commands, the user would type the following:

```

officer SEC> SHOW ALIAS

```

The output is an alphabetized list of alias names and their corresponding substitution strings:

```

---Alias command strings -----
Alias      String

```

```
-----
rebootact          reboot card actcfc force
rebootinact       reboot card inactcfc force
-----
```

To view a specific alias command, the user would type the following:

```
offi cer SEC> SHOW ALIAS=rebootact
```

The output is a display of the given alias' substitution string:

```
rebootcfc="reboot card actcfc force"
```

4. Creating an alias command with input variables

It is possible to create an alias that has placeholders for input data. An example of an alias that accepts input data:

```
offi cer SEC> CREATE ALIAS=shcard STRING="SHOW CARD=$1"
```

Then, to use this alias to view the information about card 4:

```
offi cer SEC> shcard 4
```

3.8.5.5 Multiple Command Stringing

Multiple commands can be strung together on the command line using the “;”. For example, the following commands can be entered as illustrated and the responses will be returned in the order of command entry.

```
SHOW SESSION; SHOW USER; SHOW TRANSFER ALL
```

Example:

```
offi cer SEC>> SHOW SESSION; SHOW USER; SHOW TRANSFER ALL
--- Active (logged in) Users -----
ID User          Port   Location      Status Login Time      Deact
-----
0  offi cer       Console local      CONN  2004-06-17
                               11: 18: 54      -
-----

--- User Authentication Database -----
Username:  offi cer      Description:  Securi ty Offi cer User
Pri vi lege. . . . . SECURITY Status. . . . . Enabled  Logi ns. . . . . 2
OFFICER
Tel net User. . . Yes      Last Logi n. . . 2004-06-17  Fai ls. . . . . 0
                               11: 18: 54
                               Lockouts. . . 0
-----

No Transfer in progress
-----
```

TABLE 3-10 Command Alias Commands

Object	Verb	Syntax	Description
ALIAS	CREATE	CREATE ALIAS=aliasname STRING=substitution	Creates a command alias.
	DESTROY	DESTROY ALIAS={ aliasname-list ALL }	Deletes a command alias.
	SHOW	SHOW ALIAS [= { aliasname-list ALL }]	Displays a command alias.
	SETDEFAULTS	SETDEFAULTS ALIAS	Resets alias defaults.

3.8.6 Provisioning the Login Banner

The login banner appears as the first system output presented to a user when they log into the fMAP system. The user has the ability to provision or customize the system login banner. The banner could be changed to present a message to all users or a *message of the day*.

Usage notes follow:

- The login banner is a text message presented to a user after successful login authentication.
- The login banner may be the same for all users or be different based upon privilege level (USER, MANAGER, or SECURITY OFFICER).
- This functionality supports up to three(3) different text entries for login banner.
- Login banner text may either be specified by directly entering the text using CLI commands or added using a script file that contains the desired text.
- The login banner entry may be up to **255** characters long.

Note that changes to the banner persist over software upgrades, over both active and inactive CMs in a duplex system, and over restarts.

Note: Only the Security Officer can change the login banner.

Following is an example of provisioning the login banner.

Set the login banner default:

```
officer SEC>> SETDEFAULTS LOGINBANNER ALL
Info (010017): Operation Successful
Set the login banner:
officer SEC>> SET LOGINBANNER STRING="fMAP System Testing"
Info (010017): Operation Successful
```

```

Display the login banner:
officer SEC>> SHOW LOGINBANNER
--- Login Banner Settings ---
Privilege Level: USER, MANAGER, SECURITY OFFICER
fMAP System Testing
-----
    
```

Logout and log back into the system to see the changed banner:

User Access Verification

Username: officer

Password: *****

fMAP System Testing

officer SEC>>officer SEC>>

TABLE 3-11 Provisioning the Login Banner Commands

Object	Verb	Syntax	Description
LOGINBANNER	SETDEFAULTS	SETDEFAULTS LOGINBANNER [{ USER MANAGER SECURITYOFFICER ALL }]	Sets the login banner to the default.
	SET	SET LOGINBANNER { FILE=filename STRING=string } [{ USER MANAGER SECURITYOFFICER ALL }]	Allows the Security Officer user to change the login banner.
	SHOW	SHOW LOGINBANNER	Display the current login banner text.

3.8.7 Customizing the CLI Prompt

When the user logs into the fMAP system, a default CLI prompt is provided as displayed here:

Username: officer

Password:

officer SEC>> <----- CLI prompt

The user has the ability to provision or customize the system CLI prompt. The changes to the CLI prompt affect all user sessions immediately after the prompt settings are modified.

Usage notes:

- The CLI prompt is a text message presented to a user after successful login authentication.
- The CLI prompt is the same for all users.

The user can change the format of the CLI prompt. Aside from plain text, the CLI prompt can contain any of the following formats:

- Device IP (%i)
- System name (%n)
- User name (%u)
- Date (%d)
- Time (%t)
- Security level (%s)

Note that changes to the CLI prompt persist over software upgrades, over both active and inactive CMs in a duplex system, and over restarts.

Following are examples of provisioning the CLI prompt. First, change the CLI prompt:

```
officer SEC>> SET PROMPT "Testing the CLI User Prompt"
Info (010017): Operation Successful
Testing the CLI User Prompt>>
Reset the CLI prompt:
Testing the CLI User Prompt>> SETDEFAULTS PROMPT
Info (010017): Operation Successful
officer SEC>>
Set the CLI prompt to the system IP:
officer SEC>> SET PROMPT="%i"
Info (010017): Operation Successful
172.16.66.71>>
172.16.66.71>>
Reset the CLI prompt:
172.16.66.71>> SETDEFAULTS PROMPT
Info (010017): Operation Successful
Set the CLI prompt to the system name:
officer SEC>> SET PROMPT="%n"
Info (010017): Operation Successful
Lab System 42>>
Reset the CLI prompt:
Lab System 42>> SETDEFAULTS PROMPT
Info (010017): Operation Successful
Set the CLI prompt to the system user:
officer SEC>> SET PROMPT="%u"
Info (010017): Operation Successful
LabUser>>
Reset the CLI prompt:
LabUser>> SETDEFAULTS PROMPT
Info (010017): Operation Successful
officer SEC>>
Set the CLI prompt to the system date:
officer SEC>> SET PROMPT="%d"
Info (010017): Operation Successful
2004-05-12>>
2004-05-12>>
Reset the CLI prompt:
2004-05-12>> SETDEFAULTS PROMPT
```

```

Info (010017): Operation Successful
Set the CLI prompt to the system date:
officer SEC>> SET PROMPT="%t"
Info (010017): Operation Successful
15:40:24>>
15:40:25>>
Reset the CLI prompt:
15:40:26>> SETDEFAULTS PROMPT
Info (010017): Operation Successful
Set the CLI prompt to the Security Officer:
officer SEC>> SET PROMPT="%s"
Info (010017): Operation Successful
SEC>>
SEC>>
Reset the CLI prompt:
SEC>> SETDEFAULTS PROMPT
Info (010017): Operation Successful
officer SEC>>
officer SEC>>

```

TABLE 3-12 Provisioning the CLI Prompt

Object	Verb	Syntax	Description
PROMPT	SETDEFAULTS	SETDEFAULTS PROMPT	Sets the CLI prompt to the default.
	SET	SET PROMPT=string or Device IP (%i) System name (%n) User name (%u) Date (%d) Time (%t) Security level (%s)	Allows the user to change the CLI prompt. Note that the parameter <i>string</i> must be enclosed in double quotes “ <i>string</i> ”.

3.8.8 Setting Up User Accounts, Profiles, and Sessions

The security levels required for the management interface follows these rules:

- Security Officers can add, remove, or modify other user accounts.
- Users can only change the attributes of their own accounts.
- Managers can control various aspects of a User’s account, such as showing all active sessions or showing statistics.

The specific tasks can be divided as follows. For details on these commands, refer to [Table 3-13](#).

- User Management
 - Add a User
 - Delete a User
 - Lockout a User
 - Change User attributes

- Change password
- Show current (logged in) users
- Show or reset user statistics
- Session Management
 - Display active sessions
 - Send messages between sessions
 - Deactivate session

3.8.9 Command Summary for User Administration

Table 3-13 lists the commands for User Administration

TABLE 3-13 Summary of Commands for User Administration

Object	Verb	Syntax	Description
USER PASS- WORD LOGIN	ADD	<pre> ADD USER=login-name PASSWORD=password [FORMAT={ CLEARTEXT MD5 }] [DESCRIPTION=description] [PRIVILEGE={ USER MANAGER SECURITYOFFICER }] [LOGIN={ TRUE FALSE ON OFF YES NO }] [TELNET={ YES NO }] </pre>	Adds a user and all of its attributes.

TABLE 3-13 Summary of Commands for User Administration (Continued)

Object	Verb	Syntax	Description
USER	SET	<pre> SET USER=login-name [PASSWORD=password [FORMAT={ CLEARTEXT MD5 }]] [DESCRIPTION=description] [PRIVILEGE={ USER MANAGER SECURITYOFFICER }] [LOGIN={ TRUE FALSE ON OFF YES NO }] [TELNET={ YES NO }] </pre>	Changes settings for a specific user.
	RESET	<pre> RESET USER [=login-name] [COUNTER [={ ALL GLOBAL USER }]] </pre>	Resets the User Authentication Database counters for one or all users, or resets global counters for the User Authentication Facility. Statistics about users are shown with the SHOW USER and SHOW SYSTEM USERCONFIG commands.
	DIS-ABLE	<pre> DISABLE USER=login-name </pre>	Locks out a user. The account is still present, but the user that owns the account is unable to login
	ENABLE	<pre> ENABLE USER=login-name </pre>	Re-enables an account that was previously disabled.
	DELETE	<pre> DELETE USER=login-name </pre>	Deletes a single user.
	PURGE	<pre> PURGE USER </pre>	Deletes all users from the database and recreates the default Security Officer user.
PASS-WORD	SET	<pre> SET PASSWORD (See SET USER) </pre>	Allows Users to change their password at anytime. The command prompts for the old password and asks to reconfirm the new password.

TABLE 3-13 Summary of Commands for User Administration (Continued)

Object	Verb	Syntax	Description
MORE	ENABLE	ENABLE MORE	Stops the terminal output at the end of a window and displays --MORE --. Press return to continue the output
	DIS- ABLE	DISABLE MORE	Disables the --More-- output and lets fMAP product output continue to run past the end of the window.
SYSTEM USER- CONFIG	SET	SET SYSTEM USERCONFIG [LOGINFAIL=1..10] [LOCKOUTPD=0..30000] [MANPWDFAIL=1..5] [SECUREDELAY={ OFF 0 1..90 }] [MINPWDLEN=1..23] [PERSISTTIMER=1..1440]	Change global user settings.
		SET SYSTEM USERCONFIG { MANAGERPASSWORD={ password NONE } SECURITYOFFICERPASSWORD={ password NONE } } [FORMAT={ CLEARTEXT MD5 }]	Change Manager or Security Officer password.
	SHOW	SHOW SYSTEM USERCONFIG	Shows the global user settings as a result of the SET SYSTEM USER-CONFIG command.
SESSIONS	SHOW	SHOW SESSIONS	Displays the following information about the users logged in.
MES- SAGE SESSION	SEND	SEND MESSAGE=message-text SESSION={ session-list ALL }	Allows a user to send messages through their login session.

TABLE 3-13 Summary of Commands for User Administration (Continued)

Object	Verb	Syntax	Description
SESSION	DEACTIVATE	<pre> DEACTIVATE SESSION={ session-list ALL } [{ CANCEL [MESSAGE=message-text] [DELAY=1..600] }] </pre>	Forces a user off of the system.
None	LOGOFF	--	Logs off the system. The user can also use LOGOUT or EXIT.
None	HELP	--	Refer to 3.8.13

3.8.10 TACACS+ and RADIUS Authentication

Terminal Access Controller Access Control System Plus (TACACS+) and Remote Authentication Dial In User Service (RADIUS) Authentication give the user the ability to keep a centralized database of login IDs and passwords for users. This allows the user to manage a single user authentication database over a large network; eliminating the requirement to manage many user databases over a potentially large network of devices.

RADIUS is a client/server protocol for performing network-based user authentication, authorization and accounting. RADIUS is defined by RFCs 2865, 2866, 2867, 2868, and 3575.

TACACS+ is a proprietary access control protocol as described in RFC 1492.

TACACS+ and RADIUS authentication operates by using an external server as a means to authenticate logins to the system. When a user attempts to login to the system, the system sends the request to the configured TACACS+ or RADIUS server which then processes the attempt. If the attempt is successful, the user is logged in. If the attempt fails, the system prevents the user from logging in.

The system supports:

- Up to 8 servers of type TACACS+ and RADIUS.
- Dual challenge authentication.
- Vendor-specific Attribute Value (AV) pairs for automatic assignment of security level.
- Displays PASSCODE instead of PASSWORD to inform the user they are logging in through an external service.
- A local login of “last resort” when no external authentication servers are reachable.

By default, RADIUS and TACACS+ authentication is *disabled*. When a user first adds a RADIUS or TACACS+ server, future login attempts use that server to authenticate user logins. Authentication against the local user database is disabled. Local user logins are then only allowed if none of the configured RADIUS and TACACS+ servers are reachable or if all RADIUS and TACACS+ servers are deleted.

If multiple RADIUS and/or TACACS+ servers are defined, each server is contacted in turn. First, the configured RADIUS servers are contacted, then TACACS+ servers. If a server returns an authentication failure, a request is sent to the next server. This process continues until a server returns *authentication success* or until all the servers have been contacted and returned failure.

3.8.11 Radius Details and Commands

In release 6.0, the UDP port can be set as ON or OFF to either:

AUTHENTICATION

ACCOUNTING

RADIUS user command examples:

```
officer SEC> SHOW RADIUS
--- RADIUS Servers -----
No RADIUS servers configured.
```

No RADIUS servers are provisioned.

```
officer SEC> ADD RADIUS SERVER=radi us-1. supersecure. com SECRET=new2day
Warning (020127): Parameter SERVER, could not resolve hostname "radi us-
1. supersecure. com"
Info (010017): Operation Successful
```

Note: The Warning indicates that this server has not yet been configured in the DNS server. The system adds the RADIUS server anyway.

```
officer SEC> SHOW RADIUS
--- RADIUS Servers -----
Hostname/IP Address          Status      Port      Retries  Timeout
-----
radi us-1. supersecure. com   Enabled    1812      3         5
```

Status of Enabled, Port of 1812, Retries of 3, and Timeout of 5 are all default settings.

```
officer SEC> SET RADIUS SERVER=radi us-1. supersecure. com PORT=1645
Warning (020127): Parameter SERVER, could not resolve hostname "radi us-
1. supersecure. com"
Info (010017): Operation Successful
```

Changed the port from 1812 to 1645.

```
officer SEC> SHOW RADIUS
--- RADIUS Servers -----
Hostname/IP Address          Status      Port      Retries  Timeout
-----
radi us-1. supersecure. com   Enabled    1645      3         5
```

```
officer SEC> DELETE RADIUS SERVER=radi us-1. supersecure. com
Warning (020127): Parameter SERVER, could not resolve hostname "radi us-
1. supersecure. com"
Info (010017): Operation Successful
```

```
officer SEC> SHOW RADIUS
--- RADIUS Servers -----
```

Deleted the RADIUS server **radi us-1. supersecure. com**

TABLE 3-14 RADIUS commands

Object / Key Word(s)	Verb	Syntax	Description
RADIUS SERVER	ADD	<pre> ADD RADIUS SERVER={ ipaddress-list hostname-list } SECRET=secret [AUTHPORT=1..65535] [ACCTPORT=1..65535] [RETRIES=0..10] [TIMEOUT=1..60] [AUTHENTICATION={ ON OFF }] [ACCOUNTING={ ON OFF }] </pre>	<p>The ADD RADIUS SERVER command allows the user to set up a RADIUS server for user authentication purposes. One or more IP addresses or hostnames plus a shared secret are required parameters. Users may optionally adjust the UDP port number where the RADIUS requests should be directed (port 1812 by default), the number of times a request should be retried (3 by default) and the timeout in seconds for each request (5 seconds by default).</p> <p>In release 6.0, the port type can be set to AUTHENTICATION or ACCOUNTING, and the number of that type of port is set.</p> <p>Defaults are:</p> <p>Enabled</p> <p>Port = 1812</p> <p>Retries = 3</p> <p>Timeout = 5</p>
RADIUS SERVER	DELETE	<pre> DELETE RADIUS SERVER={ ipaddress-list hostname-list ALL } </pre>	<p>The DELETE RADIUS SERVER command is used to remove RADIUS servers from the system. Once removed, user authentication requests are no longer sent to those servers.</p>
RADIUS SERVER	DISABLE	<pre> DISABLE RADIUS SERVER={ ipaddress-list hostname-list ALL } </pre>	<p>The DISABLE RADIUS SERVER command disables one or more RADIUS servers for use in user authentication requests. Once disabled, the RADIUS server(s) are not used for user authentication requests.</p>

TABLE 3-14 RADIUS commands

Object / Key Word(s)	Verb	Syntax	Description
RADIUS SERVER	ENABLE	<pre> ENABLE RADIUS SERVER={ ipaddress-list hostname-list ALL } </pre>	The ENABLE RADIUS SERVER command is used to enable one or more RADIUS servers for use in user authentication requests. Once enabled, the RADIUS server(s) are used for future user authentication requests.
RADIUS AUTHMODE	SET	<pre> SET RADIUS AUTHMODE={ LOGIN COMMAND } </pre>	The SET RADIUS AUTHMODE command is used to change the authentication mode for RADIUS servers. When the RADIUS authentication mode is set to LOGIN, the user will be logged in with the privilege level assigned by the RADIUS server. If the authentication mode is set to COMMAND, then the user is always logged in at USER privilege level and must run the ENABLE {MANAGER SECURITYOFFICER} command to request increased privilege. For RADIUS, the privilege level is determined by examining the Service-Type attribute in the Access-Accept packet returned by the RADIUS server. A Service-Type of NAS-Prompt or Login is equivalent to USER level privilege. A Service-Type of Administrative equates to MANAGER or SECURITY-OFFICER privilege.

TABLE 3-14 RADIUS commands

Object / Key Word(s)	Verb	Syntax	Description
RADIUS SERVER	SET	<pre>SET RADIUS SERVER={ ipaddress-list hostname-list ALL } [SECRET=secret] [AUTHPORT=1..65535] [ACCTPORT=1..65535] [RETRIES=0..10] [TIMEOUT=1..60] [AUTHENTICATION={ ON OFF }] [ACCOUNTING={ ON OFF }]</pre>	<p>The SET RADIUS SERVER command allows the user to change the settings of one or more existing configured RADIUS servers. Users can change the servers' shared secret, port number, retries and timeout values.</p> <p>In release 6.0, the port type can be set to AUTHENTICATION or ACCOUNTING, and the number of that type of port is set.</p>
RADIUS	SHOW	<pre>SHOW RADIUS</pre>	<p>The SHOW RADIUS command displays a table containing information regarding the RADIUS server configuration. The information includes each RADIUS server's hostname or IP address, status (enabled or disabled), port, retries, and timeout values. The shared secret is not displayed for security reasons.</p>

TABLE 3-15 RADIUS commands parameters

Object / Key Word(s)	Verb	Syntax	Description
SERVER	ADD DELETE DISABLE ENABLE SET	<pre>SERVER={ ipaddress-list hostname-list ALL }</pre>	<p>The SERVER parameter is used to specify one or more IP addresses or hostnames to send RADIUS authentication requests to.</p>
SECRET	ADD SET	<pre>secret</pre>	<p>The SECRET parameter is used to specify the secret that is shared with the RADIUS server for use in authentication requests. The secret must be an alphanumeric string of 64 characters or less in length.</p>

TABLE 3-15 RADIUS commands parameters

Object / Key Word(s)	Verb	Syntax	Description
Port	ADD SET	PORT=1..65535	The PORT parameter specifies the UDP port to which RADIUS authentication requests should be directed to on the RADIUS server. If not specified, the default is port 1812, the IANA-assigned port for RADIUS.
Retries	ADD SET	RETRIES=0..10	The RETRIES parameter specifies the number of times a user authentication request should be retried. Once the maximum number of retries has been reached without a response from the RADIUS server, the next RADIUS or TACACS+ server or local database is consulted to determine the validity of the authentication attempt.
Timeout	ADD SET	TIMEOUT=1..60	The TIMEOUT parameter is used to specify the number of seconds to wait for a response back from the RADIUS server. If no response is received within the timeout period either the request is retried if there are retries remaining, the next RADIUS or TACACS+ server is contacted or authentication is attempted against the local user database.
AUTHMODE	SET	AUTHMODE={ LOGIN COMMAND	The AUTHMODE parameter is used to specify how user privilege level is assigned when user authentication is done using a RADIUS server. When the RADIUS authentication mode is set to LOGIN, the user will be logged in with the privilege level assigned by the RADIUS server. If the authentication mode is set to COMMAND, then the user is always logged in at USER privilege level and must run the ENABLE {MANAGER SECURITYOFFICER} command to request increased privilege.

3.8.12 TACACS+ Details and Commands

In release 6.0, the TACACS+ server can be set as ON or OFF to either:

AUTHENTICATION -

AUTHORIZATION -

ACCOUNTING -

TACACS+ user commands example:

```
officer SEC> ADD TACPLUS SERVER=172.16.5.5 KEY=tea2Go4Me
Info (010017): Operation Successful
officer SEC> SHOW TACPLUS
--- TACACS+ Servers -----
Hostname/IP Address          Status      Port    Retries Timeout
-----
172.16.5.5                   Enabled    49      3        5
```

Add a TACACS+ server. Status of Enabled, Port of 49, Retries of 3, and Timeout of 5 are all default settings.

```
officer SEC> SET TACPLUS SERVER=172.16.5.5 KEY=cU2Morrow PORT=31337 TIMEOUT=30
Info (010017): Operation Successful
```

Set the key, port, and timeout values.

```
officer SEC> DELETE TACPLUS SERVER=172.16.5.5
Info (010017): Operation Successful
officer SEC> SHOW TACPLUS
--- TACACS+ Servers -----
No TACACS+ servers configured.
```

Delete the TACACS+ server.

Note: TACACS+ users have a security level of 1 to 15; for the fMAP, 1 corresponds with User, 7 with Manager, and 15 with Security Officer. Therefore, the fMAP users "manager" and "securityofficer" correspond to the TACACS+ users "\$enab7\$" and "\$enab15\$". If these are not configured in TACACS+, then TACACS+ will refuse the authentication when a user tries the command "Enable manager/securityofficer". Also, the syntax of this command depends on the level of the user; someone at the User level will have both manager/securityofficer as options while someone at the Manager level will have only securityofficer available (if they are data filled at the TACACS+ server).

TABLE 3-16 TACACS+ commands

Object / Key Word(s)	Verb	Syntax	Description
TACPLUS SERVER	ADD	<pre> ADD TACPLUS SERVER={ ipaddress-list hostname-list } KEY=key [PORT=1..65535] [RETRIES=0..10] [TIMEOUT=1..60] [AUTHENTICATION={ ON OFF }] [AUTHORIZATION={ ON OFF }] [ACCOUNTING={ ON OFF }] </pre>	<p>The ADD TACPLUS SERVER command allows the user to set up a TACACS+ server to be used for user authentication purposes. One or more IP addresses or hostnames plus a shared key are required parameters. Users may optionally adjust the TCP port number to which the TACACS+ requests should be directed (port 49 by default), the number of times a request should be retried (3 by default) and the timeout in seconds for each request (5 seconds by default).</p> <p>In release 6.0, the following are added:</p> <p>AUTHENTICATION AUTHORIZATION ACCOUNTING</p> <p>Defaults are:</p> <p>Enabled Port = 49 Retries = 3 Timeout = 5</p>
TACPLUS SERVER	DELETE	<pre> DELETE TACPLUS SERVER={ ipaddress-list hostname-list ALL } </pre>	<p>The DELETE TACPLUS SERVER command is used to remove TACACS+ servers from the system. Once removed, user authentication requests are no longer sent to those servers.</p>
TACPLUS SERVER	DISABLE	<pre> DISABLE TACPLUS SERVER={ ipaddress-list hostname-list ALL } </pre>	<p>The DISABLE TACPLUS SERVER command disables one or more TACACS+ servers for use in user authentication requests. Once disabled, the TACACS+ server(s) are not used for user authentication requests.</p>

TABLE 3-16 TACACS+ commands

Object / Key Word(s)	Verb	Syntax	Description
TACPLUS SERVER	ENABLE	<pre> ENABLE TACPLUS SERVER={ ipaddress-list hostname-list ALL } </pre>	<p>The ENABLE TACPLUS SERVER command is used to enable one or more TACACS+ servers for use in user authentication requests. Once enabled, the TACACS+ server(s) are used for future user authentication requests.</p>
TACPLUS AUTHMODE	SET	<pre> SET TACPLUS AUTHMODE={ LOGIN COMMAND } </pre>	<p>The SET TACPLUS AUTHMODE command is used to change the authentication mode for use with TACPLUS servers.</p> <p>When the TACACS+ authentication mode is set to LOGIN or COMMAND, the USER privilege level is requested by the client and either accepted or rejected by the server. The user is always logged in at the USER privilege level and must run the ENABLE {MANAGER SECURITYOFFICER} command to request increased privilege.</p>

TABLE 3-16 TACACS+ commands

Object / Key Word(s)	Verb	Syntax	Description
TACPLUS SERVER	SET	<pre>SET TACPLUS SERVER={ ipaddress-list hostname-list ALL } [KEY=key] [PORT=1..65535] [RETRIES=0..10] [TIMEOUT=1..60] [AUTHENTICATION={ ON OFF }] [AUTHORIZATION={ ON OFF }] [ACCOUNTING={ ON OFF }]</pre>	<p>The SET TACPLUS SERVER command allows the user to change the settings of one or more existing configured TACACS+ servers. Users can change the servers' shared secret, port number, retries and timeout values.</p> <p>In release 6.0, the following are added:</p> <p>AUTHENTICATION AUTHORIZATION ACCOUNTING</p>
TACPLUS	SHOW	<pre>SHOW TACPLUS</pre>	<p>The SHOW TACPLUS command displays a table containing information regarding the TACACS+ server configuration. The information includes each TACACS+ server's hostname or IP address, status (enabled or disabled), port, retries, and timeout values. The shared secret is not displayed for security reasons.</p>

TABLE 3-17 TACACS+ commands parameters

Object / Key Word(s)	Verb	Syntax	Description
SERVER	ADD DELETE DISABLE ENABLE SET	<pre>SERVER={ ipaddress-list host- name-list ALL }</pre>	<p>The SERVER parameter is used to specify one or more IP addresses or hostnames to send TACACS+ authentication requests to.</p>
KEY	ADD SET	<pre>KEY=key</pre>	<p>The KEY parameter is used to specify the key that is shared with the TACACS+ server for use in authentication requests. The secret must be an alphanumeric string of 64 characters or less in length.</p>

TABLE 3-17 TACACS+ commands parameters

Object / Key Word(s)	Verb	Syntax	Description
Port	ADD SET	PORT=1..65535	The PORT parameter specifies the TCP port to which TACACS+ authentication requests should be directed to on the TACACS+ server. If not specified, the default is TCP port 49, the IANA-assigned port for TACACS/TACACS+.
Retries	ADD SET	RETRIES=0..10	The RETRIES parameter specifies the number of times a user authentication request should be retried. Once the maximum number of retries has been reached without a response from the TACACS+ server, the next TACACS+ server or local database is consulted to determine the validity of the authentication attempt.
Timeout	ADD SET	TIMEOUT=1..60	The TIMEOUT parameter is used to specify the number of seconds to wait for a response back from the TACACS+ server. If no response is received within the timeout period either the request is retried if there are retries remaining, the next TACACS+ server is contacted or authentication is attempted against the local user database.
AUTHMODE	SET	AUTHMODE={ LOGIN COMMAND	The AUTHMODE parameter is used to specify how user privilege level is assigned when user authentication is done using a TACACS+ server. When the TACACS+ authentication mode is set to LOGIN, the user will be logged in with the privilege level assigned by the TACACS+ server. If the authentication mode is set to COMMAND, then the user is always logged in at USER privilege level and must run the ENABLE {MANAGER SECURITYOFFICER} command to request increased privilege.

3.8.13 Using Help

Online help is available for all fMAP product commands. There are two types of online help:

1. For command string help, type in the start of a command and enter a space and a “?” at the end of the line. The fMAP product will display a list of possible parameters. After entering a parameter and a “?”, online help provides an explanation of the parameter. Entering a “?” alone will display all of the verbs available.
2. For complete online help, type HELP and the command. If the command is incomplete, there is an error message. Entering a space and a “?” will show the next valid parameter. When the command is complete, a complete description of the command is displayed.

Following are examples of using HELP for a command.

```
officer SEC> HELP SHOW
% Invalid or incomplete command
officer SEC> HELP SHOW SNMP
9400 Help version 2.0.0 - English Version
```

SYNTAX:

```
SHOW SNMP
```

DESCRIPTION:

The SHOW SNMP command displays information about the device's SNMP agent.

The following is example output from the SHOW SNMP command:

```
SNMP configuration:
-----
Status ..... Enabled
Authentication failure traps ... Enabled
Community ..... public
Access ..... read-only
Status ..... Enabled
Traps ..... Enabled
Open access ..... Yes
Community ..... Administration
Access ..... read-write
Status ..... Disabled
Traps ..... Disabled
Open access ..... No

SNMP counters:
-----
inPkts ..... 0 outPkts .....
(list of counters continues)
Parameter          Meaning

Status              The status of the SNMP agent or the specified
                    community (ENABLED or DISABLED)
                    (list and description of configuration parameters continues)
```

SEE ALSO:

```
ADD SNMP COMMUNITY
CREATE SNMP COMMUNITY
(list of related commands continues)
```

```
officer SEC>
```

```
officer SEC> HELP SHOW
% Invalid or incomplete command
officer SEC> HELP SHOW SNMP
9400 Help version 2.0.0 - English Version
```

SYNTAX:

```
SHOW SNMP
```

DESCRIPTION:

The SHOW SNMP command displays information about the device's SNMP agent.

The following is example output from the SHOW SNMP command:

SNMP configuration:

```
-----
Status ..... Enabled
Authentication failure traps .... Enabled
Community ..... public
Access ..... read-only
Status ..... Enabled
Traps ..... Enabled
Open access ..... Yes
Community ..... Administration
Access ..... read-write
Status ..... Disabled
Traps ..... Disabled
Open access ..... No
```

SNMP counters:

```
-----
inPkts ..... 0 outPkts .....
(list of counters continues)
Parameter      Meaning

Status          The status of the SNMP agent or the specified
                  community (ENABLED or DISABLED)
(list and description of configuration parameters continues)
```

SEE ALSO:

```
ADD SNMP COMMUNITY
CREATE SNMP COMMUNITY
(list of related commands continues)
```

```
officer SEC>
```

3.8.14 Display and Clear ARP Table

3.8.14.1 Overview

Address Resolution Protocol (ARP) is used to obtain the MAC address of an interface given the IP address for that interface. The system keeps an ARP cache and an associated ageing timer for each entry in the ARP table. Whenever a packet needs to be sent to a destination IP, the system checks the ARP cache to see if a matching entry exists for the specified IP address. If it does and the entry has not aged out, the MAC address from that entry is assumed to be valid and is used for sending the packet. Conversely, if the entry is not found or the entry

has aged out (or cache entry is dirty), the system performs an ARP broadcast for that IP address to obtain the corresponding MAC address. It will then update the ARP cache with the appropriate timestamp.

The user can display the ARP table for the MGMT port and inband management interface and manually remove unneeded entries from the table. The following commands are available.

TABLE 3-18 ARP commands

Object / Key Word(s)	Verb	Syntax	Description
IP ARP	SHOW	<pre> SHOW IP ARP [={ ipaddress-list ALL }] [INTERFACE={ type:id-range ifname-list MGMT ALL }] </pre>	<p>Displays entries in the system's ARP cache in a tabular format. To display all entries in the ARP cache, the user can type either SHOW IP ARP ALL or simply SHOW IP ARP. Specifying an IP address or a comma separated list of IP addresses displays the ARP entries for IP addresses that are valid in the supplied list. Being a SHOW operation the above SHOW IP ARP command can be executed by users with any privilege level.</p>
IP ARP	CLEAR	<pre> CLEAR IP ARP [={ ipaddress-list ALL }] [INTERFACE={ type:id-range ifname-list MGMT ALL }] [FORCE] </pre>	<p>Allows users to delete one or more entries from the ARP cache. To completely purge the ARP cache, the user specifies ALL or no option for ARP. In other words, CLEAR IP ARP ALL and CLEAR IP ARP are equivalent. This command prompts for a [yes/no] user confirmation, with "NO" being the default. The command only executes if the user types "YES". To override the yes/no confirmation prompt, the user must type the word "FORCE".</p> <p>Example, to delete all entries in the ARP cache w/o confirmation, type CLEAR IP ARP FORCE or CLEAR IP ARP ALL FORCE.</p> <p>This command can be executed only by users with a SECURITYOFFICER privilege level or higher.</p>

3.8.15 Filtering on the MGMT interface

Traffic filters can be configured and applied to the MGMT and inband interface. For more information refer to section Traffic Management [15.1](#).

3.8.16 Telnet Client (telnet to another device from CLI)

The fMAP provides support for a telnet client, such that from a CLI session, the user may TELNET to another device.

TABLE 3-19 Telnet Client commands

Object / Key Word(s)	Verb	Syntax	Description
TELNET	SET	<pre> SET TELNET [TERMTYPE=termstring] [INSERTNULL={ ON OFF }] </pre>	<p>Allows the user to set the system-wide settings of the telnet client configuration, including the TERMTYPE and INSERTNULL data. The TERMTYPE string is the string that will be sent to a remote telnet server during the negotiation of the telnet connection. The default value is XTERM.</p> <p>The terminal identification is usually used by the remote system to set the terminal attributes for the Telnet session. The INSERTNULL parameter, when set to ON, specifies that a NULL character should be inserted after each CR sent to the remote system. The default is OFF.</p>
TELNET	SHOW	<pre> SHOW TELNET [{ SERVER SESSIONS }] </pre>	Shows the existing telnet connections.
	TELNET	<pre> TELNET={ ipaddress hostname } </pre>	<p>Attempts to open a Telnet connection to a Telnet host at the specified IP address or with the specified name. If the command is successful then the login prompt for the remote system is displayed and the user may then log in. A maximum of 5 Telnet Client connections may be created at any one time. When the user logs off from the remote system the connection is terminated and the original prompt reappears. The Telnet session can also be terminated by pressing [Ctrl/D].</p>

Note: When using the telnet client feature, if the other device fails or locks out, the tenet session to the original device is locked out as well. Use the Ctrl-D key sequence to forcibly close any open telnet sessions on the fMAP.

Note: Up to

4. Provisioning Network, Service, Control, and Resource Modules

4.1 Overview

Provisioning for any fMAP product means to query and control the configuration database, and involves the following:

- **Provisioning Data** - The provisioning data itself, which consists of:
 - States - These determine whether the card or port can be placed in service and if so whether it can process data.
 - Attributes - These are the characteristics of the card or port, usually to optimize the processing of data.
 - Management Configuration - These are all the settings that allow the fMAP product to communicate to management interfaces, and have been described in Section 3.

The provisioning data is stored in the CFC and can be retrieved and backed up when necessary, usually during a software upgrade, described in Section 5.

- **Persistence** - This is the ability of the provisioning data to survive changes such as a reboot of the shelf or the removal of a card.
- **Pre-provisioning** - The user has the option of creating a card and having it in the database prior to inserting the card.

4.1.1 Provisioning in Release 8.0 for fMAP Products

As explained in section 2.1, the components for fMAP products interfaces are divided into Service Modules, Network Modules, and Control Modules. In release 8.0 there is also the Resource Module (RM) slot, which

[Table 4-1](#) lists which of these components are available for each product in this Release.

Note: This table includes the ATN code where applicable to specify the card. For more information on these cards, refer to the fMAP Component Specification.

Controlling these components is done through the use of profiles, operational states, and provisioning modes, as explained below.

Note: The 7100 is provisioned using the automatic mode (PROVMODE=AUTO). Manual provisioning of the 7100 is not supported, but user-create Profiles are supported.

TABLE 4-1 Product/Component Compatibility for Commercial fMAP Products in Release 8.0

Category	Type	Component	9100 (CFC12)
Service Modules (SM) 4.2	Fast Ethernet 5.	FE10 (TN-102-A) ^a	x
		FX10FX (TN-104-A)	x
		FX10LX (TN-107-A)	x
		FX10BX (TN-109-A)	x
	ADSL - 6.	ADSL24A (TN-121-A) - Annex A ^b	x
	CES - 9.1	CES8 (TN-119-A)	x
Network Modules (NM) 4.3	GE8	GE8 (TN-117-A) - in Service Module slot ^{c d}	x
	Gigabit Ethernet - 5.	GE4 (CFC12)	x
		GE2RJ (CFC12)	x
Control Modules (CM) 4.5	CFC - 4.5.2	GE8 (TN-117-A) - in Resource Module slot ^e	
		CFC12 (on 9100)	x

- a. The FE/FX10 card can also be used as an upstream interface.
- b. Supports Annex M.
- c. When GE1 interface is set for customer, the card supports customer features (on port basis) at 1G rate.
- d. In 9700 simplex, cannot be plugged into the unused CFC slot (8 or 12).
- e. When plugged into Resource Module (slots 6/7 in 9700 and slots 6/8 in 9400), can provide high capacity ring distribution.

4.1.2 Provisioning Mode (Manual and Automatic)

4.1.2.1 Manual Provisioning Mode (PROVMODE = MANUAL)

In this mode, commands are used to create, modify, or delete the provisioning data. The data is persistent over reboots and restarts of the fMAP system and the removal of the card. (To delete a card, the user must explicitly do so with the DESTROY CARD command.)

Important to note is that insertion of a card when in the Manual Provisioning Mode does **not** create/provision the card in the database; this must be done using the CREATE command.

4.1.2.2 Automatic Provisioning Mode (PROVMODE = AUTO)

In the AUTO mode, hardware is discovered in a slot where there is no prior provisioning and the cards and ports are automatically provisioned. This discovery occurs when:

- The card is inserted into a slot (this would not apply to a CM in a simplex system since it is in simplex mode).
- The Network or Service Module is already inserted and the following occurs:
 - The Control Module powers up
 - The Control Module reboots
- The system mode is changed from manual to automatic.

Similar to the Manual Provisioning Mode, commands are used to create, modify, or delete provisioning data, and data is persistent over reboots/restarts of the system and the removal of the card.

Note: The default mode for the fMAP Series products is Automatic Provisioning Mode (PROVMODE=AUTO), and the mode can be changed through commands

Note: Once the user has set the PROVMODE to MANUAL, the user must explicitly provision fMAP Series modules and ports using CLI commands. It is recommended that the default AUTO mode be used.

4.1.3 The AUTOPROV Profile

When the system is first initialized, the system's PROVMODE is set to AUTO, and all modules come up with the profile name AUTOPROV.

Note: Modification of a profile does not change the attributes of a card/port that has already been provisioned.

4.1.4 User Created Profiles (Starting in Release 6.0)

4.1.4.1 Overview

In release 5.0, the user could modify the AutoProv profile that was provided for each **card type** or **port type** (called a managed entity), but the AutoProv profile was the only profile that existed. In release 6.0, user-created profiles are possible, and these profiles have the following attributes:

- Profile Creation

Profiles are created with the `CREATE PROFILE` command.

Profile names must be unique within a type; they are case insensitive.

- Applying Profiles to Managed Entities

Configuration settings of a Profile are applied to managed entities when requested at the CLI as long as the Profile and entities define the same type.

If a profile is applied to a managed entity and the user manually changes an attribute of the managed entity, the managed entity keeps its reference to the Profile but indicates that it no longer matches the Profile.

If a Profile is modified, all managed entities using the Profile indicate their provisioning no longer matches the profile.

An entity must be disabled before a different Profile can be applied.

A profile controls the attributes of the entity, but not the state.

- Destroying User Profiles

Any user-created Profile can be destroyed (unlike AutoProv, which can never be destroyed).

If a Profile that has been applied to managed entities is destroyed, the managed entity has **no** Profile (this shows up as `<none>`).

- Command Changes for Profile Names.

One of the main changes to existing commands is that the `SHOW PROFILE NAMES` command has changed to `SHOW PROFILE=name` for card and port types.

The set of commands used to create, change, and destroy profiles is not that large; however, for each type there are different parameters since each type has different attributes. The command subsection (refer to [4.1.4.7](#)) will have only the high level commands and a description of what they do; for a complete list of Profile commands, refer to the *fMAP Command Manual*.

4.1.4.2 System Interactions

Any profiles that are created, changed, or destroyed, are persistent; this means the following:

- Reboot - Any changes made to profiles survive a system reboot
- Redundancy - The profile settings are mirrored in both CFCs, and so survive an activity switch.
- Upgrade - The profile settings survive over an upgrade (not relevant for release 6.0).

Note: Although the user can SET a Profile, this does not mean that profile can be applied successfully. General checks are done on the profile, but some checks cannot be done until the user tries to apply the profile to an entity.

4.1.4.3 Example

Note: In the outputs below, the response may be abbreviated if it does not add to the concepts being explained. Removed output is shown with an extended dotted line (.....)

To show how profiles work, first input a `SHOW CARD` to show what cards are in the device:

```
officer SEC>> SHOW CARD
--- Card Information ---
    Prov
```


Slot	Card Type	State	Faults
0	ADSL24A	UP-UP-Online	-
1	POTS24	DN-DN-Offline	-
2	ADSL24	UP-UP-Online	Minor
3	-	-	-

Then SHOW PROFILE NAMES lists the currently used profiles. Since no profiles have been created, all card and port types have the profile AutoProv.

officer SEC>> show profile names
 --- Card Profiles ---

Name	Type
AutoProv	ADSL8S
AutoProv	ADSL16
AutoProv	CFC24
AutoProv	GE3
AutoProv	FE10
AutoProv	FX10
AutoProv	ADSL24
AutoProv	ADSL16B
AutoProv	POTS24
AutoProv	SHDSL16
AutoProv	CES8
AutoProv	ADSL24A
AutoProv	ADSL24B
AutoProv	ADSL16C
AutoProv	GE8

--- Port Profiles ---

Name	Type
AutoProv	ADSLPORT
AutoProv	GEPOR
AutoProv	FEPOR
AutoProv	FXPOR
AutoProv	POTSPOR
AutoProv	SHDSLPORT
AutoProv	DSLPORT
AutoProv	E1POR

SHOW the specific attributes for the AutoProv profile for the adslport, as follows:

officer SEC>> show prof autoprov adslport

--- ADSL Port Profiles ---

```
Name..... AutoProv
Type..... ADSLPORT
Initial Admin State..... Up
Mode..... Auto2+
Bitmap Mode..... DBM
Line Type..... Interleave
Interleave Delay..... 32 msec
Target SNR Margin..... 8 dB
Echo Cancellation..... Off
Line Quality Monitor..... Medium
VPI..... 0
```

```
VCI..... 35
ATU-C (Upstream)
  Maximum Rate..... 1024 kbps
  Minimum Rate..... 32 kbps
ATU-R (Downstream)
  Maximum Rate..... 26624 kbps
  Minimum Rate..... 32 kbps
```

SHOW the state of the interfaces on CARD 2.

```
officer SEC>>SHOW INTERFACE CARD=2
```

```
--- ADSL Interfaces ---
```

Interface	State	Connection	Upstream	Downstream	Mode
2.0	UP-DN	Handshake	- /1024	- /26624	T1.413/Auto2+
2.1	UP-DN	Handshake	- /1024	- /26624	T1.413/Auto2+
2.2	UP-DN	Handshake	- /1024	- /26624	T1.413/Auto2+
2.3	UP-DN	Handshake	- /1024	- /26624	T1.413/Auto2+
2.4	UP-DN	Handshake	- /1024	- /26624	T1.413/Auto2+
2.5	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.6	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.7	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.8	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.9	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.10	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.11	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.12	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.13	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.14	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.15	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.16	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.17	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.18	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.19	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.20	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.21	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.22	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.23	DN-DN	No RX/TX	- /1024	- /26624	- /Auto2+

Focus on the interface ADSL:2.0 and note the attributes

```
officer SEC>> SHOW INTERFACE=2.0
```

```
--- ADSL Interfaces ---
```

```
Interface..... 2.0
Type..... ADSL
State..... UP-DN-Failed
Description..... <none>
```

```
Interface Faults
```

```
  Loss of Signal..... Info (masked)
  Loss of Frame..... Info (masked)
  Loss of Link..... Info (masked)
  No Peer Present..... Info
```

```
Provisioning
```

```
  Provisioning Profile..... AutoProv
  Mode..... Auto2+
  Line Type..... Interleave
```

```

Interleave Delay..... 32 msec
Target SNR Margin..... 8 dB
Echo Cancellation..... Off
Line Quality Monitor..... Medium
.....
Untagged VLAN..... 1

```

Now create a Profile, **gold**, where all attributes are the same except the Target SNR Margin. Do a SHOW of the interface and note that the interface attributes have not changed, since there is no association between the created profile and the specific interface. However, a SHOW PROFILE NAMES does show that the Profile gold does exist for the ADSL interface.

```

officer SEC>> CREATE PROFILE=gold ADSLPORT targetsnr=7
Info (033561): Successfully created profile(s) gold
officer SEC>> sh int 2.0

```

```

--- ADSL Interfaces ---

```

```

Interface..... 2.0
Type..... ADSL
State..... UP-DN-Failed
Description..... <none>

```

```

Interface Faults

```

```

Loss of Signal..... Info (masked)
Loss of Frame..... Info (masked)
Loss of Link..... Info (masked)
No Peer Present..... Info

```

```

Provisioning

```

```

Provisioning Profile..... AutoProv
Mode..... Auto2+
Line Type..... Interleave
Interleave Delay..... 32 msec
Target SNR Margin..... 8 dB
Echo Cancellation..... Off
Line Quality Monitor..... Medium
.....
Untagged VLAN..... 1

```

```

officer SEC>> SHOW PROFILE NAMES

```

```

--- Card Profiles ---

```

Name	Type
AutoProv	ADSL8S
AutoProv	ADSL16
AutoProv	CFC24
AutoProv	GE3
AutoProv	FE10
AutoProv	FX10
AutoProv	ADSL24
AutoProv	ADSL16B
AutoProv	POTS24
AutoProv	SHDSL16
AutoProv	CES8

```
AutoProv          ADSL24A
AutoProv          ADSL24B
AutoProv          ADSL16C
AutoProv          GE8
```

--- Port Profiles ---

Name	Type
AutoProv	ADSLPORT
gold	ADSLPORT
AutoProv	GEPOR
AutoProv	FEPOR
AutoProv	FXPOR
AutoProv	POTSPOR
AutoProv	SHDSLPORT
AutoProv	DSLPORT
AutoProv	E1PORT

officer SEC>> show prof gold adslport

--- ADSL Port Profiles ---

```
Name..... gold
Type..... ADSLPORT
Initial Admin State..... Up
Mode..... Auto2+
Bitmap Mode..... DBM
Line Type..... Interleave
Interleave Delay..... 32 msec
Target SNR Margin..... 7 dB
Echo Cancellation..... Off
Line Quality Monitor..... Low
VPI..... 0
VCI..... 35
ATU-C (Upstream)
  Maximum Rate..... 1024 kbps
  Minimum Rate..... 32 kbps
ATU-R (Downstream)
  Maximum Rate..... 26624 kbps
  Minimum Rate..... 32 kbps
```

To associate the interface with the profile, disable the interface, then SET the interface with the Profile. A SHOW reveals that the interface now has the attributes of the profile.

```
officer SEC>> DISABLE INTERFACE2.0 FORCE
Info (039512): Operation Successful (ADSL24 Slot 02 Port 00)
officer SEC>> set int 2.0 prof gold
Info (033721): BITMAPMODE is ignored for these interface(s) 2.0
Info (020186): Successfully modified interface(s) 2.0
officer SEC>> sh int 2.0
```

--- ADSL Interfaces ---

```
Interface..... 2.0
Type..... ADSL
State..... DN-DN-Offline
Description..... <none>
```

```

Provisioning
  Provisioning Profile..... gold
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 7 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low
.....
  Untagged VLAN..... 1

```

Now change the profile so that the target SNR is no longer 7 (in this case, to 6). There is now a Profile mismatch, and this is shown by the (*) next to the profile name.

```

officer SEC>> set PROFILE=gold ADSLPORT targetsnr=6
Info (033560): Successfully modified profile(s) gold

```

```

officer SEC>> sh prof gold adslport

```

```

--- ADSL Port Profiles ---

Name..... gold
Type..... ADSLPORT
Initial Admin State..... Up
Mode..... Auto2+
Bitmap Mode..... DBM
Line Type..... Interleave
Interleave Delay..... 32 msec
Target SNR Margin..... 6 dB
Echo Cancellation..... Off
Line Quality Monitor..... Low
VPI..... 0
VCI..... 35
ATU-C (Upstream)
  Maximum Rate..... 1024 kbps
  Minimum Rate..... 32 kbps
ATU-R (Downstream)
  Maximum Rate..... 26624 kbps
  Minimum Rate..... 32 kbps

```

```

officer SEC>> SHOW INTERFACE=2.0

```

```

--- ADSL Interfaces ---

Interface..... 2.0
Type..... ADSL
State..... DN-DN-Offline
Description..... <none>

Provisioning
  Provisioning Profile..... gold (*)
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 7 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low

```

```
.....
      Untagged VLAN..... 1
```

To resolve the mismatch, the user can either change the Profile to match the changed value(s) on the interface, or change the interface attributes to match the Profile. Here the interface is changed to match the Profile, and the result is there is no longer a mismatch and no longer a (*).

```
officer SEC>> SET INTERFACE=2.0 adsl targetsnr=6
Info (020186): Successfully modified interface(s) 2.0
officer SEC>>
officer SEC>> sh int 2.0
```

```
--- ADSL Interfaces ---

Interface..... 2.0
Type..... ADSL
State..... DN-DN-Offline
Description..... <none>

Provisioning
  Provisioning Profile..... gold
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 6 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low
.....
      Untagged VLAN..... 1
```

Next the interface is changed to create a mismatch with the Profile, and a (*) is next to the profile.

```
officer SEC>> SET INTERFACE=2.0 adsl targetsnr=7
Info (020186): Successfully modified interface(s) 2.0
officer SEC>> sh int 2.0
```

```
--- ADSL Interfaces ---

Interface..... 2.0
Type..... ADSL
State..... DN-DN-Offline
Description..... <none>

Provisioning
  Provisioning Profile..... gold (*)
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 7 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low
.....
      Untagged VLAN..... 1
```

To resolve the mismatch, the profile is changes so the attributes match.

```
officer SEC>> set PROFILE=gold adslport targetsnr=7
Info (033560): Successfully modified profile(s) gold
officer SEC>> sh prof gold adslport
```

```
--- ADSL Port Profiles ---
```

```

Name..... gold
Type..... ADSLPORT
Initial Admin State..... Up
Mode..... Auto2+
Bitmap Mode..... DBM
Line Type..... Interleave
Interleave Delay..... 32 msec
Target SNR Margin..... 7 dB
Echo Cancellation..... Off
Line Quality Monitor..... Low
VPI..... 0
VCI..... 35
ATU-C (Upstream)
  Maximum Rate..... 1024 kbps
  Minimum Rate..... 32 kbps
ATU-R (Downstream)
  Maximum Rate..... 26624 kbps
  Minimum Rate..... 32 kbps

```

officer SEC>> SHOW INTERFACE=2.0

```

--- ADSL Interfaces ---

Interface..... 2.0
Type..... ADSL
State..... DN-DN-Offline
Description..... <none>

Provisioning
  Provisioning Profile..... gold
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 7 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low
.....
  Untagged VLAN..... 1

```

4.1.4.4 Example - Destroying a Profile

If desired, a Profile can be destroyed, even while associated with an entity. Continuing the example, the Profile gold is destroyed. It is therefore removed from the list of profile names. Finally, the interface that was associated with the destroyed Profile now has **<none>** for the profile.

```

officer SEC>> DESTROY PROFILE=gold adslport
Info (033571): Successfully destroyed profile(s) gold
officer SEC>> SHOW PROFILE NAMES
--- Card Profiles ---

```

Name	Type
AutoProv	ADSL8S
AutoProv	ADSL16
AutoProv	CFC24
AutoProv	GE3
AutoProv	FE10
AutoProv	FX10
AutoProv	ADSL24

```

AutoProv          ADSL16B
AutoProv          POTS24
AutoProv          SHDSL16
AutoProv          CES8
AutoProv          ADSL24A
AutoProv          ADSL24B
AutoProv          ADSL16C
AutoProv          GE8

```

```
--- Port Profiles ---
```

Name	Type
AutoProv	ADSLPORT
AutoProv	GEPOR
AutoProv	FEPOR
AutoProv	FXPOR
AutoProv	POTSPOR
AutoProv	SHDSLPORT
AutoProv	DSLPORT
AutoProv	EIPOR

```
officer SEC>> SHOW INTERFACE=2.0
```

```
--- ADSL Interfaces ---
```

```

Interface..... 2.0
Type..... ADSL
State..... DN-DN-Offline
Description..... <none>

```

```
Provisioning
```

```

  Provisioning Profile..... <none>
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 7 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low

```

```
.....
  Untagged VLAN..... 1
```

4.1.4.5 Example - Creating an Entity

When creating a card, the user has the option to CREATE the card and associate it with a Profile. If the user does not include a Profile (that exists), the card and its associated entities (interfaces) will have <none> as the associated Profile.

```

officer SEC>> CREATE CARD=3 adsl24
Info (038012): Successfully created card(s) 3
officer SEC>> SHOW INTERFACE=3.0

```

```
--- ADSL Interfaces ---
```

```

Interface..... 3.0
Type..... ADSL
State..... UP-DN-Dependency
Description..... <none>

```

```
Provisioning
```

```

  Provisioning Profile..... <none>
  Mode..... Auto2+

```



```

Line Type..... Interleave
Interleave Delay..... 32 msec
Target SNR Margin..... 8 dB
Echo Cancellation..... Off
.....
Untagged VLAN..... 1

```

The user could now create a new Profile (silver) with changed attributes, and SET the interface (3.0) to this Profile. As a result, the interface will be associated with the silver profile and have its attributes.

```

officer SEC>> CREATE PROFILE=silver adslport targetsnr=9
Info (033561): Successfully created profile(s) silver
officer SEC>> sh profile names

```

```

--- Card Profiles ---

```

Name	Type
AutoProv	ADSL8S
AutoProv	ADSL16
AutoProv	CFC24
AutoProv	GE3
AutoProv	FE10
AutoProv	FX10
AutoProv	ADSL24
AutoProv	ADSL16B
AutoProv	POTS24
AutoProv	SHDSL16
AutoProv	CES8
AutoProv	ADSL24A
AutoProv	ADSL24B
AutoProv	ADSL16C
AutoProv	GE8

```

--- Port Profiles ---

```

Name	Type
AutoProv	ADSLPORT
silver	ADSLPORT
AutoProv	GEPOR
AutoProv	FEPOR
AutoProv	FXPOR
AutoProv	POTSPOR
AutoProv	SHDSLPORT
AutoProv	DSLPORT
AutoProv	E1PORT

```

officer SEC>> DISABLE INTERFACE=3.0 FORCE
Info (039512): Operation Successful (ADSL24 Slot 03 Port 00)
officer SEC>> set int 3.0 prof silver
Info (033721): BITMAPMODE is ignored for these interface(s) 3.0
Info (020186): Successfully modified interface(s) 3.0
officer SEC>> SHOW INTERFACE=3.0

```

```

--- ADSL Interfaces ---

```

```

Interface..... 3.0
Type..... ADSL
State..... DN-DN-Dependency
Description..... <none>

```

```

Provisioning
  Provisioning Profile..... silver
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 9 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low
.....
  Untagged VLAN..... 1

```

4.1.4.6 Setting an Interface to No Profile

If the user wishes to disassociate an entity with any Profile, two double quotes are used. Whatever profile the entity was associated with is dropped and the entity has <none> for a Profile association. (In the previous example, a <none> was done by creating a Profile-Entity mismatch.)

```
officer SEC>>set int 3.0 profile ""
```

```
officer SEC>> sh int 3.0
```

```

--- ADSL Interfaces ---

Interface..... 3.0
Type..... ADSL
State..... DN-DN-Dependency
Description..... <none>

Provisioning
  Provisioning Profile..... <none>
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 9 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Low
.....
  Untagged VLAN..... 1

```

4.1.4.7 Command Summary

Note: This table only includes generic commands for card and interface types. For attributes for specific card and interface types, refer to subsequent sections.

TABLE 4-2 User-Created Profile Commands

Object / Key Word(s)	Verb	Syntax	Description
PROFILE	SHOW	<pre>SHOW PROFILE [={ name-list NAMES ALL }] [{ card or port type }] [FULL]</pre>	Displays a summary of profiles including their card/interface types. If only a single profile is given or the FULL option is supplied, the details of the profile are displayed.
PROFILE card or port type	CREATE	<pre>CREATE PROFILE=name card or port type [parameter={value}]</pre>	Creates a profile with the specified name. Included are the type and type specific attributes for the profile.
PROFILE	SET	<pre>SET PROFILE=name <card_type> [PREFLOAD=filename] [ADMINSTATE={ UP DOWN }] or SET PROFILE=name <port_type> [attribute=value]</pre>	Changes the attributes for the Profile. For most card types, the profile attributes are the PREFLOAD and ADMINSTATE. For port types, these depend on the type of card. Refer to the other sections of this Guide for these attributes.
INTERFACE PROFILE	SET	<pre>SET INTERFACE={ type: type:id-range id-range ifname-list ALL } PROFILE=name</pre>	Applied the Profile to the interface. If double quotes are used after PROFILE, the interface has no Profile.

TABLE 4-2 User-Created Profile Commands

Object / Key Word(s)	Verb	Syntax	Description
CARD	CREATE	<pre>CREATE CARD=slot cardtype [[[ADMINSTATE={ UP DOWN }] PROFILE=name]]</pre>	<p>Creates a card in which the AutoProv, user-created, or no Profile is applied.</p> <p>This command syntax has not changed except for the use of user-created Profiles</p>
PROFILE	DESTROY	<pre>DESTROY PROFILE=name card or port type</pre>	<p>Destroys the user-created Profile.</p> <p>Any managed entity that had a Profile applied is set to <none> (no profile association).</p>

4.1.5 Administrative and Operational States

Administrative and Operational States determine whether the card or port is available for service and, if available for service, whether it is being provided:

- The **Administrative State** is controlled by the user and can be set to either UP (available for service) or DOWN (Not available for service). Control of this state is through the ENABLE/DISABLE command.
- The **Operational State** is either UP (providing service) or DOWN (not providing service). This state is not user controllable but does depend on the Administrative State:
 - If the Administrative State of a card is UP, the Operational State will be UP if the card/port can provide service.
 - If the Administrative State is DOWN, the Operational State will always be DOWN.

Note: The exception to these rules is the FM7 fan module. This module cannot be disabled, but the alarms can be masked while the fan module is removed.

4.1.6 Software Loads

When creating or changing the attributes of a card, the type of load must be considered, especially when doing a software upgrade (covered in detail in Section 11.).

4.1.6.1 Card Load Preferences

Once a software load is present in the control module FLASH file system, it can be designated as the target software load for one or more cards using the parameters of the SET CARD command. One or more of the following types of designations can be set for a card:

- Preferred - selected using the **PREFLOAD** parameter. A load designated as **PREFLOAD** indicates that this is the primary load that the specified card will load from. For system integrity reasons, load files designated as **PREFLOAD** cannot be renamed or deleted. Any changes made in load designations, for a system configured for duplex operation, while the system is operating in sync, will be reflected on both the ACTCFC and INACTCFC.
- Alternate - (**Control Module Only**) selected using the **ALTLOAD** parameter. A load designated as **ALTLOAD** indicates that this is the alternate load that the specified CM will load from. The **ALTLOAD** is used when a redundant copy of the preferred load file is made on the CM FLASH file system; it specifies an alternate load preference for the redundant file. Establishing an alternate load provides a backup in the unlikely event that the preferred load file cannot boot. For a duplex system configuration, any changes made in the **ALTLOAD** designation apply to both the active (ACTCFC) and inactive (INACTCFC) control modules.

Note: This parameter is not supported for the service modules because the copy of the service module load stored on the control module FLASH file system is the alternate by default (the preferred is the copy located in the service module flash memory).

- Temporary - selected using the **TEMPLOAD** parameter. A load designated as **TEMPLOAD** indicates that this is the load that the specified card will load from, *one time*, during the *next* loading process. The **TEMPLOAD** designation is used during the software upgrade procedure. A load designated as **TEMPLOAD** indicates this is the load that the specified card will load from, one time, during the next loading process.

TEMPLOAD designation results in two things. First, if for any reason the new load file is unusable, the system will erase the designation of **TEMPLOAD** for the new file and revert back to using its original load, allowing the system to automatically recover from an initialization failure of the **TEMPLOAD**. Second, setting a load as **TEMPLOAD** puts the configuration into the upgrade mode. For upgrade purposes, changes made to the designation of temporary are independent of system synchronization status.

Note that load preferences for the CM(s) are stored in the non-volatile RAM (NVRAM) of each module, while load preferences for the SMs are stored in the configuration database, explained in [10.4](#).

Note: An Inconsistent Load Minor alarm will be posted against any service module whose running major and minor software load version does not match the preferred major and minor software load version of the active control module. The alarm is raised whenever the CM is taken out of upgrade mode. This is intended to maintain consistency of load versions throughout the system.

4.1.7 Provisioning Data at Startup

When the system is first brought up, it is configured as follows:

- The Provisioning Mode is set to AUTO (PROVMODE=AUTO)
- All modules and ports use the AUTOPROV profile
- The AUTOPROV profile is set to the factory defaults.

- The Administrative State of all modules and ports is UP, and the Operational State is set to UP if the module/port can process data.

4.1.8 Provisioning Mode (SHOW SYSTEM PROVMODE)

Use this command to view whether the Provisioning Mode is MANUAL or AUTO.

```
officer SEC> SHOW SYSTEM PROVMODE
System is in AUTO provisioning mode
```

4.2 SM Category Attributes

4.2.1 SM Card

The attributes for a SM are shown in the display for the **SHOW CARD <slot number>** command.

Following is the output for the command. [Table 4-3](#) describes the attributes and states that are common for ADSL cards

```
officer SEC>> SHOW CARD=2

--- Card Information ---

Slot..... 2
Type..... ADSL24
State..... DN-DN-Offline
Provisioning Profile..... AutoProv (*)

Hardware
  Model Number (Revision)..... TN-112-A (Rev X8)
  Serial Number..... ATNLAB4030200740
  CLEI Code..... <none>

Software
  Running Load..... 6.0.0.GAMMA.20050306
  Preferred Load..... ads124_6.0.0.GAMMA.20050306.tar
  Temporary Load..... <none>
```

TABLE 4-3 Common SM Card Attributes - Defaults are in Bold)

Card Attribute	Values / Range	Description
Slot	Slot Number	The slot number occupied by the card
Type	Depends on the card type	The type of card

TABLE 4-3 Common SM Card Attributes - Defaults are in Bold> (Continued)

Card Attribute	Values / Range	Description
State	Three attributes: - Admin State - Operational State - Status	<p>These three attributes determine the state of the card; whether it is capable of carrying traffic and the status (Implied Operational Status)</p> <p>ONLINE - Card is configured and can provide service. (UP)</p> <p>DEGRADED - There is a fault but the card can still provide service (UP)</p> <p>OFFLINE - The normal status when a card is in a DOWN state. The card requires a routine operation to place it ONLINE and available for service. (DOWN)</p> <p>FAILED - The card has detected a hardware or software fault that makes it unable to provide service. (DOWN)</p> <p>NOT INSTALLED - Card is provisioned in software (CREATE) but not physically present (DOWN)</p> <p>RESET - transient state as card resets (DOWN)</p> <p>LOADING - The software load is being transferred from the CFC to the flash memory in the card. (DOWN)</p> <p><i>Note: In release 8.0, a percentage number for loading is added. Once at 100%, there may still be a delay so that the transfer of software to the card is complete.</i></p> <p>BOOTING - The software load is being copied from the flash memory into its RAM memory. (DOWN)</p> <p>IN TEST - Card is running diagnostics (DOWN)</p> <p>CONFIGURING - Provisioning data for the card is being copied from the CFC to the RAM memory on the card. (DOWN)</p> <p>TERMINATING - The card is performing an operation in preparing to go out of service. (UP or DOWN)</p>
Provisioning Profile	Profile that has been applied to the card and if there is a Profile mismatch.	If there is a status mismatch, a (*) appears next to the Profile Name. Refer to 4.1.4 .

TABLE 4-3 Common SM Card Attributes - Defaults are in Bold) (Continued)

Card Attribute	Values / Range	Description
Hardware	Model Number	The TN number for card type
	Serial Number	The unique serial number for the card
	CLEI Code	The CLEI code, if the card has one.
Preferred SW Load	Running	Refer to 4.1.6 .
	Preferred SW Load	
	Temporary SW Load	

4.2.2 SM Interface

The attributes for a SM port are shown in the display for the **SHOW INTERFACE <interface>** command.

The following output shows samples for the ADSL and FE interfaces.

Note: For an overview of interfaces, refer to [4.4](#). For an explanation of parameters and values for specific interfaces, refer to the other Sections.

```
officer SEC>> SHOW INTERFACE ADSL:2.0

--- ADSL Interfaces ---

Interface..... 2.0
Type..... ADSL
State..... UP-DN-Dependency
Description..... <none>

Provisioning
  Provisioning Profile..... AutoProv
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 8 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Medium
  ATU-C (Upstream)
    Maximum Rate..... 1024 kbps
    Minimum Rate..... 32 kbps
  ATU-R (Downstream)
    Maximum Rate..... 26624 kbps
    Minimum Rate..... 32 kbps
  Performance Monitoring..... Off
  Remote Monitoring..... Off

Actual
  Connection State..... No RX/TX
  Direction..... Customer
  Physical Address..... 00:0C:25:00:0A:D4

VC Information
  VC Identifier..... 0
```



```

VPI..... 0
VCI..... 35
Service Category..... UBR
TX Peak Cell Rate..... Maximum cps

VLAN Information
VC Identifier..... 0
Acceptable Frame Types..... All
Ingress Filtering..... On
TPID..... 0x8100
TAGALL..... Off
Dynamic MAC Learning Limit.... Off
Untagged VLAN..... 105

```

officer SEC>> SHOW INTERFACE ETH:4.0

--- FE Interfaces ---

```

Interface..... 4.0
Type..... FE
State..... UP-DN-Dependency
Description..... <none>

```

Provisioning

```

Provisioning Profile..... AutoProv
Direction..... Customer
Auto Negotiation..... On
Speed..... Auto
Duplex..... Auto
Flow Control..... Auto
Remote Monitoring..... Off

```

Actual

```

Direction..... Customer
Physical Address..... <none>

```

VLAN Information

```

Acceptable Frame Types..... All
Ingress Filtering..... On
TPID..... 0x8100
TAGALL..... Off
Dynamic MAC Learning Limit.... Off
Untagged VLAN..... 1

```

Table 4-4 describes these attributes and states

Note: Only the common attributes for interfaces are listed here. For interface specific attributes, refer to the appropriate section.

TABLE 4-4 Common SM Interface Attributes - Default is in Bold

SM Port Attribute	Values / Range	Description
Interface	Number of the interface	The identifying number of the interface.
Type	Interface Type	The type of interface, such as FE or ADSL
State	Three attributes: - Admin State - Operational State - Status	<p>These three attributes determine the state of the card; whether it is capable of carrying traffic and the status (Implied Operational Status)</p> <p>ONLINE - Port is configured and can provide service. (UP)</p> <p>DEGRADED - There is a fault but the port can still provide service (UP)</p> <p>OFFLINE - The normal status when a port is in a DOWN state. The card requires a routine operation to place it ONLINE and available for service. (DOWN)</p> <p>FAILED - The port has detected a hardware or software fault that makes it unable to provide service. (DOWN)</p> <p>DEPENDENCY - The port cannot provide service because the card on which it depends is unavailable. (DOWN)</p> <p><i>Note: This status is important to understand when reviewing alarms. Refer to 18.3.</i></p> <p>CONFIGURING - Provisioning data for the port is being copied from the CFC to the RAM memory on the card. (DOWN)</p> <p>TERMINATING - The port is performing an operation in preparing to go out of service. (UP or DOWN)</p>
Description	Text	This is an attribute for a user-created description.
Provisioning	Profile that has been applied to the card and if there is a Profile mismatch.	<p>If there is a status mismatch, a (*) appears next to the Profile Name. Refer to 4.1.4.</p> <p>Other attributes are determined by the Interface Type. Refer to the appropriate Section.</p>
Actual	Attributes measured when the interface is operationally UP.	Other attributes are determined by the Interface Type. Refer to the appropriate Section.

TABLE 4-4 Common SM Interface Attributes - Default is in Bold (Continued)

SM Port Attribute	Values / Range	Description
VC Information	Virtual Channel Attributes	This applies to ADSL card types. Refer to 12.4 .
VLAN Information	VLAN attributes	The attributes for the VLAN over the interface.

4.3 NM Category Attributes

4.3.1 NM Card

The only common attribute for NM cards is the ADMIN STATE (UP or DOWN).

Note: If the user sets the Administrative State of the NM card to DOWN (using the DISABLE command with the FORCE option) and there is only one NM provisioned, upstream data service is lost for the fMAP product.

4.3.2 NM Interface

Following are the outputs for the Network Interfaces. (For details on the DIRECTION attribute, refer to [14.2](#).)

```
officer SEC>> SHOW INTERFACE DIRECTION=NETWORK
```

```
--- GE Interfaces ---
```

```
Interface  State Autonegotiate Flow Control
-----
10.0      DN-DN On           Off
10.1      UP-DN On           Off
10.2      UP-DN On           Off
11.0      UP-DN On           Off
11.1      UP-DN On           Off
11.2      UP-DN On           Off
```

```
--- General Interfaces ---
```

```
Interface      State Name
-----
ETH:0          UP-UP MGMT
VLAN:402.0    UP-UP <none>
```

```
officer SEC>> SHOW INTERFACE ETH:10.0
```

```
--- GE Interfaces ---
```

```
Interface..... 10.0
Type..... GE
State..... DN-DN-Offline
Description..... <none>
```

```

Provisioning
  Provisioning Profile..... AutoProv
  Auto Negotiation..... On
  Flow Control..... Off
  Remote Monitoring..... Off

Actual
  Direction..... Network
  Physical Address..... 00:0C:25:00:06:1A
  Port Speed..... 1 Gbps

VLAN Information
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Untagged VLAN..... 1
  Tagged VLAN(s)..... 100, 105, 200, 205, 300, 305, 400, 402,
                        405, 500, 505, 512, 600, 605, 700, 705,
                        800, 805, 900, 905, 1000, 1005, 1100,
                        1105, 1200, 1205, 1300, 1305, 1400, 1405,
                        1500, 1505, 1600, 1605
    
```

Table 4-5 lists the Port attributes for the NM port.

TABLE 4-5 NM Port Attributes (Default value is in Bold)

NM Port Attribute	Values / Range	Description
Interface	Number of the interface	The identifying number of the interface.
Type	Interface Type	The type of interface, such as ETH.
State	Three attributes: - Admin State - Operational State - Status	These three attributes determine the state of the card; whether it is capable of carrying traffic and the status. For status refer to Table 4-4 .
Description	Text	This is an attribute for a user-created description.
Provisioning	Provisioning Profile	The Profile that is being applied to this interface. If there is a Profile mismatch a (*) appears.

TABLE 4-5 NM Port Attributes (Default value is in Bold)

NM Port Attribute	Values / Range	Description
	Auto Negotiation	<p>Specifies whether automatic negotiation of transmission parameters for the ports is allowed.</p> <p>If ON, the port has increased flexibility to communicate with the remote peer. The port has the ability to advertise flow control and to provide single direction fault coverage. The port will drive the link state up and down based on the ability to communicate with the remote peer, triggering on both transmit and receive failures Loss of Signal (LOS).</p> <p>If OFF, the port state is driven by receive failure (LOS). Flow control is still provided as long as the FLOWCONTROL parameter is ON.</p> <p>The default value is ON.</p>
	Flow Control	<p>Specifies whether flow control is enabled.</p> <p>If ON, the port can generate and respond to pause signals with the remote peer.</p> <p>If OFF, pause is ignored and not generated, and potential for packet loss is increased.</p> <p>FLOWCONTROL is independent of AUTO-NEGOTIATION.</p> <p>The default value is OFF.</p>
	Remote Monitoring	
Actual	Direction	CUSTOMER or NETWORK (14.2)
	Physical Address	
	Port Speed	Port speed (such as 1 Gigabit)
VLAN Information	VLAN attributes	Includes a listing of the untagged VLAN and all the tagged VLANs associated with this interface.

4.4 Provisioning Interfaces

4.4.1 Overview

An **interface** is a capability associated with a physical port. The interface, therefore, provides a logical representation of one or many physical ports. A specific instance of an interface has an identifier which can be used when configuring these capabilities.

In earlier releases, the relationship between interfaces and physical ports was one-to-one. However, the relationship of interfaces to ports will evolve towards a many-to-many relationship. This means that one port can have more than one interface type and an interface type can use more than one port.

4.4.2 Interface Provisioning

For the fMAP series products, there are several interface types:

- Ethernet - with each instance having a Type of ETH, an ID of the port number, and no name. The management interface has an interface of ETH:0 and has a category of General.
- LAG - which occurs when a LAG is created, and has a Type of LAG, an ID of 0.0 (first), and associated ports that depend on the ports that have been chosen to be part of the LAG group.
- VLAN - support of the INBAND interface, as well as the CES8 IP interfaces
- ADSL - support of the ADSL interface.
- DS1 - support of DS1 interfaces for circuit emulation services (CES) and PPP / MLPPP services.
- E1 - support of E1 interfaces for circuit emulation services (CES) and PPP / MLPPP services.
- PSPAN - support of Pseudo-span interfaces for circuit emulation services (CES)
- PPP - This is the protocol used for Ethernet packet transport over the DS1/E1 interface.
- EPON - support of the EPON card

Following are some examples of setting up interfaces across a set of ports.

```
officer SEC>> show interface ETH: 11.0 FULL

--- GE Interfaces ---

Interface..... 11.0
Type..... GE
State..... UP-UP-Online
Description..... <none>

Provisioning
  Provisioning Profile..... AutoProv
  Auto Negotiation..... On
  Flow Control..... Auto
  Remote Monitoring..... On

Actual
  Direction..... Network
  Physical Address..... 00:0C:25:00:05:F9
  Port Speed..... 1 Gbps

VLAN Information
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Untagged VLAN..... 1

Packet Statistics

```

	Input	Output
Octets.....	13404736	1772822
Unicast Packets.....	0	153

Discarded Packets.....	110198	0
Errored Packets.....	0	0
Unknown Proto Packets.....	110198	N/A

Command examples using interfaces are included in this Guide; those commands that can use a Interface or Port will use only Interface, since Port is deprecated.

4.4.3 Displaying ADSL Interface information

The fMAP system provides the user with the ability to display the ATU-C (Near End) and ATU-R (Far End) ADSL values for a subscriber. The values returned by the command are expressed in dBm. This functionality is implemented as part of the SHOW PORT command. Values will change depending on the increase and decrease in loop length. Values vary from 7 dBm or so (full power cutback) to 20 dBm (no cutback). Following is a sample display of a port with the ATU-C and ATU-R fields highlighted.

```
officer SEC>> SHOW INTERFACE ADSL:2.0

--- ADSL Interfaces ---

Interface..... 2.0
Type..... ADSL
State..... UP-DN-Dependency
Description..... <none>

Provisioning
  Provisioning Profile..... AutoProv
  Mode..... Auto2+
  Line Type..... Interleave
  Interleave Delay..... 32 msec
  Target SNR Margin..... 8 dB
  Echo Cancellation..... Off
  Line Quality Monitor..... Medium
  ATU-C (Upstream)
    Maximum Rate..... 1024 kbps
    Minimum Rate..... 32 kbps
  ATU-R (Downstream)
    Maximum Rate..... 26624 kbps
    Minimum Rate..... 32 kbps
  Performance Monitoring..... Off
  Remote Monitoring..... Off

Actual
  Connection State..... No RX/TX
  Direction..... Customer
  Physical Address..... 00:0C:25:00:0A:D4

VC Information
  VC Identifier..... 0
  VPI..... 0
  VCI..... 35
  Service Category..... UBR
  TX Peak Cell Rate..... Maximum cps

VLAN Information
  VC Identifier..... 0
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Dynamic MAC Learning Limit.... Off
```

Untagged VLAN..... 105

The state of an interface can be displayed at the CLI using the SHOW INTERFACE command. [Table 4-6](#) includes the syntax (note the use of the wildcard *).

officer SEC>> SHOW INTERFACE ADSL: 2. * STATE ALL

--- ADSL Interfaces ---

Interface	State	Connection	Upstream	Downstream	Mode
2.0	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.1	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.2	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.3	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.4	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.5	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.6	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.7	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.8	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.9	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.10	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.11	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.12	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.13	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.14	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.15	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.16	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.17	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.18	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.19	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.20	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.21	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.22	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.23	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+

officer SEC>> SHOW INTERFACE STATE DOWN

--- ADSL Interfaces ---

Interface	State	Connection	Upstream	Downstream	Mode
2.0	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.1	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.2	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.3	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.4	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.22	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+
2.23	UP-DN	No RX/TX	- /1024	- /26624	- /Auto2+

--- FE Interfaces ---

Interface	State	Autonegotiate	Flow Control	Duplex	Speed	Direction
4.0	UP-DN	On	-	-	-	Customer
4.1	UP-DN	On	-	-	-	Customer
4.8	UP-DN	On	-	-	-	Customer
4.9	UP-DN	On	-	-	-	Customer
16.0	UP-DN	On	-	-	-	Customer
16.1	UP-DN	On	-	-	-	Customer
16.9	UP-DN	On	-	-	-	Customer

--- GE Interfaces ---

Interface	State	Autonegotiate	Flow Control
10.0	DN-DN	On	Off
10.1	UP-DN	On	Off
10.2	UP-DN	On	Off
11.0	UP-DN	On	Off
11.1	UP-DN	On	Off
11.2	UP-DN	On	Off

(some output omitted)

The following table lists the command set for Interfaces.

TABLE 4-6 SHOW INTERFACE command

Object	Verb	Syntax	Description
INTER-FACE	SHOW	<pre> SHOW INTERFACE [={ type: type:id-range id-range ifname-list ALL }] [CARD=slot-list] [STATE={ UP DOWN ALL }] [DIRECTION={ NETWORK CUSTOMER INTERNAL }] [FULL] </pre>	<p>Displays provisioning attributes for the interface specified.</p> <p>STATE - The STATE parameter is used to specify that only interfaces in the given state should be displayed.</p> <p>For interface DIRECTION, refer to 14.2.</p>
INTER-FACE	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } <Interface-type> <Parameters> </pre>	<p>Adds or changes the attribute for the interface. The user can identify the interface type through the INTERFACE value or the Interface-type value. These commands are the same:</p> <ul style="list-style-type: none"> - SET INTERFACE=20.0 ADSL DESC="test" - SET INTERFACE ADSL:20.0 DESC="test"

4.4.4 Interface States

The SHOW INTERFACE command provides a STATE attribute. STATE is a combination of Administrative and Operational states as illustrated in the following table.

TABLE 4-7 STATE table (Boolean AND)

Admin State	Operational State	STATE
DOWN	DOWN	DOWN
UP	DOWN	DOWN
UP	UP	UP
DOWN	UP	DOWN

4.4.5 Port Attributes no Longer Supported

Prior to release 7.0, the use of both PORT and INTERFACE was supported, but the response for a port-based command was that the command would be rejected in a future release. In 7.0, the command is rejected, with a message to use the INTERFACE parameter instead.

For example, if the user wishes to add / change a description to the slot.port 1.1, and uses the SET PORT command, the following occurs:

```
officer SEC>> SET PORT=1.1 ADSL DESCRIPTION="this is a test"
Error (010033): Command has been removed.
Use 'SET INTERFACE=1.1 ADSL DESCRIPTION="this is a test"' instead.
```

All commands that use port are treated similarly, with the response including the correct syntax that uses INTERFACE.

If the user tries to define attributes for an interface that is not a type for the specified port, there is an error message:

```
officer SEC>> SET INTERFACE=20.0 FE DESCRIPTION="test"
Error (033965): The command given does not apply to interface(s) 20.0
```

4.4.6 VC Configuration

Accompanying the concept of a VC is the concept of a subinterface. For example, a user may input a command at the CLI such as:

```
<VERB> <NOUN> INTERFACE=ETH:1.0.2 <parameter> <parameter>
```

The “2” in this command line would be considered a subinterface by the system referring to VC number 2 configured on Ethernet interface 1.0. For commands related to ADSL cards, as well as RMON, STP, and IGMP functionality, whether a subinterface is included by the user on the command line or not, the system uses the interface. For example, in the command illustrated above, *ETH:1.0.2* would translate internally to *ETH:1.0*.

4.4.7 EPON/ONU Configuration

For the EPON/ONU interfaces, the EPON interface is slot.port (such as epon:4.1). When the user creates an ONU on the EPON port, an ONU interface is created (onu:4.1.0). The ETH interface is also created (eth:4.1.0). The ONU interface is used for commands to query/control the ONU, while the ETH interface is used to query/change the attributes of the ONU services. Refer to [8.2](#).

4.5 Control Module Provisioning Data

4.5.1 Overview (Simplex versus Duplex)

For the fMAP product in **simplex** mode there is one Control Module, and it is called the active CFC: it has the only copy of the configuration database, and if the CFC restarts, service is temporarily lost.

When a fMAP product is in **duplex** mode, one of the CFCs is redundant. The system is usually brought up in duplex mode by first configuring the shelf in simplex mode and then inserting the second CFC. Because the default for the shelf is AutoProv, the newly inserted CFC is discovered and provisioned and comes up as the inactive CFC. The system is now equipped with an **active** and **inactive** CFC; referred to on the system as **ACTCFC** and **INACTCFC**, respectively, and the inactive keeps a copy of the persistent data and dynamic data, as well as the software load on the active CFC. Moreover, the inactive receives incremental updates from the active CFC, called **data synchronization**, or data sync. This ensures there is matching data (called **data mirroring**) in both CFCs.

Because of this mirroring, the inactive CFC can take over the shelf if there is a fault in the active CFC. This is called a **swap activity** or a **swap**; all persistent and transient data is retained, so the Allied Telesis duplex system can continue to process subscriber services as well as receive requests and produce outputs to the management interfaces, including the alarms associated with the swap.

The above explanation of the fMAP in duplex mode assumes that both CFCs are functioning normally prior to the degradation of the active CFC, that all data and software loads have synched, and that the inactive CFC has been successfully mirroring data up to the swap. There are situations where these conditions do not exist, and these can be created by the user (such as disabling the inactive CFC), or autonomously. These will be explained below.

The following shows the output for the `SHOW CARD ACTCFC` and `SHOW CARD INACTCFC` commands.

```
officer SEC>> SHOW CARD ACTCFC
--- Card Information ---
Slot..... 13
Type..... CFC24
State..... UP-UP-Online (Active)
Provisioning Profile..... AutoProv
```

```

Hardware
  Model Number (Revision)..... TN-401-A0 (Rev X3)
  Serial Number..... ATNLAB4030200441
  CLEI Code..... <none>

Software
  Running Load..... cfc24_6.0.0.DBH.20050315.tar
  Preferred Load..... cfc24_6.0.0.DBH.20050315.tar
  Temporary Load..... None
  Alternate Load..... None

Software Build Information
  Load File..... cfc24_6.0.0.DBH.20050315.tar
  Build Name..... 24G Central Fabric Controller
  Build Type..... Lab-Only Build
  Revision..... 6.0.0.GAMMA.20050315
  Built On..... Wed 03/16/2005 at 10:20a
  Built By..... ccadmin
  Environment..... dhays_R6.0_XSTP_EPSR_Interaction
  Baseline..... R6.0_Nightly_03_15_05.1953
  Boot ROM Build Name..... 24G Central Fabric Controller
  Boot Loader..... Boot Loader
  Boot ROM Version..... 3.0.g.1

Card Type Specific Information
  Timing Reference..... N/A

officer SEC>> SHOW CARD INACTCFC

--- Card Information ---

Slot..... 9
Type..... CFC24
State..... UP-UP-Online (Inactive)
Provisioning Profile..... AutoProv

Hardware
  Model Number (Revision)..... TN-401-A0 (Rev X3)
  Serial Number..... ATNLAB4030200447
  CLEI Code..... <none>

Software
  Running Load..... cfc24_6.0.0.DBH.20050315.tar
  Preferred Load..... cfc24_6.0.0.DBH.20050315.tar
  Temporary Load..... None
  Alternate Load..... None

Software Build Information
  Boot ROM Build Name..... 24G Central Fabric Controller
  Boot Loader..... Boot Loader
  Boot ROM Version..... 3.0.g.1

Card Type Specific Information
  Timing Reference..... N/A

```

4.5.2 CFC Card Attributes and States (SHOW CARD ACTCFC)

Table 4-8 describes the common attributes and states for the CFC card.

Note: For completeness, this table includes all attribute values. Some attributes are explained in more detail in Section 10.: **File Management** and Section 18.: **Alarms and Troubleshooting**, since they are relevant in software upgrade and troubleshooting scenarios.

TABLE 4-8 CFC Card Attributes - Defaults are in Bold

CFC Card Attribute	Values / Range	Description
Slot	Slot Number of the CFC card	The identifying slot of the CFC card. For a duplex, the slots are 9 and 13.)
Type	Interface Type	The type of interface, such as FE or ADSL
State	Admin State	Refer to 4.1.5 for an overview. Active: The Admin state must be UP so the CFC can have its Op Status as UP. Inactive: This state by default is UP, since it must be capable of having an Op Status of UP and be able to take over activity.
	Operational State	Refer to 4.1.5 for an overview. Active: The default status is UP, since the active CFC must be UP to process subscriber services. If the status were to change to DOWN, a swap (or attempted swap) would occur. Inactive: The default status is UP, since the inactive CFC must be ready to take over.
	- Status No Faults Simplex Card not Present Card Mismatch Death of non-critical task Heartbeat Failed Power Fuse Failed RTC Battery Power 18 Failed Power 25 Adj. Failed	Any fault condition on the card. No Faults means the card is error free Simplex - The shelf is configured as a duplex shelf but only one CFC is capable of providing service, and may be the result of a swap of activity. The inactive CFC is in an Admin State of UP but an Op State of DOWN. This is an alarm, and a CFC needs to be brought back into service. Card not Present - For the Inactive CFC only, the card is provisioned and enabled, but the card is not present. For a description of the other values, See the fMAP Log Manual .
Provisioning Profile	Profile that has been applied to the card and if there is a Profile mismatch.	If there is a status mismatch, a (*) appears next to the Profile Name. Refer to 4.1.4.

TABLE 4-8 CFC Card Attributes - Defaults are in Bold (Continued)

CFC Card Attribute	Values / Range	Description
Hardware	Model Number	The TN number for card type
	Serial Number	The unique serial number for the card
	CLEI Code	The CLEI code, if the card has one.
Software	Running	Refer to 4.1.6 .
	Preferred SW Load	
	Temporary SW Load	
	Alternate Load	
Software Build Information	Boot ROM Build Name	
	Boot ROM Version	

TABLE 4-8 CFC Card Attributes - Defaults are in Bold (Continued)

CFC Card Attribute	Values / Range	Description
Card Type Specific Information	Timing Reference	This applies only to the CFC6 card with the mezzanine card.
Status QUESTION ON RELATION IN OUT- PUT	<p>ONLINE</p> <p>IN TEST</p> <p>FAILED</p> <p>POWER OFF</p> <p>OFFLINE</p> <p>DEGRADED</p> <p>NOT INSTALLED</p> <p>NOT PROVISIONED</p> <p>INITIALIZING</p>	<p>The provisioning or procedural status of the card. (Implied Operation status)</p> <p>Active: The default is ONLINE - The card is controlling the shelf and providing services.</p> <p>Inactive: The default is ONLINE - The CFC6 is configured and can provide service. (UP)</p> <p>IN TEST - Inactive CFC6 Card is being tested and may be in or out of service (UP or DOWN).</p> <p>FAILED - A fault on the inactive CFC6 has made the card unable to provide service and take over the shelf if necessary (DOWN).</p> <p>POWER OFF - The inactive CFC6 is powered down or in a low power state and cannot provide service (DOWN)</p> <p>OFFLINE - The inactive status when the card is in a DOWN state (DOWN).</p> <p>DEGRADED - There is a fault but the card can still provide service (UP). If this is the status of the active CFC6 and the inactive is in an ONLINE status, a swap will occur.</p> <p>NOT INSTALLED - For the inactive CFC6 card, it is provisioned in software (CREATE) but not physically present (DOWN)</p> <p>NOT PROVISIONED - For the inactive CFC6 card, the card may be physically installed but has not been provisioned in software (CREATE).</p> <p>INITIALIZING - Card is being initialized as part of attempt to restore it to service (DOWN)</p>

4.5.3 Changing the Administrative State of the Inactive CFC

When both CFC cards are functioning properly, the inactive CFC has an Admin and Operational Status of UP and data is synched so it can take activity if necessary. However, it is possible to change the state of the inactive CFC so that normal duplex operation does not take place and the inactive CFC cannot assume activity. There are two ways to do this:

1. Disable the inactive CFC (i.e. `DISABLE CARD=12`)

The inactive CFC now cannot take over activity, and the card goes through the following changes:

- Admin State - DOWN, since the CFC is now disabled.
 - Operational State - DOWN, since the Operational State follows the Admin State.
 - STATUS - OFFLINE, since the card is disabled and is no longer synching with the active CFC.
 - The OK to Pull LED in the inactive CFC is on, since the card can no longer take over the shelf.
2. Press the SWACT button on the Inactive CFC. The same changes occur as when inputting the DISABLE command above.

Note: The Admin State and Op State of the active CFC will always be UP.

When the inactive CFC is disabled, data sync is no longer occurring between the two CFCs. To bring the inactive (and *disabled*) CFC back into service and have it receive the data stored on the active CFC so it can take over the shelf, enable the card.

The inactive CFC now reboots, and goes through the following steps before it can return to normal operation and can take over the shelf.

1. Status sequence - This is the most important attribute, and it is a sequence that shows the progress of data sync with the active CFC. During data sync, the status is *Initializing*, and the Admin State and Operational State continue to be DOWN, since the card cannot provide service yet. Once data sync is complete, the inactive card states change to normal:
2. Admin State - UP, since the card now is synched to the active CFC and is enabled.
3. Status - IN TEST, since the inactive CFC is running diagnostics on itself to ensure it has no faults and is ready to go into service
4. Operational State - UP, since the CFC will now be able to take over the shelf if it can sync with the active CFC and has no faults.

Note: Whenever the inactive CFC reboots, it assumes that the active CFC is functioning normally during the reboot process; this ensures the inactive will not come up as the newly active CFC if the currently active CFC were to fail during the reboot. Moreover, in assuming the active CFC is functioning normally during the reboot, the inactive CFC will wait for the active CFC to establish communications and begin the data sync. However, if the inactive card detects that the active card is no longer providing service (rebooting, failed, removed), the inactive card will start a 5 minute timer. If the timer expires, the inactive CFC will try to come into service as the active using its own (possibly not current) data.

4.5.4 Manual versus Automatic Swaps

Having the inactive CFC take over the system from the active can be done by system software (automatic) or by user action (manual).

- Manual Swap - There are two ways to manually perform a swap
 1. Enter the command **SWAP ACTIVITY [FORCE]**

Without the `FORCE` option, there is a prompt that a swap will occur, and system software will check that a swap can occur without loss of service. If a swap will affect service, the swap will not occur and the reason why will be displayed.

With the `FORCE` option, there is no prompt and no checks; there is only a warning that this is an uncontrolled Swap and can lead to loss of service, as shown below:

SWAP ACTIVITY FORCE

Warning (030106): UNCONTROLLED switchover in progress. No safeguards are used to guarantee sane switchover.

Note: After the command is entered, the user's terminal session will be terminated. The user should reconnect to the system.

2. Press the `SWACT` button on the Active CFC. This is the same as the `SWAP ACTIVITY`, but there will be no output if the swap could not be done.
- Automatic Swap - If there is a fault on the active CFC and the Inactive CFC can take over service, the system software will initiate a swap.

4.6 Provisioning Scenarios for Control Modules

4.6.1 Overview

The following procedures are used when changing the `fMAP` mode from simplex to duplex (one always active CFC to an active and inactive CFC) or duplex to simplex (an active and inactive CFC to only one always active CFC). Both types of procedures can be done in `AutoProv` or `Manual` mode.

Note: To minimize the possibility of loss of service, all procedures to change the CFC configuration involve inserting or removing the inactive CFC. Ensure that all of the commands used in this section apply to the currently inactive CFC. If the slots the user wishes to provision involve the CFC that is currently active, perform a Swap Activity to make it the inactive CFC.

4.6.2 Simplex to Duplex (AutoProv Mode)

When the `fMAP` system is in simplex mode, one of the double slots (8/9, 12/13) can contain one Service Module card and a face plate full (FPF). Changing the mode from simplex to duplex is done as follows:

1. If a Service Module card was installed, follow antistatic procedures and remove the Service Module card from either Slot 8 or 12.
2. Deprovision the Service Module card if it occupies a slot for the new CFC card.
3. Follow antistatic procedures and remove any Filler Plate Full card(s) in slots 8/9 or 12/13.
4. Remove the new CFC card from its antistatic container.

5. Following antistatic procedures, insert the CFC card into the available double slot. Refer to the Installation Guide for details.
6. Because the CFC cards are in AutoProv mode, the active CFC will detect the newly inserted CFC and try to provision it, synch all of its data with the inactive CFC (bulk sync), and then bring the CFC into service (an Operational State of UP, a status of ONLINE, and a status of INACTIVE).

4.6.3 Duplex to Simplex (AutoProv Mode)

In AutoProv mode, whenever a CFC card is physically present, the shelf will enable that CFC (to bring its Admin Status to UP) and then try to bring the CFC card into service (perform data sync). Therefore, the card must be physically removed to prevent this. Follow these steps:

1. Disable the inactive CFC (**DI SABLE CARD=8 INACTCFC** or **DI SABLE CARD=12 INACTCFC**).
2. Follow antistatic procedures and remove the inactive CFC from its double slot and place in an antistatic container.
3. With the inactive CFC card now removed, delete the card from the configuration database (**DESTROY CARD=8** or **DESTROY CARD=12**).
4. If desired, provision the now NOT PROVISIONED slot with a Service Module card.
5. Fill any empty slots with Filler Plates.

4.6.4 Simplex to Duplex (Manual Mode)

When the fMAP system is in simplex mode, one of the double slots (8/9, 12/13) can contain one Service Module card (such as an ADSL16) and a face plate full (FPF). Moreover, because the CFC is in manual mode, the CFC must be explicitly created and enabled, as shown in the following steps:

1. Deprovision the Service Module card if it occupies a slot for the new CFC card.
2. If a Service Module card was installed, follow antistatic procedures and remove the Service Module card from either Slot 8 or 12.
3. Follow antistatic procedures and remove any Filler Plate Full card(s) in slots 8/9 or 12/13.
4. Remove the CFC card from its antistatic container.
5. Follow antistatic procedures and insert the CFC card into the available double slot. Refer to the Installation Guide for details.
6. Because the CFC cards are in Manual mode, the slot(s) will have a status of NOT PROVISIONED. The card must therefore be provisioned using the command:

```
CREATE CARD=8 <CFC6 or CFC24> or CREATE CARD=12 <CFC6 or CFC24>
```

7. The card is now provisioned in the configuration database, but must change its Admin State to UP to sync with the active CFC, and to try to go into service. The default for the manual mode is when cards are created, **ENABLED=UP**, so no additional actions are needed. If this has been changed, enter the **ENABLE CARD=INACTCFC** command.

8. The newly enabled CFC will boot, sync all of its data with the active CFC (bulk sync), and then bring the CFC into service (an Operational State of UP, a status of ONLINE, and a state of INACTIVE).

4.6.5 Duplex to Simplex (Manual Mode)

In Manual mode, the CFC card can be physically present and the shelf will not try to create and enable the card. Therefore, the card does not need to be removed before deleting it from the database. Follow these steps:

1. Disable the inactive CFC (`DISABLE CARD=I NACTCFC`).
2. Delete the card from the configuration database (`DESTROY CARD=8` or `DESTROY CARD=12`).
3. Follow antistatic procedures and remove the inactive CFC from its double slot and place in an antistatic container.
4. Provision the now NOT PROVISIONED slots with Filler Plates or a Service Module card.

4.7 Provisioning the 9100 (GE Settings)

4.7.1 Overview

The fMAP operates only in the **Automatic Provisioning Mode** (PROVMODE = AUTO). The **Manual Provisioning Mode** (PROVMODE = MANUAL) is not supported. Upon start-up, the CFC12 detects the existing hardware, automatically provisions it using a default AUTOPROV profile, and attempts to bring it into service. Commands can then be used to modify provisioning data. When the system is first initialized, the system's PROVMODE is set to AUTO, and all configurable components come up with the profile name AUTOPROV.

Provisioning data is persistent over reboots/restarts of the CFC12.

The 9100 shares most of the same attributes as the 9400 system, since the 9100 is a simplex system that supports many of the same network and service modules as the 9400. Commands show the same type of output, with the exceptions highlighted below.

- SHOW SYSTEM - Shows generic system information as well as a CLEI code
- SHOW CARD - Displays card configuration and their states
- SET INTERFACE GE - The 9100 has SPEED and DIRECTION as provisionable values. Refer to [4.7.2](#).

Note: The SET PORT command is deprecated in 7.0

- SET PROFILE - The 9100 has SPEED as a provisionable value. Refer to [4.7.2](#)

4.7.2 GE Speed Settings are Configurable

For the GE2RJ interface, SPEED is provisionable (10, 100, or 1000); this parameter is ignored with a warning on other GE interface types.

For the GE4 interface, when copper SFP support, the provisionable speeds are 100 or 1000.

For both the GE2RJ and GE4 interfaces, DIRECTION can be set to NETWORK or CUSTOMER.

When setting a profile for the GEPORT, SPEED can be set, If the GE port profile is applied to an interface that does not support the speed setting, the speed setting is ignored.

Refer to [5.3](#) for sample interfaces.

4.8 Command Summary for Provisioning

[Table 4-9](#) lists all the commands and parameters used to provision (and deprovision) the fMAP product components.)

Note: This list only includes commands that are generic across card_types and interface_types. For specific information on card and interface attributes, refer to the next Sections. To see where command have been added, deleted, or changed, refer to the Command Handbook

TABLE 4-9 Command Summary for Provisioning

Object	Verb	Syntax	Description
CARD	CREATE	<pre>CREATE CARD=slot <card_type> [{ [PREFLOAD=filename] [ADMINSTATE={ UP DOWN }] PROFILE=name }]</pre>	<p>For each card type, creates the software provisioning for a card in a specific slot.</p> <p><i>Note:</i> CREATE CARD CFC4 <i>and</i> CREATE CARD GE2 <i>are not supported.</i></p> <p>ADMINSTATE - Refer to 4.1.5.</p> <p>PROFILE - Specifies the name of the profile used to provision the card. The contents of a profile can be displayed (SHOW PROFILE) and changed (SET PROFILE).</p>
CARD	DESTROY	<pre>DESTROY CARD=slot-list [FORCE]</pre>	Removes software provisioning for the specified card or list of cards.
CARD	DISABLE	<pre>DISABLE CARD={ slot-list INACTCFC } [FORCE]</pre>	Takes a card out-of-service and sets the card's administrative state to DOWN.
CARD	ENABLE	<pre>ENABLE CARD={ slot-list INACTCFC } [NODIAGS] [VERBOSE]</pre>	<p>Changes the administrative state of the specified card or set of cards to UP, making it available for service.</p> <p>VERBOSE lists the change in card status as the card is enabled. (Logs, however, are always produced even if this option is not used.)</p>

TABLE 4-9 Command Summary for Provisioning (Continued)

Object	Verb	Syntax	Description
CARD	RESET	<pre> RESET CARD={ slot-list ACTCFC INACTCFC ALL } { CPUSTATS } </pre>	
CARD	RESTART	<pre> RESTART CARD={ slot-list INACTCFC ACTCFC } [COLD] [FORCE] </pre>	Performs a restart of the software running on the specified card.
SYSTEM	RESTART	<pre> RESTART SYSTEM [FORCE] </pre>	Allows the user to restart an entire duplex system; avoiding the requirement to restart both the ACTCFC and INACTCFC.
CARD	SET	<pre> SET CARD={ slot-list ACTCFC INACTCFC } [PREFLOAD={ filename NONE }] [ALTLOAD={ filename NONE }] [TEMPLOAD={ filename NONE }] </pre>	Modifies the provisioning attributes for the specified card or list of cards.
CARD PRO-FILE	SET	<pre> SET CARD=slot-list PROFILE=name </pre>	Changes the attributes of a card by modifying them with a profile.
SYSTEM PROV-MODE	SET	<pre> SET SYSTEM [PROVMODE={ MANUAL AUTO }] </pre>	Changes the system provisioning mode from AUTO to MANUAL or vice-versa.
SYSTEM PROV-MODE	SHOW	<pre> SHOW SYSTEM PROVMODE </pre>	

TABLE 4-9 Command Summary for Provisioning (Continued)

Object	Verb	Syntax	Description
CARD	SHOW	<pre> SHOW CARD [= { slot-list ACTCFC INACTCFC ALL }] [{ INVENTORY SOFTWARE FULL }] </pre>	
CARD	SHOW	<pre> SHOW CARD [= { slot-list ACTCFC INACTCFC ALL }] { CPUSTATS [TASKS] MEMORY { HEAP MESSAGEBUFFERS QUICKHEAP } } </pre>	<p>Displays various information about the provisioned card in the specified slot(s)</p> <p>Output for <code>SHOW CARD=ACTCFC</code> is in 4.5.2.</p>
CARD PORTS	SHOW	<pre> SHOW CARD={ slot-list ALL } PORTS </pre>	Show for the specified card(s) the associated ports and their status (State, Connection, Mode, etc.)
PORT	ENABLE	<pre> ENABLE PORT=port-list </pre>	Places the port(s) in the UP administrative state and attempts to make the port in-service.
INTERFACE	ENABLE	<pre> ENABLE INTERFACE={ type:id-range id-range ifname-list } </pre>	Places the interface in the UP administrative state and attempts to make the port in-service.
INTERFACE	DISABLE	<pre> DISABLE INTERFACE={ type:id-range id-range ifname-list } [FORCE] </pre>	Places the interface(s) in the DOWN administrative state; FORCE will do this even is the interface is operationally UP.
PORT	DISABLE	<pre> DISABLE PORT=port-list [FORCE] </pre>	Places the port(s) in the DOWN administrative state; FORCE will do this even is the interface is operationally UP.

TABLE 4-9 Command Summary for Provisioning (Continued)

Object	Verb	Syntax	Description
INTERFACE	SET	<pre> SET INTERFACE={ type:id-range id-range ifname-list ALL } [ACCEPTABLE={ ALL VLAN HVLAN }] [INFILTERING={ OFF ON }] [TAGALL={ ON OFF }] [TPID=tpidvalue] [LEARNLIMIT={ 1..64 OFF }] </pre>	<p>Changes a variety of attributes for the interface:</p> <p>ACCEPTABLE -</p> <p>INFILTERING -</p> <p>TAGALL -</p> <p>TPID -</p> <p>LEARNLIMIT -</p>
INTERFACE DESCRIP- TION	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } DESCRIPTION={ description NONE } </pre>	Associates a description with the interface(s).
INTERFACE PROFILE	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } PROFILE=name </pre>	Associates the interface type with the Profile name.
PORT PROFILE	SET	<pre> SET PORT=port-list PROFILE=name </pre>	Associates the port(s) with the Profile name.

TABLE 4-9 Command Summary for Provisioning (Continued)

Object	Verb	Syntax	Description
INTERFACE	SHOW	<pre> SHOW INTERFACE [={ type: type:id-range id-range ifname-list ALL }] [CARD=slot-list] [STATE={ UP DOWN ALL }] [DIRECTION={ NETWORK CUSTOMER INTERNAL }] [FULL] </pre>	Shows the attributes of the interface(s). What is shown depends on the type of interface.
PORT	SHOW	<pre> SHOW PORT [={ port-list ALL }] [DIRECTION={ NETWORK CUSTOMER INTERNAL }] </pre>	Shows the attributes of the interface(s). What is shown depends on the type of interface.

5. Provisioning Ethernet Interfaces

5.1 Overview

Ethernet is a widely-installed local area network (LAN) technology, that is specified in standard IEEE 802.3, and is a physical and data link layer protocol when used for LAN connections. fMAPs use Ethernet to deliver carrier-class services such as video, data, and voice.

5.1.1 7000/9000

fMAP uplink connections are made through gigabit ethernet Network Modules (NM). Each fMAP is configured with at least one NM. Subscriber access to the fMAP, and eventually the uplink, is made through a Service Module (SM). Service modules using Ethernet as their connection technology, are shown in [Figure 5-1](#):

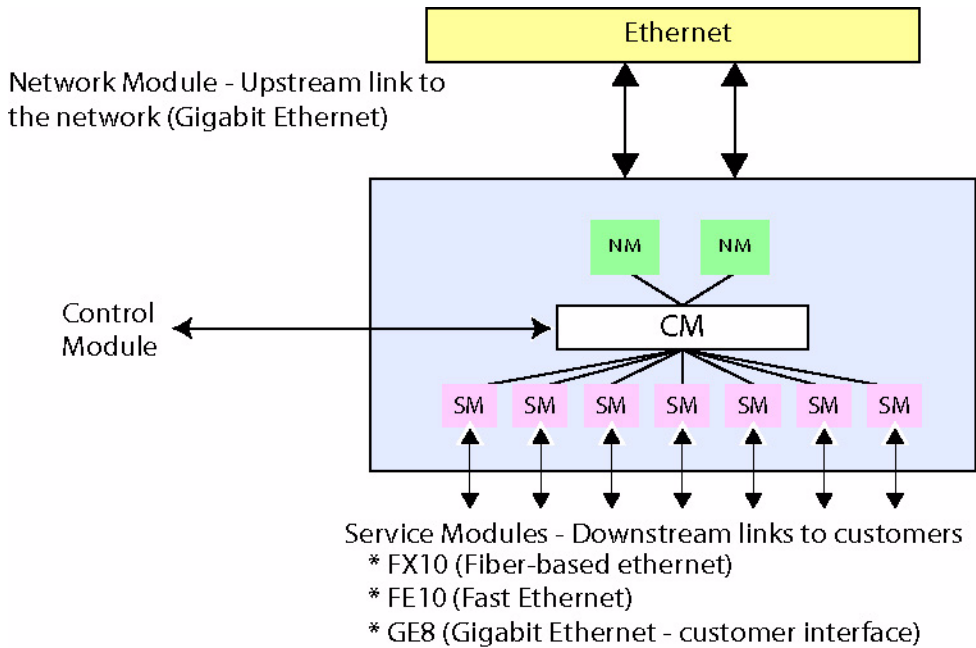


FIGURE 5-1 7000 and 9000 fMAP NM, CM, and SM for Ethernet Interfaces

5.2 Fast Ethernet and Fiber-based Interfaces

5.2.1 Overview

This section describes provisioning for the FE and FX card and interfaces. For more information on provisioning features provided by these cards, see the **Allied fMAP Service guide**.

5.2.2 FE10 Interface

The following shows the output for the **SHOW INTERFACE** command for the FE10. For a description of the General attributes, refer to Section 4. For a description of the VLAN attributes, refer to Section 14.

```
officer SEC>> SHOW INTERFACE ETH:20.0

--- FE Interfaces ---

Interface..... 20.0
Type..... FE
State..... UP-DN-Dependency
Description..... <none>

Provisioning
  Provisioning Profile..... AutoProv
  Direction..... Customer
  Auto Negotiation..... On
  Speed..... Auto
  Duplex..... Auto
  Flow Control..... Auto
  Remote Monitoring..... Off

Actual
  Direction..... Customer
  Physical Address..... 00:0C:25:00:05:C2

VLAN Information
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Dynamic MAC Learning Limit..... Off
  Untagged VLAN..... 1
```

```
officer SEC>> SHOW INTERFACE CARD 20

--- FE Interfaces ---

Interface   State Autonegotiate Flow Control Duplex Speed   Direction
-----
20.0        UP-DN On          -          -          -       Customer
20.1        UP-DN On          -          -          -       Customer
20.2        UP-DN On          -          -          -       Customer
20.3        UP-DN On          -          -          -       Customer
20.4        UP-DN On          -          -          -       Customer
20.5        UP-DN On          -          -          -       Customer
20.6        UP-DN On          -          -          -       Customer
20.7        UP-DN On          -          -          -       Customer
20.8        UP-DN On          -          -          -       Customer
20.9        UP-DN On          -          -          -       Customer
```

Note: The attribute `Direction` is added in release 5.0, which controls whether the port uses the upstream (`NETWORK`) or downstream (`CUSTOMER`) direction. The `NETWORK` value is used for the FE10 upstream interface feature. For details on this attribute and examples.

5.2.3 FX10 Interface

The following shows the output for the **SHOW INTERFACE** command for the FX10. This is similar to the FE interface.

```
officer SEC>> SHOW INTERFACE 17.0

--- FX Interfaces ---

Interface..... 17.0
Type..... FX
State..... DN-DN-Offline
Description..... <none>

Provisioning
  Provisioning Profile..... AutoProv
  Direction..... Customer
  Auto Negotiation..... Off
  Flow Control..... Off
  Remote Monitoring..... Off

Actual
  Direction..... Customer
  Physical Address..... 00:0C:25:00:06:7C

VLAN Information
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Dynamic MAC Learning Limit..... Off
  Untagged VLAN..... 1

officer SEC>> SHOW INTERFACE CARD 17

--- FX Interfaces ---

Interface  State Flow Control Duplex Speed  Direction
-----
17.0      DN-DN Auto          -      -      Customer
17.1      DN-DN Auto          -      -      Customer
17.2      DN-DN Auto          -      -      Customer
17.3      DN-DN Auto          -      -      Customer
17.4      DN-DN Auto          -      -      Customer
17.5      DN-DN Auto          -      -      Customer
17.6      DN-DN Auto          -      -      Customer
17.7      DN-DN Auto          -      -      Customer
17.8      DN-DN Auto          -      -      Customer
17.9      UP-UP Off            Full   100 Mbps Customer
```

5.3 9100 Ethernet Interfaces

A key difference between the 9100 and other fMAP devices is that like the 7100, the 9100 does not have physical GE cards but has 2 **virtual** network/service modules; these are identified as GE4 and GE2RJ. The associated hardware is located on the CFC12, but is represented to the user as two virtual card types.

To control the attributes of these interfaces, the following command is used:

```
SET
INTERFACE={ type:
             type:id-range
             id-range
             ifname-list
             ALL
            }

GE
[ AUTONEGOTIATION={ ON
                   OFF
                  } ]
[ SPEED={ AUTONEGOTIATE
          10
          100
          1000
        } ]
[ DUPLEX={ AUTONEGOTIATE
           FULL
           HALF
         } ]
[ FLOWCONTROL={ AUTONEGOTIATE
                ON
                OFF
               } ]
[ DIRECTION={ NETWORK
               CUSTOMER
             } ]
[ DESCRIPTION=description ]
[ FORCE ]
```

The attribute **SPEED** is only provisionable on the GE2RJ interfaces, and is ignored with a warning on other GE interface types that do not support this provisioning (GE1, GE2, GE3, GE4). When copper SFP support is added to the GE4, a provisionable speed of 100 or 1000 will be allowable for its interfaces. **SPEED** is also an attribute that can be set for the GE interface profile. Note that for the other GE interfaces, this attribute is ignored.

The attribute **DIRECTION** is also provisionable on the GE2RJ and GE4.

The following illustrates the output from the **SHOW INTERFACE** as it relates to GE detailed output. The example shows possible output for the GE2RJ (new items in the output shown in bold).

```
officer SEC>> SHOW INTERFACE 5.0
--- GE Interfaces ---

Interface..... 4.0
Type..... GE
```

State..... UP-UP-Online
Description..... <none>

Provisioning

Provisioning Profile..... AutoProv
Auto Negotiation..... On
Port Speed..... Auto
Flow Control..... Auto
Remote Monitoring..... Off

Actual

Direction..... Customer
Physical Address..... 00:0C:25:00:00:1D
Port Speed..... 100 Mbps (based on the far-end rate)

VLAN Information

Acceptable Frame Types..... All
Ingress Filtering..... On
TPID..... 0x8100
TAGALL..... Off
Untagged VLAN..... 1

6. Provisioning ADSL

6.1 Overview

Asymmetric digital subscriber line (ADSL) is a DSL technology that allows data flow in the downstream direction at faster speeds than the upstream direction. ADSL service modules (SM) provide a IP/Ethernet-based access solution that features advanced ADSL technologies to deliver voice, video, and data services over copper wire. ADSL SMs come in several versions, as illustrated below

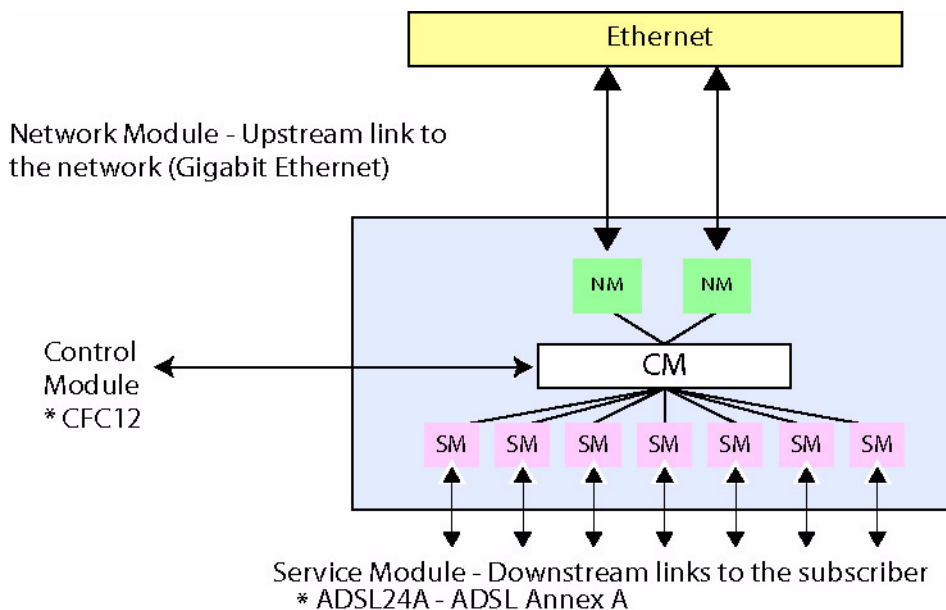


FIGURE 6-1 fMAP 9000 ADSL SMs

6.2 ADSL Card

This section describes provisioning for the ADSL card and interfaces. For more information on provisioning features provided by ADSL, see the **Allied fMAP Service guide**.

The following shows the output for the **SHOW CARD** command for an ADSL type.

```
officer SEC>> show card 0
```

```
--- Card Information ---
```

```
Slot..... 0
Type..... ADSL24A
State..... UP-UP-Online
Provisioning Profile..... AutoProv

Hardware
  Model Number (Revision)..... TN-123-A (Rev X4)
  Serial Number..... ATNLAB4040302469
  CLEI Code..... <none>

Software
  Running Load..... 6.0.0.ALPHA.20050124
  Preferred Load..... <none>
  Temporary Load..... <none>
```

```
officer SEC>> show card 2
```

```
--- Card Information ---
```

```
Slot..... 2
Type..... ADSL24
State..... UP-UP-Online
Provisioning Profile..... AutoProv

Card Faults
  Inconsistent Load..... Minor

Hardware
  Model Number (Revision)..... TN-112-A (Rev X8)
  Serial Number..... ATNLAB4030200696
  CLEI Code..... <none>

Software
  Running Load..... 5.0.1
  Preferred Load..... <none>
  Temporary Load..... <none>
```

The ADSL16, ADSL8S, ADSL16B, and ADSL24 cards all have similar attributes.

6.3 ADSL Interface Attributes

6.3.1 Output of SHOW INTERFACE Command

The following shows the output for the **SHOW PORT** command for the ADSL16. Attributes for the ADSL8S, ADSL16B, ADSL24, ADSL24A, and ADSL24B are similar.

```
officer SEC> officer SEC>> show interface 0.0
```

```
--- ADSL Interfaces ---
```

```
Interface..... 0.0
Type..... ADSL
State..... UP-DN-Failed
Description..... <none>

Interface Faults
  Loss of Signal..... Info (masked)
  Loss of Frame..... Info (masked)
  Loss of Link..... Info (masked)
  No Peer Present..... Info
```


Provi si oni ng	
Provi si oni ng Profi le.....	AutoProv
Mode.....	Auto2+
Li ne Type.....	Interleave
Interleave Del ay.....	32 msec
Target SNR Margin.....	8 dB
Echo Cancell ation.....	Off
Li ne Quali ty Moni tor.....	Low
ATU-C (Upstream)	
Maxi mum Rate.....	1024 kbps
Mi ni mum Rate.....	32 kbps
ATU-R (Downstream)	
Maxi mum Rate.....	26624 kbps
Mi ni mum Rate.....	32 kbps
Performance Moni tori ng.....	Off
Remote Moni tori ng.....	Off
Actual	
Connecti on State.....	Handshake
Di recti on.....	Customer
Physi cal Address.....	00: 0C: 25: 03: 90: 48
VC I nformati on	
VC I denti fier.....	0
VPI.....	0
VCI.....	35
Service Category.....	UBR
TX Peak Cell Rate.....	Maxi mum cps
VLAN I nformati on	
VC I denti fier.....	0
Acceptable Frame Types.....	All
Ingress Fi lteri ng.....	On
TPI D.....	0x8100
TAGALL.....	Off
Dynami c MAC Learni ng Li mi t.....	Off
Untagged VLAN.....	1

6.3.2 ADSL Mode Selection

6.3.2.1 Mode Selection up to Release 4.0

Up to and including release 4.0, fMAP's line of ADSL service modules employed three common Physical-Layer modulation schemes to provide data access over customer telephone loops:

- ITU G.992.1 (G.DMT, ADSL), circa June 1999. For the fMAP, this equates to the **GDMT** mode for ADSL ports.
- ITU G.992.2 (GLITE, ADSL-Lite), circa June 1999. For the fMAP, this equates to the **GLITE** mode for ADSL ports.
- G.span (ADSL+), [vendor-proprietary]. G.span was in effect the Alpha version of ITU G.992.5. For the fMAP, this equates to the **ADSL2+** mode for ADSL ports.

6.3.2.2 Mode Selection

Industry requirements for higher bandwidth, better management, and support for a broader range of premise loops, have spawned the evolution of "ADSL2". ADSL2 is, effectively, an extension to ADSL, providing a new

family of ITU-ratified modulations, intended to address limitations of the original specifications. Several potential improvements have been identified in areas such as data rate versus loop reach performance, loop diagnostics, deployment from remote cabinets, spectrum control, power control, robustness against loop impairments and RFI, and operations and maintenance.

ADSL2 is defined by:

- ITU G.992.3 (G.DMT.bis, ADSL2), circa July 2002
- ITU G.992.3 Annex L (Reach Extended ADSL2)
- ITU G.992.4 (G.LITE.bis, ADSL2-Lite), circa July 2002
- ITU G.992.5 (ADSL2plus), circa May 2003

Beginning with release 5.0, the fMAP supports the new ADSL modulations (except for ITU G.992.4) and their value-added features on all Annex A-supporting service modules, including ADSL16, ADSL8S and ADSL24. There is no change to configuration and support of higher layer services that run on ADSL interfaces. After upgrading the CFC and ADSL cards to release 5.0, G.Span is no longer supported and is replaced by ITU G.992.5 when running ADSL2+ mode.

Note: This functionality requires an ADSL2 supporting modem and release 5.0 software running in both the CM and SM to utilize the new modes.

6.3.2.3 Configurable Modes from 5.0

The MODE attribute is expanded to include the following choices:

- **AUTO2+** - an expanded multimode, where the choice of all old and new modes (excluding G.span) is auto negotiated with the remote peer. This is the default mode beginning with release 5.0. Note that if the CFC is running release 5.0 and the ADSL card is running release 4.0.x, then AUTO2+ has the behavior of release 4.0.x ADSL2+ (G.span and fallback).

The following existing modes are still supported but have different functionality:

- **ADSL2+** - ITU G.992.5 instead of G.span when running release 5.0 in the CFC.

The following pre-existing modes are not changed:

- **T1.413** - ANSI T1.413 issue 2
- **GDMT** - ITU G.992.1
- **GLITE** - ITU G.992.2
- **AUTO** - multimode, where the choice mode is auto negotiated with the remote peer. The choices include the pre-existing modes only (T1.413, GDMT, GLITE).

Changes are also made to the functionality of the following attributes:

- **EC** (Echo Cancellation) - prior to release 5.0 this was only allowed to be ON, if the mode was set to GDMT. In release 5.0, EC is allowed to be set to ON if the mode is set to either GDMT or any of the new modes. The default is OFF.

- **MINDOWNSTREAMRATE** - The range is from 32Kb to 16128Kb for all new modes. Default is 32Kb. No change to range and default for pre-existing modes.
- **MAXDOWNSTREAMRATE** - The range is from 32Kb to 16128Kb for all new modes. Default is 10016Kb. No change to range and default for pre-existing modes.
- **MINUPSTREAMRATE** - The range is from 32Kb to 1024Kb for all new modes. Default is 32Kb. No change to range and default for pre-existing modes.
- **MAXUPSTREAMRATE** - The range is from 32Kb to 1024Kb for all new modes. Default is 1024Kb. No change to range and default for pre-existing modes.
- **ENCAPSULATIONTYPE** - Up to and including release 4.0, allowed values were LLCSNAP and VCMUX, although VCMUX was not commonly used. The default is LLCSNAP. In release 5.0, support for VCMUX is removed for all pre-existing and new modes.

All other ADSL port attributes are unchanged in release 5.0 and are supported for all modes. These attributes include:

- INTERLEAVEDELAY
- TARGETSNRMARGIN
- LINEQUALITYMONITOR
- LINETYPE
- VPI
- VCI
- DESCRIPTION

6.3.2.4 Handshake and Fallback Behavior

During handshake protocol with the CPE modem, the ADSL ports are capable of *fallback* in certain modes. Fallback refers to the ability of the port to auto-negotiate mode with the remote peer (the CPE modem) and pick the mode that delivers optimum performance.

In release 4.0.x, only AUTO and ADSL2+ supported fallback behavior. In release 5.0, only AUTO and AUTO2+ support fallback behavior. Because of this, during upgrade from release 4.0.x to 5.0, all ADSL2+ mode settings are automatically migrated to AUTO2+. If the CFC is running release 5.0 and the ADSL card is running 4.0.x, then AUTO2+ has the behavior of release 4.0.x ADSL2+ (G.span and fallback).

The following tables illustrate the mode support and fallback behavior of the ADSL ports based on the mode attribute setting and the load version running in the CFC and ADSL card. The fallback paths are from right to left.

TABLE 6-1 CFC and ADSL SM both running release 4.0.x software

	CPE mode support (fallback from left to right)			
CLI Mode Setting	G.Span	G.DMT (ITU G992.1)	G.LITE (ITU G992.2)	T1.413
T1.413	No	No	No	Yes
GLITE	No	No	Yes	No
GDMT	No	Yes	No	No
AUTO	No	Yes	Yes	Yes
ADSL2+	Yes	Yes	Yes	Yes

TABLE 6-2 CFC running release 5.0 or later, ADSL SM running release 4.0.x software

	CPE mode support (fallback from left to right)						
CLI Mode Setting	ADSL2+ (ITU G.992.5)	ADSL2 (ITU G992.3)	READSL2 (ITU G992.3 Annex L)	G.span	G.DMT (ITU G992.1)	G.LITE (ITU G992.2)	T1.413
T1.413	No	No	No	No	No	No	Yes
GLITE	No	No	No	No	No	Yes	No
GDMT	No	No	No	No	Yes	No	No
AUTO	No	No	No	No	Yes	Yes	Yes
ADSL2	No	No	No	No	No	No	No
ADSL2+	No	No	No	No	No	No	No
AUTO2+	No	No	No	Yes	Yes	Yes	Yes

TABLE 6-3 CFC and ADSL SM running release 5.0 or later software

CLI Mode Setting	CPE mode support (fallback from left to right)						
	ADSL2+ (ITU G.992.5)	ADSL2 (ITU G992.3)	READSL2 (ITU G992.3 Annex L)	G.span	G.DMT (ITU G992.1)	G.LITE (ITU G992.2)	T1.413
T1.413	No	No	No	No	No	No	Yes
GLITE	No	No	No	No	No	Yes	No
GDMT	No	No	No	No	Yes	No	No
AUTO	No	No	No	No	Yes	Yes	Yes
ADSL2	No	Yes	Yes	No	No	No	No
ADSL2+	Yes	No	No	No	No	No	No
AUTO2+	Yes	Yes	Yes	No	Yes	Yes	Yes

6.3.2.5 Runtime Attributes

In release 5.0, modifications are made to the display of the runtime attributes in response to the SHOW PORT command, as follows:

- Actual Line Standard - can be one of the following:
 - GLITE - ITU G.992.2
 - GDMT - ITU G.992.1
 - T1413 - ANSI T1.413
 - ADSL2+ (ITU GSPAN) - G.span (only if running release 4.0 or earlier software in the ADSL card)
 - ADSL2 - ITU G.992.3
 - ADSL2+ - ITU G.992.5
 - READSL2 - ITU G.992.3 Annex L
- Actual Upstream Rate - In release 4.0 and earlier, all rates were reported as numbers that were multiples of 32 kbps. With release 5.0 running in the ADSL card, rates are reported as numbers that are multiples of 4 kbps.
- Actual Downstream Rate - In release 4.0 and earlier, all rates were reported as numbers that were multiples of 32 kbps. With release 5.0 running in the ADSL card, rates are reported as numbers that are multiples of 4 kbps.

In addition, new runtime attributes are added, as follows:

- Trellis Coding - values displayed include:

- Active
- Not Active
- Unknown (only if the ADSL card is running release 4.0 or earlier software or the modem doesn't support ADSL2)
- Max Attainable Rate (ATUC) - from 0 to 26000 kbps. A value of 0 is displayed if the port is not in Show-time. A value of Unknown is displayed if the ADSL card is running release 4.0 or earlier software or the modem doesn't support ADSL2.
- Max Attainable Rate (ATUR) - from 0 to 26000 kbps. A value of 0 is displayed if the port is not in Show-time. A value of Unknown is displayed if the ADSL card is running release 4.0 or earlier software or the modem doesn't support ADSL2.

All other runtime attributes remain the same as the release 4.0 or earlier software.

6.3.3 Annex A versus Annex B

The Annex A provides ADSL service and POTS as described in ITU-T Recommendation G.992.

Annex B allows ISDN Basic Rate Access and ADSL services to co-exist over the same subscriber line. GLITE, GDMT, and AUTO are supported for Annex B, and the new ADSL2+ mode (as well as T1.413) is supported for Annex B in release 6.0 as well.

6.3.4 ADSL Cards Starting with Release 6.0 (ADSL24A)

In release 6.0, the new technology versions of both cards are introduced and **supported in 9x00 product lines**. Because it's expected that the old and new versions of ADSL24 might coexist in the same shelf and have different capabilities and require different loads, the new versions have different card types, as follows:

- **ADSL24A** (TN-121-A, the annex A version)

The cards are provisioned and managed like all other ADSL variant cards. All the same services are supported on the new cards as was supported on the old cards, with exceptions in the following areas:

- Classifier support is limited on these cards - Refer to [15.1](#).
- QOS queue support is increased (8 queues instead of 4)
- VLAN-VC mapping is supported on the new cards (it was not supported on the older ADSL24 versions) - Refer to [12.4](#).

6.3.5 Listing of ADSL Interface Attributes

[Table 6-4](#) lists the ADSL Interface attributes, and summarizes much of the information presented above.

TABLE 6-4 ADSL Interface Attributes - Default is in Bold

ADSL Port Attribute	Values / Range	Description
AdslLineType	FAST INTERLEAVE	<p>The LINETYPE parameter specifies the ADSL line type as per ITU G.992. Allowed values are FAST and INTERLEAVE, although FAST is not allowed if the MODE is GLITE.</p> <p>The default is INTERLEAVE.</p> <p>The FAST parameter specifies the ADSL line type as using the fast path as described in ITU G.992. The fast path provides low latency.</p> <p>The INTERLEAVE parameter specifies the ADSL line type as using the interleaved path as described in ITU G.992. The interleaved path provides a low error rate but greater latency than the fast path.</p>
AdslMode	GLITE GDMT T1.413 ADSL2 ADSL2+ AUTO AUTO2+ ADSL2M ADSL2+M	<p>Specifies the ADSL line mode standard.</p> <p>GLITE: GLite (ITU G.992.2)</p> <p>GDMT: G.DMT (ITU G.992.1)</p> <p>T1.413: ANSI T1.413 Issue 2</p> <p>AUTO: multimode, where the choice of G.lite, G.dmt, or T1.413 is auto negotiated with the remote peer.</p> <p>ADSL2+ - This is for higher rates, and can support more than two Set Top Boxes (STB). When this mode is chosen, the Echo Cancellation (EC) is set to OFF, and the MAXDOWNSTREAMRATE is set at 26.</p> <p>ADSL2M/2+M - For Annex M.</p> <p>The default for new (release 5.0 and later) card provisioning is AUTO2+.</p> <p><i>Note: Only with AUTO2+, ADSL2, and ADSL2+ can the Data Boost feature work.</i></p>
Max Interleave Delay	1..64	<p>Specifies the maximum interleave delay in milliseconds used when the ADSL linetype is set to INTERLEAVE. The valid range is from 1 to 64.</p> <p>32 is the default.</p>

TABLE 6-4 ADSL Interface Attributes - Default is in Bold (Continued)

ADSL Port Attribute	Values / Range	Description
Echo Cancel	ON OFF	<p>Specifies whether echo cancellation is utilized on ADSL ports running G.DMT mode as per ITU-T Recommendation G.992.1.</p> <p>If ON the port uses overlapping spectrum operation to more effectively use bandwidth between the upstream and downstream frequencies, thus boosting the connect rate.</p> <p>This parameter is only allowed to be set to ON if the MODE is set to GDMT. The default value is OFF.</p>
Encapsulation Type	LLCSNAP VCMUX	<p>Specifies the ATM data encapsulation protocol used by an ADSL port, as defined in IETF RFC-1483.</p> <p>LLC/SNAP - (Logical Link Control with Subnetwork Attachment Point)</p> <p>VCMUX - VC Mux (Virtual Circuit Based Multiplexing)</p> <p>The default is LLCSNAP.</p>
VPI	0..255	<p>Specifies the value for the ATM virtual path identifier on an ADSL port.</p> <p>The default is 0.</p>
VCI	32..65535	<p>Specifies the value for the ATM virtual channel identifier on an ADSL port. The valid range for this parameter is from 32 to 65535.</p> <p>The default is 35.</p>
Maximum Upstream Rate	32..1024 (Kb)	<p>Specifies the maximum upstream bit rate to attain for an ADSL port. The valid range for this parameter is from 32Kb to 1024Kb for all line modes except GLITE. If the MODE is set to GLITE, then the valid range is from 32Kb to 512Kb.</p> <p>The default is 1024Kb.</p>

TABLE 6-4 ADSL Interface Attributes - Default is in Bold (Continued)

ADSL Port Attribute	Values / Range	Description
Minimum Upstream Rate	32..1024 (Kb)	<p>Specifies the minimum upstream bit rate to attain for an ADSL port. The valid range for this parameter is from 32Kb to 1024Kb for all line modes except GLITE. If the MODE is set to GLITE, then the valid range is from 32Kb to 512Kb.</p> <p>The default is 32Kb.</p> <p>The MINUPSTREAMRATE must be equal or less than the MAXUPSTREAMRATE.</p>
Maximum Downstream Rate	32..26624 (Kb)	<p>Specifies the maximum downstream bit rate to attain for an ADSL port. The valid range for this parameter is from 32Kb to 16128Kb for all line modes except GLITE. If the MODE is set to GLITE, then the valid range is from 32Kb to 1536Kb.</p> <p>The default is 26624.</p> <p><i>Note: In release 6.0, changing the mode of a line does not change a previously set rate.</i></p>
Minimum Downstream Rate	32..26624 (Kb)	<p>Specifies the minimum downstream bit rate to attain for an ADSL port. The valid range for this parameter is from 32Kb to 26624Kb for all line modes except GLITE. If the MODE is set to GLITE, then the valid range is from 32Kb to 1536Kb.</p> <p>The default is 32Kb.</p> <p>The MINDOWNSTREAMRATE must be less than the MAXDOWNSTREAMRATE.</p> <p>This value is especially important for configuring Set Top Boxes (STB), and should have the following minimum values:</p> <p>One STB - 5000 Kb Two STBs - 9000 Kb For information on configuring three STBs.</p>
Target SNR Margin	0..15 (dB)	<p>Specifies the target signal-to-noise ratio (in dB) to achieve on an ADSL port.</p> <p>The default value is 8.</p>

TABLE 6-4 ADSL Interface Attributes - Default is in Bold (Continued)

ADSL Port Attribute	Values / Range	Description
Line Quality Monitor	Medium	<p>The LINEQUALITYMONITOR parameter specifies a level of quality (amount of errors) acceptable for the ADSL port. If an unacceptable number of errors occurs, the ADSL port will automatically retrain. Allowed values include:</p> <ul style="list-style-type: none"> - Low: Port is monitored for Data applications (approx. 10⁻⁷ bit errors) - Medium: Port is monitored for Video applications (approx. 10⁻⁹ bit errors) - High: Port is monitored for High Quality applications (approx. 10⁻¹⁰ bit errors). <p>The default is Medium. The user can set this parameter only when the port is disabled (See DISABLE PORT) This parameter is only applicable to ADSL ports.</p>
Connection State	SHOWTIME HANDSHAKE IDLE	<p>This is the state of the data connection. Values are:</p> <ul style="list-style-type: none"> • SHOWTIME - Modems are synchronized and data is exchanged. • HANDSHAKE - Negotiation is in progress and the are synchronizing. • IDLE - No data is being exchanged.
Dynamic Attributes	Actual Line Standard Actual Line Type Actual Upstream Rate Actual Downstream Rate Actual SNR (ATUC-Near End) Actual Attenuation (ATUC-Near End) Actual Output Power (ATUC-Near End) Actual SNR (ATUR-Far End) Actual Attenuation (ATUR-Far End) Actual Output Power (ATUR-Far End)	<p>These are the actual runtime attributes of the port after showtime status has been achieved. They are based upon port provisioning</p> <p>If the AdslMode is Auto, the actual mode chosen is negotiated, and the actual values may or may not match how the card is provisioned (what is contained in the AUTOPROV profile).</p> <p>Ensure that any rates set are at the minimum allowed for a service so that the actual rates do not drop below the set level.</p>

6.3.6 Provisioning Scenarios for ADSL Cards

6.3.6.1 Add an ADSL Card

Empty service module slots can be provisioned with additional ADSL cards. It is recommended that SM slots be filled in numerical order 0-*n*. In this example an ADSL card is being provisioned in slot 10.

1. When it is time to install the card, follow antistatic procedures and remove the Filler Plate Full (FPF).
2. Insert in its place the ADSL card. Refer to the Installation Guide for details.
3. The card will provision itself and come up in an Administration State of UP and an Operational State of DOWN, since there is no data cable attached yet. An alarm is produced.
4. Disable the card. This will disable the alarm on the card as well as the associated ports.

DI SABLE CARD=10

5. Connect the RJ21 data cables and configure them for the customer installation (such as splitters). Refer to the Installation Guide.
6. Use the ENABLE card to run diagnostics on the card, and if the diagnostics pass, to bring the card to the Operational State of UP. The ports will come up as well and the card can now begin data processing.

ENABLE CARD=10

Note: It is not necessary to ENABLE the ADSL ports, since the card, once enabled, will put the ports in service if they have no faults.

7. Enable any features that require separate commands for the ADSL card, such as associating a MAC/STB with the ADSL port.

6.3.6.2 Deprovision an ADSL Card

To deprovision the Service Module card, the reverse sequence is followed:

1. Delete any features that explicitly use the Service Module card. For example, delete the MAC/STB associations for video service.
2. Use the DISABLE command to disable the ports so that they will no longer provide data service. The ADMINSTATE of the port will change to DOWN, and the OPERATIONAL STATE will follow to DOWN.

DI SABLE PORT=10. 0-10. 15 (This is the port range for the ADSL16)

3. Use the DISABLE command to disable the ADSL card.

DI SABLE CARD=10

Note: Since the port has already been disabled, there is no need to use the FORCE option.

4. Remove the data cable.
5. Following antistatic procedures, remove the card from its slot and place in an antistatic container. Refer to the Installation Guide, since it shows how to use the latch-locks.
6. Replace the now empty slot with a filler card.
7. The card can now be destroyed. Use the DESTROY command as follows:

DESTROY CARD=10

6.4 Command Summary

TABLE 6-5 Command Summary for Provisioning ADSL

Object	Verb	Syntax	Description
PROFILE ADSLPORT	CREATE	<pre> CREATE PROFILE=name ADSLPORT [MODE={ GLITE GDMT T1.413 ADSL2 ADSL2+ AUTO AUTO2+ ADSL2M ADSL2+M }] [BITMAPMODE={ FBM DEM }] [LINETYPE={ FAST INTERLEAVE }] [INTERLEAVEDELAY=1..64] [ECHOCANCELLATION={ ON OFF }] [DATABOOST={ ON OFF }] [MAXUPSTREAMRATE=32..3072] [MINUPSTREAMRATE=32..3072] [MAXDOWNSTREAMRATE=32..26624] [MINDOWNSTREAMRATE=32..26624] [TARGETSNRMARGIN=0..15] [MAXSNRMARGIN={ OFF 1..30 }] [LINEQUALITYMONITOR={ LOW MEDIUM HIGH }] [VPI=0..4095] [VCI=32..65535] [ADMINSTATE={ UP DOWN }] </pre>	<p>Creates the Profile for an ADSL port that can then be applied to multiple ports.</p> <p>Refer to 6.3 for a description of these attributes.</p>

TABLE 6-5 Command Summary for Provisioning ADSL (Continued)

Object	Verb	Syntax	Description
INTERFACE ADSL	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } ADSL [MODE={ GLITE GDMT T1.413 ADSL2 ADSL2+ AUTO AUTO2+ ADSL2M ADSL2+M }] [BITMAPMODE={ FBM DBM }] [LINETYPE={ FAST INTERLEAVE }] [INTERLEAVEDELAY=1..64] [ECHOCANCELLATION={ ON OFF }] [DATABOOST={ ON OFF }] [MAXUPSTREAMRATE=32..3072] [MINUPSTREAMRATE=32..3072] [MAXDOWNSTREAMRATE=32..26624] [MINDOWNSTREAMRATE=32..26624] [TARGETSNRMARGIN=0..15] [MAXSNRMARGIN={ OFF 1..30 }] [LINEQUALITYMONITOR={ LOW MEDIUM HIGH }] [VPI=0..4095] [VCI=32..65535] [DESCRIPTION=description] </pre>	<p>Sets the attributes for the ADSL interface.</p> <p>Refer to 6.3 for a description of these attributes.</p> <p>For Data Boost.</p>

7. Provisioning Services with DS1/E1

7.1 Overview

7.1.1 CES8 Configuration

Circuit Emulation Service (CES) over Ethernet is provided by the CES8 card on the fMAP. Here is a diagram of a sample customer network employing CES. See [Figure 7-1](#).

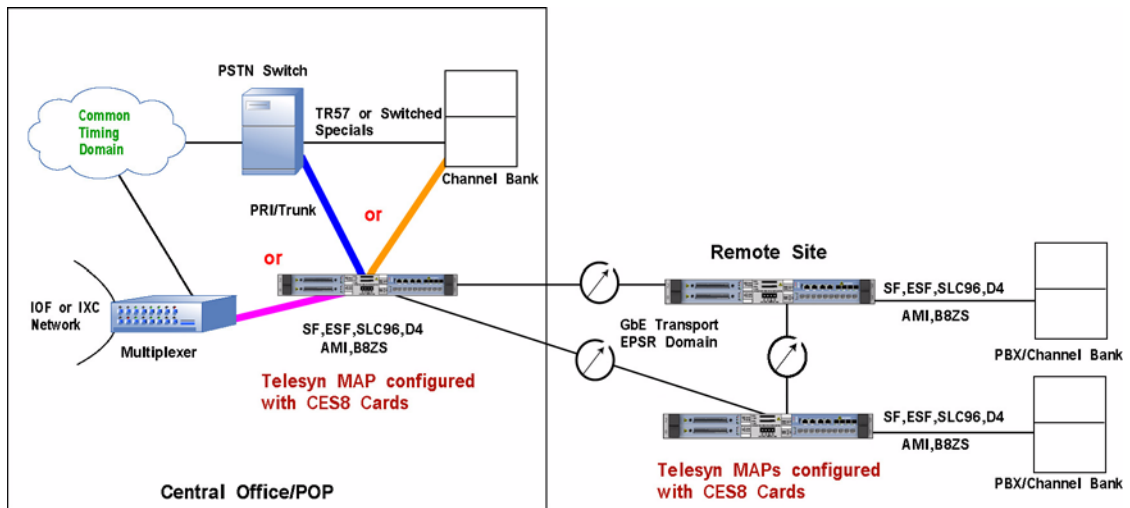


FIGURE 7-1 Sample Customer Network using CES8 cards

The CES8 card allows users of the fMAP to continue to use their legacy DS1/E1 interfaces and provide DS1/E1-based services to their subscribers; while shifting to an Ethernet-based infrastructure for the rest of their network. This requires an “interworking function” between the Ethernet world and the Digital Carrier world. The CES8 card provides interworking functionality.

7.2 Concepts and Terms

The following concepts and terms will be used when describing how the CES8 and NTE8 provides services that include DS1/E1

7.2.1 Common Terms

- **DS1/E1 Port** - The physical, external port on the DS1-related card. The format (DS1 or E1) is configured at a card level. Generally this models things at the line-level parameters on this entity.
- **DS0** - A timeslot within the DS1/E1 path that transports 64Kbps of bandwidth.
- **Reference Clock** - The signal the system used for TDM timing synchronization. Regardless of the input frequency, it is divided down into an 8000Hz timing signal in the fMAP. Reference clock signals may originate from a physical Digital Carrier Port or may be derived from a Pseudo-span using Adaptive Timing. If required, there is also an internal oscillator that provides a timing reference, a Stratum 4 interface.
- **FDL Facility Data Link** - The 4000 bps overhead channel embedded in an Extended Super Frame (ESF) format DS1 signal. It can be used to convey both bit-oriented and message-oriented signals between DS1 nodes.
- **Loopback**: A configuration of the Framer and/or Line Interface Unit hardware where the data sent in one direction is “looped back” in the direction from which it was received. This can be done at the path level or the line level. This is typically used to diagnose / troubleshoot problems with an interface.
- **AIS** - Alarm Indication Signal. The DS1 AIS is an **unframed** “all ones” signal.
- **RAI** - Remote Alarm Indication. When a DS1 sink detects LOS or it receives a signal for which framing cannot be found (e.g. AIS), it forces a zero into the second bit of each channel of the 24 channel structure. RAI is a **framed** signal.

7.3 Provisioning CES8

7.3.1 CES8 Terms

- **Pseudo-span** - A logical interface that emulates a bi-directional TDM circuit using a stream of packets. The stream is transmitted to another node in the network where it is converted back into a TDM signal. This is sometimes also referred to as a “pseudo-wire”. The Pseudo-span can be one of two types with respect to knowledge of the framing structure of the content being carried:
 - Unstructured - no knowledge of frame format. Also called Structure Agnostic
 - Structured - has knowledge of the frame format. Also called Structure Aware (Not supported in this release.)
- **IP interface** - Each CES8 card supports **one** VLAN for transmitting Pseudo-span packets up into the network. When the user interacts with this entity from the CLI, it will be seen as an interface with a name, for example, “vlan:403.0”. For an IP interface to communicate on the network it requires an IP address. The association of an IP address with an IP interface allows the CES8 card to use CLI tools such as PING, ARP, and TRACEROUTE.
- **DS1 Cross-Connect** (i.e. *pass-thru connection*) - The connection between a DS1 or E1 port and an unstructured Pseudo-span. In this case, the content of the framed DS1 is not interpreted. The bits are simply passed through to the destination.
- **Digital Carrier Port** - The physical, external port on the CES8 card. At the CLI, the user will see this as a DS1 or E1 port (depending on a card-level mode setting).

Note: Note that the term DS1 port will be used throughout the remainder of this section for simplicity. All references, except where noted, apply to E1 port as well.

7.3.2 CES Connectivity in Release 6.0

In release 6.0, “Pass-thru” Circuit Emulation Service for both E1 and DS1 circuits is supported on the CES8 SM. In Pass-Thru, the whole DS1/E1 frame is passed. The fMAP does not terminate/interpret the FDL in Pass-Thru. The only layer terminated at the fMAP is the line layer. See [Figure 7-2](#).




FIGURE 7-2 CES8 Point to Point Switching

7.3.3 Provisioning Model

Referring to [Figure 7-3](#), these are the provisioning entities involved in establishing a CES connection in release 6.0. Each entity has specific attributes that are parameters used in CLI commands that must be configured correctly in order for CES to function properly.




FIGURE 7-3 CES8 Provisioning Model

7.3.4 PSPAN

IP-based encapsulation, following the IETF Structure-Agnostic TDM over Packet (SAToP) standard is supported, including the user-provisioned option of running Real-Time Protocol (RTP) in addition to the regular UDP/IP header. The SAToP packet stream is referred to as a “pseudo-span” (PSPAN). The user creates, and can delete, up to 8 pseudo-span interfaces on a CES8 card. Each pseudo-span interface is hosted by (layered on) an IP interface.

Note: In summary, one IP Interface on one VLAN per CES8 card is supported in release 6.0.

Two packet formats are supported in release 5.0:

- SAToP over IPv4/UDP with RTP
- SAToP over IPv4/UDP without RTP

The user may provision the pseudo-span to include an RTP header. The RTP header is defined in detail in RFC3550. RTP is required in order to support adaptive timing so that a pseudo-span may be used as a timing reference.

Variable-sized payloads are supported. The user provisions the number of bytes of payload that the system collects before sending a packet.

To ensure proper packet prioritization, the user may configure the pseudo-span's 802.1p VLAN prioritization bit (p-bit, VPRIORITY) value and the IP DSCP value (when using IP encapsulation). **Default** values for DS1 are:

VPRIORITY=6, IPDSCP=46, JITTERBUFFER=6000, NUMBYTES=192, RTP=ON

For an E1, the only difference is NUMBYTES=256.

Note that these values are used only for the TDM payload packets. Other packets that are generated by the card (e.g. for ARP, PING, TRACEROUTE) will have p-bit (VPRIORITY) and DSCP set to 0.

UDP port numbers are used to uniquely identify the pseudo-span when received at the IP Interface on the card. The values are assigned by the user and are in the range 49152-65535. This is the dynamic range defined by Internet Assigned Numbers Authority (IANA).

Note: The fMAP does not interpret the content of the framed DS1. The bits are simply passed-through to the destination end of the emulated circuit.

Table 7-1 lists the pseudo-span attributes.

TABLE 7-1 Pseudo-span Attributes - Default is in Bold

Attribute	Values / Range	Description
Pseudo-span ID	0 to 127 (Only 8 can be created)	Id within the IP interface.
Interface Name	n/a	The name of this interface.
Admin State	UP DOWN	Refer to 4.1
Encapsulation	n/a	In release 5.0, only SAToP is supported.
INTERFACE	VLAN:id or id or ifname	IPADDRESS (vs. INTERFACE) can only be used if the IP address is unique. Note that the interface id of the pspan and the interface name can be used interchangeably. The format of the pspan's interface id is "pspan:vid.x.y" where "vid.x" is the id of the IP interface that the pspan is stacked on (e.g. vlan:402.0) and "y" is the PSPANID from the CREATE PSPAN command.
IPADDRESS	n/a	Matches the ipaddress for this CES8 card.
UDPPORT	udp-port 49152..65535 in Release 6.0	UDPPORT - The UDP port of the near end interface, the local receive ID. Must be unique within an IP address on a card. Selected by the user. This value is placed in the UDP source port for packets that are transmitted and is expected in the UDP destination port for packets that received for this pseudo-span.
PEERIPADDRESS	ipaddress	PEERIPADDRESS - Matches the IP address of the IP interface of the PEER's CES8 card.

TABLE 7-1 Pseudo-span Attributes - Default is in Bold (Continued)

Attribute	Values / Range	Description
PEERUDPPORT	udp-port 49152..65535 in Release 6.0	PEERUDPPORT - Matches with the peer's UDP-PORT attribute, the local transmit ID. Must be unique within the IP address on the PEER's card. Selected by the user. This value is placed in the UDP destination port for packets that are transmitted and is expected in the UDP source port for packets that are received for this pseudo-span.
Actual Received Indication Actual Transmitted Indication		Status of the remote/local port the PSPAN is connected to. For example, Local Loss Carrier could mean the local port is unplugged. Refer to 18.14 .
NUMBYTES	16..1023	NUMBYTES - The number of bytes per packet, DS1 - 192, E1 - 256 .
JITTERBUFFER	value 6000 usec (default)	JITTERBUFFER - The packet delay variation that the PSPAN should be able to absorb. A value of 6000 usec means the PSPAN can nominally absorb +/- 6000 usec of jitter. (The size of the jitter buffer is usually larger because the hardware works in integral numbers and rounds up to the power of 2. The actual size of the configured jitter buffer using SHOW PSPAN.) <i>Note: In certain cases, the average fill level may be too low or high and the PSPAN should be disabled/enabled. Refer to 9.6 for more on jitter buffer statistics.</i>
TIMINGREFERENCE	SELF or CONNECTION or CARD	TIMINGREFERENCE - This parameter controls the interface timing reference the PSPAN will use.. The card-level timing reference is a common reference for all interfaces on the card.
RTP	ON or OFF	RTP - This parameter determines whether RTP is ON or OFF. This determines whether the RTP header is included or not.
VPRIORITY	0..7 (6)	VPRIORITY - 802.1p priority bit setting. <i>Note: The default (6) should not be changed, since CES traffic should be a high priority. (EPSR traffic is given a level of 7.)</i>
IPDSCP	0..63 (46)	IPDSCP - Specifies the DSCP (Differentiated Services Code Point) value.

Note: Critical values must be matched on both sides of the PSPAN in order for the circuit to operate correctly. They are:

- *PEERIPADDRESS* - The IP address of the distant end of the emulated circuit.
- *PEERUDPPORT* - The UDP Port of the distant end of the emulated circuit.
- *RTP* - Whether the RTP header is included or not at both ends of the emulated circuit.
- *NUMBYTES* - The number of bytes per packet.

7.3.5 CES8 Card Attributes

CES8 card attributes, such as Admin State, Preferred Load, etc., are similar to other Service Modules. See the section 4.1 for more information.

The following table lists card attributes that are specific to the CES8 card.

TABLE 7-2 CES8 Card Attributes - Default is in Bold

Attribute	Values / Range	Description
PORTTYPE	DS1 E1	<p>PORTTYPE - DS1 or E1. The PORTTYPE parameter specifies the type of port, either DS1 or E1, to be utilized for the card. All ports on the card will be of the same type.</p> <p>Either DS1 or E1.</p> <ul style="list-style-type: none"> - DS1 and E1 ports cannot be provisioned on the same card - CES8 SMs provisioned to support DS1 and E1 can be configured in the same fMAP shelf. - The user must DISABLE the card to change the mode. - Changing the PORTTYPE effectively destroys the card and creates a new card with the new port types.
TIMINGREFERENCE	type:id or ifname or INTERNAL	TIMINGREFERENCE - A common reference for all interfaces on the card that are configured with TIMINGREFERENCE=CARD.

To modify the PORTTYPE, the user must DISABLE the card, then change the mode using the SET CARD command. Changing the PORTTYPE effectively destroys the card and creates a new card with the new port types. Due to the critical nature of changing the PORTTYPE, the command provides a clear warning to the user that all of the provisioning on the specified card (e.g. physical ports, pseudo-spans, services, PM thresholds, etc.) will be lost when the mode is changed.

7.3.6 CES8 Port

The port on the CES8 card is the connecting point for the DS1 or E1 link. 8 ports are provided on the card. Each port is an individual provisionable entity. The following table lists DS1PORT or E1PORT attributes.

TABLE 7-3 DS1 or E1 Port Attributes - Default is in Bold

Attribute	Values / Range	Description
Admin State	UP DOWN	Refer to 4.1
TIMINGREFERENCE	SELF or CONNECTION or CARD	TIMINGREFERENCE - This parameter controls the interface-level timing reference the port will use. The card-level timing reference is a common reference for all interfaces on the card
LINEENCODING	B8ZS or AMI or HDB3	LINEENCODING - This parameter specifies the line encoding scheme to be employed for this port. DS1 - B8ZS or AMI E1 - AMI or HDB3
LINEBUILDOUT	LONGHAUL = 0.0DB -7.5DB -15.0DB -22.5DB SHORTHAUL = 133FT 266FT 399FT 533FT 655FT	LINEBUILDOUT - This parameter indicates that the port line build out is being set. - LONGHAUL - This parameter specifies that the LINEBUILDOUT being specified will be for a longer length and will be specified in DB. - SHORTHAUL - This parameter specifies that the LINEBUILDOUT being specified will be for a shorter length and will be specified in FT. Default is 0.0DB . <i>Note: These settings are only used for the DS1.</i>
LOOPBACK	NONE INWARD LINE	LOOPBACK - This parameter specifies the loop-back setting for the port.
DESCRIPTION	description	A text description of the port.

7.3.7 IP Interface (IP Address/VLAN)

In release 6.0, each CES8 card can be assigned to one VLAN. Also, each card is assigned one IP address. The association of the IP address with the VLAN results in an entity referred to as an *IP Interface*. The following table lists IP Interface attributes.

TABLE 7-4 IP Interface Attributes - Default is in Bold

Attribute	Values / Range	Description
IP Source Address	ip address range	Used as the Source address for CES.
VLAN interface	n/a	For example, VLAN:402.0. <i>Note: These VLANs must be STD and not UFO.</i>
Interface Name	n/a	The interface name.
Slot number	n/a	The slot number occupied by the card.
Subnet Mask	n/a	Subnet Address.
Default Gateway	ip address range	The IP address of a gateway device.

Note: In release 6.0, there is one IP Interface configured per card. Since each card can have only one IP address, configuring two IP Interface on the same CES8 card (to create a hairpin connection on one card) is not supported. For an example of two cards being used in one device, refer to [Figure 7-5](#), connection B. At the PSPAN level, the IP address (IPSA) and peer IP address (IPDA) cannot be the same. Finally, the address 127.0.0.1 (loopback network connection) cannot be used.

7.3.8 DS1/E1 Connection (PSPAN<->DS1)

A connection is the link between a DS1 or E1 physical port and an unstructured Pseudo-span. [Table 7-5](#) list connection attributes:

TABLE 7-5 DS1/E1 Connection Attributes - Default is in Bold

Attribute	Values / Range	Description
Interface ID for DS1/E1 port	n/a	The interface ID.
Interface ID for Pseudo-span	Name or ID	Must exist on the same slot as the DS1/E1 port.

7.3.9 8KHz Timing References

As indicated in the [Figure 7-4](#) below, the timing source for the PBX/Channel Banks served by the Remote Site fMAPs is the PSTN Switch.

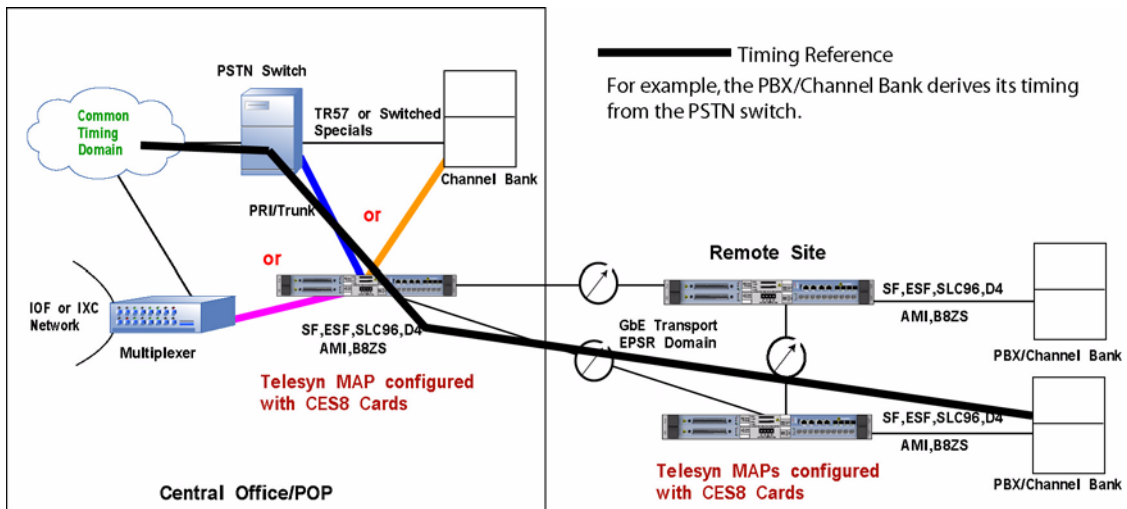


FIGURE 7-4 Timing relationship

Each interface, DS1/E1 port or Pseudo-span, must be driven by a timing reference which provides the clocking signal for receiving and transmitting data.

In release 6.0, two types of timing reference settings are supported:

- a “card-level” timing reference setting, and
- a “per-interface” timing reference setting.

7.3.9.1 Card Level

The “**card-level**” timing reference is a common reference for all interfaces on the card. One timing reference, any interface on the CES8 card, can be designated as the common timing reference and can be distributed to any other interface on the CES8 card. There is only one card-level timing reference per CES8 card.

The card-level timing reference may be one of the following sources:

- A “self-timed” DS1/E1 port physical interface
- A “self-timed,” connected Pseudo-span (using RTP-based derived, adaptive timing)
- The internal oscillator (locked to a timing signal from the active CFC)

Note that the internal oscillator is a 10ppm oscillator that provides the required line frequency stability of +/-32 ppm for DS1 and +/-50 ppm for E1.

The card-level timing reference default is the internal oscillator. The internal oscillator will be locked to a system-wide clock reference provided by the active CFC. This ensures that the internal oscillator(s) are in sync across all cards in the system. CES8 failure detection circuitry will detect whether the clock is present. If the

clock is not present then the CES8 internal oscillator will run free and attempt to preserve the last good frequency it was locked to.

7.3.9.2 Per-Interface

Each interface has a designated “**per-interface**” timing reference. An interface can have one of 3 sources for its timing reference:

- Itself
- The interface to which it is connected
- The “card-level” timing reference (Explained in [7.3.9.1](#))

Note: A pseudo-span must be connected to a DS1/E1 port, using the CONNECT command, and running RTP protocol, to be used as a timing reference.

By default, a DS1/E1 port derives a timing reference from the Card level.

By default, a PSPAN derives a timing reference from the Card level.

Note: This is a change from release 6.0, when the default for the DS1/E1 port was itself, and the default for the PSPAN was its connection. Existing values, however, will not be affected.

In release 5.0, there is no support for a secondary timing reference assignment. The interface will always fall back to the card-level timing reference if the designated interface is no longer an appropriate timing reference (e.g. because a DS1 interface detects LOS).

There are restrictions on the combinations of interface-level timing references that are allowed on a CES8 card. Below is the allowable timing configurations of connected interfaces:

TABLE 7-6 Allowable timing configuration of connected interfaces

DS1	PSPAN
Self	Connected Interface
Connected Interface	Self
Card	Card

The timing reference of a connected interface cannot be changed. The connection to the specified interface must first be removed. Then the timing reference configuration can be changed.

Timing references are not entities, but simply attribute(s) on interface(s). As such they have neither operational state nor admin state.

Failure of a provisioned timing reference will result in a **degraded** alarm on all interfaces timed from that reference. If the internal oscillator fails, all interfaces timed from the internal oscillator will have a **failed** alarm.

7.3.10 Commands and Usage Notes

7.3.10.1 SHOW PSPAN

The following system output shows the results of the `SHOW PSPAN` command for PSPAN c19p0. Note that the PSPAN has a Timing-reference degraded alarm. Due to the alarm, it's Status is degraded.

```
officer SEC>> SHOW PSPAN=c19p0
```

```
-----
Fault:                                     Severi ty
-----
PSPAN Timing-reference degraded           Minor

Interface..... PSPAN: [100. 19. 19]
Name..... c19p0
Description..... <none>
Admin State..... UP
Operational State..... UP
Status..... Degraded
Timing Reference..... CONNECTION
Encapsulation..... SAToP over IPv4
  IP Address..... 2. 2. 2. 3
  UDP Port..... 61000
  Peer IP Address..... 2. 2. 2. 2
  Peer UDP Port..... 61000
  Bytes per packet..... 256
  RTP..... ON
  Requested Jitter Buffer Size..... 6000
  VLAN Priority..... 6
  IP DiffServ Code Point..... 46
Actual Received Indication(s)..... Remote Loss Carrier
Actual Transmitted Indication(s)..... Local Loss Carrier
Jitter Buffer Statistics
  Average Fill Level..... 5500
  Actual Jitter Buffer Depth..... 11000
  Minimum Fill Level..... 5500
  Maximum Fill Level..... 5500
-----
```

7.3.10.2 Connection-related Configuration CLI Commands

```
CONNECT INTERFACE=if-name TO=if-name
```

Connects a PSPAN to a DS1/E1 port. Note that the connection must be present before the `SHOW CONNECTIONS` command will display the connection.

`DISCONNECT INTERFACE=if-name` removes a connection between a PSPAN and a DS1/E1 port.

`SHOW CONNECTIONS [INTERFACE=type:id-list|id-name-list|ALL]` shows connections among interfaces in the system. You can show it by either end of the connection. All current connection types are bi-directional.

```
officer SEC>> SHOW CONNECTIONS
```

```
--- Connections -----
DS1: 6. 0          <----->          PSPAN: 100. 6. 0  c6p0
E1: 19. 0         <----->          PSPAN: 100. 19. 19 c19p0
-----
```

```

officer SEC>> SHOW CONNECTIONS | INTERFACE=DS1: 6. 0
--- Connections -----
DS1: 6. 0                <----->                PSPAN: 100. 6. 0   c6p0
-----
    
```

7.3.10.3 Existing Commands

Command	Description
SHOW PORT	This shows the DS1/E1 port-specific information, such as line build-out, associated timing reference, etc. The command syntax is the same as the existing SHOW PORT command, but the output is different. Note that if it was an E1PORT, it would indicate that in the header. Also, some attributes (e.g. LBO) will be different. <i>Note: In release 5.0, line type cannot be configured; therefore, it is not displayed.</i>
SHOW INTERFACE	The SHOW INTERFACE command shows the DS1/E1 and PSPAN specific information. The command syntax is the same as the current SHOW INTERFACE, but the output is different.
SHOW ALARMS	The SHOW ALARMS command shows kinds of alarms that are DS1/E1 specific.
SHOW PROFILE	The SHOW PROFILE shows DS1/E1 profile information.
SET CARD	The SET CARD command supports provisioning a PORTTYPE of DS1 or E1 of the CES8 card.
SET PORT	The SET PORT command supports the provisioning of DS1/E1 port-specific information.
SET INTERFACE	The SET INTERFACE command allows the user to turn counters on/off and set PM threshold values.
SET PROFILE	The SET PROFILE allows the user to set DS1/E1 profile information.

7.4 Example CES8 Configurations

Figure 7-5 shows a sample configuration that shows how DS1 circuits can be configured. Included are the values the PSPAN and IP Interfaces.

This figure includes two possible CES8 configurations:

- A connection within the same shelf, with timing off of the Card (internal oscillator). This is the connection labeled B in the figure.
- A connection across a network, with timing off an external source. This is the connection labeled A in the figure.

The provisioning sequences for these two scenarios are:

- **Card** (especially for port-type, in this case DS1)
- **Port** - This uses the attributes that apply for DS1.
- **VLAN/IP Address** - This sets up the VLAN used (110) and the IP Interfaces for each card.
- **PSPAN** - This sets up the PSPAN for each end of the circuit and the Peer (**other end**) IP Interfaces and UDP ports. **These must match each other for the circuit to work properly.**
- **Connection** - This ties together the PSAN and the DS1 so the circuit can begin to carry traffic.

Since Connection B is the simpler configuration, this is shown first.

7.4.1 Same Shelf, Card Timing (Internal Oscillator)

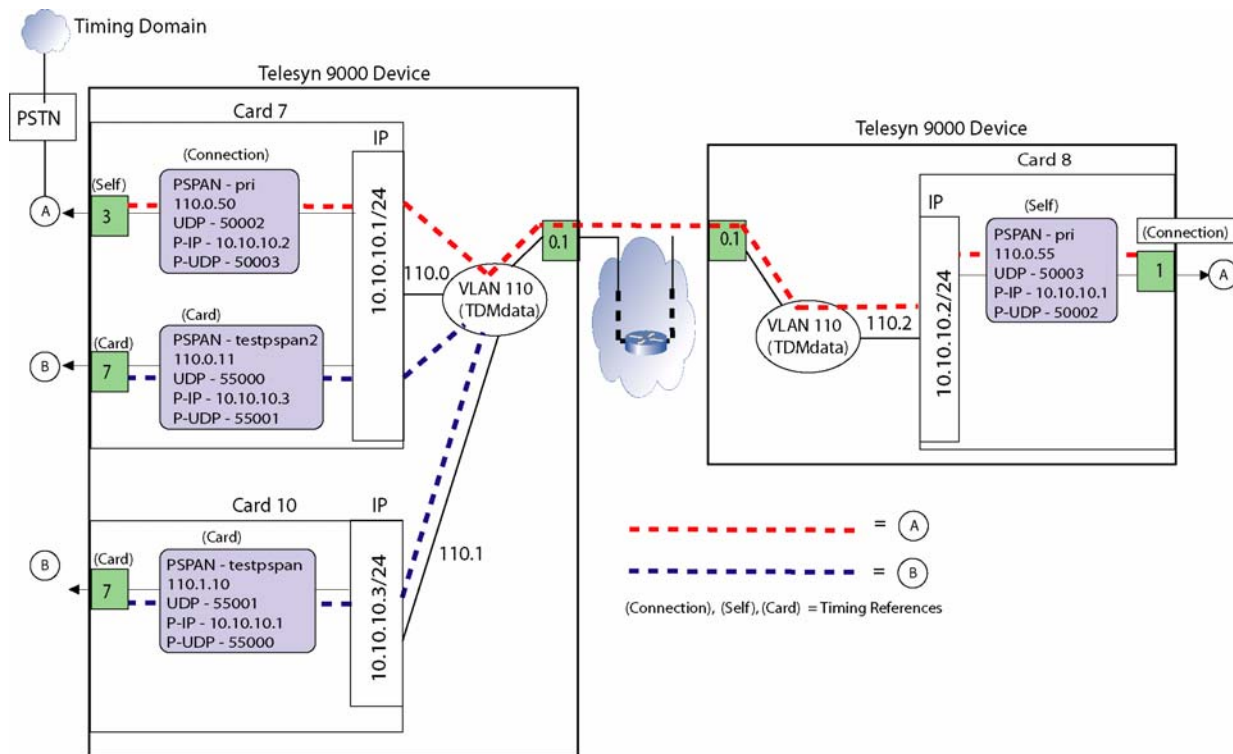


FIGURE 7-5 Sample CES Configuration

7.4.1.1 Card

The CES8 cards at each end are provisioned, using the CREATE CARD command. Attributes such as the load and profile are included to make the card ready to provide service.

Note: As mentioned in 7.3.5, to modify the PORTTYPE, the user must DISABLE the card, then change the mode using the SET CARD command. Since the default is DS1, this does not need to be changed.

The user could set the timing reference for one of the cards to the internal oscillator (INTERNAL). This means that for the PORT and PSPAN interfaces, the only timing reference is the internal oscillator.

```
SET CARD=7 CES8 TIMINGREFERENCE=INTERNAL
SET CARD=10 CES8 TIMINGREFERENCE=INTERNAL
```

7.4.1.2 Port

The user could configure the DS1 ports (for each end) as:

```
SET PORT=7.7 DS1 TIMINGREFERENCE=CARD LINEENCODING=B8ZS LINEBUILDOUT LONGHAUL=-7.5DB
SET PORT=10.7 DS1 TIMINGREFERENCE=CARD LINEENCODING=B8ZS LINEBUILDOUT LONGHAUL=-7.5DB
```

At this point the DS1 ports can report faults and a loopback could be set up at each end using the SETLOOPBACK parameter.

7.4.1.3 VLAN/IP Address

Creating the VLAN (in this case 110) and associating it with an IP Interface on the CES8 card are done using the following commands:

Note: The creation of the VLAN in the fMAP(s) that are used for the DS1 connection can be done in advance, and in fact this is encouraged since configuring network VLANs helps the user plan out the network and know which nodes associated with the VLAN need to be configured.

For each node in the figure, the IP Interface would use these commands:

```
ADD IP INTERFACE=vlan:110.0 IPADDR=10.10.10.1 Subnet=255.255.255.0 CARD=7
IFNAME=IP_End_7
ADD IP INTERFACE=vlan:110.1 IPADDR=10.10.10.3 Subnet=255.255.255.0 CARD=10
IFNAME=IP_End_10
```

Note: If the IP interfaces are on the same shelf, each IFNAME must be unique.

7.4.1.4 PSPAN

Referring to [Figure 7-5](#), a PSPAN could be created by using the following commands, one for each node:

```
CREATE PSPAN=testspan2 PSPANID=11 SATOP INTERFACE=vlan:110.0 UDPPORT=55000 PEER-  
ADDRESS=10.10.10.3 PEERUDPPORT=50001 NUMBYTES=193 JITTERBUFFER=6000 TIMINGREFER-  
ENCE=CARD RTP=ON VPRIORITY=6 IPDSCP=46
```

```
CREATE PSPAN=testspan PSPANID=10 SATOP INTERFACE=vlan:110.1 UDPPORT=55001 PEER-  
ADDRESS=10.10.10.1 PEERUDPPORT=55000 NUMBYTES=193 JITTERBUFFER=6000 TIMINGREFER-  
ENCE=CARD RTP=ON VPRIORITY=6 IPDSCP=46
```

Note: The **VPRIORITY** and **IPDSCP** values are the defaults and therefore do not need to be entered. As mentioned in [Table 7-1](#), these values should not be changed.

7.4.1.5 Connection (PSPAN/DS1)

The user could connect the PSPAN and DS1 interface for each node as follows:

```
CONNECT INT DS1:7.7 TO testspan2  
CONNECT INT DS1:10.7 TO testspan
```

7.4.2 Across a Network, External Timing

In this configuration, an external timing reference is used and is connected through port 7.3 in the network end 9000 device. To have the entire path follow this timing, the timing references are as follows:

- The network-end port (7.3) that is connected to the clocking reference is set to **SELF**, since that will be the reference for the rest of the path.
- The PSPAN associated with 7.3 is set to **CONNECTION**, meaning it will get its clocking from the 7.3 port.
- The PSPAN on the subscriber end (associated with port 8.1) will have its timing set to **SELF** since it receives the RTP from the network end and so needs to be the source clock for the subscriber end.
- The associated port (8.1) is set to **CONNECTION**, since it receives its timing from the PSPAN.

These timing reference values are included in [Figure 7-5](#).

7.4.2.1 Card

The CES8 card at the network end is set with a timing reference of **INTERNAL**; in this case this will be the backup timing reference if the clocking to the external reference fails.

(network-side device)

```
SET CARD=7 CES8 TIMINGREFERENCE=INTERNAL
```

(subscriber-side device)

```
SET CARD=8 CES8 TIMINGREFERENCE=INTERNAL
```

7.4.2.2 Port

The user would configure the DS1 ports as:

(network-side device)

```
SET PORT=7.3 DS1 TIMINGREFERENCE=SELF LINEENCODING=B8ZS LINEBUILDOUT LONGHAUL=-7.5DB
```

(subscriber-side device)

```
SET PORT=8.1 DS1 TIMINGREFERENCE=SELF LINEENCODING=B8ZS LINEBUILDOUT LONGHAUL=-7.5DB
```

7.4.2.3 VLAN/IP Address

In this case, the VLAN 110 is created on both the network-side and subscriber-side devices; moreover, the VLAN must be associated with the GE ports as well as the CES8 cards.

Note: Since this is a VLAN that crosses devices, it is termed a Network VLAN. Network VLANs should be planned out, and the Allied Telesis NMS provides an easy method to provision these and associate them with the correct ports. Refer to the Allied Telesis NMS Administration Guide for more details.

For each node in the figure, the IP Interface would use these commands:

(network-side device)

```
ADD IP INTERFACE=vlan:110.0 CARD=7 IFNAME=IP_End_PRI IPADDR=10.10.10.1 Sub-
net=255.255.255.0
```

(subscriber-side device)

```
ADD IP INTERFACE=vlan:110.1 CARD=8 IFNAME=IP_End_PRI IPADDR=10.10.10.2 Sub-
net=255.255.255.0
```

7.4.2.4 PSPAN

Referring to [Figure 7-5](#), a PSPAN could be created by using the following commands, one for each node:

(network-side device)

```
CREATE PSPAN=pri PSPANID=50 SATOP INTERFACE=vlan:110.0 UDPPORT=50002 PEERIPAD-
DRESS=10.10.10.2 PEERUDPPORT=50003 NUMBYTES=193 JITTERBUFFER=6000 TIMINGREFER-
ENCE=CARD RTP=ON VPRIORITY=6 IPDSCP=46
```

(subscriber-side device)

```
CREATE PSPAN=pri PSPANID=55 SATOP INTERFACE=vlan:110.0 UDPPORT=50003 PEERIPAD-
DRESS=10.10.10.1 PEERUDPPORT=50002 NUMBYTES=193 JITTERBUFFER=6000 TIMINGREFER-
ENCE=CARD RTP=ON VPRIORITY=6 IPDSCP=46
```

Note: The VPRIORITY and IPDSCP values are the defaults and therefore do not need to be entered. As mentioned in [Table 7-1](#), these values should not be changed.

7.4.2.5 Connection (PSPAN/DS1)

The user could connect the PSPAN and DS1 interface for each node as follows:

(network-side device)

```
CONNECT INT DS1:7.3 TO pri
```

(subscriber-side device)

```
CONNECT INT DS1:8.1 TO pri
```

7.5 Engineering the Packet Network for CES

Many characteristics of packet networks are not compatible to constant bit rate services, and so engineering of the components in the packet network is important. Packet network congestion, blocking, QoS prioritization, and multiple paths with varying latency can lead to:

- jitter
- lost packets
- duplicate packets
- packets out of sequence.

All of these error conditions must be engineered into the CES8 configuration; some of these are done automatically by the card, while some are done through configurable parameters.

Refer to the *fMAP Services Guide* for more information on these engineering issues.

7.6 Command Summary for CES8 Card/Interface

Note: The commands listed here apply specifically to CES8 components. For commands that apply to components common to CES8 and NTE8, refer to .

TABLE 7-7 CES8/CES Provisioning Commands

Object	Verb	Syntax	Description
CARD CES8	CREATE	<pre> CREATE CARD=slot CES8 [{ [PREFLOAD=filename] [ADMINSTATE={ UP DOWN }] [PORTTYPE={ DS1 E1 }] PROFILE=name }] </pre>	<p>Creates the CES8 card.</p> <p>Refer to Table 7-2 for more information on these attributes.</p>
CARD	SET	<pre> SET CARD={ slot-list ACTCFC INACTCFC } { PREFLOAD={ filename NONE } ALTLOAD={ filename NONE } TEMPLOAD={ filename NONE } CES8 [PORTTYPE={ DS1 E1 }] [TIMINGREFERENCE={ type:id ifname INTERNAL }] } </pre>	<p>Allows the user to modify card attributes, here the ones for the CES8 are set.</p> <p>Refer to Table 7-2 for more information on these attributes.</p>

TABLE 7-7 CES8/CES Provisioning Commands (Continued)

Object	Verb	Syntax	Description
PROFILE	SET	<pre>SET PROFILE=name CES8 [PREFLOAD=filename] [ADMINSTATE={ UP DOWN }] [PORTTYPE={ DS1 E1 }]</pre>	<p>Allows the user to modify the card profile. For the CES8, this changes the PORTTYPE.</p> <p>Refer to Table 7-2 for more information on these attributes.</p>
PROFILE	SHOW	<pre>SHOW PROFILE=name { CES8 DS1PORT E1PORT }</pre>	<p>Allows the user to display the card profile for the CES8 card(s) or the CS* card port type.</p>
PROFILE NAMES	SHOW	<pre>SHOW PROFILE NAMES [{ CES8 DS1PORT E1PORT }]</pre>	<p>Allows the user to display card profile names associated with the CES8 card or CES8 port type.</p>
PSPAN	CREATE	<pre>CREATE PSPAN=pspanname PSPANID=0..127 SATOP { INTERFACE={ VLAN:id id ifname } IPADDRESS=ipaddress } { UDPPORT=49152..65535 } { PEERIPADDRESS=ipaddress } { PEERUDPPORT=49152..65535 } [NUMBYTES=16..1023] [JITTERBUFFER=value] [TIMINGREFERENCE={ SELF CONNECTION CARD }] [RTP={ ON OFF }] [VPRIORITY=0..7] [IPDSCP=0..63]</pre>	<p>Allows the user to create a Pseudo-span and provision its attributes.</p> <p>Refer to Table 7-1 for more information on PSPAN attributes.</p>

TABLE 7-7 CES8/CES Provisioning Commands (Continued)

Object	Verb	Syntax	Description
PSPAN	DESTROY	<pre>DESTROY PSPAN={ pspanname-list ALL } [INTERFACE={ type:id-range ifname-list ALL }] [FORCE]</pre>	<p>Allows the user to destroy a Pseudo-span and all its provisioning.</p> <p>Note: FORCE will automatically disconnect any connections to the pseudo-span and clear any use of the pseudo-span as a timing reference. Then it will destroy the pseudo-span without any confirmation.</p>
PSPAN	DISABLE	<pre>DISABLE PSPAN={ pspanname-list ALL } [{ INTERFACE={ type:id-range ifname-list ALL } CARD={ slot-list ALL } }]</pre>	<p>Allows the user to disable a Pseudo-span.</p>
PSPAN	ENABLE	<pre>ENABLE PSPAN={ pspanname-list ALL } [{ INTERFACE={ type:id-range ifname-list ALL } CARD={ slot-list ALL } }]</pre>	<p>Allows the user to enable a Pseudo-span.</p>

TABLE 7-7 CES8/CES Provisioning Commands (Continued)

Object	Verb	Syntax	Description
PSPAN	SET	<pre> SET PSPAN={ pspanname-list ALL } SATOP [UDPPORT=49152..65535] [PEERIPADDRESS=ipaddress] [PEERUDPPORT=49152..65535] [NUMBYTES=16..1023] [JITTERBUFFER=value] [TIMINGREFERENCE={ SELF CONNECTION CARD }] [RTP={ ON OFF }] [VPRIORITY=0..7] [IPDSCP=0..63] </pre>	<p>Allows a user to modify Pseudo-span attributes.</p> <p>Refer to Table 7-1 for more information on PSPAN attributes.</p>

TABLE 7-7 CES8/CES Provisioning Commands (Continued)

Object	Verb	Syntax	Description
PSPAN	SHOW	<pre> SHOW PSPAN [= { pspanname-list ALL }] [{ INTERFACE= { type:id-range ifname-list } CARD=slot-list }] [{ JITTERBUFFER FULL }] </pre>	<p>Allows the user to display a Pseudo-span. SHOW PSPAN will allow the user to display one or more pseudo-spans by:</p> <ul style="list-style-type: none"> - name (unique in the system) - interface (may return multiple pseudo-spans) - source IP (may return multiple pseudo-spans) <p>This does not allow the user to display them by destination, encapsulation type, support of RTP, etc.</p> <p>By default, the command shows the attributes of the pseudo-span.</p> <p>JITTERBUFFER shows the current value in usec.</p> <p>Use the SHOW CONNECTIONS command to show connections on the pspan.</p> <p><i>Note: If two users try to access PSPAN attributes at the same time, there may be an error message. Refer to 9.6.3.</i></p>
INTERFACE	CONNECT	<pre> CONNECT INTERFACE= { type:id ifname } TO= { type:id ifname } </pre>	<p>Allows the user to connect two interfaces.</p>
INTERFACE	DISCONNECT	<pre> DISCONNECT INTERFACE= { type:id-range ifname-list ALL } </pre>	<p>Allows the user to disconnect previously connected interfaces.</p>

TABLE 7-7 CES8/CES Provisioning Commands (Continued)

Object	Verb	Syntax	Description
CONNECTIONS	SHOW	<pre> SHOW CONNECTIONS [INTERFACE={ type:id-range ifname-list ALL }] [FULL] </pre>	Allows the user to display connected interfaces.

7.7 Common Command Summary for DS1/E1

TABLE 7-8 Common DS1/E1 Provisioning Commands

Object	Verb	Syntax	Description
PROFILE DS1PORT	CREATE	<pre> CREATE PROFILE=name DS1PORT [ADMINSTATE={ UP DOWN }] [TIMINGREFERENCE={ SELF CONNECTION CARD }] [LINEENCODING={ B8ZS AMI }] [LINEBUILDOUT { LONGHAUL={ 0.0DB -7.5DB -15.0DB -22.5DB } SHORThAUL={ 133FT 266FT 399FT 533FT 655FT } }] [FRAMING={ UNFRAMED SF ESF STANDARD }] </pre>	<p>Creates a Profile for the DS1PORT.</p> <p>The default (AUTOPROV) profile attributes apply to the NTE8 and are:</p> <pre> TIMINGREFERENCE = SELF LINEENCODING = B8ZS LINEBUILDOUT= 0.0DB FRAMING = ESF </pre> <p>Note: - When the AUTOPROV DS1PORT profile is applied to the CES, where FRAMING=UNFRAMED, the system forces the FRAMING attribute to UNFRAMED.</p> <p>If a profile for one type of service is applied to another type, the default values for that other type are used.</p> <p>For example if a profile for the DS1PORT is for CES and is applied to an NTE8, the FRAMING is set for the default NTE8 value (ESF).</p>

TABLE 7-8 Common DS1/E1 Provisioning Commands (Continued)

Object	Verb	Syntax	Description
PROFILE E1PORT	CREATE	<pre> CREATE PROFILE=name E1PORT [ADMINSTATE={ UP DOWN }] [TIMINGREFERENCE={ SELF CONNECTION CARD }] [LINEENCODING={ HDB3 AMI }] [FRAMING={ UNFRAMED E1 E1CRC STANDARD }] </pre>	<p>Creates a Profile for the E1PORT</p> <p>The default (AUTOPROV) profile attributes apply to the NTE8 and are:</p> <p>TIMINGREFERENCE = SELF LINEENCODING = AMI FRAMING = E1CRC</p> <p>Note: - When the AUTOPROV E1PORT profile is applied to the CES, where FRAMING=UNFRAMED, the system forces the FRAMING attribute to UNFRAMED.</p>
PROFILE DS1PORT	SET	<pre> SET PROFILE=name DS1PORT [ADMINSTATE={ UP DOWN }] [TIMINGREFERENCE={ SELF CONNECTION CARD }] [LINEENCODING={ B8ZS AMI }] [LINEBUILDOUT { LONGHAUL={ 0.0DB -7.5DB -15.0DB -22.5DB } SHORThAUL={ 133FT 266FT 399FT 533FT 655FT } }] [FRAMING={ UNFRAMED SF ESF STANDARD }] </pre>	<p>Changes the attributes of the DS1PORT profile.</p> <p>Refer to 4.1.4 on the relationship between components and the changing of attributes for a Profile.</p> <p>Refer to Table 7-3 for more information on DS1 port attributes.</p>

TABLE 7-8 Common DS1/E1 Provisioning Commands (Continued)

Object	Verb	Syntax	Description
PROFILE E1PORT	SET	<pre> SET PROFILE=name E1PORT [ADMINSTATE={ UP DOWN }] [TIMINGREFERENCE={ SELF CONNECTION CARD }] [LINEENCODING={ HDB3 AMI }] [FRAMING={ UNFRAMED E1 E1CRC STANDARD }] </pre>	<p>Changes the attributes of the E1PORT profile.</p> <p>Refer to 4.1.4 on the relationship between components and the changing of attributes for a Profile.</p> <p>Refer to Table 7-3 for more information on E1 port attributes.</p>

TABLE 7-8 Common DS1/E1 Provisioning Commands (Continued)

Object	Verb	Syntax	Description
INTERFACE DS1	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } DS1 [TIMINGREFERENCE={ SELF CONNECTION CARD }] [LINEENCODING={ B8ZS AMI }] [LINEBUILDOUT { LONGHAUL={ 0.0DB -7.5DB -15.0DB -22.5DB } SHORThAUL={ 133FT 266FT 399FT 533FT 655FT } }] [FRAMING={ UNFRAMED SF ESF STANDARD }] [DIRECTION={ NETWORK CUSTOMER }] [DESCRIPTION=description] [FORCE] </pre>	<p>Changes the attributes for the DS1 interface.</p> <p>Refer to 4.1.4 on the relationship between profiles and the changing of attributes for a specific interface.</p>

TABLE 7-8 Common DS1/E1 Provisioning Commands (Continued)

Object	Verb	Syntax	Description
INTERFACE E1	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } E1 [TIMINGREFERENCE={ SELF CONNECTION CARD }] [LINEENCODING={ HDB3 AMI }] [FRAMING={ UNFRAMED E1 E1CRC STANDARD }] [DIRECTION={ NETWORK CUSTOMER }] [DESCRIPTION=description] [FORCE] </pre>	<p>Changes the attributes for the E1 interface.</p> <p>Refer to 4.1.4 on the relationship between profiles and the changing of attributes for a specific interface.</p>
		<pre> LOOPBACK INTERFACE={ type: type:id-range id-range ifname-list ALL } { NEAREND FAREND } { INWARD PAYLOAD LINE PACKET NONE } </pre>	<p>Sets the loopback state of the interface. Loop backs are removed by setting the value to NONE.</p> <p>Defaults to the near-end of the interface. Release 7.0 will only support near-end requests</p> <p>PACKET loopback type is not applicable for the DS1/E1 interface types.</p>

8. Provisioning EPON

8.1 Overview

8.1.1 EPON Configuration

The Ethernet Passive Optical Network (EPON) has the following characteristics:

- It is a point-to-multipoint protocol, so one physical port (on the fMAP) emits packets to multiple Ethernet ports (at the CPE-side) without active electronics in between.
- In the downstream direction all the CPE devices on the network can see the signal transmitted by the fMAP, so the CPE must select only the content they are supposed to receive.
- In the upstream direction, only the fMAP port can see the signal transmitted by the CPE, so the fMAP must coordinate the CPE so that they share the medium efficiently.
- Because it is a shared medium, bandwidth has to be explicitly provisioned for each service and each subscriber on the shared medium.

In an example configuration, the EPON2 card is used with the iMG/RG to provide Triple Play. See [Figure 8-1](#).




FIGURE 8-1 Sample Customer Network using EPON cards

8.2 Concepts and Terms

8.2.1 Common Terms

The following concepts and terms will be used when describing how the EPON and iMG/RG work together.

- **OLT (Optical Line Termination)** - This is a piece of equipment located at the central-office end of the EPON. It transmits and receives packets to/from multiple ONUs on the EPON. It also controls bandwidth allocation and OAMP for the ONUs. This is part of the EPON2 card in the fMAP.
- **ONU (Optical Network Unit)** - This is a piece of equipment located at the end-user end of the EPON. It terminates the EPON from the OLT and outputs packets to the user/network interface (UNI). This is part of the iMG/RG.
- **Logical Link** - This is an entity in the system representing a stream of packets on the EPON. The stream may be destined for a unique ONU (point-to-point), or a collection of ONUs (point-to-multipoint). The packets in the logical link are identified by a Logical Link ID (LLID), with a tag placed in the preamble before each frame on the EPON.
- **Upstream / Downstream Data Flows** - Depending on the type of packet, a packet stream (a logical link) may be upstream/downstream or downstream only. The direction determines which attributes apply for the level of service.
- **SLA - (Service Level Agreement)** - In the CLI syntax this is called a **QOSPOLICY**. This is an agreed-upon level of service (min/max bandwidth and delay sensitivity) that must be provided for each traffic flow on the EPON. This traffic flow is on a VLAN basis, so a user specifies an SLA for a particular ONU per VLAN.

8.2.2 EPON Connectivity in Release 8.0

In release 8.0, the EPON card can be used with the fMAP to provide Triple-Play. See [Figure 8-2](#).

- iMG/RG - The RG provides the customer connections and is configured using the AlliedView NMS.
- ONU - The ONU is co-located in an enclosure with the iMG/RG, so for the customer it is one piece of equipment.

Note: In release 8.0, the number of ONUs per EPON2 port is 16. This will be increased to 32.

- VLANs - To provide customer as well as configuration services, usually five VLANs are used. (A sixth VLAN is available for an additional customer service):
 - An untagged VLAN for the iMG/RG to boot
 - A tagged VLAN for iMG/RG to receive configuration
 - One tagged VLAN for each of the 3 service types (voice, video, and data)
- IP Subnets - Each of the services is provided by a subnet that must be configured correctly so that the service can work correctly.
- UPS - To provide service in case of a power outage at the customer site, a UPS can be connected.

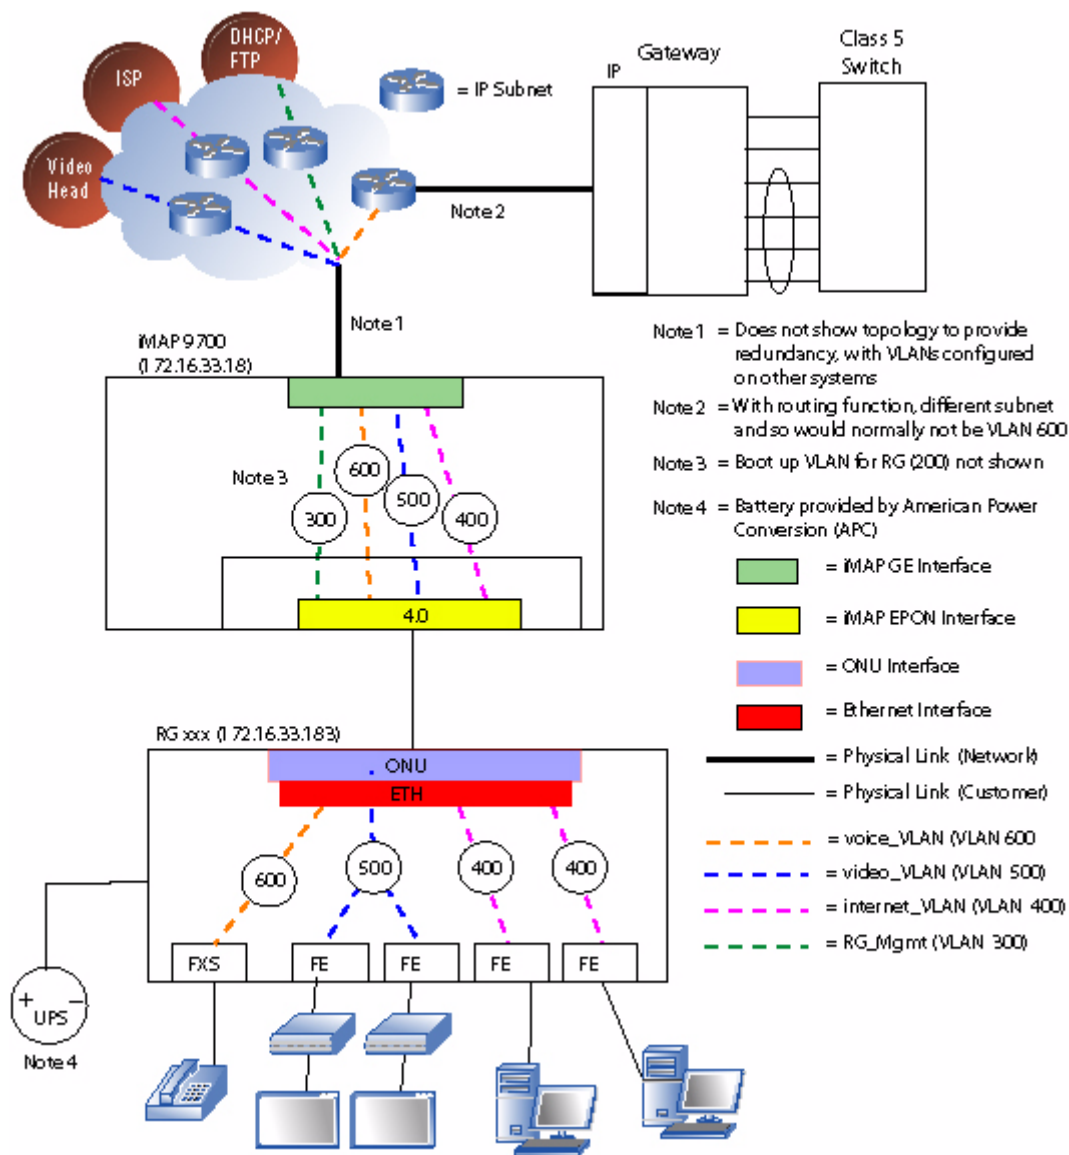


FIGURE 8-2 The EPON in a Triple Play Configuration

8.2.3 Provisioning Components

8.2.3.1 EPON Interface

The EPON interface is one-to-one with the physical EPON port on the card. Like other interfaces the system can raise alarms on it, collect statistics on it, enable/disable it, etc. It will host ONU interfaces but it does not support any ETH interfaces directly.

The EPON interface is always present when the card is present (i.e. they are not created/deleted by the user). The EPON interface is identified as an interface like `epon:4.1` with the slot and physical port as the indices.

The EPON interface has IGMP-specific attributes for video multicasting:

- The VLAN for IP Multicast
- IP Source Address for IGMP Proxy (refer to [8.8](#) for address to use)

Once the IP Multicast (IPMC) VLAN is associated with the EPON interface, IGMP Proxy is enabled as well.

All EPON interfaces are considered customer-facing and thus have `DIRECTION=Customer`.

8.2.3.2 ONU (Included with iMG/RG)

The ONU interface is a dynamically created interface on an EPON interface. Although the physical ONU is a separate piece of hardware, it is provisioned as an extension of the fMAP. It will behave similarly to other interfaces in that the system can raise alarms on it, collect statistics on it, enable/disable it, etc.

The ONU interface supports/hosts one (and only one) ETH interface, which is actually an Ethernet port inside the iMG/RG.

The ONU interface is created/deleted by the user as required. The ONU interface is identified like “`onu:4.1.7`” with the EPON's slot and physical port as the first two indices. The last index is a logical identifier and identifies the ONU interface for CLI commands.

The ONU is provisioned with these attributes:

- A logical interface ID. This is added to the EPON interface (the EPON slot.port) to make up the complete ONU interface (`onu:4.1.0`), which is used for command query/control.
- The provisioned MAC address. This is used to correlate the physical ONU to the logical interface once the ONU is discovered on the EPON.

For example, to create an ONU on EPON port 4.1 the command would be:

```
CREATE ONU=onu_example ONUID=0 INT=epon:4.1 MAC=00:aa:00:08:07:06
```

This creates the ONU interface “`onu:4.1.0`” with the following attributes:

- A created ETH interface “`eth:4.1.0`”.
- The ETH interface associated with the default VLAN (VID 1)
- An SLA (which is the ONU/VLAN association) which is the default SLA, “NONE”.

Note: See the next subsection for a description of the SLA.

To configure/change the attributes of the ONU services (association with VLANs, IGMP configuration, IP Filter settings, etc.), the commands refer to the ETH interface (eth:4.1.0). The user can also use the ONU name that has been configured (`onu_example`) in entering ONU-related commands.

All ONU interfaces are considered customer-facing and thus have `DIRECTION=Customer`.

Note: You can use the `SHOW INTERFACE` on the EPON interface to display the MAC for ONUs that have been discovered on that EPON but have not yet been configured.

8.2.3.3 ONU compatible with EPON2 Card

For release 8.0, there are two ONUs that are compatible with the EPON2 card:

- iMG646PX-ON - This is an iMG in which the ONU is housed within the iMG unit
- Media Converter - on the Media Converter the UNI port is exposed.

8.2.3.4 SLA / QOSPOLICY (VLAN basis)

As mentioned in [8.2.1](#), the SLA is defined at the CLI as a QOSPOLICY, which provides attributes to ensure that a traffic flow is given adequate bandwidth to support a service on an ONU. Since the service may involve downstream only or upstream/downstream data flows, the QOSPOLICY has both upstream and downstream attributes.

The QOSPOLICY is associated with a VLAN as well, and so to configure the QOSPOLICY, the user must understand the following:

- The VLANs associated with a service.
- The logical links, which are the traffic flows (upstream/downstream and downstream only) for each type of service.

Note: For details on provisioning each service, refer to [8.3](#)

There are two types of traffic flows on which QOSPOLICYs are configured:

1. Upstream/Downstream Links

- There is one or more per ONU
- Each one carries **one** VLAN to **one** ONU.
- Downstream, they carry known unicast packets to the ONU
- Upstream, they carry unicast, broadcast, multicast, and unknown MAC packets.
- The first one provisioned on the ONU carries some control and management traffic upstream

2. Downstream Only Link

- One is for all ONUs
- Carries two types of traffic, with each having a separately defined SLA (and therefore QOSPOLICY)

- Multicast traffic for only the IP Multicast (IPMC) VLAN
- Broadcast, Unknown Unicast, and flooded Multicast (BRUUM). This downstream link is shared for all VLANs on all ONUs on the EPON.

To understand these relationships, the QOSPOLICYs are highlighted for each type of service, since each type will need to have particular QOSPOLICYs defined for its specific types of traffic.

8.3 Provisioning Models

8.3.1 Video

Refer to [Figure 8-3](#) while reading the following:

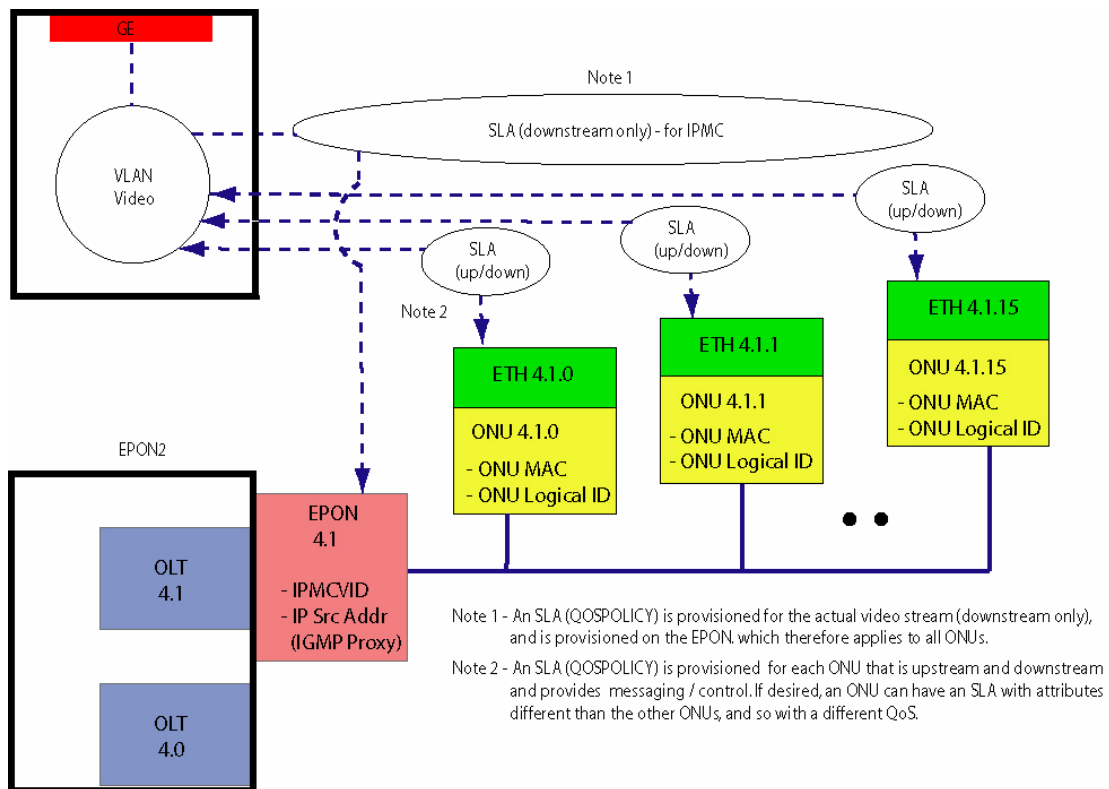


FIGURE 8-3 SLA Provisioning Model for Video Traffic

As mentioned in [8.2.3.1](#), the EPON card is provisioned with two attributes necessary to process Video:

- IPMC VID - This is the VLAN for IP Multicast that carries the actual video traffic (downstream only)
- IP Source Address - This is configured so the EPON2 card can perform IGMP proxy functions

For video, two types of SLAs are configured:

1. SLA for the video stream (downstream only) - This is the IPMC SLA attribute on the EPON interface and has the attributes that control the downstream attributes for the IP Multicast VID. (Any upstream attributes are ignored.) This SLA therefore applies to all the ONUs on the EPON2 interface.
2. SLA for the Video VLAN (upstream/downstream) - This is the **bi-directional** attribute on the ONU/VLAN association. This is for the same VID as the downstream-only video stream, but it applies to all upstream traffic from that ONU, and known/learned unicast downstream traffic to that ONU. The ONU/VLAN association and corresponding SLA must be provisioned if unicast or broadcast traffic is required for operation (e.g. DHCP is used for IP address assignment).

Table 8-1 lists the SLA attributes as they apply to a video service example (150 channel IPTV service)..

TABLE 8-1 SLA attributes for example Video

Attribute	Values / Range	Notes for Video
Description	text	Appropriate description for function of SLA.
MAXUPSTREAMRATE	100Mbps	For downstream SLA, ignored.
MAXDOWNSTREAMRATE	1000M	The maximum allowable bandwidth for downstream traffic on this traffic flow.
MINUPSTREAMRATE	0Mbps	For downstream SLA, ignored.
MINDOWNSTREAMRATE	650M	For upstream/downstream SLA, the minimum, guaranteed (but not reserved) bandwidth for downstream traffic on this traffic flow.
UPBURSTSIZE	100Mbps	For downstream SLA, ignored.
DOWNBURSTSIZE	30K	For upstream/downstream SLA, the maximum burst size that may briefly exceed the max rate for downstream traffic on this traffic flow.
UPDELAYSENSITIVITY	Tolerant <i>Note - Delay-sensitive traffic must have MAX and MIN rates equal.</i>	For downstream SLA, ignored.
DOWNDELAYSENSITIVITY	NO See note above	Indicates that the upstream traffic in this traffic flow is sensitive to delay.

8.3.2 Data, Voice

Refer to Figure 8-4 while reading the following:

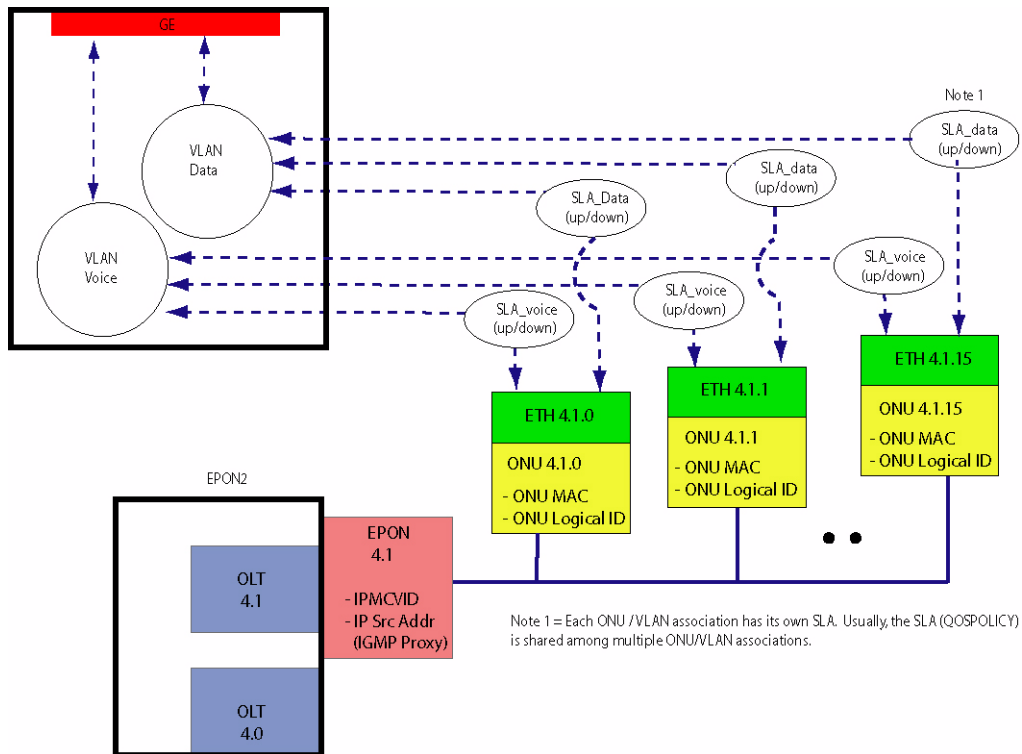


FIGURE 8-4 SLA Provisioning Model for Data, Voice Traffic

Only bi-directional SLAs are configured for either data or voice service:

- SLA for the Data VLAN (upstream/downstream) - This SLA associates its attributes with the Data VID and a specific ONU. This SLA must therefore be provisioned for each ONU, and so an ONU that did not support data would not have an association to the Data VLAN provisioned.
- SLA for the Voice VLAN (upstream/downstream) - This SLA associates its attributes with the Voice VID and a specific ONU. This SLA must therefore be provisioned for each ONU, and so an ONU that did not support voice would not have an association to the Voice VLAN provisioned.

Note: Since each bi-directional SLA is configured on a per-ONU/VLAN association basis, the user could, if desired, configure an ONU/VLAN with a different SLA and therefore with a different QoS. There is, however, a limit to how much this can be done (refer to 8.5).

Table 8-2 lists the SLA attributes as they apply to an example voice service (VoIP Service with 2 handsets, assuming 10ms packetization)..

TABLE 8-2 SLA attributes for Example Voice Service

Attribute	Values / Range	Notes for Data, Voice
Description	text	
MAXUPSTREAMRATE	256K	For downstream SLA, ignored. For upstream/downstream SLA, the maximum allowable bandwidth for upstream traffic on this traffic flow.
MAXDOWNSTREAMRATE	256K	The maximum allowable bandwidth for downstream traffic on this traffic flow.
MINUPSTREAMRATE	256K	For downstream SLA, ignored. For upstream/downstream SLA, the minimum, guaranteed (but not reserved) bandwidth for upstream traffic on this traffic flow.
MINDOWNSTREAMRATE	256K	For upstream/downstream SLA, the minimum, guaranteed (but not reserved) bandwidth for downstream traffic on this traffic flow.
UPBURSTSIZE	6K	For downstream SLA, ignored. For upstream/downstream SLA, the maximum burst size that may briefly exceed the max rate for upstream traffic on this traffic flow
DOWNBURSTSIZE	6K	For upstream/downstream SLA, the maximum burst size that may briefly exceed the max rate for downstream traffic on this traffic flow.
UPDELAYSENSITIVITY	YES	For downstream SLA, ignored. For upstream/downstream SLA, indicates that the upstream traffic in this traffic flow is sensitive to delay.
DOWNDELAYSENSITIVITY	YES	Indicates that the upstream traffic in this traffic flow is sensitive to delay.

Table 8-3 lists the attributes for a typical data service.

TABLE 8-3 SLA attributes for Data Service

Attribute	Values / Range	Notes for Data, Voice
Description	text	
MAXUPSTREAMRATE	100Mbps	For downstream SLA, ignored. For upstream/downstream SLA, the maximum allowable bandwidth for upstream traffic on this traffic flow.
MAXDOWNSTREAMRATE	100Mbps	The maximum allowable bandwidth for downstream traffic on this traffic flow.
MINUPSTREAMRATE	0Mbps	For downstream SLA, ignored. For upstream/downstream SLA, the minimum, guaranteed (but not reserved) bandwidth for upstream traffic on this traffic flow.
MINDOWNSTREAMRATE	0Mbps	For upstream/downstream SLA, the minimum, guaranteed (but not reserved) bandwidth for downstream traffic on this traffic flow.
UPBURSTSIZE	100Mbps	For downstream SLA, ignored. For upstream/downstream SLA, the maximum burst size that may briefly exceed the max rate for upstream traffic on this traffic flow
DOWNBURSTSIZE	100Mbps	For upstream/downstream SLA, the maximum burst size that may briefly exceed the max rate for downstream traffic on this traffic flow.
UPDELAYSENSITIVITY	Tolerant Sensitive <i>Note - Delay sensitive traffic must have MAX = MIN values.</i>	For downstream SLA, ignored. For upstream/downstream SLA, indicates that the upstream traffic in this traffic flow is sensitive to delay.
DOWNDELAYSENSITIVITY	Tolerant Sensitive See note above	Indicates that the upstream traffic in this traffic flow is sensitive to delay.

8.4 Provisioning Model - BRUUM

Refer to [Figure 8-5](#) while reading the following:

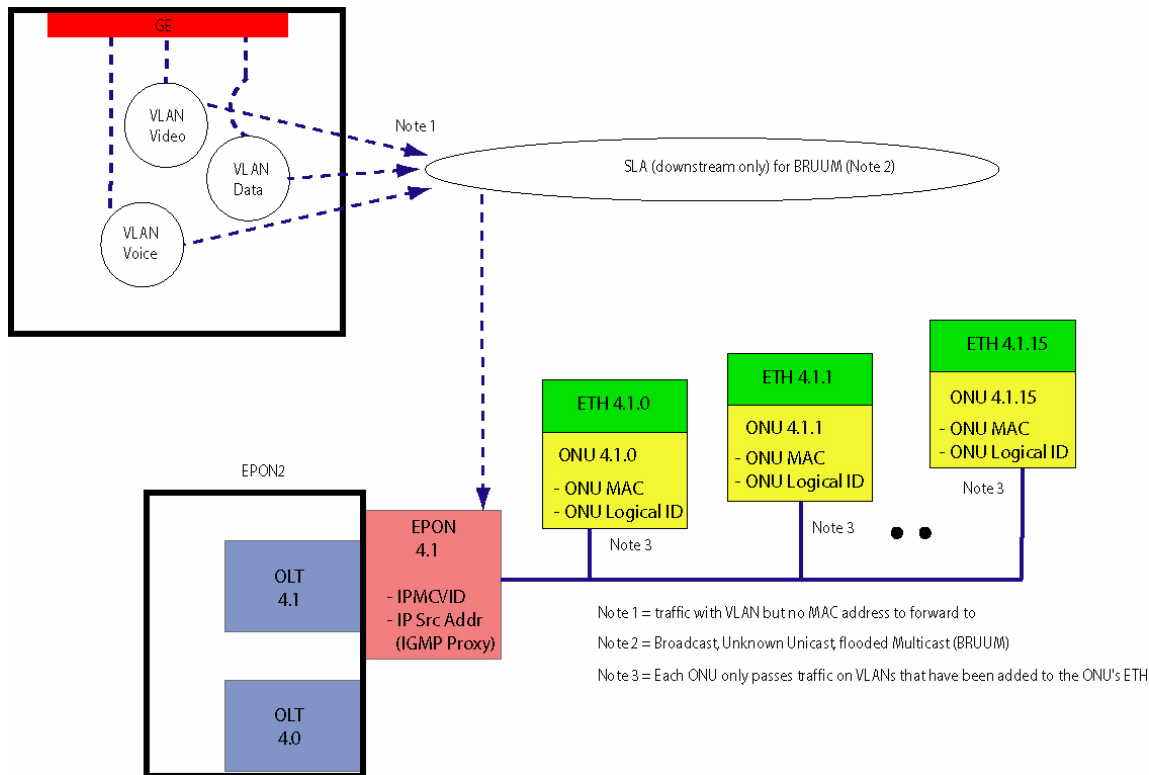


FIGURE 8-5 SLA Provisioning Model for BRUUM Traffic

For BRUUM, only one downstream-only SLA is configured, the SLA for all service VLANs (downstream only). This SLA associates its attributes with the BRUUM traffic. This is traffic that has a VLAN but no MAC address to be forwarded to. Therefore, all ONUs receive the traffic and each ONU decides whether to forward or discard based on whether it matches a configured VLAN.

For the SLA attributes, all of the upstream values are ignored.

8.5 Traffic Management

For an overview of all aspects of traffic management, refer to Section 15. This subsection highlights how traffic management is handled for the EPON configuration.

8.5.1 Classifiers

Filtering is based on VLAN and IPSOURCE address.

IPSOURCE filtering can be done:

- Statically, using ACL or user classifiers
- Dynamically, using the Auto_IP filtering option on DHCP relay.

Note the following rules for classification; these are enforced by the CLI:

- The VID match rule is required on all IP filters.
- Ingress metering is not supported.
- The user can do only FORWARD and DROP actions (e.g. can't COUNT, remark, etc.).
- There is no support for IP address masks .
- All FORWARD actions precede any DROP actions, and all DROP actions follow any FORWARD actions.

Note: These attributes are included in the summary tables in Section 15.1.

8.5.2 QoS (Traffic Queues/Priorities)

With the introduction of the SLA model, there is a change in how traffic is prioritized as it flows upstream and downstream. Between the OLT and ONU, traffic management is done as follows:

- The SLA provides traffic management per VLAN
- The SLA must take into account all traffic on that VLAN (service, ping, DHCP, etc.)

At points outside the OLT-ONU, p-bits/classifiers may still be used at various points; moreover, these are passed through the OLT-ONU. This has the following results:

- In the upstream direction, traffic is passed with no controls from the UNI to the ONU. From the OLT to the ONU, the SLA is used to prioritize traffic flows per VLAN as follows:
 - High - UPDELAYSENSITIVITY=Sensitive
 - Medium - UPDELAYSENSITIVITY=Tolerant, MINUPSTREAMRATE not=0
 - Low - MINUPSTREAMRATE=0

At the EPON interface, p-bits may be used with the VPRIORITY setting to separate and prioritize traffic for up to 8 queues.

Note: The user must be sure that there are no conflicts between the flows set up by the SLA and those by the p-bit settings, since they are separate traffic management tools.

- In the downstream direction, p-bits/classifiers are set at the EPON interface. For known unicast traffic (non-video and non-BRUUM), the SLA per VLAN is used to prioritize the data flows. At the ONU, the p-bit/classifier settings are passed down to the ONU and are actually applied to the traffic flows.

Note: Traffic management is therefore not performed by the EPON but by the ONU, which has the interface to the UNI. The ONU is modeled as an fMAP extension.

8.5.3 Connection Admission Control (CAC)

The CAC function is to ensure the hardware can provide the guarantees configured by the SLAs. There are two type of CAC check on the EPON2 interface:

1. Sum of Minimum Bandwidths - The sum of provisioned minimum bandwidths for the QOSPOLICYs of all logical links on an EPON port must not exceed the bandwidth capacity of that port in either upstream or downstream direction. This limit is slightly below 1G due to administrative overhead (REPORTs, GRANTs, OAMPDU, etc.). This function is performed in OAM so that CAC can be enforced even if the EPON2 card is not physically present (i.e. when pre-provisioning).
2. Availability of priority categories - This differentiates among three categories of traffic at the queues from the OLT upstream to the EPON switching fabric. There is a number of links allowed in each category at initialization time. The following table shows these categories and the number of links allowed:

TABLE 8-4 CAC for the EPON2

Priority Level	Correlation to SLA	Type of Traffic	Number of Links	Total
0	DELAY=SENSITIVE AND Min = Max	Traffic that is sensitive to jitter (e.g. Voice VLAN)	2 per ONU	64
1	DELAY=TOLERANT AND Min > 0	Traffic that can tolerate some jitter, and has some guaranteed bandwidth (e.g. Video VLAN)	2 per ONU, plus 15 additional	64 +15 (79)
2	DELAY=TOLERANT AND Min = 0	Best Effort” traffic (e.g. Internet VLAN, RG Boot VLAN, RG Mgmt VLAN)	3 per ONU	96
				239

Note: The user must be aware of this allocation when choosing SLAs. Attempts to assign SLAs that exceed the number that are allocated for that type of SLA will be rejected at the CLI.

8.6 Feature Interaction

Following are the feature interactions/limitations for the EPON2:

- The EPON2 does not support the MAC limiting feature.
- The EPON2 does not support the STB Mobility feature.
- ONU switching from ONU to ONU connected to the same EPON2 is not supported in this release.

- Each ONU supports up to 6 VLANs, and up to 24 different VIDs among all the ONUs that are connected to the EPON2.
- The EPON2 card does not support FLOODUNKNOWN=ON the same way as other cards/interfaces. Refer to 14.3.2.
- Upgrading the EPON2 load upgrades the software on all the ONUs associated with the EPON2 automatically, to ensure they are in synch with the OLT configuration on the EPON2 card.

8.7 Example Configuration

8.7.1 Example Figures

The following figures show an example configuration and highlight the parameters that are datafilled. Refer to these figures while reading the steps and associated commands.

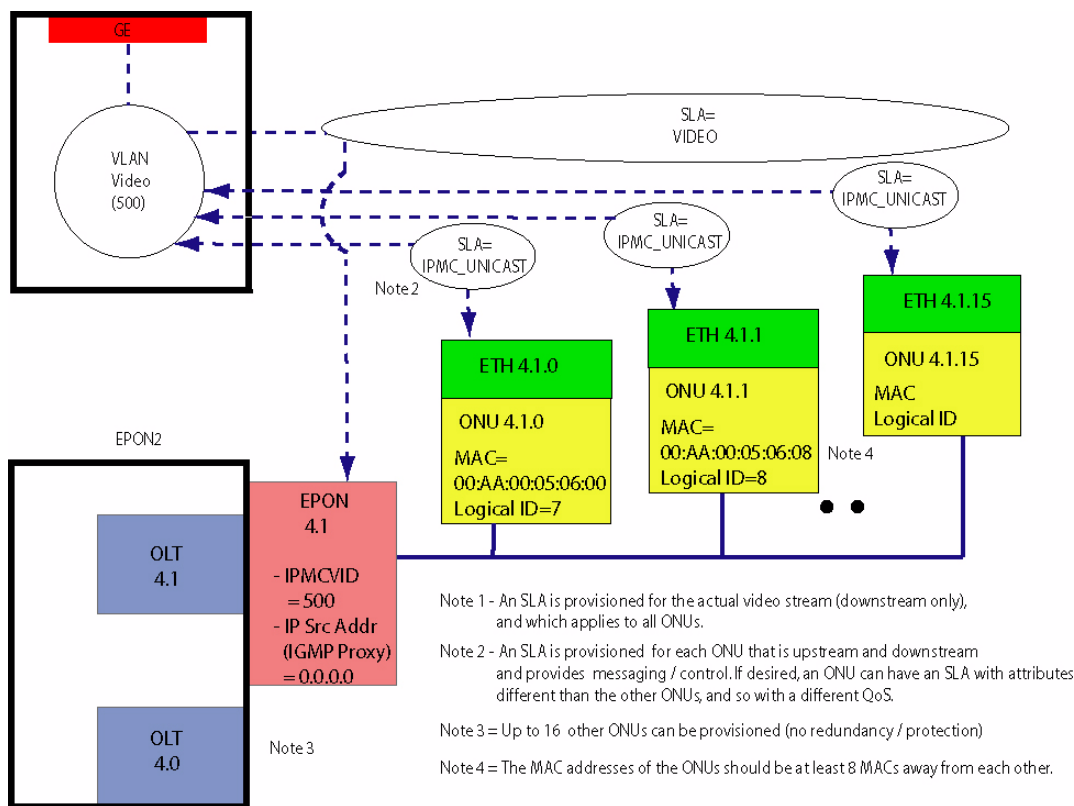


FIGURE 8-6 Video Provisioning Model - Example




FIGURE 8-7 Data, Voice Provisioning Model - Example

8.7.2 Pre-provisioning Tasks

Before provisioning the OLT/ONU and the associated SLAs, the following VLANs must already have been created on the fMAP

- Bootstrap VLAN - This is the VLAN 200 in the example
- RG Mgmt VLAN - This is VID=300 in the example
- Data - This is the VLAN 400 in the example
- Video - This is the VLAN 500 in the example
- Voice - This is the VLAN 600 in the example

Note that once a VLAN is created and associated with a QOSPOLICY, the VLAN cannot be destroyed until all associations with the ETH interface for an ONU are destroyed.

Note: Once the EPON and ONU have been provisioned, the RG is configured using the AlliedView NMS. The iMG/RG is power cycled, and using DHCP the RG gets the latest files for the RG, reboots, and when finished the RG can be controlled by the AlliedView NMS. Refer to the AlliedView NMS Administration Guide for details.

8.7.3 Major Tasks

The creation of the EPON card and its connection to the ONU can be divided into these main tasks

1. Configure the cards / ports to put the card in operation

When this is done, there are still no ONUs to support service, but you can enable/disable the EPON card and interface.

2. Create the Triple-play service to an ONU
3. Show the configuration results
4. Show the forwarding database

8.7.3.1 Configure the cards / ports to put the card in operation

1. **Create the card, and the EPON interfaces are automatically provisioned.**

This can be done:

- automatically, by inserting the card in AutoProv mode
- manually, by entering the command:

```
CREATE CARD 4 EPON2
```

Note the following:

- For each EPON interface, the BRUUM Logical Link is also automatically created, since it will be used for any ONU on the EPON.
- The BRUUM is created with default QOSPOLICY, NONE. Since this policy is unrestricted (no guarantee and no limits), this will need be changed, as shown below.

2. **Configure the IP address for the EPON interface, using the following command:**

```
SET INTERFACE 4.1 EPON IPADDRESS=0.0.0.0
```

This value identifies the EPON as the IGMP entity that the Multicast router will request reports from.

Note: In most cases, the default 0.0.0.0 should work, If 0.0.0.0 does not work, then the user should set IPADDRESS to an address that is appropriate for the IPMC subnet.

3. **Restrict the (initially unrestricted) QOSPOLICY on the BRUUM logical link for the EPON interface**

This is done to limit broadcast storms. The policy BC_LIMIT is created as follows:

```
>CREATE QOSPOLICY BC_LIMIT DOWNMAX=1M DOWNBURST=16K
>ADD QOSPOLICY=BC_LIMIT INT epon:4.1 BRUUM
```

Note that the UPMAX/UPMIN parameters are ignored for this path.

4. **Configure the IPMC VID and QOSPOLICY**

This is done to support the IPMC logical link (downstream) for video. Note that the IP multicast VID and QOSPOLICY have automatically been set from the EPON “port” profile. This is just to change them.

```
>SET INT 4.1 EPON IPMCVLAN 201
>CREATE QOSPOLICY VIDEO DOWNMIN=200M DOWNBURST=256K
>ADD QOSPOLICY=VIDEO INT epon:4.1 IPMC
```

Note: Configuring the DOWNMIN value is an engineering decision and is determined by the total bandwidth needed to support the video configuration. For example, with 3 STBs per residence with a fully configured EPON interface (16 ONUs), and 4Mbps per stream, the value could be 200M. Refer to the Services Guide for other examples.

8.7.3.2 Create the Triple-play service to an ONU

1. Create an ONU on the EPON interface, which adds the ETH interface to the untagged VLAN (VID 1)

```
>CREATE ONU my_onu ONUID=7 INT=epon:4.1 MAC=00:AA:00:05:06:07
```

This creates an ONU with the id `onu:4.1.7`, auto-creates an ETH with id `eth:4.1.7`, and automatically adds it to the untagged VLAN on the iMG/RG.

At this point the QOSPOLICY for the ONU's VID=1 traffic is “NONE”.

Note: For RG discovery, the user should use the AlliedView NMS and its tools. Refer to the AlliedView NMS Administration Guide for information on how the untagged VLAN is used for initial bootup/configuration of the iMG/RG.

2. Associate the RG Mgmt VLAN with the ETH Interface and Define a QOSPOLICY

```
>ADD VLAN 300 INT 4.1.7 FRAME TAGGED
>CREATE QOSPOLICY RG_MGMT UPMIN=512K UPMAX=1M DOWNMIN=64K DOWNMAX=1M
>ADD QOSPOLICY=RG_MGMT INT=4.1.7 BIDIR VID=300
```

The ADD VLAN creates another Logical Link and starts sending any learned VID=300 traffic over it. You only need to create the QOSPOLICY one time.

Note: For RG discovery, the user should use the AlliedView NMS and its tools which include the use of the RG Mgmt VLAN. Refer to the AlliedView NMS Administration Guide for information on how the default VLAN is used for initial bootup/configuration of the iMG/RG.

3. Add the data VLAN (400)

```
>ADD VLAN 400 INT 4.1.7 FRAME TAGGED
>CREATE QOSPOLICY DATA_GOLD UPMIN=5M UPMAX=10M DOWNMIN=5M DOWNMAX=10M
>ADD QOSPOLICY=DATA_GOLD INT=4.1.7 BIDIR VID=400
```

4. Add the voice VLAN (500)

```
>ADD VLAN 500 INT 4.1.7 FRAME TAGGED
>CREATE QOSPOLICY VOIP_2PHONE UPMIN=512K UPMAX=1M DOWNMIN=1M DOWNMAX=10M
>ADD QOSPOLICY=VOIP_2PHONE INT=4.1.7 BIDIR VID=500
```

5. Add the video VLAN (600)

```
>CREATE QOSPOLICY IPMC_UNICAST UPMIN=512K UPMAX=512K DOWNMIN=512K DOWN-
MAX=512K
>ADD VLAN 600 INT 4.1.7 FRAME TAGGED
>ADD QOSPOLICY=IPMC_UNICAST INT=4.1.7 BIDIR VID=600
>ENABLE IGMP INT 4.1.7 (This is enabled by default)
```

Note: The QOSPOLICY configured here is just the one for the unicast path (upstream and downstream). The downstream multicast QOSPOLICY covers the actual video traffic.

8.7.3.3 Show the Configuration Results

If the above steps have been done and the ONUs are enabled and operationally UP, then the SHOW INTERFACE for the EPON and ONU would show the following:

```
officer SEC>> SHOW INTERFACE=4.1
```

```
--- EPON Interfaces ---

Interface..... 4.1
Type..... EPON
State..... UP-UP-Online
Description..... epon port 4.1

Provisioning
  Provisioning Profile..... AutoProv (*)
  IPMC Source Address..... 192.200.2.0
  IPMC VLAN..... 512
  Performance Monitoring..... On

Actual
  Direction..... Customer

QoS Policy Information
  Broadcast Link (BRUUM)..... bruum_limit
  IP Multicast Link (IPMC)..... ipmc_multicast

Discovered ONU Information
  00:0D:DA:04:40:00..... onu_2p0_01
  54:4B:37:02:10:A0..... onu_2p0_04
  54:4B:37:02:10:B0..... onu_2p0_10
  54:4B:37:02:10:C0..... onu_2p0_14
  54:4B:37:02:10:D0..... onu_2p0_05
  54:4B:37:02:10:F0..... onu_2p0_09
  54:4B:37:02:11:10..... onu_2p0_11
  54:4B:37:02:11:20..... onu_2p0_13
```

```

54:4B:37:02:11:30..... onu_2p0_12
54:4B:37:02:11:40..... onu_2p0_00
54:4B:37:02:11:50..... onu_2p0_06
54:4B:37:02:11:60..... onu_2p0_03
54:4B:37:02:11:70..... onu_2p0_02
54:4B:37:02:11:80..... onu_2p0_07
54:4B:37:02:11:90..... onu_2p0_08

```

VLAN Information

```

Acceptable Frame Types..... All
Ingress Filtering..... On
TPID..... 0x8100
TAGALL..... Off
Untagged VLAN..... 1

```

Packet Statistics

	Input	Output
	-----	-----
Octets.....	1408	0
Unicast Packets.....	0	0
Discarded Packets.....	0	0
Errored Packets.....	0	0
Unknown Proto Packets.....	0	N/A

--- ONU Interfaces ---

Interface	State	EPON	ID	MAC Address
-----	-----	-----	---	-----
onu_2p0_00	UP-UP	2.0	0	54:4B:37:02:11:40
onu_2p0_01	UP-UP	2.0	1	00:0D:DA:04:40:00
onu_2p0_02	UP-UP	2.0	2	54:4B:37:02:11:70
onu_2p0_03	UP-UP	2.0	3	54:4B:37:02:11:60
onu_2p0_04	UP-UP	2.0	4	54:4B:37:02:10:A0
onu_2p0_05	UP-UP	2.0	5	54:4B:37:02:10:D0
onu_2p0_06	UP-UP	2.0	6	54:4B:37:02:11:50
onu_2p0_07	UP-UP	2.0	7	54:4B:37:02:11:80
onu_2p0_08	UP-UP	2.0	8	54:4B:37:02:11:90
onu_2p0_09	UP-UP	2.0	9	54:4B:37:02:10:F0
onu_2p0_10	UP-UP	2.0	10	54:4B:37:02:10:B0
onu_2p0_11	UP-UP	2.0	11	54:4B:37:02:11:10
onu_2p0_12	UP-UP	2.0	12	54:4B:37:02:11:30
onu_2p0_13	UP-UP	2.0	13	54:4B:37:02:11:20
onu_2p0_14	UP-UP	2.0	14	54:4B:37:02:10:C0

```
officer SEC>> SHOW INTERFACE=ONU:4.1.7
```

```
--- ONU Interfaces ---
```

```
Interface..... onu_4p1_07
Type..... ONU
State..... UP-UP-Online
Description..... <none>

Provisioning
  ID..... 1
  Associated EPON..... 4.1
  Auto Negotiation..... On
  Speed..... Auto
  Duplex..... Auto
  Flow Control..... Auto
  MAC Addresses
    MAC 1 (user setting)..... 00:0D:DA:04:40:00
    MAC 2 (generated)..... 00:0D:DA:04:40:01
    MAC 3 (generated)..... 00:0D:DA:04:40:02
    MAC 4 (generated)..... 00:0D:DA:04:40:03
    MAC 5 (generated)..... 00:0D:DA:04:40:04
    MAC 6 (generated)..... 00:0D:DA:04:40:05
    MAC 7 (generated)..... 00:0D:DA:04:40:06
    MAC 8 (generated)..... 00:0D:DA:04:40:07
  Performance Monitoring..... On

Actual
  Product ID..... Unknown
  Serial Number..... Unknown
  Speed..... 100 Mbps
  Duplex..... Half
  Flow Control..... On

QoS Policy Information
  VID 1..... NONE

VLAN Information
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Untagged VLAN..... 1
```

```

officer SEC>> SHOW INTERFACE=2.0.1

--- ONU Interfaces ---

Interface..... onu_2p0_01
Type..... ONU
State..... UP-UP-Degraded
Description..... <none>

Interface Faults
  Remote Loss of Signal (UNI)..... Info

Provisioning
  ID..... 1
  Associated EPON..... 2.0
  Auto Negotiation..... On
  Speed..... Auto
  Duplex..... Auto
  Flow Control..... Auto
  MAC Addresses
    MAC 1 (user setting)..... 00:0D:DA:04:40:00
    MAC 2 (generated)..... 00:0D:DA:04:40:01
    MAC 3 (generated)..... 00:0D:DA:04:40:02
    MAC 4 (generated)..... 00:0D:DA:04:40:03
    MAC 5 (generated)..... 00:0D:DA:04:40:04
    MAC 6 (generated)..... 00:0D:DA:04:40:05
    MAC 7 (generated)..... 00:0D:DA:04:40:06
    MAC 8 (generated)..... 00:0D:DA:04:40:07
  Performance Monitoring..... On

Actual
  Product ID..... Unknown
  Serial Number..... Unknown
  Speed..... 100 Mbps
  Duplex..... Half
  Flow Control..... On

QoS Policy Information
  VID 1..... NONE

VLAN Information
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Untagged VLAN..... 1

```

8.7.4 Existing Commands

Table 8-5 lists the existing commands that include the EPON card/interface.

TABLE 8-5 EPON/ONU for Existing Commands

Object(s)	Verb	Syntax	Description
CARD EPON2	CREATE	<pre>CREATE CARD=slot EPON2 [{ [PREFLOAD=filename] [ADMINSTATE={ UP DOWN }] PROFILE=name }]</pre>	<p>Creates the card, which automatically provisions the EPON interfaces and the EPON interfaces are automatically provisioned.</p> <p>For each EPON interface, the BRUUM Logical Link is also automatically created, since it will be used for any ONU on the EPON. The BRUUM is created with the default QOSPOLICY, NONE.</p>
PROFILE EPON2	CREATE	<pre>CREATE PROFILE=name EPON2 [PREFLOAD=filename] [ADMINSTATE={ UP DOWN }]</pre>	Creates a profile for the card type.
PROFILE EPONPORT	CREATE	<pre>CREATE PROFILE=name EPONPORT [ADMINSTATE={ UP DOWN }] [IPMCVLAN={ vlanname vid }] [IPADDRESS=ipaddress]</pre>	Creates a profile for the port type.
PROFILE EPON2 EPONPORT	SHOW	<pre>SHOW PROFILE [={ name-list NAMES ALL }] { EPON2 EPONPORT } [FULL]</pre>	Shows the profile attributes for the card or the port type.

8.8 Command Summary for EPON/ONU

8.8.1 Command Set

TABLE 8-6 EPON/ONU Provisioning Commands

Object(s)	Verb	Syntax	Description
INTERFACE EPON	SET	<pre>SET INTERFACE={ type: type:id-range id-range ifname-list ALL } EPON [IPMCVLAN={ vlanname vid }] [IPADDRESS=ipaddress] [DESCRIPTION=description]</pre>	<p>Configures for the EPON interface the IP address, the IPMC VID, and a text description.</p> <p>The IP address identifies the EPON as the IGMP entity that the Multicast router will request reports from</p> <p><i>Note:</i> In most cases, the default 0.0.0.0 should work. If 0.0.0.0 does not work, set this to an address that is appropriate for the IPMC subnet.</p>
CARD	CREATE	<pre>CREATE ONU=onuname ONUID=0..15 INTERFACE={ type:id id ifname } MACADDRESS=macaddress</pre>	<p>Creates an ONU interface and its associated ETH interface.</p> <ul style="list-style-type: none"> - Adds the ETH interface to the default VLAN, untagged. - ONU interfaces share the same namespace as other interface names. <p><i>Note:</i> The MAC addresses of the ONUs should be at least 8 MACs away from each other. The links are (internally) identified by a MAC address calculated from the base MAC.</p>
ONU	SET	<pre>SET ONU=onuname [MACADDRESS=macaddress]</pre>	<p>Requests to modify an ONU interface(s).</p> <p>The ONU name is unique to reflect that there should be no duplicate MACs.</p>
ONU	RENAME	<pre>RENAME ONU=onuname TO=onuname</pre>	<p>RENAME ONU</p> <p>Requests to change the name of an ONU interface(s).</p>

TABLE 8-6 EPON/ONU Provisioning Commands (Continued)

Object(s)	Verb	Syntax	Description
INTERFACE ONU	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } ONU [AUTONEGOTIATION={ ON OFF }] [SPEED={ AUTONEGOTIATE 10 100 1000 }] [DUPLEX={ AUTONEGOTIATE FULL HALF }] [FLOWCONTROL={ AUTONEGOTIATE ON OFF }] </pre>	<p>Sets the ONU ethernet provisioning. For the iMG646PX-ON (where the the ethernet port is internal to the device) the values are ignored, with actual values reflecting the real state of the port provisioning.</p> <p>For ONUs that support auto-negotiation, the values can change depending on the far-end capabilities.</p>
ONU	DESTROY	<pre> DESTROY ONU={ onuname-list ALL } [INTERFACE={ type: type:id-range id-range ifname-list ALL }] [FORCE] </pre>	<p>Requests to destroy ONU interface(s).</p> <p>Destroys one or more ONU interfaces and its associated ETH interface.</p> <p>- FORCE does not prompt the user, and does not require that the ONU interface be disabled</p>
ONU	SHOW	<pre> SHOW ONU [={ onuname-list ALL }] [ONUID={ 0..15 ALL }] [INTERFACE={ type: type:id-range id-range ifname-list ALL }] [MACADDRESS={ macaddress ALL }] [FULL] </pre>	<p>Requests to display ONU interface(s).</p> <p>- Shows details of the ONU interface. You can ID by interface name/id or by MAC address.</p> <p>- INTERFACE should be able to SHOW based on the ONU or EPON interface.</p>

TABLE 8-6 EPON/ONU Provisioning Commands (Continued)

Object(s)	Verb	Syntax	Description
QOSPOLICY	CREATE	<pre> CREATE QOSPOLICY=policyname [DESCRIPTION=text] [MAXUPSTREAMRATE={ bits-per-second MAX }] [MAXDOWNSTREAMRATE={ bits-per-second MAX }] [MINUPSTREAMRATE={ bits-per-second MIN }] [MINDOWNSTREAMRATE={ bits-per-second MIN }] [UPBURSTSIZE={ 1..256 MAX }] [DOWNBURSTSIZE={ 1..256 MAX }] [UPDELAYSENSITIVITY={ SENSITIVE TOLERANT }] [DOWNDELAYSENSITIVITY={ SENSITIVE TOLERANT }] </pre>	<p>Requests to create a QOSPOLICY.</p> <p>There is a default policy NONE that matches the default values.</p> <p>The bits-per-second values can be entered with a “units” postfix (e.g. “512K” means 512Kbps, while “10M” means 10000Kbps).</p> <p>The bits-per-second values can be between 0 and 1G.</p> <p>The values of “MIN” and “MAX” will be appropriate to the interface. MIN is the bandwidth that is guaranteed for the flow.</p> <p>Default burst values are all 100K.</p>
QOSPOLICY	ADD	<pre> ADD QOSPOLICY=policyname INTERFACE={ type:id-range id-range ifname-list } { BRUUM IPMC BIDIRECTIONAL VLAN={ vllanname-list vid-range ALL } }] </pre>	<p>Requests to add a QOSPOLICY to an interface.</p> <ul style="list-style-type: none"> - BRUUM and IPMC options are only available on the EPON interfaces. - BIDIRECTIONAL option is only available on an ONU interface (or the ONU’s ETH) - Adding the “NONE” QOSPOLICY is effectively the same as DELETE QOSPOLICY.

TABLE 8-6 EPON/ONU Provisioning Commands (Continued)

Object(s)	Verb	Syntax	Description
QOSPOLICY	DELETE	<pre>DELETE QOSPOLICY={ policynome-list ALL } INTERFACE={ type:id-range id-range ifname-list ALL } { BRUUM IPMC BIDIRECTIONAL VLAN={ vlannome-list vid-range ALL } ALL }</pre>	<p>Requests to delete QOSPOLICY from an interface.</p> <p>- This sets the QOSPOLICY to “NONE”.</p>
QOSPOLICY	DESTROY	<pre>DESTROY QOSPOLICY={ policynome-list ALL } [FORCE]</pre>	<p>Requests to destroy a QOSPOLICY.</p> <p>- Can’t destroy the QOSPOLICY of “NONE”.</p> <p>- By default, this command will be rejected if the QOSPOLICY is in use.</p>
QOSPOLICY	RENAME	<pre>RENAME QOSPOLICY=policynome TO=policynome</pre>	<p>Requests to change the name of a QOSPOLICY.</p> <p>The policy NONE cannot be renamed.</p>

TABLE 8-6 EPON/ONU Provisioning Commands (Continued)

Object(s)	Verb	Syntax	Description
QOSPOLICY	SET	<pre> SET QOSPOLICY={ policyname-list ALL } [DESCRIPTION=text] [MAXUPSTREAMRATE={ bits-per-second MAX }] [MAXDOWNSTREAMRATE={ bits-per-second MAX }] [MINUPSTREAMRATE={ bits-per-second MIN }] [MINDOWNSTREAMRATE={ bits-per-second MIN }] [UPBURSTSIZE={ 1..256 MAX }] [DOWNBURSTSIZE={ 1..256 MAX }] [UPDELAYSENSITIVITY={ SENSITIVE TOLERANT }] [DOWNDELAYSENSITIVITY={ SENSITIVE TOLERANT }] </pre>	<p>Requests to modify a QOSPOLICY.</p> <p>For parameters, see CREATE QOSPOLICY.</p>
QOSPOLICY	SHOW	<pre> SHOW QOSPOLICY [={ policyname-list ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] [{ BRUUM IPMC BIDIRECTIONAL [VLAN={ vlannamelist vid-range ALL }] ALL }] [FULL] </pre>	<p>Requests to display QOSPOLICY info.</p> <ul style="list-style-type: none"> - INTERFACE should accept EPON, ONU or ETH. - The command shows QOSPOLICYs and their usage in VLAN / interface associations. - ETHs that do not support QOSPOLICYs will not be displayed, or will indicate “N/A”.

9. Configuring Performance Monitoring

9.1 Overview

Performance Management is the collection of traffic statistics over the interfaces (usually ports) over a specified time period (called the interval). During this period, if the value for a certain statistic crosses a threshold value, it is noted and a log or alarm may be produced.

Note: By default, ports are disabled for statistics and must be explicitly enabled. Enabled ports for statistics are persistent over reboots/restarts.

Note: Monitor interfaces must be specifically enabled in order for counts to increment.

Figure 9-1 shows the performance management configuration for the fMAP product.

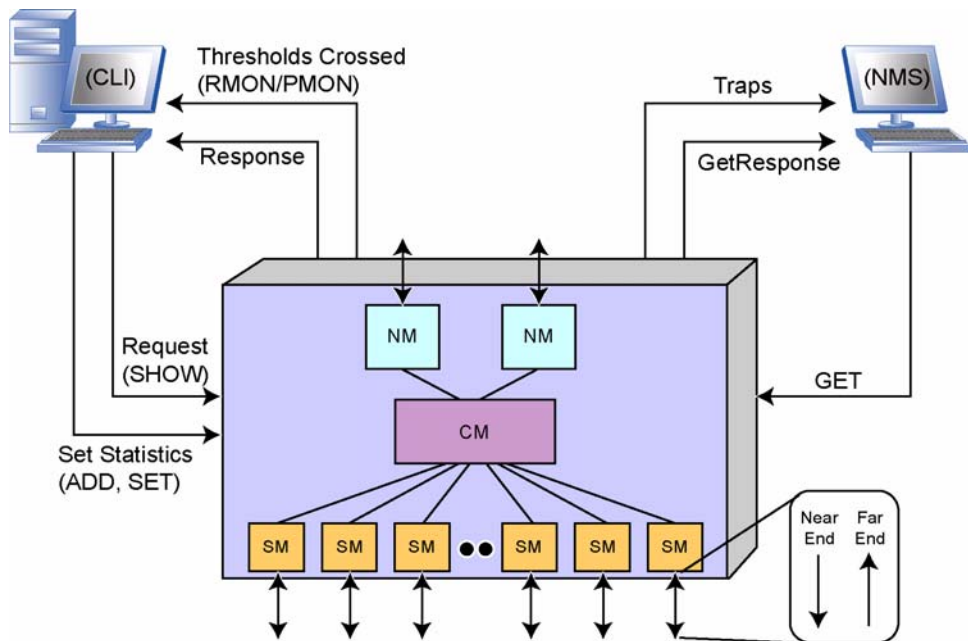


FIGURE 9-1 Performance Management Configuration for the 9000

Performance monitoring is based on two types of statistics set by commands:

1. **RMONALERT (based on RFC2819, RFC3273)**, which is for *Ethernet-type statistics*, has the following attributes:
 - Interval times are controllable (defaults are set at 30 seconds and 30 minutes).
 - Two thresholds can be set:
 - Rising - When a statistical value rises above a threshold, a log occurs, as well as a trap to an SNMP browser
 - Falling - When a statistical value drops below a threshold, an associated log/trap is produced.
 - During the time interval, once the falling threshold is crossed, the rising threshold trap/log is produced again (this is explained in [Figure 9-2](#)).
2. **PMONALERT (based on RFC2662, RFC3440)**, which is for *ADSL statistics*, has the following attributes:
 - Interval times are set (15 minutes).
 - When a threshold is crossed, a log occurs, as well as a trap to an SNMP browser, but that is the only time the alert is produced during the 15 minute or 24 -hour period.
 - At the end of the time interval, the statistical value is reset to zero.

Note: When Performance Measurements are set to on, they cannot generate traps or logs unless a threshold is explicitly set.

Note: *Once ports are enabled and thresholds are set, management logs can be produced. These management logs have the category of RMON (for ethernet-based logs) or ADSL (for PMON-based logs). Regardless of the type of threshold, these logs have a severity level of NONE.*

For the fMAP series, both PMON and RMON statistics are monitored, since there is an ADSL interface on the network or subscriber side. For the list of the RMON and PMON statistics supported by the fMAP products for this release, refer to [Section 19.1](#) (RFCs can be found on www.ietf.org.)

9.2 RMON (Ethernet-Based) Statistics

9.2.1 Overview

[Table 9-1](#) lists the RMON statistics supported by the fMAP product.

Note: For the ports on the GE3 card, the following count both egress and ingress packets:

- etherStatsPkts65to127Octets
- etherStatsPkts128to255Octets
- etherStatsPkts256to511Octets
- etherStatsPkts512to1023Octets

- etherStatsPkts1024to1518Octets

TABLE 9-1 RMON Statistics

Statistic	Description
etherStatsDropEvents	Change alert setting for dropped packet events
etherStatsOctets	Change alert settings for octets
etherStatsPkts	Change alert setting for packets
etherStatsBroadcastPkts	Change alert setting for broadcast packets
etherStatsMulticastPkts	Change alert setting for multicast packets
etherStatsCRCAlignErrors	Change alert setting for CRC alignment errors
etherStatsUndersizePkts	Change alert settings for undersize packets
etherStatsOversizePkts	Change alert setting for oversize packets
etherStatsFragments	Change alert setting for fragmented packets
etherStatsJabbers	Change alert setting for jabbers
etherStatsCollisions	Change alert setting for packet collisions
etherStatsPkts64Octets	Change alert settings for packets that are up to 64 octets long
etherStatsPkts65to127Octets	Change alert settings for packets that are 65 to 127 octets long
etherStatsPkts128to255Octets	Change alert settings for packets that are 128 to 255 octets long
etherStatsPkts256to511Octets	Change alert settings for packets that are 256 to 511 octets long
etherStatsPkts512to1023Octets	Change alert settings for packets that are 512 to 1023 octets long
etherStatsPkts1024to1518Octets	Change alert settings for packets that are 1024 to 1518 octets long

Figure 9-2 shows how the rising and falling thresholds work for the RMON statistics.

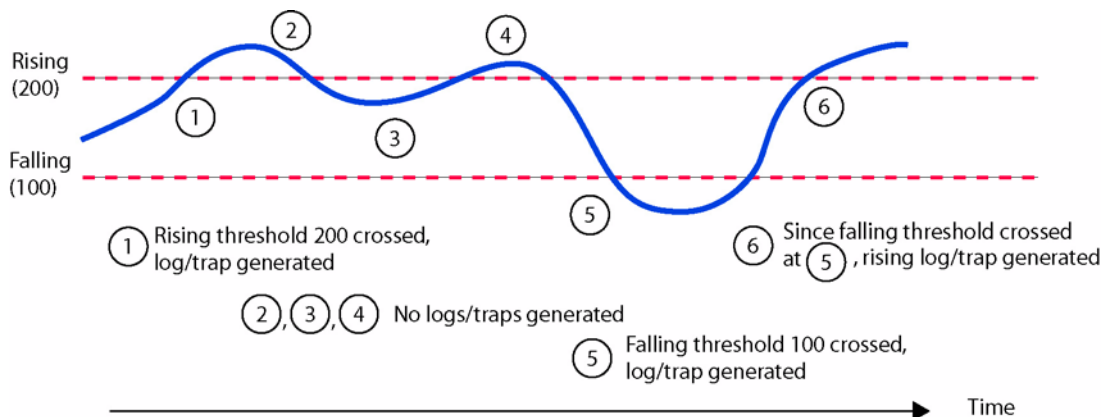


FIGURE 9-2 Thresholds for RMON Statistics

9.2.2 Example of Configuring RMON Statistics

Note: After a CFC Swap of activity on a duplex system or a SM Restart on a simplex or duplex system, disregard the first interval of RMON statistics.

An example of enabling the RMON (ethernet-based) statistic follows:

- **SHOW CARD=ALL** - This lists the cards for the device and shows which interfaces (eth) can be configured.
- **SET INTERFACE=10.0 COUNTER ON** - This would enable the 10.0 interface for Performance Management for the ethernet statistics.
- **SHOW INTERFACE=10.0 COUNTER STATUS** - This queries the current counters.

Note that since no thresholds have been set, no traps or logs can be generated. Following is an example of the use of these commands:

```
officer SEC>> SET INTERFACE=11.0 COUNTER ON
Info (032016): Counters turned on for interfaces ETH: 11.0
Info (010017): Operation Successful
officer SEC>> SHOW INTERFACE=ETH: 11.0 COUNTER
--- Ether Statistics -----
Interface: ETH: 11.0
State      : On (Collecting)

DropEvents... 0          BroadcastPkts. 0          MulticastPkts. 0
CRCAlignErrs.. 0          UndersizePkts. 0          OversizePkts.. 0
Fragments..... 0          Jabbers..... 0          Collisions.... 0

Name                               Current Counts      Overflow
Counters                          Counts
-----
Packets                             0                   0
Octets                              0                   0
Packets 64 Octets                    0                   0
Packets 65 to 127 Octets             0                   0
Packets 128 to 255 Octets            0                   0
Packets 256 to 511 Octets            0                   0
Packets 512 to 1023 Octets           0                   0
Packets 1024 to 1518 Octets          0                   0
```

Below is an example of enabling the RMON (ethernet-based) thresholds for PACKETS. The PACKETS parameter indicates that the rising/falling threshold values are to be used for the packets statistical counter. The user sets the rising and falling thresholds for each parameter. The user would use the ADD INTERFACE command to provision each of the required statistics in a similar manner. See the **fMAP Command Handbook** for more information about the ADD INTERFACE command. The sample commands are:

- **ADD INTERFACE=ETH: 11.0 RMONALERT PACKETS ABSOLUTE INTERVAL=5 RISINGTHRESHOLD=100000 FALLINGTHRESHOLD=10000** - This adds the PACKET RMONALERT threshold.
- **SHOW INTERFACE=11.0 COUNTER STATUS** - Since a threshold has been set, the alarm table for the statistic(s) is shown.

```
officer SEC> ADD INTERFACE=ETH: 11.0 RMONALERT PACKETS ABSOLUTE INTERVAL=5 RISINGTHRESHOLD=100000 FALLINGTHRESHOLD=10000
Info (032008): ADD operation for PACKETS affected interfaces ETH: 11.0
Info (010017): Operation Successful
officer SEC>> SHOW INTERFACE=ETH: 11.0 COUNTER
```

```
--- Ether Statistics -----
Interface: ETH: 11.0
State      : On (Collecting)

DropEvents... 0          BroadcastPkts. 0          MulticastPkts. 0
CRCAlignErrs.. 0          UndersizePkts. 0          OversizePkts.. 0
Fragments..... 0          Jabbers..... 0          Collisions.... 0

Name                      Current Counts          Overflow Counts
-----
Packets                    0                        0
Octets                    0                        0
Packets 64 Octets         0                        0
Packets 65 to 127 Octets  0                        0
Packets 128 to 255 Octets 0                        0
Packets 256 to 511 Octets 0                        0
Packets 512 to 1023 Octets 0                        0
Packets 1024 to 1518 Octets 0                        0

Name                      Sample Type  Interval  Rising Threshold  Falling Threshold
-----
Packets                    Absolute     5         100000            10000
```

When thresholds for RMON statistics are no longer needed, they should be deleted, in this example with:

```
officer SEC>> DELETE INTERFACE=ETH: 11.0 RMONALERT PACKETS
Info (032008): DELETE operation for PACKETS affected interfaces ETH: 11.0
Info (010017): Operation Successful
```

```
officer SEC>> SHOW INTERFACE=ETH: 11.0 COUNTER
--- Ether Statistics -----
Interface: ETH: 11.0
State      : On (Collecting)

DropEvents... 0          BroadcastPkts. 0          MulticastPkts. 0
CRCAlignErrs.. 0          UndersizePkts. 0          OversizePkts.. 0
Fragments..... 0          Jabbers..... 0          Collisions.... 0

Name                      Current Counts          Overflow Counts
-----
Packets                    0                        0
Octets                    0                        0
Packets 64 Octets         0                        0
Packets 65 to 127 Octets  0                        0
```

Packets 128 to 255 Octets	0	0
Packets 256 to 511 Octets	0	0
Packets 512 to 1023 Octets	0	0
Packets 1024 to 1518 Octets	0	0

Note that the PACKET RMONALERT threshold has been deprovisioned.

The user can stop the collection of packet statistics using the SET INTERFACE=<interface identifier> COUNTER OFF command:

```
officer SEC>> SET INTERFACE=ETH: 11.0 COUNTER OFF
Info (032016): Counters turned off for interfaces ETH: 11.0
Info (010017): Operation Successful
officer SEC>> SHOW INTERFACE=ETH: 11.0 COUNTER
```

```
--- Ether Statistics -----
Interface: ETH: 11.0
State      : Off (Not Collecting)

DropEvents... -      BroadcastPkts. -      MulticastPkts. -
CRCAlignErrs.. -      UndersizePkts. -      OversizePkts.. -
Fragments..... -      Jabbers..... -      Collisions.... -

Name                               Current Counts      Overflow
-----
Packets                             -                   -
Octets                              -                   -
Packets 64 Octets                    -                   -
Packets 65 to 127 Octets              -                   -
Packets 128 to 255 Octets              -                   -
Packets 256 to 511 Octets              -                   -
Packets 512 to 1023 Octets             -                   -
Packets 1024 to 1518 Octets            -                   -
```

9.2.3 Sample Management Logs for RMON (Ethernet-Based) Thresholds

The following log examples show the management logs produced when the set thresholds (in this case, etherStatsBroadcastPkts for a falling value of 10 and a rising value of 100) are crossed. Refer to [Figure 9-2](#) to understand how these thresholds work.

- Management log produced when rising threshold occurs:

```
RMON003 2004-06-16 13:27:07 6071 INFO
Location : ETH: 6.0
Description: etherStatsBroadcastPkts Rising Alert
Type : Absolute
Threshold : 100
Value : 0
```

- Management log produced when falling threshold occurs:

```
RMON004 2004-06-16 13:34:54 6080 INFO
Location : ETH: 6.0
Description: etherStatsBroadcastPkts Falling Alert
Type : Absolute
Threshold : 10
Value : 0
```

9.3 PMON (ADSL Port) Statistics

9.3.1 Overview

Table 9-2 lists the PMON statistics supported by the fMAP. In this table statistic names use the following naming conventions. For a complete description of these statistics, refer to **RFC2662** and **RFC3440**. Also, the RateUp and RateDown thresholds only apply to rate-adaptive ADSL mode, which is not currently supported.

- **Atuc** and **Atur** - Atuc refers to the upstream (near end) direction of a particular channel; Atur refers to downstream (far end).
- **Fast** and **Interleave** - These are the channel interfaces that can exist on a physical channel.
- “Loss of” names are Loss of Frame (**LOF**), Signal (**LOS**), Link (**LOL**), and Power (**LPR**). LPR is used by T1E1, so this is used for consistency (rather than LOP).
- **ES** - ES means Errored Second, any second containing one or more CRC anomalies, or one or more LOFs or Severely Errored Frame defects.
- **SesL** and **UasL** - SesL is Severely Errored Seconds-Line. UasL is Unavailable Seconds-Line.

Note: For certain modems/firmware, the ATUR LOFs and LOSs performance monitoring counts may not be correct.

TABLE 9-2 ADSL (PMON) Statistics

Statistic	Value	Description
adslAtucThresh15MinLOFs	0-900	The number of Loss of Frame Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAtucThresh15MinLOSs	0-900	The number of Loss of Signal Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAtucThresh15MinLOLs	0-900	The number of Loss of Link Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAtucThresh15MinLPRs	0-900	The number of Loss of Power Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAtucThresh15MinESs	0-900	The number of Errored Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAtucInitFailureTrapEnable	1-2	Enables and disables the InitFailureTrap.

TABLE 9-2 ADSL (PMON) Statistics (Continued)

Statistic	Value	Description
adslAturThresh15MinLOFs	0-900	The number of Loss of Frame Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAturThresh15MinLOSs	0-900	The number of Loss of Signal Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAturThresh15MinLPRs	0-900	The number of Loss of Power Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAturThresh15MinESs	0-900	The number of Errored Seconds encountered by an ADSL interface within any given 15 minutes performance data collection period.
adslAtucThresh15MinSesL	0-900	The first time the value of the corresponding instance of adslAtucPerf15MinSesL reaches or exceeds this value within a given 15-minute performance data collection period.
adslAtucThresh15MinUasL	0-900	The first time the value of the corresponding instance of adslAtucPerf15MinUasL reaches or exceeds this value within a given 15-minute performance data collection period.
adslAturThresh15MinSesL	0-900	The first time the value of the corresponding instance of adslAturPerf15MinSesL reaches or exceeds this value within a given 15-minute performance data collection period.
adslAtucSesLThreshTrap	0-900	Severely errored seconds-line 15-minute threshold reached
adslAtucUasLThreshTrap	0-900	Unavailable seconds-line 15-minute threshold reached.

9.3.2 Example of Configuring PMON Statistics (ADSL)

The following shows an example of setting PMONALERT (ADSL-based) statistics on. The sample commands are:

- **SHOW CARD=ALL** - This lists the cards for the product and shows which cards have ports that can be configured.
- **SET INTERFACE=0.4 COUNTER ON** - This enables interface (port) 0 on SHDSL card 16 for Performance Management.
- **SHOW INTERFACE=0.4 COUNTER** - This queries the current counters on interface 16.0.

Note that since no thresholds have been set, no traps or logs can be generated.

officer SEC>> **SHOW INTERFACE=0.4 COUNTER**

```

--- ADSL Statistics -----
Interface: ADSL:0.4
State      : On (Collecting)

Valid Intervals.... 32          Curr 15Min Elapsed.. 72
Invalid Intervals... 0          Curr 1Day Elapsed... 72
                                   Prev Day Elapsed.... 0
    
```

ATU-C (Receive)

Name	15Min Thresh	Curr 15Min	Curr 1Day	Prev Day	Failure Count
LOFs	-	0	0	0	0
LOSSs	-	0	0	0	0
LOLS	-	0	0	0	0
LPRs	-	0	0	0	0
ES	60	0	0	0	0
SES	-	0	0	0	0
UAS	-	0	0	0	0
Ini ts	N/A	0	0	0	0
FastRetry	N/A	0	0	0	0
Fai l FastRetry	-	0	0	0	0

ATU-R (Transmit)

Name	15Min Thresh	Curr 15Min	Curr 1Day	Prev Day	Failure Count
LOFs	-	0	0	0	0
LOSSs	-	0	0	0	0
LPRs	-	0	0	0	0
ES	30	0	0	0	0
SES	N/A	0	0	0	0
UAS	N/A	0	0	0	0

--- Ether Statistics -----

```

Interface: ETH:[0.4.0]
State      : On (Collecting)
    
```

Name	Current Counts	Overflow Counts
Packets		0 0
Octets		0 0
Broadcast Packets		0 N/A
Mul ti cast Packets		0 N/A

The following shows an example of setting the PMONALERT (ADSL-based) thresholds. The sample commands are:

- **SHOW INTERFACE=0.4 COUNTER** - This is to show the threshold information. Since all thresholds are at 0, no thresholds will be logged or trapped.

The example command below sets the PMON thresholds for ES on this interface. The ES parameter is used to set a limit on the number of allowed errored seconds over a fifteen minute interval. The user would use the SET INTERFACE command to provision each of the required PMON statistics in a similar manner. See the **fMAP Command Handbook** for more information about the SET INTERFACE command. The sample commands are:

- **SET INTERFACE=0.4 PMONALERT ATUC ES 60** - This sets the PMON threshold for ATUC ES (Errored Seconds) at 60 on ADSL interface 0.4.
- **SET INTERFACE=0.4 PMONALERT ATUR ES 30** - This sets the PMON threshold for ATUR ES (Errored Seconds) at 30 on ADSL interface 0.4.
- **SHOW INTERFACE=0.4 COUNTER** - Since a threshold has been set for ADSL interface 0.4, this is now shown.

```

officer SEC>> SET INTERFACE=0.4 PMONALERT ATUC ES 60
Info (010017): Operation Successful
officer SEC>> SET INTERFACE=0.4 PMONALERT ATUR ES 30
Info (010017): Operation Successful
officer SEC>> SHOW INTERFACE=0.4 COUNTER
--- ADSL Statistics -----
Interface: ADSL: 0.4
State      : On (Collecting)

Valid Intervals.... 32          Curr 15Min Elapsed.. 138
Invalid Intervals... 0         Curr 1Day Elapsed... 18139
                                   Prev Day Elapsed.... 86403
    
```

ATU-C (Receive)

Name	15Min Thresh	Curr 15Min	Curr 1Day	Prev Day	Failure Count
LOFs	-	0	0	0	0
LOSs	-	0	0	0	0
LOLs	-	0	0	0	0
LPRs	-	0	0	0	0
ES	60	0	0	0	0
SES	-	0	0	0	0
UAS	-	0	0	0	0
Inits	N/A	0	0	0	1
FastRetry	N/A	0	0	0	0
FailFastRetry	-	0	0	0	0

ATU-R (Transmit)

Name	15Min Thresh	Curr 15Min	Curr 1Day	Prev Day	Failure Count
LOFs	-	0	0	0	0
LOSs	-	0	0	0	0
LPRs	-	0	0	0	0
ES	30	0	0	1	3
SES	N/A	0	0	0	0
UAS	N/A	0	0	0	0

--- Ether Statistics -----

```

Interface: ETH: [0.4.0]
State      : On (Collecting)
    
```

Name	Current Counts	Overflow Counts
Packets	24087378	0
Octets	31828641226	0
Broadcast Packets	0	N/A
Multicast Packets	24087378	N/A

When thresholds for a PMON statistic is no longer needed, it should be set to 0, in this example with:

```
officer SEC>> RESET INTERFACE=0.4 COUNTER
Reset statistical counts (Y/N)? Y
Info (032009): Reset statistics for interface ADSL: [0.4]/ETH: [0.4.0]
Info (010017): Operation Successful
```

Once all the thresholds have been reset, the response to the **SHOW INTERFACE=0.4 COUNTER** will show all threshold values back at 0 and if the counter is set to ON (using the **SET INTERFACE=0.4 COUNTER ON** command), the registers will start accumulating statistics again. This is indicated in the Ether Statistics from the system output below.

```
officer SEC>> SHOW INTERFACE=0.4 COUNTER
--- ADSL Statistics -----
Interface: ADSL: 0.4
State      : On (Collecting)

Valid Intervals.... 32          Curr 15Min Elapsed.. 72
Invalid Intervals... 0          Curr 1Day Elapsed... 72
                               Prev Day Elapsed.... 0

ATU-C (Receive)
-----
Name          15Min Thresh  Curr 15Min   Curr 1Day   Prev Day   Failure
-----
LOFs          -                0           0           0           0
LOSSs        -                0           0           0           0
LOLs         -                0           0           0           0
LPRs         -                0           0           0           0
ES           60                0           0           0           0
SES          -                0           0           0           0
UAS          -                0           0           0           0
Inits        N/A                0           0           0           0
FastRetry    N/A                0           0           0           0
FailFastRetry -                0           0           0           0

ATU-R (Transmit)
-----
Name          15Min Thresh  Curr 15Min   Curr 1Day   Prev Day   Failure
-----
LOFs          -                0           0           0           0
LOSSs        -                0           0           0           0
LPRs         -                0           0           0           0
ES           30                0           0           0           0
SES          N/A                0           0           0           0
UAS          N/A                0           0           0           0

--- Ether Statistics -----
Interface: ETH: [0.4.0]
State      : On (Collecting)

Name          Current Counts   Overfl ow
-----
Packets          5201           0
Octets         6871966         0
Broadcast Packets 0             N/A
Multicast Packets 5201          N/A
```

9.3.3 ADSL24 Egress Queue Counts are not supported

The system does not support egress queue counts for the ADSL24 SM card. Following is an example. Card 9 is an ADSL24 card.

Note: Egress queue counts are supported for the ADSL24A/ADSL24B card, similar to the ADSL16 card.

officer SEC>> SHOW INTERFACE ALL QUEUE STATUS

--- Egress Queue Statistics Summary -----

Interface	State	Status	Dropped Packets			
			Queue 3	Queue 2	Queue 1	Queue 0
ETH: 0	-	Not Supported	-	-	-	-
ETH: 5. 0. 0	Off	Not Collecting	-	-	-	-
ETH: 5. 1. 0	Off	Not Collecting	-	-	-	-

(Some text omitted)

ETH: 7. 6. 0	Off	Not Collecting	-	-	-	-
ETH: 7. 7. 0	Off	Not Collecting	-	-	-	-
ETH: 9. 0. 0	-	Not Supported	-	-	-	-
ETH: 9. 1. 0	-	Not Supported	-	-	-	-
ETH: 9. 2. 0	-	Not Supported	-	-	-	-
ETH: 9. 3. 0	-	Not Supported	-	-	-	-
ETH: 9. 4. 0	-	Not Supported	-	-	-	-
ETH: 9. 5. 0	-	Not Supported	-	-	-	-
ETH: 9. 6. 0	-	Not Supported	-	-	-	-
ETH: 9. 7. 0	-	Not Supported	-	-	-	-
ETH: 9. 8. 0	-	Not Supported	-	-	-	-
ETH: 9. 9. 0	-	Not Supported	-	-	-	-
ETH: 9. 10. 0	-	Not Supported	-	-	-	-
ETH: 9. 11. 0	-	Not Supported	-	-	-	-
ETH: 9. 12. 0	-	Not Supported	-	-	-	-
ETH: 9. 13. 0	-	Not Supported	-	-	-	-
ETH: 9. 14. 0	-	Not Supported	-	-	-	-
ETH: 9. 15. 0	-	Not Supported	-	-	-	-
ETH: 9. 16. 0	-	Not Supported	-	-	-	-
ETH: 9. 17. 0	-	Not Supported	-	-	-	-
ETH: 9. 18. 0	-	Not Supported	-	-	-	-
ETH: 9. 19. 0	-	Not Supported	-	-	-	-
ETH: 9. 20. 0	-	Not Supported	-	-	-	-
ETH: 9. 21. 0	-	Not Supported	-	-	-	-
ETH: 9. 22. 0	-	Not Supported	-	-	-	-
ETH: 9. 23. 0	-	Not Supported	-	-	-	-

9.3.4 Sample Management Logs for PMON (ADSL-Based) Thresholds

Below is a management log produced when the threshold set in this example (adslatucthresh15minlofs) is exceeded, meaning the LOF condition has created a trap for more than 10 seconds of the 900 second period.

```
(MANAGEMENT LOG PRODUCED WHEN Condition has occurred more than set value)
ADSL031 2004-06-17 10:37:02 3285 OTHER
Location : ADSL:7.0
Description: ATU-C Unavailable Seconds(UAS) TCA
Threshold : 10
Value : 72
```

Note: Refer to the fMAP Log / Troubleshooting Manual for descriptions of log messages.

9.3.5 Commands for Setting Up Performance Management

TABLE 9-3 Performance Management Commands - Set Up

Object	Verb	Syntax	Description
INTERFACE RMONALERT INTERVAL RISINGTHRESH- OLD FALLINGTHRESH- OLD	ADD	ADD INTERFACE={ type:id-range id-range ifname-list ALL } RMONALERT { DROPEVENTS OCTETS PACKETS BROADCAST MULTICAST UNDERSIZE OVERSIZE CRCALIGN FRAGMENTS JABBERS COLLISIONS PKTS64OCTETS PKTS65TO127OCTETS PKTS128TO255OCTETS PKTS256TO511OCTETS PKTS512TO1023OCTETS PKTS1024TO1518OCTETS } { ABSOLUTE CHANGE } { INTERVAL=2..3600 RISINGTHRESHOLD=threshold FALLINGTHRESHOLD=threshold }	Configures an interface for RMON statistics and adds an interface to the alarmTable (RFC2819). RMONALERT- Refer to Table 9-1 .

TABLE 9-3 Performance Management Commands - Set Up (Continued)

Object	Verb	Syntax	Description
INTERFACE RMONALERT	SET	<pre> SET INTERFACE={ type:id-range id-range ifname-list ALL } RMONALERT { DROPEVENTS OCTETS PACKETS BROADCAST MULTICAST UNDERSIZE OVERSIZE CRCALIGN FRAGMENTS JABBERS COLLISIONS PKTS64OCTETS PKTS65TO127OCTETS PKTS128TO255OCTETS PKTS256TO511OCTETS PKTS512TO1023OCTETS PKTS1024TO1518OCTETS } [{ ABSOLUTE CHANGE }] [INTERVAL=2..3600] [RISINGTHRESHOLD=threshold] [FALLINGTHRESHOLD=threshold] </pre>	<p>Changes the attributes of the RMON statistic for the interface(s).</p> <p>The range for RISINGTHRESHOLD and FALLINGTHRESHOLD is 0..2147483647.</p> <p>If ABSOLUTE is chosen, the statistic must be reset before the threshold can be crossed again and a log produced.</p> <p>For RMONALERT refer to Table 9-1.</p>

TABLE 9-3 Performance Management Commands - Set Up (Continued)

Object	Verb	Syntax	Description
INTERFACE PMONALERT ADSL	SET	<pre> SET INTERFACE={ type:id-range id-range ifname-list ALL } PMONALERT ADSL { ATUC [LOFS=0..900] [LOSS=0..900] [LPRS=0..900] [ES=0..900] [SES=0..900] [UAS=0..900] [LOLS=0..900] [FAILEDFASTRETRAIN=threshold] ATUR [LOFS=0..900] [LOSS=0..900] [LPRS=0..900] [ES=0..900] </pre>	<p>Changes the attributes of the PMON statistic for the interface(s).</p> <p>For PMONALERT, refer to Table 9-2.</p> <p><i>Note</i> - This is for an ADSL interface. For other interfaces, there are other keywords (SHSDL, PSPAN, PPP, etc.) than ADSL. These are explained in other subsections. Also, the command without an interface keyword is no longer used.</p>
INTERFACE COUNTER	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } COUNTER { ON OFF } </pre>	For the collection of statistics, activates the interface(s).
	RESET	<pre> RESET INTERFACE={ type: type:id-range id-range ifname-list ALL } COUNTER [FORCE] </pre>	Resets the counters to 0 and restarts the time interval for the interface(s).

TABLE 9-3 Performance Management Commands - Set Up (Continued)

Object	Verb	Syntax	Description
INTERFACE RMONALERT	DELETE	<pre>DELETE INTERFACE={ type:id-range id-range ifname-list ALL } RMONALERT { DROPEVENTS OCTETS PACKETS BROADCAST MULTICAST UNDERSIZE OVERSIZE CRCALIGN FRAGMENTS JABBERS COLLISIONS PKTS64OCTETS PKTS65TO127OCTETS PKTS128TO255OCTETS PKTS256TO511OCTETS PKTS512TO1023OCTETS PKTS1024TO1518OCTETS ALL} [FORCE]</pre>	Deletes the RMON statistic from the interface(s)
INTERFACE COUNTER	SHOW	<pre>SHOW INTERFACE={ type: type:id-range id-range ifname-list ALL } COUNTER [{ STATUS FULL }]</pre>	Show the statistics for the interface(s). COUNTER shows the thresholds that have been set.

TABLE 9-3 Performance Management Commands - Set Up (Continued)

Object	Verb	Syntax	Description
INTERFACE FAULTCOUNT	SHOW	<pre> SHOW INTERFACE={ type: type:id-range id-range ifname-list ALL } FAULTCOUNT </pre>	Display the fault counts for the interface(s).
INTERFACE QUEUECOUNT	SHOW	<pre> SHOW INTERFACE={ type: type:id-range id-range ifname-list ALL } QUEUECOUNT [STATUS] </pre>	Display the queue counts for the interface(s).

9.4 Monitoring Performance Management

9.4.1 Overview

To view a history of performance management statistics, an interface or set of interfaces is associated with the following:

- **Interval** - This is the number of seconds in which the bucket gathers the statistics until the next bucket begins collecting data.
- **Bucket** - This is the container or collection of the historical data.

Note: For historical data to be collected, the interface(s) must be enabled to collect the data.

The total amount of data collected depends on the intervals and the number of buckets configured. As shown in Figure 9-3, when bucket 1 is full (a bucket being full when the interval has occurred), it becomes bucket 2, and the new bucket 1 contains the most recent historical data. When the last bucket (in this case bucket 4) is full and new data arrives, the data in bucket 3 becomes bucket 4, and the data that was in bucket 4 is discarded.




FIGURE 9-3 Data Collection Using Buckets (4) and Intervals (60)

Note: Buckets are a system wide resource. For RMON data, the maximum number of buckets that can be configured for a fMAP product is 2700. If the user tries to add a history that will go over the 2700 limit, the system will allocate what is available and produce a message saying how many buckets were granted.

9.4.2 RMON Collection

9.4.2.1 Overview

To collect RMON data, the user must explicitly associate an interface or set of interfaces with the number of

buckets and the interval each bucket collects data.

To do this, the user creates an interface history and adds an interval and number of buckets to an interface using the `ADD INTERFACE= <interface ID> COUNTER HISTORY INTERVAL=interval -list [BUCKETS=1..2700]`.

After associating an interface with an interval and buckets, the user can modify the number of buckets if needed using the `SET INTERFACE=<interface ID> COUNTER HISTORY [INTERVAL={interval -list|ALL}][BUCKETS=1..2700]`.

To modify the interval, the user must delete the interface history and read it with the modified interval using the `DELETE INTERFACE=<interface ID> COUNTER HISTORY [INTERVAL={interval -list|ALL}]`.

Finally, the RMON history for an interface or set of interfaces can be displayed with the `SHOW` command.

9.4.2.2 Sample Configuration

Figure 9-4 shows a possible traffic monitoring configuration that has the following:

- Traffic is being collected for one of the ports on the GE3 card.
- Each bucket will collect 1800 seconds (30 minutes) of data
- There are 48 buckets, so data older than 24 hours will be discarded.




FIGURE 9-4 Data Collection Using Buckets (48) and Intervals (1800 seconds)

9.4.2.3 Adding/Modifying RMON History Data for Interfaces

To configure data monitoring for the set of interfaces, use the command `ADD HISTORY INTERFACE` with the `INTERVAL` and `BUCKETS` parameters.

To modify the number of buckets:

1. Input the `SHOW` command for the interfaces.
2. Use the `SET INTERFACE` command to modify the number of buckets.

Note: To modify the interval, the HISTORY INTERFACE must be deleted (DELETE INTERFACE HISTORY) and then a new HISTORY INTERFACE must be added.

```
officer SEC> ADD INTERFACE=10.0 COUNTER HISTORY INTERVAL=1800 BUCKETS=48
Info (033865): ADD History operation successful for interfaces ETH:10.0
Info (010017): Operation Successful
```

```
officer SEC> SHOW INTERFACE=10.0 COUNTER HISTORY STATUS
```

```
--- Ether History -----
Buckets Available:2654

Interface C Int  Requested  Granted  Coll
-----
ETH:10.0  Y 1800      48       48      0
-----
```

9.4.2.4 Viewing RMON Data

The command to view data, `SHOW INTERFACE=Interface-Id COUNTER HISTORY`, includes RMON data.

9.4.3 PMON Collection

9.4.3.1 Overview

For PMON collection, the buckets and interval are active once an interface is enabled.

Note: The maximum number of buckets in the RFC 2662 is 96, or 24 hours of data collection. For this release, the interval is set at 15 minutes (900 seconds), and the maximum number of buckets is set at 32, so the system can hold 8 hours of PMON data for each interface.

9.4.3.2 Showing PMON History for Interfaces

The **SHOW INTERFACE=interface-id COUNTER HISTORY** command will show PMON as well as RMON data.

```

officer SEC>> SHOW INTERFACE=ADSL: 9.8 COUNTER HISTORY
----- ADSL History -----
Interface      C Coll Bkt  Interval      ATU-C Statistics      ATU-R Statistics
-----
Interface      C Coll Bkt  Start         ES   SES   UAS         ES   SES   UAS
-----
ADSL: 9.8      Y   0   1   -             -   -   -             -   -   -
-----

officer SEC>> SHOW INTERFACE=9.8 COUNTER
----- ADSL Statistics -----

Interface: ADSL: 9.8
State      : On (Collecting)

Valid Intervals.... 0          Curr 15Min Elapsed... 203
Invalid Intervals... 0          Curr 1Day Elapsed... 203
Prev Day Elapsed.... 0

ATU-C (Receive)
Name          15Min Thresh Curr 15Min  Curr 1Day  Prev Day  Failure
-----
LOFs          -             0         0         0         0         0
LOSSs         -             0         0         0         0         0
LOLs          -             0         0         0         0         0
LPRs          -             0         0         0         0         0
ES            -             0         0         0         0         0
SES           -             0         0         0         0         0
UAS           -             0         0         0         0         0
Inits         N/A           1         1         0         0         1
FastRetry     N/A           0         0         0         0         0
FailFastRetry -             0         0         0         0         0

ATU-R (Transmit)
Name          15Min Thresh Curr 15Min  Curr 1Day  Prev Day  Failure
-----
LOFs          -             0         0         0         0         0
LOSSs         -             0         0         0         0         0
LPRs          -             0         0         0         0         0
ES            -             1         1         0         0         1
SES           N/A           0         0         0         0         0
UAS           N/A           0         0         0         0         0

----- Ether Statistics -----

Interface: ETH: [9.8.0]
State      : On (Collecting)

Name          Current Counts      Overflow
-----
Packets       0                   0
Octets        0                   0
Broadcast Packets 0                   N/A
Multicast Packets 0                   N/A
-----
    
```

9.4.3.3 History for SHDSL Interfaces

Historical display shows a summary regardless of whether a single or multiple interfaces are requested. For example, the display for a specific interface for SHDSL looks like the following:

```
officer SEC>> SHOW INTERFACE=5.0-5.2 COUNTER HISTORY
--- SHDSL History ---
-----
Interface      C Coll Bkt   Interval          STU-C Stats      STU-R Stats
-----
                C Coll Bkt   Start          ES      SES   UAS      ES      SES   UAS
-----
SHDSL:5.0      Y   2   1 15:15:00         0       0   0         0       0   0
                2 15:00:00         0       0   0         0       0   0
SHDSL:5.1      N   1   1 15:00:00         0       0   0         0       0   0
SHDSL:5.2      Y   2   1 15:15:00         0       0   0         0       0   0
                2 15:00:00         0       0   0         0       0   0
-----
--- Ether History ---
No Ether History Data information to display
```

A detailed presentation is also available for history through use of the FULL option. The following is an example of the full display for SHDSL history:

```
officer SEC>> SHOW INTERFACE=5.0 COUNTER HISTORY FULL
--- SHDSL History ---
-----
Interface      C Coll   Valid   Invalid
                C Coll   Intervals Intervals
-----
SHDSL:5.0      Y   2       2       0
STU-C (Customer Side)
Name           Bucket: 1      Bucket: 2
-----
IntervalStart 15:15:00      15:00:00
LOSW           0              0
CRCAnomalies  0              0
ES             0              0
SES           0              0
UAS           0              0
STU-R (Network Side)
Name           Bucket: 1      Bucket: 2
-----
LOSW           0              0
CRCAnomalies  0              0
ES             0              0
SES           0              0
UAS           0              0
-----
--- Ether History ---
No Ether History Data information to display
```

The FULL display contains all of the statistics available unlike the standard display which includes just the high-runner statistics which are of most interest to a user. Though the interval start value is listed in the STU-C statistics collection, its value applies to both STU-C and STU-R relative to each bucket.

9.4.4 Commands for Collecting Performance Management Data

TABLE 9-4 Performance Management Commands -History

Object / Key Word(s)	Verb	Syntax	Description
INTERFACE COUNTER HISTORY INTERVAL BUCKETS	ADD	ADD INTERFACE={ type:id-range id-range ifname-list ALL } COUNTER HISTORY INTERVAL=interval-list [BUCKETS=1..2700]	Specifies data collection history for RMON statistics for an interface(s). A BUCKET is a collection of historical data. An INTERVAL is the time in seconds a bucket can collect data. The maximum number of buckets allowed for all data collection (RMON only) is 2700.
INTERFACE COUNTER HISTORY INTERVAL	DELETE	DELETE INTERFACE={ type:id-range id-range ifname-list ALL } COUNTER HISTORY [INTERVAL={ interval-list ALL }]	Deletes an interface from the interval association.

TABLE 9-4 Performance Management Commands -History (Continued)

Object / Key Word(s)	Verb	Syntax	Description
INTERFACE COUNTER HISTORY INTERVAL BUCKETS	SET	<pre>SET INTERFACE={ type:id-range id-range ifname-list ALL } COUNTER HISTORY [INTERVAL={ interval-list ALL }] [BUCKETS=1..2700]</pre>	Adds an interface to the bucket-interval association.
INTERFACE COUNTER HISTORY INTERVAL BUCKET	SHOW	<pre>SHOW INTERFACE={ type:id-range id-range ifname-list ALL } COUNTER HISTORY [STATUS] [INTERVAL={ interval-list ALL }] [BUCKET={ bucket-list ALL }] [FULL]</pre>	Allows the user to view data collection entries for Remote Monitoring (RMON) as well as bucket data collected for both RMON and Performance Monitoring (PMON).

9.5 Retrieving IP Statistics

Subsection 9.4 describes how to gather and retrieve statistical information. In addition, there are a number of statistics associated with SNMP MIBs that can be retrieved from both the CLI and an SNMP-enabled interface. Table 9-5 lists the CLI commands used and the associated MIB information.

Note: For command output, commands that provide the option of displaying one record will show two columns, with the first column listing the fieldnames and the second column listing the

associated values. When a command allows all records to be displayed, each record's data will be contained in one row. When a record exceeds an 80-column width, the data will wrap.

TABLE 9-5 SNMP Statistics

Command	MIB File	SNMP MIB
SHOW IP ROUTE[={ipaddress-list ALL}]	ipRouteEntry	rfc1213.mib
SHOW IP COUNTER={TCP UDP ICMP}	tcp	rfc1213.mib
SHOW IP CONNECTIONS [= {TCP UDP}]	tcpConnEntry	rfc1213.mib
SHOW IP COUNTER={TCP UDP ICMP}	udp	rfc1213.mib
SHOW IP CONNECTIONS [= {UDP TCP}]	udpEntry	rfc1213.mib
SHOW IP COUNTER={TCP UDP ICMP}	icmp	rfc1213.mib
SHOW IP [INTERFACE[={MGMT type:id-range if- name-list ALL}]] [FULL]	ifEntry	rfc1573.mib

9.5.1 IP Routing

The **SHOW IP ROUTE** command lists the routing table data, which could have been obtained from static routing data (manually entered) or dynamic routing data (using routing protocols such as OSPF or RIP).

When the command is entered with no values, the output is the routing table for all interfaces, and does not include routing metrics. Each row displays the next hop to use for the destination IP, type of routing, the routing protocol, the mask, and the age. In general, a route is chosen based on the row whose ageing timer has not expired (non-dirty) and has the largest match with the destination IP address. When a route is chosen the packet is forwarded to the next hop router entry in the row. If no match is found, the packet is forwarded to a default router.

When the command include an IP address, the specific row is listed, and includes the routing metrics. The following shows an example of this command and includes descriptions of specific columns.

```
officer SEC> SHOW IP ROUTE
```

```
--- IP Routing Table ---
-----
```

Intf Id	Dest Id	Next Hop	Type	Proto	Mask	Age
1	0.0.0.0	172.16.66.1	4	1	255.255.255.0	9969
2	127.0.0.1	127.0.0.1	3	2	255.255.255.0	10115
1	172.16.18.40	172.16.66.1	4	2	255.255.255.0	253
1	172.16.66.0	172.16.66.155	3	2	255.255.255.0	10115

9.5.2 TCP

The following commands list TCP-related information:

- **SHOW IP COUNTER=TCP** - This command displays overall TCP-related information for all TCP connections on the fMAP product.
- **SHOW IP TCP** - This command displays all information in the TCP table.

The following shows examples of these commands.

```
officer SEC> SHOW IP COUNTER=TCP
```

```
--- TCP Statistics Available on the Device ---
-----
```

TimeOut Determination Algorithm.....	4
Minimum Retransmission Timeout (in msec).....	1000
Maximum Retransmission Timeout (in msec).....	64000
Maximum allowed TCP connections.....	-1
Active Opens.....	0
Passive Opens.....	8
Attempt Failures.....	0
Connection Resets.....	0
Established Connections.....	1
Received Segments (incl. errors).....	2921
Transmitted Segments (incl. errors).....	3350
Retransmitted Segments.....	0
Received Error Packets.....	0
Transmitted Error Packets with RST flag.....	1

```
officer SEC> SHOW IP TCP
```

```
--- TCP System Data ---
-----
```

Id	State	Local IP	Local Port	Remote IP	Remote Port
0	2	0.0.0.0	23	0.0.0.0	0
1	5	172.16.66.55	23	172.16.18.40	4074

9.5.3 UDP

The following commands list UDP-related information:

- **SHOW IP COUNTER=UDP** - This command displays overall UDP-related information for all UDP connections on the fMAP product.
- **SHOW IP UDP** - This command displays all information in the UDP table.

The following shows examples of these commands.

```
officer SEC>SHOW IP COUNTER=UDP
```



```

--- UDP Statistics Available on the Device -----
Datagrams Received..... 4           Datagrams Transmitted..... 4
Error Datagrams Received 0         Datagrams with invalid port..... 4

```

```
officer SEC> SHOW IP UDP
```

```

--- UDP System Data -----
Local IP                               Local Port
-----
0.0.0.0                                0
0.0.0.0                                161
0.0.0.0                                1025

```

9.5.4 ICMP

The following command lists ICMP-related information:

- **SHOW IP COUNTER=ICMP** - This command displays the information in the Internet Control Message Protocol (ICMP) table.

The following shows an example of this command.

```
officer SEC> SHOW IP COUNTER=ICMP
```

```

--- ICMP Statistics Available on the Device -----
Received Packet Stats
-----
ICMP Packets..... 4
Error Packets..... 0
Time Exceeded Packets..... 4
Parameter Problem Messages..... 0
ICMP Source Quench..... 0
ICMP Redirect Messages..... 0
ICMP Echo Request Messages..... 0
ICMP Echo Reply Messages..... 0
ICMP Timestamp Request Messages..... 0
ICMP Timestamp Reply Messages..... 0
ICMP Address Mask Request Messages..... 0
ICMP Address Mask Reply Messages..... 0

Transmitted Packet Stats
-----
ICMP Packets..... 4
Error Packets..... 0
Time Exceeded Packets..... 4
Parameter Problem Messages..... 0
ICMP Source Quench..... 0
ICMP Redirect Messages..... 0
ICMP Echo Request Messages..... 0
ICMP Echo Reply Messages..... 0
ICMP Timestamp Request Messages..... 0
ICMP Timestamp Reply Messages..... 0
ICMP Address Mask Request Messages..... 0
ICMP Address Mask Reply Messages..... 0

```

9.5.5 Access of MIB Statistics Using an SNMP Browser

With an appropriate v1/v2c browser and with the relevant MIBs loaded, these same MIB statistics can be retrieved one at a time (using GET or GET-NEXT) or collectively (using GET-BULK). The following shows an example, where the GET operation is used repeatedly to get output for the entire tcpConnEntry row.

```

tcpConnState. 0. 0. 0. 0. 23. 0. 0. 0. 0. 0: -->listen(2)
tcpConnState. 90. 0. 0. 1. 23. 90. 0. 0. 254. 2526: -->established(5)
tcpConnLocal Address. 0. 0. 0. 0. 23. 0. 0. 0. 0: -->0. 0. 0. 0
tcpConnLocal Address. 90. 0. 0. 1. 23. 90. 0. 0. 254. 2526: -->90. 0. 0. 1
tcpConnLocal Port. 0. 0. 0. 0. 23. 0. 0. 0. 0: -->23
tcpConnLocal Port. 90. 0. 0. 1. 23. 90. 0. 0. 254. 2526: -->23
tcpConnRemAddress. 0. 0. 0. 0. 23. 0. 0. 0. 0: -->0. 0. 0. 0
tcpConnRemAddress. 90. 0. 0. 1. 23. 90. 0. 0. 254. 2526: -->254. 0. 0. 90
tcpConnRemPort. 0. 0. 0. 0. 23. 0. 0. 0. 0: -->0
tcpConnRemPort. 90. 0. 0. 1. 23. 90. 0. 0. 254. 2526: -->2526

```

9.6 Performance Monitoring for CES

As explained in Section 9., CES includes the following components:

- IP Address/VLAN - This makes up the IP Endpoint or IP Interface of the card (such as VLAN:402.0).
- DS1/E1 Port - The connecting point for the DS1 or E1 link.
- PSPAN - The IP-based encapsulation of the DS1 packets. This is the IP Endpoint and the unique PSPAN ID (such as 402.0.1).

For each one of these, various RMON and PMON statistics are available.

Note: With Circuit Emulation, far end statistics are not available. (In circuit emulation, the Line and not the Path of the DS1/E1 is terminated; since the far end statistics are embedded in the DS1/E1 frame, the fMAP device cannot access this information.)

9.6.1 IP Interface (RMON)

9.6.1.1 Overview

As explained in Section 9.4.2.3, an IP endpoint (the association of the VLAN and the IP address), is created using the ADD IP INTERFACE command, such as:

```
ADD IP INTERFACE=vlan:110.0 CARD=7 IFNAME=IP_End IPADDR=10.10.10.1 Sub-
net=255.255.255.0 GATEWAY=10.10.5.1
```

Once this is created and the connection is set up and running, RMON statistics can be measured for the interface using the SHOW INTERFACE <if_name> COUNTER ON command (the ON is the option that activates the RMON counters): Refer to the following:

```

officer SEC>> SET INTERFACE=vlan:400.1 COUNTER ON
Info (032016): Counters turned on for interfaces VLAN:[400.1]
Info (010017): Operation Successful
officer SEC>> SHOW INTERFACE=vlan:400.1 COUNTER
--- Ether Statistics -----
Interface: VLAN:[400.1]
State      : On (Collecting)
DropEvents.... 0          BroadcastPkts. 0          MulticastPkts. 0
RCAlignErrs... 0          UndersizePkts. 0          OversizePkts.. 0

```

```

Fragments..... 0          Jabbers..... 0          Collisions.... 0
                                Overflow
Name                          Current Counts      Counts
-----
Packets                        0                  0
Octets                         0                  0
Packets 64 Octets              0                  0
Packets 65 to 127 Octets      0                  0
Packets 128 to 255 Octets     0                  0
Packets 256 to 511 Octets    0                  0
Packets 512 to 1023 Octets   0                  0
Packets 1024 to 1518 Octets  0                  0

```

9.6.1.2 Threshold Setting

The IP interface also supports rising/falling alerts. When set, the alert configuration is visible in the SHOW INTERFACE COUNTER output. The following shows Octets being given threshold values.

```

officer SEC>> SET INTERFACE=vlan:400.1 RMONALERT OCTETS ABSOLUTE INTERVAL=900 RIS-
INGTHRESHOLD=4000 FALLINGTHRESHOLD=1000

```

```

Info (032008): SET operation for OCTETS affected interfaces VLAN:[400.1]
Info (010017): Operation Successful

```

```

officer SEC>> SHOW INTERFACE=vlan:400.1 COUNTER

```

```

--- Ether Statistics -----
Interface: VLAN:[400.1]
State      : On (Collecting)
DropEvents... 0          BroadcastPkts. 0          MulticastPkts. 0
CRCAAlignErrs.. 0          UndersizePkts. 0          OversizePkts.. 0
Fragments..... 0          Jabbers..... 0          Collisions.... 0

Name                          Current Counts      Overflow
-----
Packets                        0                  0
Octets                         0                  0
Packets 64 Octets              0                  0
Packets 65 to 127 Octets      0                  0
Packets 128 to 255 Octets     0                  0
Packets 256 to 511 Octets    0                  0
Packets 512 to 1023 Octets   0                  0
Packets 1024 to 1518 Octets  0                  0

Name                          Sample      Rising      Falling
                               Type        Interval   Threshold  Threshold
-----
Octets                        Absolute    900        4000       1000

```

Note: As with thresholds for other RMON data, management logs and SNMP traps are generated whenever the thresholds are crossed.

9.6.1.3 History Collection

Historical data collection is also the same as with other RMON collection.

9.6.2 DS1/E1 Port Performance (PMON)

9.6.2.1 Overview

The following line level statistics are collected:

- Line Coding Violations (CV-L) - A count of coding violations on the line
- Line Errored Seconds (ES-L) - The number of seconds with one or more coding violation or LOS defects

Note: When there are no packets received, the user will see an "Egress Path Failed". As the system begins to receive packets, the user will see Loss of Packet Stream (LOPS).

- SES-L : Line Severely Errored Seconds (there is no RFC mapping, this is part of the fMAP Enterprise MIB.)

Note: In Release 5.0 for the CES8, DS1 interfaces used "dsx1IntervalSESSs," a MIB variable, and mapped to the Line level. For the NTE8, "dsx1IntervalSESSs" should represent the Path-level SES rather than Line level SES. The CES8 also now maps to the Path-level SES in release 7.0.

- Line LOS Seconds (LOSS-L) - The number of seconds in which there was one or more LOS defects
- Line UAS (UAS-L) - The number of seconds in which the line was unavailable

Collection of BER is not supported.

The user can provision the thresholds and collection of all values.

Note: Path-level detection and reporting of AIS is not supported.

Thresholds, as defined in T1.231, are supported and accessible from both the CLI and SNMP Enterprise MIB.

E1 specific performance monitoring statistics ITU G.826, are collected. Certain statistics, specific to the E1 standard (e.g. Line Background Block Errors - BBE-L), will be collected, but can be displayed only for E1PORT instances.

All of the above statistics are managed using the `SET INTERFACE <interface> COUNTER {ON|OFF}` command. Statistics are similar to ADSL and SHDSL. When looking at multiple DS1/E1 interfaces, a summary view is provided. Refer to the following.

```
officer SEC>> SET INTERFACE=9.* COUNTER ON
Info (032016): Counters turned on for interfaces DS1:[9.0-7]
Info (010017): Operation Successful
```

```
officer SEC>> SHOW INTERFACE=9.* COUNTER
```

```
--- DS1/E1 Statistics Summary -----
Interface      C          LOSS          ES          SES          UAS          CV
-----
```

DS1:9.0	Y	0	0	0	9	0
DS1:9.1	Y	0	9	0	9	0
DS1:9.2	Y	0	9	0	9	0
DS1:9.3	Y	0	9	0	9	0
DS1:9.4	Y	0	9	0	9	0
DS1:9.5	Y	0	9	0	9	0
DS1:9.6	Y	0	0	0	10	0
DS1:9.7	Y	0	10	0	10	0

Detailed information is available by looking at an individual interface or by performing a `SHOW INTERFACE <interface> COUNTER FULL`. The following shows an example output:

```
officer SEC>> SHOW INTERFACE=DS1:9.0 COUNTER FULL
```

```
--- DS1/E1 Statistics -----
Interface: DS1:9.0
State      : On (Collecting)
  Valid Intervals..... 0
  Invalid Intervals... 0
Near End
  Name      15Min Thresh  Curr 15Min  Curr 1Day
  ----      -
  LOSS      N/A          0           0
  ES        -            0           0
  SES       -            0           0
  UAS       -            221        221
  CV        -            0           0
```

9.6.2.2 Threshold Setting

The user can set threshold crossing alerts that are invoked when an associated statistics count exceeds the user provided threshold. The following shows an example command:

```
officer SEC>> SET INTERFACE=DS1:9.0 PMONALERT DS1 NEAREND LINE CV=150 ES=75 SES=15
UAS=10
```

```
Info (010017): Operation Successful
```

```
officer SEC>> SHOW INTERFACE=DS1:9.0 COUNTER FULL
```

```
--- DS1/E1 Statistics -----
Interface: DS1:9.0

State      : On (Collecting)
  Valid Intervals..... 0
  Invalid Intervals... 0

Near End

  Name      15Min Thresh  Curr 15Min  Curr 1Day
  ----      -
  LOSS      N/A          0           0
  ES        75          0           0
  SES       15          0           0
  UAS       10          387        387
```


- BytesSent (Current counts only) - This is used to verify the expected number of bytes are being sent for this PSPAN while other PSPANs are also in operation.
- JitterBuffFillMin (Current, including the last 512 packets and 15-minute interval. See Note below.) - The minimum fill level of the jitter buffer. If this goes to 0, packet loss will occur.
- JitterBuffFillMax (Current, including the last 512 packets and 15-minute interval. See Note below.) - The maximum fill level of the jitter buffer. If this goes to the size of the jitter buffer, packet loss will occur.
- JitterBuffFillAvg (Current, including the last 512 packets and 15-minute interval. See Note below.) - The average fill level of the jitter buffer.

The jitter buffer values shown are from the most recent 15 minute interval. The max value is the max during the last 15 minute interval. The min is the min over the last 15 minute interval. And the average is the average over the last 15 minute interval. If there is packetization of more than 1 millisecond, then every packet is considered in those stats. If it is less than 1 millisecond, then the statistics are a sampling.

Note: For the 1Day value, since jitter buffer is a gauge, not a counter, the value within the last 512 packets is updated every 15 minutes to keep an historical record for the values.

Caution: Any time the user performs a configuration change (changing of timing or a loopback), the jitter buffer statistics should be checked to see if the actual values are far from the value set (JITTERBUFFER). If they are, the user should disable/enable the PSPAN. The jitter buffer levels should also be checked if there is an ES threshold being crossed and a log produced. In this case as well, if the actual values are far from the value set (JITTERBUFFER), the user should disable/enable the PSPAN.

```
officer SEC>> SHOW INTERFACE PSPAN:* COUNTER
```

```
--- PSpan Statistics Summary -----
```

Interface	C	LOPSS	ES
PSPAN:400.1.2	Y	0	0
PSPAN:400.1.3	Y	0	0

```
officer SEC>> SHOW INTERFACE pwe3a COUNTER
```

```
--- PSpan Statistics -----
```

```
Interface: PSPAN:400.1.2
State      : On (Collecting)
```

```
Valid Intervals..... 0
```

```
Invalid Intervals... 0
```

Name	15Min Thresh	Curr 15Min	Curr 1Day
LOPSS	-	0	0
ES	-	0	0
LatePkts	-	0	0
EarlyPkts	-	0	0
LostPkts	-	0	0
PacketsReceived	N/A	N/A	0
BytesReceived	N/A	N/A	0
PacketsSent	N/A	N/A	0
BytesSent	N/A	N/A	0
JitterBuffFillMin	N/A	N/A	0
JitterBuffFillMax	N/A	N/A	0
JitterBuffFillAvg	N/A	N/A	0

9.6.3.2 Threshold Setting

Thresholds can be set for the following:

- LOPSS
- ES
- LatePkts
- EarlyPkts
- LostPkts

The following shows example thresholds being set.

```
officer SEC>> SET INTERFACE=PSPAN:400.1.2 PMONALERT PSPAN SATOP ES=75 LOPSS=10
LATEPACKETS=250 EARLYPACKETS=150 LOSTPACKETS=50
```

```
Info (010017): Operation Successful
```

```
officer SEC>> SHOW INTERFACE pwe3a COUNTER
```



```

--- PSpan Statistics -----
Interface: PSPAN:400.1.2
State      : On (Collecting)

Valid Intervals..... 0
Invalid Intervals... 0

Name          15Min Thresh  Curr 15Min  Curr 1Day
-----
LOPSS                10           0           0
ES                   75           0           0
LatePkts            250           0           0
EarlyPkts           150           0           0
LostPkts             50           0           0
PacketsReceived     N/A          N/A          0
BytesReceived        N/A          N/A          0
PacketsSent          N/A          N/A          0
BytesSent            N/A          N/A          0
JitterBuffFillMin   N/A          N/A          0
JitterBuffFillMax   N/A          N/A          0
JitterBuffFillAvg   N/A          N/A          0
    
```

Like other PMON-oriented statistics, these values are handled as part of a 15-minute collection window.

Note: The packet and byte counts along with Jitter Buffer statistics are only collected as a day total and have no associated threshold in Release 5.0.

9.6.3.3 History Collection

History collection of PSPAN statistics is similar to other types.

Note: LOPSS and ES are shown if the FULL option is not included. If the FULL option is included, the other types are included.

The following is an example of PSPAN historical output

```
officer SEC>> show interface pspan:400.1.2 counter history
```

```

--- PSpan History -----
Interface      C Coll Bkt  Start      ES  LOPSS
-----
PSPAN:400.1.2  Y    7    1 21:45:00    0   0
                2 21:30:00    0   0
                3 21:15:00    0   0
                4 21:00:00    0   0
                5 20:45:00    0   0
                6 20:30:00    0   0
                7 *          *   *
    
```

Note: If two users are accessing information on the PSPAN at the same time, the following type message may appear:

```
officer SEC>> sh pspan all
Error (040909): The following Pspans failed
Pspan-IDs: c9p0, c9p1, c9p2, c9p3, c9p4, c9p5, c9p6, c9p7
reason: Maintenance in progress

Error (010009): Operation Failure
officer SEC>>
```

If a message like the above one appears, simply wait a retry the command.

9.7 Performance Monitoring for EPON

9.7.1 Overview

TABLE 9-6 EPON and ONU Statistics

Statistic	EPON Ingress ^a	EPON Egress (a)	ONU Ingress ^b	ONU Ingress (b)
RMON-like statistics common to all Ethernet interfaces (running count and user defined history)				
Total Octets	x	x	x	x
Total Frames	x	x	x	x
Unicast Frames	x	x	x	x
Broadcast Frames	x	x	x	x
Multicast Frames	x	x	x	x
CRC-32 Frames	x		x	
Undersize Frames	x		x	
Oversize Frames	x		x	
Collisions	x			
Size-based Packet Counts ^c			x	x
RMON-like statistics specific to EPON card interfaces (running count and user defined history)				
Frames Dropped ^d		x		x
Octets Dropped		x		x
CRC-8 (preamble) errors	x		x	
Octets Delayed (units of 100 us)		x		x
Octets Granted	x			x
Octets Granted but not used				x
PMON-like stats (current count, daily count, 15 min. interval history)				
Line-code Violations	x		x	
Errored Frame Seconds			x	

- a. EPON Interface Statistics are for the aggregate of links to all ONUs
- b. ONU statistics are for the aggregate of links and are collected on the EPON side
- c. 64-octet packets, 65-127 octet packets, etc.
- d. Counted only when they occur because of queue overflow conditions

9.7.2 Example Outputs

Note: The user should be aware that when displaying counters, there can be a large number for the octet count (especially ingress) while little actual traffic is being passed. For the egress octet count, this includes OAM packets, which are the heartbeat packets. For ingress, the octet count is showing OAM as well as the Multi-Point Control Protocol (MPCP) packets, which include reports from the ONUs. Therefore the ingress octet count may be much higher.

Note: The output for stats on the ONU are in relation to the EPON, so they show the EPON port side of the ONU. Ingress is therefore downstream from the EPON, and egress is upstream to the EPON. The output for the ETH interface shows the UNI port side.

```
officer SEC>> SHOW INT 2.0 COUNTER
```

```
--- Ether Statistics ---
```

Interface	C	Packets	Octets	Broadcast Packets	Multicast Packets
ETH:2.0.0	Y	14	896	0	14
ETH:2.0.1	Y	14	896	0	14
ETH:2.0.2	Y	13	832	0	13
ETH:2.0.3	Y	14	896	0	14
ETH:2.0.4	Y	14	896	0	14
ETH:2.0.5	Y	13	832	0	13
ETH:2.0.6	Y	14	896	0	14
ETH:2.0.7	Y	14	896	0	14
ETH:2.0.8	Y	13	832	0	13
ETH:2.0.9	Y	14	896	0	14
ETH:2.0.10	Y	14	896	0	14
ETH:2.0.11	Y	13	832	0	13
ETH:2.0.12	Y	13	832	0	13
ETH:2.0.13	Y	14	896	0	14
ETH:2.0.14	Y	75	22002	61	14

```
--- EPON Statistics ---
```

Interface	C	Total Octets	Total Packets
EPON:2.0	Y	2176640097	451

```
--- ONU Statistics ---
```

Interface	C	Total Octets	Total Packets
ONU:2.0.0	Y	5568	87
ONU:2.0.1	Y	5184	81
ONU:2.0.2	Y	2112	33
ONU:2.0.3	Y	4800	75
ONU:2.0.4	Y	4544	71
ONU:2.0.5	Y	1344	21
ONU:2.0.6	Y	3648	57
ONU:2.0.7	Y	4288	67
ONU:2.0.8	Y	2176	34
ONU:2.0.9	Y	5888	92

```
ONU:2.0.10    Y           4096           64
ONU:2.0.11    Y           2560           40
ONU:2.0.12    Y           1024           16
ONU:2.0.13    Y           6080           95
ONU:2.0.14    Y           5440           85
```

officer SEC>> sh int epon:2.1 counter

--- EPON Statistics ---

Interface: EPON:2.1

State : On (Collecting)

Name	Ingress	Egress
Total Octets	7983639	24614
Total Packets	0	7
Unicast Packets	0	0
Broadcast Packets	0	2
Multicast Packets	0	-
CRC-32 Errors	0	-
Undersize Packets	0	-
Oversize Packets	0	-
CRC-8 Errors	0	-
LCVs	0	-
Collisions	0	-
Packets Dropped	-	0
Octets Dropped	-	0
Octets Delayed	-	0
Octets Granted	0	-

officer SEC>> SHOW INTERFACE 2.0.0 COUNTER

--- Ether Statistics ---

Interface: ETH:[2.0.0]

State : On (Collecting)

```
DropEvents.... 0          BroadcastPkts. 0          MulticastPkts. 553
CRCAlignErrs.. 0          UndersizePkts. 0          OversizePkts.. 0
Fragments..... 0          Jabbers..... 0          Collisions.... 0
```

Name	Current Counts	Overflow Counts
Packets	553	0
Octets	35392	0
Packets 64 Octets	553	0
Packets 65 to 127 Octets	0	0
Packets 128 to 255 Octets	0	0
Packets 256 to 511 Octets	0	0
Packets 512 to 1023 Octets	0	0
Packets 1024 to 1518 Octets	0	0

--- ONU Statistics ---

Interface: ONU:2.0.0

State : On (Collecting)

Name	Ingress	Egress
Total Octets	40512	0
Total Packets	633	0
Unicast Packets	80	0
Broadcast Packets	0	0
Multicast Packets	553	4
CRC-32 Errors	0	-
Undersize Packets	0	-
Oversize Packets	0	-
CRC-8 Errors	0	-
LCVs	0	-
EFSSs	0	-
Packets Dropped	-	0
Octets Dropped	-	0
Octets Delayed	-	0
Octets Granted	-	0
Octets Not Used	-	0
Packets 64 Octets	0	4
Packets 65 to 127 Octets	0	0
Packets 128 to 255 Octets	0	0
Packets 256 to 511 Octets	0	0
Packets 512 to 1023 Octets	0	0
Packets 1024 to 1518 Octets	0	0
Packets >= 1519 Octets	0	0

officer SEC>> sh int onu: 2.1.0 co

--- ONU Statistics ---

Interface: ONU: 2.1.0

State : On (Collecting)

Name	Ingress	Egress
Total Octets	1106	0
Total Packets	11	0
Unicast Packets	0	0
Broadcast Packets	2	0
Multicast Packets	9	0
CRC-32 Errors	0	-
Undersize Packets	0	-
Oversize Packets	0	-
CRC-8 Errors	0	-
LCVs	0	-
EFSSs	0	-
Packets Dropped	-	0
Octets Dropped	-	0
Octets Delayed	-	0
Octets Granted	-	0
Octets Not Used	-	0
Packets 64 Octets	0	0
Packets 65 to 127 Octets	0	0
Packets 128 to 255 Octets	0	0
Packets 256 to 511 Octets	0	0
Packets 512 to 1023 Octets	0	0
Packets 1024 to 1518 Octets	0	0
Packets >= 1519 Octets	0	0

officer SEC>> sh int epon:2.1 co history

--- Epon RMON History ---

Interface	C	Int	Requested	Granted	Coll
EPON: 2.1	Y	15	10	10	10

Name	Bucket: 1	Bucket: 2	Bucket: 3
Interval Start	0:02:30	0:02:25	0:02:20
Ingress			
Total Octets	2062080	2061952	2067217
Total Packets	0	0	0
Unicast Packets	0	0	0
Broadcast Packets	0	0	0
Multicast Packets	0	0	0
CRC-32 Errors	0	0	0
Undersize Packets	0	0	0
Oversize Packets	0	0	0
CRC-8 Errors	0	0	0
Collisions	0	0	0
Octets Granted	0	0	0

10. File Management and Software Release Upgrade

10.1 Overview

A software release is the set of executable binary code that runs on system cards. Software releases are delivered in the form of executable files, or load files. Depending on the card, some will require a load, others may not. Card load files and the system configuration database are normally stored on the control module card. New functionality and feature content for the system is delivered in the form of new software releases. Users perform software upgrades to load new releases. Management of system files is important to maintaining optimum operational performance from the system. System file management, load management, database management, and the software upgrade process will be described in this section. Detailed software upgrade procedures for the fMAP is included in section [12.1](#) and [11.1](#), respectively.

fMAP offers two system configurations, duplex or simplex. Software upgrades for both configurations are described below. Some differences between a simplex and duplex system are:

- Simplex systems are equipped with a single control module, the active CFC (ACTCFC)
- Duplex systems are configured with two control modules, the active CFC (ACTCFC) and the inactive CFC (INACTCFC), providing a hot standby control module

For more information on simplex and duplex systems, refer to Section [4.5.1](#)

[Table 10-1](#) lists the main tasks to perform for file management and upgrade.

TABLE 10-1 Tasks for File Management and Upgrade

Task	Description	Section
File Management	Copying, transferring, renaming, and deleting files	10.2
Software Load Management	Designating files as software loads for cards	10.3
Boot Server	Designating a network server as the boot server	10.3.3
Database Management	Backing up, restoring, and purging the configuration database	10.4
Configuration Files	Backing up and restoring the database configuration using a text-based file	10.5

TABLE 10-1 Tasks for File Management and Upgrade

Task	Description	Section
Software Compatibility	Ensuring cards with different loads can communicate and downloading is possible.older and newer loads can	10.6
Software Upgrade Process	Upgrading the loads on the fMAP product cards	10.7

10.2 File Management

10.2.1 File Names

Load file names describe both the hardware type that uses the software and the version identification for the release. There is internal meta data imbedded within the file that contains the same information. This external and internal data is used during the loading and installation processes to verify and assure correct software installation for appropriate hardware. For example, the meta data insures that a control module load file will only load into and execute on a control module card.

These are examples of software release file:

- CFC24 card load file name: **CFC24_5.0.1.tar**
- FE10 card load file name: **FE10_5.0.1.tar**

Consider the load file names listed above. The load release is subdivided into three levels. The release level will be important and will be referred to during a software upgrade. They are as follows:




FIGURE 10-1 Release levels

Note: Software load files are delivered as tar files, and the ".tar" extension must be retained.

10.2.2 File Storage

Load files can be stored in numerous locations. For example, loads will be stored in FLASH on the control module, service module, and could be stored in the users network on a TFTP server, ZMODEM, or FTP server, or CFLASH unit.

FLASH memory is a nonvolatile, reusable memory device that allows storage of large volumes of data. RAM is volatile, dynamic memory that contains the executable software. Software loads are stored in FLASH and RAM

memory on the control module and service module. The primary function of FLASH memory on the system is to store software releases, simplifying the servicing and maintenance requirements of the system and reducing recovery time during system restorations. Control modules and service modules are shipped with release files already present in FLASH memory.

Note: The files on these control and service modules may or may not have the latest release.

For a duplex system, both control modules have the same files on their respective FLASH file systems. When both control modules are UP and are in sync, all file operations are automatically applied to both the active and inactive control module. If the inactive control module is DOWN, it is out-of-sync and file operations are not performed on it. When the inactive control module is enabled, it performs a bulk sync, and once complete, the two control modules have file operations applied to both.

Control modules have enough FLASH memory to store two release files for each card type, allowing older and new releases to coexist in FLASH memory during upgrade procedures. The configuration database is also stored in FLASH on the control modules. Service modules have enough memory to store a single copy of its software release file.

Note: When downloading files to the CFC24, a 1 meg buffer of memory should be allowed.

When performing an upgrade, management commands retrieve the new release files from a network host and load them into FLASH memory on the control module and service module.

10.2.3 Compact FLASH

10.2.3.1 Overview

Compact Flash (CFLASH) is a removable memory card that will insert into the faceplate of CFC24 Control Module. This extra memory can be used to store release load files, databases, and captured log files. 256MB of memory is available in CFLASH. Compact FLASH is a small portable card that the user inserts into the CFC24 similar to FLASH on a digital camera or other device.

CFLASH is simply a local device that performs the same actions as a network server, with exceptions listed below:

1. CFLASH cannot be used as a bootserver.
2. CFLASH does not support rebooting. The system will not reboot a system card from a load file located on CFLASH even the load is designated as PREFLOAD.
3. Storing a backup database on the CFC's own flash is NOT recommended.

10.2.3.2 Inserting CFLASH

The following system actions occur when CFLASH is inserted into a CFC24:

- The system will automatically create the CFLASH media object when the insertion is detected.

- The CFC24 will verify the type of device detected (media). The CFC24 will acquire data from the CFLASH to determine if the device is compatible with the system software.
- If the CFLASH is not supported by the software currently running on the system, a system log will be generated and the device will be left out of service.
- If the CFLASH is supported by the current system software load the device will be tested.
- If the CFLASH fails the testing, a minor alarm will be raised on the CFC24 and the device will be placed out of service.
- If the CFLASH test passes, the device will be brought into a valid service state consistent with the database.
- If the system is configured to log to the CFLASH this service will be enabled and the logging tasks will be notified that it may open (reopen) a file and start logging to CFLASH.
- A management log will be generated indicating the insertion of a CFLASH device.

The CFLASH card must be formatted and activated in order to use it. To format and activate the card, ensure that it is inserted into one of the control modules, then execute the `DEACTIVATE MEDIA` command, the `FORMAT MEDIA` command, and then the `ACTIVATE MEDIA` command. An example follows:

```
DEACTIVATE MEDIA=CFLASH13 FORCE
```

Where `CFLASH13` is the unit on the control module in slot 12/13. Note that the unit must be deactivated before it can be formatted.

```
FORMAT MEDIA=CFLASH13
```

Where `CFLASH13` is the unit on the control module in slot 12/13.

```
ACTIVATE MEDIA=CFLASH13
```

Where `CFLASH13` is the unit on the control module in slot 12/13.

10.2.3.3 Storing system load files on CFLASH

In order to store files onto CFLASH, use the `GET FILE` command, for example:

```
GET FILE=CFC24_2.0.0.tar TFTP SERVER=<ip address> TO=CFLASH13
```

This example gets a load file, `CFC24_5.0.0.tar`, from a TFTP server and puts it onto CFLASH on the control module in slot 12/13.

10.2.3.4 Back up and restore databases on CFLASH

Database management commands are available for the CFLASH. To support database management commands, the CFLASH must be in service with no alarms.

An example of backing up the database would be:

```
BACKUP DATABASE FILE=CFLASH13:DBBackup TFTP SERVER=Database_Server
```

Backups the current system database from the active control module to CFLASH on the control module in slot 12/13. The command will be rejected if the CFLASH is not in service or there is insufficient space in the CFLASH.

An example of restoring a database from CFLASH follows:

```
RESTORE DATABASE FILE=CFLASH13: DBBackup TFTP SERVER=Database_Server
```

Restores the system database from CFLASH on the control module in slot 12/13 to the active control module. The system will attempt to verify that the file being transferred is actually a database. The system will not transfer a non-database into CFC FLASH. Finally, this command results in a reboot of the control module.

10.2.3.5 Streaming system logs to CFLASH

System logs can be streamed to CFLASH. This is done by specifying CFLASH in the CREATE LOG and SET LOG commands. Examples follow:

```
CREATE LOG FILTER=Sys_Logs
```

Creates a log filter and names it Sys_Logs. Since no parameters were included in the command, all logs will be included in the output.

```
CREATE LOG OUTPUT=Sys_Logs FILE=CFLASH13: Logs_File
```

Creates a log file called Sys_Logs and start log streaming directly to the CFLASH on the control module in slot 12/13.

```
ADD LOG FILTER=Sys_Logs OUTPUT=Sys_Logs
```

This commands adds the log filter to the log out.

```
ENABLE LOG OUTPUT=Sys_Logs
```

This command enables a log file called Sys_Logs and starts log streaming directly to the CFLASH on the control module in slot 12/13. This command turns on and off logging. If the CFLASH attribute is added to DESTINATION, CFLASH will be added to the output devices. If the CFLASH attribute is left off, then logging to the CFLASH will terminate. If the file already exists, the system will open the file and append current system logs to it. Otherwise, the system will create the file and start the output.

10.2.3.6 Other commands

Other system file and log commands used for working with CFLASH are:

```
ACTIVATE MEDIA=unit
```

Executes diagnostics (chkdsk & other self tests) on the device, puts the device in *ready to use* state. Logging will resume if set to this device previously.

```
DEACTIVATE MEDIA=unit [FORCE]
```

where:

MEDIA=unit specifies the medium and the card slot this media is mounted on.

For example, `MEDIA=CFLASH13` indicates CFLASH mounted on control module residing in slot 12/13.

This command suspends operations to the device and prepares for extraction. It is used to suspend logging to the CFLASH. It can be used to gracefully interrupt logging to the device in preparation for removal. The command will stop logging, flush data to the CFLASH, and close the files so data loss is minimized.

Example: `DEACTIVATE MEDIA=CFLASH9`

Deactivates the CFLASH card mounted on the CFC card in slot 9.

PURGE MEDIA=unit t

Clears all files on the device. The command will warn the user that all files will be deleted and await a confirmation to delete. To execute this command, the card must be deactivated. This will ensure that all writing to the card is suspended during the clear.

SHOW MEDIA=unit t

Shows CFLASH device information such as amount of FLASH memory, format information, vendor etc.

Example of output:

```
Compact Flash..... CFLASH13
Device State..... Activated
CFC..... 12
Hardware found..... Yes
Serial Number..... AA435267711134D
Number of Sectors..... 250000
Size..... 125056 KB
Used..... 4 KB ( 9 files)
Free..... 125052 KB
```

DIAGNOSE MEDIA=unit t

This command will run `chkdsk` and other self tests on CFLASH. The device must already be deactivated.

PUT FILE={ sourcefile | unit:sourcefile } { TFTP SERVER={ ipaddress | hostname } | FTP SERVER={ ipaddress | hostname } USER=userid PASSWORD=password | ZMODEM } [TO=serverpath]

Sends a file from CFLASH to a server. *unit* specifies the shelf slot this CFLASH CFC is in. For example, `CFLASH9` specifies compact flash in CFC slot 9. *sourcefile* is the file name.

Example:

PUT FILE=CFLASH9:cfc6.tar TFTP SERVER=172.16.17.18 TO="Test1"

Transfers `cfc6.tar` file from CFLASH of CFC in slot 9 to a server with ip address 172.16.17.18. The parameter, `TO`, can be used to specify an absolute directory path on the server. In this example, the file will be sent to a directory called `Test1`.

SHOW TRANSFER={transferid|transferid-list|ALL}

Shows the status of transfer in progress between CFLASH and Servers.

STOP TRANSFER={transferid|transferid-list|ALL}

Allows the user to stop a transfer between CFLASH and server.

COPY FILE={filename|unit:filename} TO={filename|unit:filename}

This command copies files between CFLASH and local control module FLASH.

Example:

```
Copy FILE=script.txt TO=CFLASH9: scriptcopy.txt
```

This command will copy the file “script.txt” from the local CFC’s flash, to a file named “scriptcopy.txt” on CFLASH card mounted on CFC in slot 9.

```
DELETE FILE={filename|unit: filename} [FORCE]
```

This command deletes a file from the CFLASH device.

```
RENAME FILE={filename|unit: filename} TO={filename|unit: filename}
```

This command allows the user to rename a file on CFLASH. Note that the compact flash unit in source and destination must match if specified. The following command syntax is not allowed:

```
RENAME FILE=CFLASH9: script.txt TO=CFLASH13: newscrip t.txt
```

```
SHOW FILES [MEDIA=unit] [FULL]
```

This command will list all files on the CFLASH in the targeted slot.

```
PUT LOG FILE={ destinationfile | unit: destinationfile | serverpath/destinationfile } [ { TFTP  
SERVER={ ipaddress | hostname } | ZMODEM | FTP SERVER={ ipaddress | hostname } USER=userid  
PASSWORD=password } ] [ TYPE={ MGMT | ERROR | TRACE | CRASH } ] [ CARD={ ACTCFC | INACTCFC } ]
```

This command will put a log file on the active CFC’s local flash or CFLASH.

10.2.3.7 CFLASH device alarms

In all error conditions, the system will place the failed or missing device out of service and a minor alarm will be generated. Any logging to the device will be disabled until the device is placed back in service.

- **Device full**

An alarm indicating all available memory has been used will be generated when background streaming of logs to CFLASH is enabled and the system runs out of memory space in CFLASH. This alarm will be raised against the CFC card on which CFLASH is mounted. Disabling the log output will clear this alarm.

- **Device read / write failure**

The CFLASH device driver returned a read or write failure upon receiving a read or write request.

- **Device Improper Removal**

A management log will be generated indicating improper removal of CFLASH. No alarm will be generated.

10.2.4 Commands for File Management

The system user will normally work with two types of files, card load files, and script files. Commands are provided for the management of these files. The most often used file management commands are listed in [Table 10-2](#). Sample outputs are shown below.

TABLE 10-2 File Commands Summary

Noun	Verb	Syntax	Description
FILES	SHOW	SHOW FILES [FULL]	Displays all user manageable files that exist on the control module FLASH file system. Also displays file system memory consumption information.
FILES	SHOW	SHOW FILES MEDIA=unit [FULL]	Displays all user manageable files that exist on the CFLASH device.
SCRIPT	SHOW	SHOW SCRIPT= filename	Displays the contents of a CLI script file on the control module FLASH file system.
FLASH	SHOW	SHOW FLASH [INACTCFC]	Displays information about the flash memory on the control module.
FILE TO	COPY	COPY FILE={ sourcefile unit:sourcefile } TO={ destinationfile unit:destinationfile }	Makes a copy of an existing file. The new file will have a different file name.
FILE TO	RENAME	RENAME FILE={ sourcefile unit:sourcefile } TO={ destinationfile unit:destinationfile }	Renames the specified file on the control module FLASH file system or CFLASH.
FILE	DELETE	DELETE FILE={ filename unit:filename } [FORCE]	Delete a file from the control module FLASH file system or CFLASH.
NONPREF-LOADS	DELETE	DELETE NONPREFLOADS	Deletes ALL files that are stored in FLASH that are not set as PREFLOAD on any card.

TABLE 10-2 File Commands Summary (Continued)

Noun	Verb	Syntax	Description
<p>FILE <server> USER PASSWORD</p>	<p>GET</p>	<pre>GET FILE={ sourcefilename serverpath/sourcefile- name } { TFTP SERVER={ ipaddress hostname } ZMODEM FTP SERVER={ ipaddress hostname } USER=userid PASSWORD=password } [TO=unit:]</pre>	<p>Transfers a copy of the specified file from a specified network server to the control module FLASH file system or CFLASH.</p>
<p>FILE CARD</p>	<p>PUT</p>	<pre>PUT FILE=filename CARD={ slot slot-list }</pre>	<p>Transfers a copy of a specified existing file on the control module FLASH file system to a specified card or list of cards.</p>
<p>FILE <server> USER PASSWORD</p>	<p>PUT</p>	<pre>PUT FILE={ sourcefile unit:sourcefile } { TFTP SERVER={ ipaddress hostname } FTP SERVER={ ipaddress hostname } USER=userid PASSWORD=password ZMODEM } [TO=serverpath]</pre>	<p>Transfers a copy of the specified file on the control module FLASH file system or CFLASH to the specified network server.</p>

TABLE 10-2 File Commands Summary (Continued)

Noun	Verb	Syntax	Description
TRANSFER	SHOW	SHOW TRANSFER [={ transferid-list ALL }]	Displays current file transfer operations, including those in progress and those that are pending.
TRANSFER	STOP	STOP TRANSFER={ transferid-list ALL }	Aborts an in-progress or pending file transfers involving a network server. Intercard transfers cannot be aborted.
FILES	AUDIT	AUDIT FILES	Audits CRC on all *.TAR files and raises or clears file corruption alarm accordingly. Refer to 10.3.2 .

Following is a sample output of SHOW FILES command. Note that files beginning with uppercase are listed in order before files beginning with lowercase, which are also listed in order.

```
offi cer SEC> SHOW FILES
```

```
-----
File..... Size Kb
-----
Dec_01_cfg..... 7
Dec_01_out..... 16
Dec_01_out2..... 17
Dec_02_cfg..... 4
Dec_03_cfg..... 7
Dec_03_out..... 17
Nov_02_Timing_cfg..... 7
Nov_02_Timing_out..... 20
Nov_04_Timing_cfg..... 7
Nov_04_Timing_out..... 20
Nov_05_Timing_cfg..... 7
Nov_05_Timing_out..... 20
Nov_09_HalfTime_cfg..... 6
Nov_09_HalfTime_out..... 14
Nov_17_cfg..... 7
Nov_17_out..... 18
addpotsip.txt..... <1
adsl16_4_0_5.tar..... 1283
adsl16_4_1_3.tar..... 1283
adsl16_5_0_0.tar..... 1283
adsl16b_4_0_5.tar..... 1219
adsl16b_4_1_3.tar..... 1219
adsl16b_5_0_0.tar..... 1219
adsl24_4_0_5.tar..... 2563
adsl24_4_1_3.tar..... 2627
adsl24_5_0_0.tar..... 2627
adsl8_4_0_5.tar..... 1283
adsl8_4_1_3.tar..... 1283
adsl8_5_0_0.tar..... 1283
alias.scr..... <1
ces8_5_0_0.tar..... 1997
cfc24_3_0_3.tar..... 6713
cfc24_4_0_5.tar..... 6775
cfc24_4_1_3.tar..... 6793
cfc24_5_0_0.tar..... 7086
cfc24_60_pez2.tar..... 7219
cfg2.txt..... 14
del50files.txt..... <1
epsr_4_0_setup.txt..... 1
f51_4_0_cfg_dec09.txt..... 12
f51_4_0_setup.txt..... 5
f51_5_0_cfg.txt..... 14
```

```

f51_5_0_setup.txt..... 8
fe10_4_0_5.tar..... 3278
fe10_4_1_3.tar..... 3279
fe10_5_0_0.tar..... 3332
get405files.txt..... <1
get413files.txt..... <1
get50files.txt..... <1
inband.cfg..... <1
kevinConfig.txt..... 7
pots24_4_0_5.tar..... 4349
pots24_4_1_3.tar..... 4831
pots24_5_0_0.tar..... 4960
pspans.txt..... 1
set405pref.txt..... <1
set413pref.txt..... <1
set50pref.txt..... <1
shdsl16_4_0_5.tar..... 1219
shdsl16_4_1_3.tar..... 1219
shdsl16_5_0_0.tar..... 1219

Allowed KiB..... 112000
Used KiB..... 98150
Available KiB..... 13849

```

Sample output of SHOW FILE FULL command:

```
officer SEC> SHOW FILES FULL
```

```

-----
File..... v1anscript1.txt
Version.....
Model.....
Size Kb..... 3
Modified..... 11/04/2003, 18:34:08

File..... v1anscript2.txt
Version.....
Model.....
Size Kb..... 1
Modified..... 11/04/2003, 18:34:08

File..... adsl16_500_1.1.1.tar
Version..... 1.1.1
Model..... TN-ADSL-X: 01, TN-ADSL16-X: 01, TN-100-
X: 01, TN-106-X: 00
Size Kb..... 1027
Modified..... 10/21/2003, 15:59:32

```

(some text omitted)

10.3 Software Load Management

10.3.1 Card Load Preferences

Once a software load is present in the control module FLASH file system, it can be designated as the target software load for one or more cards using the parameters on the SET CARD command. The setting can be PREFLOAD, ALTLOAD, or TEMPLOAD. Load preferences are discussed in detail in [4.1.6](#).

10.3.2 Load File Verification

When software load files are created, a CRC value is calculated and written into the internal file data. Once a file has been transferred to the control module FLASH file system (and to Service Module FLASH), the file contents can be verified by recalculating the CRC value and comparing it to the internal CRC value. In the unlikely event that there is a mismatch between the value, the file is designated corrupt.

File CRC validation is performed on all load files (those with the “.tar” extension on the file name) in the control module FLASH file system a follows:

- immediately after a restart or swap of activity
- periodic audit (every 24 hours)
- whenever a user enters the AUDIT FILES command
- File CRC validation is performed on individual files when:
 - the file is being designated as a parameter on the SET CARD command
 - the file is being used during the card restart sequence, as a result of system action or manual command (RESTART CARD or ENABLE CARD)

As mentioned above, the user can audit system files using the AUDIT FILES command. Following is an example of the use of the command.

```
officer SEC> AUDIT FILES
Command has been submitted
officer SEC>
```

```
-----
File..... Audit result
-----
adsl16_3_0_0.GAMMA.20040331.tar..... Pass
adsl16_500_2_0_3.tar..... Pass
adsl16_500_2_1_0.ANNEXB_BETA.20040129. Pass
adsl16_500_2_1_1.tar..... Pass
adsl16b_3_0_0.GAMMA.20040331.tar..... Pass
adsl16b_500_2_1_1.tar..... Pass
adsl24_3_0_0.GAMMA.20040331.tar..... Pass
adsl8_3_0_0.GAMMA.20040331.tar..... Pass
adsl8_500_2_0_3.tar..... Pass
adsl8_500_2_1_1.tar..... Pass
cfc24_2_0_3.tar..... Pass
cfc24_2_1_1.tar..... Pass
cfc24_3_0_0.BETA.20040301.tar..... Pass
cfc24_3_0_0.GAMMA.20040322.tar..... Pass
cfc24_3_0_0.GAMMA.20040405.tar..... Pass
fe10_2_0_3.tar..... Pass
fe10_2_1_1.tar..... Pass
fe10_3_0_0.GAMMA.20040331.tar..... Pass
fe10_3_0_0.GAMMA.20040405.tar..... Pass
```

10.3.3 Boot Server (Control Module Only)

Users of the system have the option of storing secure copies of software release load files on their network servers. This assists the user in providing optimum network reliability. The boot server should be configured and the most current control module load file should be stored there. This insures that secure load files are always available.

Note that boot server functionality is **only** available for the active control module.

Users can also configure the system to boot from the network servers where the load files are stored using the SET BOOTSERVER. It permits users to designate a server to be the system boot server. As discussed in 10.2.2, load files are stored in FLASH memory on the control module. If the FLASH should become corrupted or the files become unusable for any reason, the system will boot from the secure load files stored on the boot server. To ensure system recovery back to its normal operating state, the card load files that are stored on the boot server must be a copy of the card load file designated as **PREFLOAD** for the control module.

Loading the system from the boot server is not intended to be the primary method for software release delivery. Rather, it is a backup or secondary method that the system will utilize if the primary method is unusable for any reason.

To configure the boot server, the flow of commands would be.

1. Get the preferred load using the GET FILE command
2. Set the preferred load using the SET CARD PREFLOAD command
3. Make a backup of preferred load using the COPY FILE command
4. Make the backup the alternate load using the SET CARD ALTLOAD command
5. Set the Bootserver using the SET BOOTSERVER command
6. Copy the current control module preferred load files, that are designated as PREFLOAD, onto the boot server using the PUT FILE command.
- 7.

TABLE 10-3 Bootserver Commands Summary

Noun	Verb	Syntax	Description
BOOTSERVER	SET	SET BOOTSERVER [IPADDRESS=ipaddress] [PATH={ pathname NONE }]	Sets the network server that stores the secure system load files. PATH is the directory path on the network server where the files are stored. If PATH is not specified, the root directory is used.
BOOTSERVER	SHOW	SHOW BOOTSERVER	Displays information on the network server that has been set as the boot device.

Figure 10-2 and Figure 10-3 illustrate the sequence the system will follow to reload cards and recover the system.

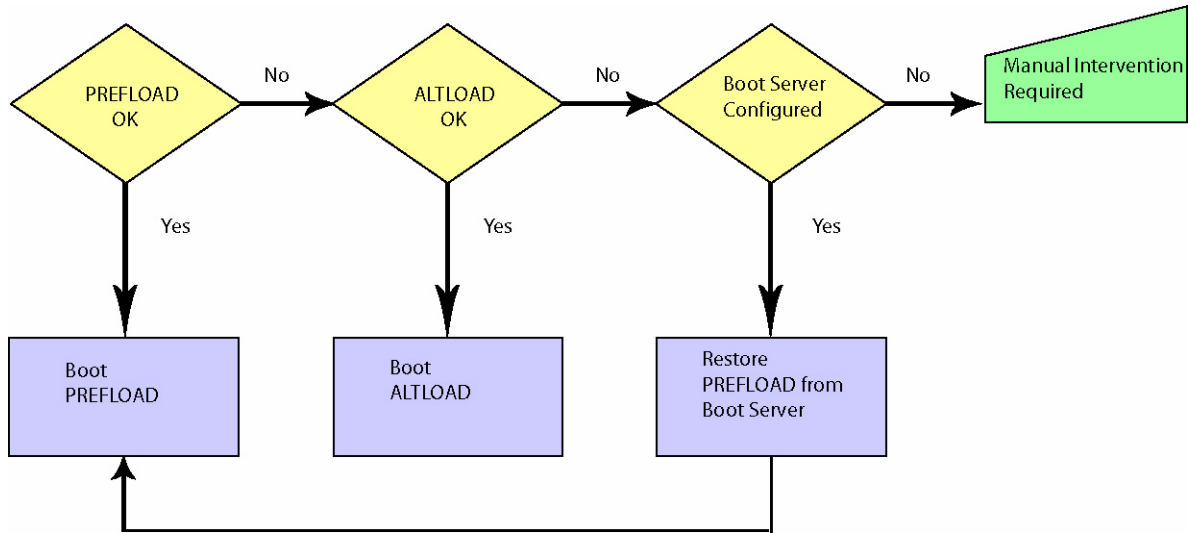


FIGURE 10-2 Control Module loading

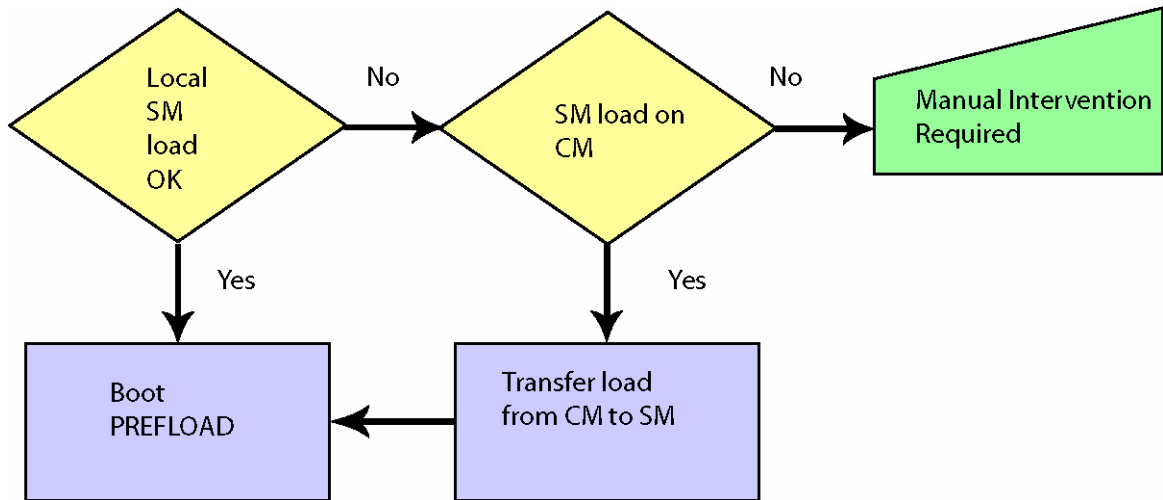


FIGURE 10-3 Service Module Loading

10.4 Database Management

10.4.1 Overview

All configuration and provisioning changes made by the system user are stored internally in the system configuration database. The database is updated dynamically any time a configuration change is made to the system, through the normal use of CLI commands. When the control module is restarted, it restores the system configuration from the database.

The database is stored in the control module FLASH memory, but is not a user manageable file in the FLASH system, so no file-related CLI commands apply to it. A separate set of CLI commands is provided to manipulate the database.

For a system configured for duplex operation, a copy of the current database is located on both the active and inactive control module. Configuration changes made on the active control module are automatically written to the database and propagated to the copy of the database located on the inactive control module. As long as the control modules are operating in synchronization, any action that results in configuration data being written to the database is reflected on both control modules. When the inactive control module is being brought online from an offline condition, its database is bulk synchronized with the active control module.

During a normal software release upgrade, commands that manage the configuration database are executed. Summaries of the commands follow. For detailed information on database commands, refer to the **Routine Administration** section [17.2](#).

10.4.1.1 Database Back Up

If required, the database containing the most recent provisioning and configuration data would be used to restore the system back to normal.

*Note: Keep a copy of the most recent database on a secure network server by performing a backup after any configuration change. Also, storing a backup database on the CFC's own flash is **not** recommended. Refer to the **Routine Administration** section.*

During a normal software release upgrade, the database is manually backed up if a downgrade is required after the new load files have been committed. If a downgrade is required, a database restoration is performed. The steps in performing a database backup include:

1. Designate a network server for secure storage of the current configuration database
2. Using the BACKUP DATABASE command, back up the current configuration database to the secure server
3. The user may execute the SHOW TRANSFER command to display the progress of the backup.

10.4.1.2 Database Purge

The PURGE DATABASE command erases the current configuration database. This command would be used if the user wanted to reconfigure the system *back to factory defaults*. When the command is entered, the system reboots and recovers with the factory defaults.

Note: *When the PURGE DATABASE command is used on a duplex configuration that is in normal (synchronized) mode, the database on the active control module and the copy on the inactive control module will be purged.*

As explained in a previous section, TELNET is disabled by default. If the user is connected through TELNET, when the database purge completes, TELNET will be back in the default disabled mode, and the user will no longer be connected to the system. The user should therefore connect and login to the CONSOLE interface of the control module **prior** to executing the PURGE DATABASE command.

Once the database is purged, and the system reboots, the system configuration database can be built by either:

- Restoring a previously backed up database (RESTORE DATABASE)
- Rebuilding the database manually using CLI commands and scripts
- Restoring a Config File

The following shows an example of the PURGE DATABASE command and the reply that the purge is completed. Note that this example was completed at the TELNET interface and is last message before the connection was lost.

```
officer SEC> PURGE DATABASE FORCE
Command has been submitted
PURGE DATABASE - success
```

Note: *Use of the PURGE DATABASE command can cause network outages.*

10.4.1.3 Restore Database

With the fMAP products, the user can use the RESTORE DATABASE command to replace the current database with a database that had been previously backed up and sent to a network server. (The user should be aware of backward compatibility criteria, as explained in [10.6](#).)

Note: *While the database transfer is occurring, the old database is still intact in FLASH memory, and the new database is written to RAM. The user can abort the database restore operation while the database transfer is still in progress. Once the database transfer is complete, the old database is erased from FLASH and the new database is written to FLASH. The control module is then automatically restarted, and the new database is used to configure the system.*

10.4.2 Database in upgrade mode

During a software upgrade, a schema migration is performed, where the configuration data read from the original database in flash memory is written to a new database. The data in the new database is converted (schema migrated) to a new format that is compatible with the load being upgraded to. During this process, the original

database is left intact in flash memory, and the new database is held in RAM memory on the control module. When in this condition, the database is considered to be in “upgrade mode”, and an alarm is raised against the control module being upgraded.

To get out of upgrade mode, the user must commit to the new load using the SET CARD command on the CM, which erases the original database in flash memory and then copies the new database from RAM to flash memory. Alternatively, during duplex upgrades only, upgrade mode can be cleared by doing an abort of the upgrade process, which erases the new database in RAM memory and causes the system to revert back to the original database still in flash memory.

10.5 Text File Configuration

10.5.1 Overview

A text configuration file is a “snapshot” of the configuration database including all of the non-default configuration commands in a text-based (rather than binary) file. (The binary configuration database and its associated commands are in the previous subsection, 10.4).

The advantage to having a text file is that it can be read (unlike a binary file), modified if necessary, saved, and then used to configure (or reconfigure) this or other systems.

Note: Since the file is in text rather than binary format, applying a configuration file will take longer (by minutes) than a binary file.

Following are the major tasks and commands used for this feature.

10.5.2 Creating a Text Configuration file

The command to create the file is:

```
BACKUP CONFIG [ FILE={ destinationfile | unit:destinationfile } ]
```

If only BACKUP CONFIG is input, the commands are written to the CONSOLE output. If the optional key word FILE is included, the command set is sent to either the local FLASH (destinationfile) or the Compact FLASH (using the format CFLASH<no.>:destinationfile).

Note: While the BACKUP CONFIG command is executing, commands that further affect the system configuration are disallowed, with the user receiving a command rejected message. Moreover, the execution of this command may take several minutes to complete.

Following is an example of this command.

```
officer SEC> BACKUP CONFIG FILE bkupcfgfl
officer SEC> Info (020139): Configuration Backup Processing...

officer SEC> SHOW CONFIG STATUS
--- Configuration File Progress -----
State..... Backup

Backup Initiated..... 2004-03-26 06:47:29
Backup Completed..... In progress...
```

```

Progress Summary..... 9 of 10 complete

Backup Configuration Progress Details
User Data..... Complete (100.00%)
OA&M Configuration..... Complete (100.00%)
DSLagTextConfig..... Complete (100.00%)
Traffic Management..... Complete (100.00%)
DsbseTextConfig..... Complete (100.00%)
DsstpTextConfig..... Complete (100.00%)
DsdhcpTextConfig..... Complete (100.00%)
DSL2vnTextConfig..... Processing (50.84%)
OampNmTextConfig..... Complete (100.00%)
Dsi gmpTextConfig..... Complete (100.00%)

officer SEC> Info (020147): Configuration file "/tffs/load/bkupcfgfl" successfully created.

officer SEC> SHOW CONFIG STATUS
--- Configuration File Progress -----
State..... Not In Progress

Previous Backup Configuration Status
-----
Backup Initiated..... 2004-03-26 06:47:29
Backup Completed..... 2004-03-26 06:49:47
Progress Summary..... 10 of 10 complete

Backup Configuration Progress Details
User Data..... Complete (100.00%)
OA&M Configuration..... Complete (100.00%)
DSLagTextConfig..... Complete (100.00%)
Traffic Management..... Complete (100.00%)
DsbseTextConfig..... Complete (100.00%)
DsstpTextConfig..... Complete (100.00%)
DsdhcpTextConfig..... Complete (100.00%)
DSL2vnTextConfig..... Complete (100.00%)
OampNmTextConfig..... Complete (100.00%)
Dsi gmpTextConfig..... Complete (100.00%)

Previous Restore Configuration Status
-----
No restore configuration information available

```

10.5.3 Restoring a Configuration Database Using a Text Configuration File

A text configuration file can be used to populate the configuration database of a device, and can be useful during system upgrades and downgrades (refer to [10.5.7](#)). The command used to execute the file and restore a configuration database is:

```
RESTORE CONFIG FILE={sourcefile | unit:sourcefile}
[OUTPUT={CONSOLE | logfile | unit:logfile}]
```

The keyword FILE requires that a sourcefile (from FLASH) or unit:sourcefile (from CFLASH) be supplied. The optional keyword OUTPUT is recommended since this can be used to capture logs that are produced by the script.

Note: *Once the text configuration file has finished running, the system will purge its current database and reboot using the configuration reflected in the text configuration file. The user can stop this from occurring using the STOP CONFIG file, as explained below.*

Following is an example of this command.

```

officer SEC> RESTORE CONFIG FILE=bkupcfgfl OUTPUT=CONSOLE
Database will be cleared and system will reboot. Restore configuration (Y/N)? Y
Command has been submitted

```

```

Info (020148): Restore configuration successfully requested.
Info (033765): Database purge succeeded

officer SEC> RESTORE CONFIG FILE=bkupcfgfl OUTPUT=restlogfl
Database will be cleared and system will reboot. Restore configuration (Y/N)? Y
Command has been submitted
Info (020148): Restore configuration successfully requested.
Info (033765): Database purge succeeded

officer SEC> SHOW CONFIG STATUS
--- Configuration File Progress -----
State..... Not In Progress

Previous Backup Configuration Status
-----
No backup configuration information available

Previous Restore Configuration Status
-----
Restore Initiated..... 2004-03-26 07:08:16
Restore Completed..... 2004-03-26 07:08:55
Restore File..... bkupcfgfl
Restore Log File..... restlogfl
Restore Progress..... 100.00% complete

```

10.5.4 Stopping a Backup/Restore in Progress

Both the BACKUP and RESTORE commands take several minutes to execute, and the user may wish to stop the command before it is complete. The command to do this is STOP CONFIG and it has the following effect:

- If entered during a BACKUP CONFIG command, STOP CONFIG throws away the configuration text file being created, and configuration commands can be input.
- If entered during a RESTORE CONFIG command, STOP CONFIG cancels the execution of the file, closes out the log file (if one is being output) and configuration commands can be input.

Note: *Stopping a RESTORE CONFIG should be done before the database is purged and the system reboots; otherwise, an incorrect configuration could be installed.*

10.5.5 Viewing the Progress of a BACKUP or RESTORE

While a BACKUP or RESTORE is in progress, the user can view the status of the file execution using the command SHOW CONFIG STATUS. The command displays the processes that are run, and as each one is completed, it will have a 100% displayed next to it.

If a STOP CONFIG is input during a BACKUP or RESTORE, the SHOW CONFIG STATUS will show which process was being run when the STOP CONFIG command was input, and what percentage of that process was completed. When the text file configuration is present on the system, the system will use the text file to create the DB which it runs from. The text file configuration will also be kept up to date as additional commands are entered into the system.

Following is an example of the SHOW CONFIG command when a STOP CONFIG command has been input during a BACKUP CONFIG. Notice that Dsl 2vnTextConfig is in the Processing state with (6.10%) complete.

```

officer SEC> SHOW CONFIG STATUS
--- Configuration File Progress -----

```

```

State..... Backup

Backup Initiated..... 2004-03-25 12:52:59
Backup Completed..... In progress...
Progress Summary..... 9 of 10 complete

Backup Configuration Progress Details
User Data..... Complete (100.00%)
OA&M Configuration..... Complete (100.00%)
DslagTextConfig..... Complete (100.00%)
Traffic Management..... Complete (100.00%)
DsbasetextConfig..... Complete (100.00%)
DsstptextConfig..... Complete (100.00%)
DsdhcpTextConfig..... Complete (100.00%)
Dsl2vntextConfig..... Processing (6.10%)
OampNmtextConfig..... Complete (100.00%)
DslgmpTextConfig..... Complete (100.00%)

```

10.5.6 Editing a Text Configuration File

Once the BACKUP command has been used to create a text file of the system configuration, the file can be edited by using the PUT FILE command to send the file to a destination (such as a server) where it can be edited. The file can then be placed back onto the FLASH or FLASH media using the GET FILE command.

Following is an example of this sequence:

// Put file onto server

```

officer SEC> PUT FILE bkupcfgfl TFTP SERVER 172.16.18.50
Command has been submitted Transfer ID: 1
officer SEC> Info (033012): Successfully transferred file: bkupcfgfl

```

// File arrives on TFTP server

```

TFTP Daemon started

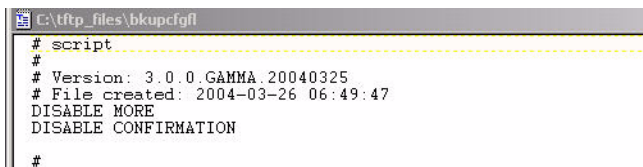
New session created
Requested to write file bkupcfgfl in format octet
Fully expanded file specification : C:\tftp_files\bkupcfgfl
Received 4486 bytes in < 1 second

```

// File arrives on server

script_file_1.txt	1 KB	Text Document	5/1/2003 6:22 AM
test_script1.txt	1 KB	Text Document	6/24/2003 6:11 AM
bkupcfgfl	5 KB	File	3/26/2004 8:05 AM

// Editing the backup file on the server



```

C:\tftp_files\bkupcfgfl
# script
#
# Version: 3.0.0.GAMMA.20040325
# File created: 2004-03-26 06:49:47
DISABLE MORE
DISABLE CONFIRMATION
#

```

// Get file from server after editing

```

officer SEC> GET FILE bkupcfgfl TFTP SERVER 172.16.18.50
Command has been submitted Transfer ID: 2
officer SEC> Info (033012): Successfully transferred file: bkupcfgfl

```

10.5.7 Using Configuration Text Files During Upgrades/Downgrades

The text config file can be useful when performing software release upgrades, especially when upgrading hardware at the same time. Refer to the software release upgrade sections for the 9000 [11.1](#) for more information.

10.5.8 Summary of Commands

TABLE 10-4 Text File Commands Summary

Noun	Verb	Syntax	Description
CONFIG	BACKUP	BACKUP CONFIG FILE={ destinationfile unit:destinationfile }	<p>Allows the user to create a configuration file which reflects current configuration of the system. This configuration file can be used to recreate the configuration on the same or similar system, using the RESTORE CONFIG command.</p> <p>Note that while the BACKUP CONFIG command is executing, commands that further affect the system configuration are disallowed, with the user receiving a command rejected message. Moreover, the execution of this command may take several minutes to complete.</p>
	RESTORE	RESTORE CONFIG FILE={ sourcefile unit:sourcefile } [OUTPUT={ CONSOLE logfile unit:logfile }]	<p>Used to populate the configuration database of a device from a text configuration file. This can be useful during system upgrades and downgrades (refer to 10.5.7).</p> <p>The keyword FILE requires that a sourcefile (from FLASH) or unit:sourcefile (from CFLASH) be supplied. The optional keyword OUTPUT is recommended since this can be used to capture logs that are produced by the script.</p>
	SHOW	SHOW CONFIG [STATUS]	<p>Allows the user to view the current configuration information or to monitor the progress of a currently running BACKUP CONFIG or RESTORE CONFIG command.</p> <p>If the STATUS parameter is provided, the current or previous status of a BACKUP CONFIG or RESTORE CONFIG is displayed.</p> <p>Without the STATUS parameter, the current configuration information is generated and displayed to the user.</p> <p>The status information is not saved over reboots of the system.</p>
	STOP	STOP CONFIG	<p>Stops the execution of a RESTORE or BACKUP.</p>

10.6 Software Compatibility

10.6.1 Overview

A software release will retain backward compatibility with certain releases that preceded it. *Backward compatible* means that configuration data that was saved in flash memory using the older software release will be restorable when the newer software release is loaded, allowing the upgrade to occur without requiring the entire system to be reconfigured. It also means that a card using a given software release can communicate with another card using an older software release.

An understanding of *software upgrades* and *interim upgrades* is important to the concept of backward compatibility. An upgrade where system load files are changed in order to add new, significant feature functionality is called a *software upgrade*. For example, an upgrade from release 4.0.0 to 5.0.0. An upgrade where system load files are changed in order to possibly add minor feature functionality and software fixes is referred to as an *interim upgrade*. For example, an upgrade from release 5.0.1 to 5.0.3. During an interim upgrade, no schema migration is to be performed on the configuration database. Also, during the interim upgrade, no Database in Upgrade alarm is raised.

Software upgrades and interim upgrades will be reemphasized in each of the system upgrade procedure sections.

Backward compatibility is unidirectional; a newer load may be able to understand data that originates from an older load, but the reverse, forward compatibility, is not necessarily true.

The rules for backward compatibility are:

- Interim-release changes have no effect on backward compatibility. For example, release “2.0.1” is fully backward compatible with release “2.0.0”. Any release that is backward compatible with release “2.0.1” is also, by extension, compatible with release “2.0.0”.
- Backward compatibility is supported between any two releases that share the same major release number. For example, release “2.1.1” is backward compatible with release “2.0.1”.
- If a software release has a minor release number of “0” (zero), it is backward compatible with any releases in the previous major release. For example, release “3.0.1” is backward compatible with release “2.1.0”.

The following figure shows how this works.

In cases where a direct upgrade is not supported, a multi-step upgrade process may be used. For example, to upgrade from release “2.1.1” to “4.0.0”, a user must first upgrade from “2.1.1” to “3.0.1”, and then incrementally upgrade to release “4.0.0”.

*Note: Only the control module card (not the service modules) has this backward compatibility logic. Therefore, a control module running a newer software load is compatible with a service module running an equivalent or older load, but a control module running an older release load is **not** compatible with a service module running a newer release load.*

Note: The third digit, for patches, is automatically compatible for upgrade. For any anomalies, refer to the specific Release Notes.




FIGURE 10-4 Compatibility Hierarchy

10.6.2 Supported Loads

Any CM load supports its major release load and the same release for NMs and SMs. Also, any CM load supports **one major release** back on NMs and SMs for the purpose of upgrading. For example:

- A Release 5.x CM load will support any Release 5.x NM and SM load.
- A Release 5.x CM load will support any Release 4.x NM and SM load for upgrading.

10.7 Software Upgrade

10.7.1 Overview

To upgrade the system to a new software release, load files can be remotely loaded into FLASH memory and then individually selected for use at runtime. Along with the user commands described earlier in this guide, there are comprehensive commands provided to examine the state of the FLASH memory and to load, view, and remove release files. There are also commands provided to install, query, and activate the software on each card, at which time the software is booted into RAM memory and executed.

10.7.2 Obtaining New Loads

The latest software release files that are available from Allied Telesis. Contact information is located in the Introduction section of this manual.

The general steps are:

1. If not already done, set up the MGMT/Inband interface and network servers. Refer to Section 3.
2. Do `SHOW SYSTEM` or `SHOW CARD ACTCFC SOFTWARE` to see what CFC load is used.

3. Compare current system loads with the latest load's information.
4. If the loads are the same, no action is required.
5. If the loads are not the latest, contact Allied Telesis K.K. to obtain them.

11. Software Release Upgrade for the fMAP

11.1 Simplex Configuration

11.1.1 Overview

A software upgrade involves obtaining new load files. Next, the loads are stored in FLASH on the CM. The boot status of the new loads is set to TEMPLOAD. The user then executes a restart on the control module. The user verifies that the new loads are working correctly. The user must commit, then downgrade if necessary.

Commit means that the user chooses to accept the new loads as the new system load files. After setting the loads as PREFLOAD, the system will boot from these loads until new loads are set as PREFLOAD.

The user can choose to perform a downgrade after a software upgrade has been completed. The user may revert back to the original load files and database, but any changes made to the new database are lost.

11.1.1.1 Software upgrade

An upgrade where system load files are changed in order to add new, significant feature functionality. For example, an upgrade from release 3.0.3 to 4.0.0.

11.1.1.2 Interim Upgrade

An upgrade where system load files are changed in order to possibly add minor feature functionality and software fixes. For example, an upgrade from release 3.0.2 to 3.0.3. During this type of upgrade, no schema migration is to be performed on the configuration database. Also, during the upgrade, no Database in Upgrade Mode alarm is raised.

Note: Refer to section *Software Compatibility*, [10.6](#), for software upgrade compatibility information.

Note: Software upgrades on a system in simplex are service affecting. Video, data, and voice traffic will be affected.

11.2 Simplex Upgrade Procedure

11.2.1 Upgrade

This is a walkthrough of a software upgrade for a simplex 9000. It is assumed that the system is running in a stable state.

TABLE 11-1 9000 Simplex Software Upgrade Steps

Step	State or Action	Details
1.	Pre-Upgrade Configuration	Checking the Allied Telesis website, find the latest loads for the hardware and software release this upgrade will support and download these to the network server so that they may be copied if necessary.
2.	Back up the current database: BACKUP DATABASE	For network reliability purposes, backup the existing configuration database to the external network server using the BACKUP DATABASE command.
3.	Retrieve the new load files for the CM and SM cards from the network server and store them in FLASH on the CM card.	Follow these steps to get the latest CM and SM load files and set them as the TEMP load. 1 Ensure the latest loads are on the network server. 2. Transfer the new loads from the server to the control module FLASH - GET FILE=<load> <server> USER=id PASSWORD=pw 3. Set the new CM load as the TEMP load for the CFC56. - SET CARD=ACTCFC TEMPLOAD=<latest CM load> 4. Set the new ADSL loads as the TEMP load for the ADSL cards - SET CARD=<SM slot-list> TEMPLOAD=<new SM load>
4.	Execute a restart on the CM card	Restart the card with the command: RESTART CARD=ACTCFC COLD The CFC24 and SM cards reboot to the new loads and recover. The database contains all the original configuration, schema-migrated to the new load, and held in CFC24 RAM during this step only. The original database is still intact in FLASH memory.

TABLE 11-1 9000 Simplex Software Upgrade Steps (Continued)

Step	State or Action	Details
5.	Run any verification tests.	<p>After step 4 is completed, the system will be in a “Database in upgrade mode” alarm condition unless this is an Interim Upgrade (See 11.1.1.2).</p> <p><i>Note: Whenever the active CFC is in upgrade mode, all commands are allowed, but a warning is displayed to remind the user that the system is in upgrade mode and changes are not being saved to FLASH memory until the upgrade is committed to (by setting the preferred load of the active CFC to the running load).</i></p>
6.	Commit to the new load: Set the CFC24 card with the new load file as the preferred load.	<p>The new load is set as PREFERRED so that on the next reboot the new load will be loaded.</p> <p>SET CARD=ACTCFC PREFLOAD=<latest CM load></p> <p>The new schema-migrated database is now written into FLASH memory. The old database is erased. The upgrade mode alarm is cleared.</p> <p>If the user is sure they are not going to perform a downgrade, the original CM loads can be deleted with the DELETE FILE command.</p>
7.	Set the latest SM load as the preferred load for all SM cards	<p>The new load is set as PREFERRED so that when the system reboots the new load will always be used.</p> <p>Immediately upgrade the remaining service modules.</p> <p>SET CARD=SM PREFLOAD=<latest SM load></p>
8.	Back up the current database: BACKUP DATABASE	<p>For network reliability purposes, backup the existing configuration database to the external network server using the command:</p> <p>BACKUP DATABASE</p>

11.2.2 Downgrade

To revert back to the original load files after they have committed to, a downgrade must be performed. The user must fully commit to the upgrade before performing a downgrade. Load preferences must be set to the new loads.

Note: Any configuration changes completed while the upgrade was in progress will not be saved to the database if a downgrade is performed.

TABLE 11-2 Simplex Software Downgrade Steps

Step	State or Action	Details
1.	Obtain the original load files.	If the original load files are not in FLASH on the CM and SMs, obtain copies and put them in FLASH on the CM.
2.	Set the original load files as PREFLOAD.	Set the original CFC load files to PREFERRED. SM load preferences are saved on the database and will be taken from it after it has been restored.
3.	Restore the original database.	Restore the original database. The user inputs the command: RESTORE DATABASE During this process, the database reverts back to its original configuration data, the CM restarts, and the CM and SMs revert back to their original loads.

12. 802.3 Ethernet, VLAN, UFO

12.1 Layer 2 Switching (Learning)

12.1.1 Overview

As a layer 2 switching device, the fMAP product ensures data packets arrive at their proper destination by using:

- **VLAN** - This is a software-defined subnetwork that allows devices to be grouped into one logical broadcast domain. Refer to [12.2](#).
- **MAC address** - The MAC address uniquely identifies each hardware device attached to the network.

The layer 2 switching process includes four separate but related processes:

- **Ingress Rules** admit or discard frames based on their VLAN tagging.
- **Learning Process** learns the MAC addresses for each VLAN as frames are admitted to each interface.
- **Forwarding Process** determines which interfaces the frames are forwarded to.
- **Egress Rules** determine for each frame whether VLAN tags are included in the Ethernet frames that are Transmitted.

Since this is layer 2, the learning process assumes that each host on the extended LAN has a unique data link layer address, and all data link layer frames have a header which includes the source (sender) MAC address and destination (receiver) MAC address.

12.1.2 Ingress Rules

When a frame first arrives at a port, the Ingress Rules for the port check the VLAN tagging in the frame to determine whether it will be discarded or forwarded to the Learning Process.

Every frame received by the switch must be associated with a VLAN. If a received frame is untagged, then the port's untagged VLAN Identifier (VID) will be associated with the received frame. Since every port belongs to one or more VLANs, every incoming frame will have a VID to indicate which VLAN it belongs to.

The Ingress Rule will check whether the port, in which the frame was received on, belongs to the VLAN indicated by the received frame's VID. If the port is not a member of the VLAN, then the frame will be discarded; otherwise, the frame will be passed on to the Learning Process.

12.1.3 Learning Process

A layer 2 ethernet switch, when it first receives frames, floods the switch with data packets. The Learning Process uses an adaptive learning algorithm, sometimes called backward learning, to discover the location (port) of each host on the extended LAN and ensure frames are set to their destination as efficiently as possible.

All frames admitted by the Ingress Rules on any port are passed on to the Learning Process, where the frame's source MAC address and numerical (VID) are compared with entries in the Forwarding Database for the VLAN (also known as a MAC address table, or a forwarding table) maintained by the switch. The Forwarding Database contains one entry for every unique host MAC address the switch knows in each VLAN.

If the frame's source address is not already in the Forwarding Database for the VLAN, the address is added (MAC address and VLAN ID) and an ageing timer for that entry is started. If the frame's source address is already in the Forwarding Database, the ageing timer for that entry is restarted.

By default, switch learning is enabled, and it can be disabled or enabled using the commands:

```
DISABLE SWITCH LEARNING
```

```
ENALBLE SWITCH LEARNING
```



If the Learning Process is disabled, MAC addresses are no longer added to the forwarding database, and as the ageingtimer (discussed next) expires and frames with their source addresses and VLAN IDs are no longer learned, the fMAP product will slowly depopulate its forwarding database.

If the ageing timer for an entry in the Forwarding Database expires before another frame with the same source address and VID is received, the entry is removed from the Forwarding Database. This prevents the Forwarding Database from being filled up with information about hosts that are inactive or have been disconnected from the network, while ensuring that entries for active hosts are kept alive. By default, the ageing timer is enabled, and it can be disabled or enabled using the commands

```
ENABLE SWITCH AGEINGTIMER
```

```
DISABLE SWITCH AGEINGTIMER
```

By default, the ageing time is set to a value of 300 second (5 minutes), and can be modified.

The Forwarding Database relates a host's (source) address to a port on the switch, and is used by the switch to determine from which port (if any) to transmit frames with a destination MAC address matching the entry in the host map.

To display the contents of the Forwarding Database, use the command `SHOW SWITCH FDB`.

Note: *If an ADSL line is incurring noise, it is possible that bogus MAC addresses could be learned by the system. These bogus addresses will eventually age out, but they may appear in the SHOW FDB command output and may also interfere with MAC limiting by consuming valid learned-MAC slots.*

12.1.4 Forwarding Process

The Forwarding Process forwards received frames that are to be relayed to other ports in the same VLAN. If a frame is received on the port for a destination in a different VLAN, it will need to be serviced by a Layer 3 switch/router external to the fMAP product.

The Spanning Tree Protocol (STP) can impact the forwarding process by disabling forwarding; see [13.2](#) for more information.

12.1.5 Clearing of FDB (Selective or Global)

12.1.5.1 Overview

The user can clear the switch forwarding database both on a selective and global basis. The following commands are available.

TABLE 12-1 Switch FDB clearing commands

Object	Verb	Syntax	Description
SWITCH FDB HVLAN	CLEAR	<pre> CLEAR SWITCH FDB [INTERFACE={ type:id-range id-range ifname-list ALL }] [ADDRESS=macaddress] [HVLAN={ hvlanname vid }] </pre>	Clear the switch FDB for HVLAN
SWITCH FDB VLAN	CLEAR	<pre> CLEAR SWITCH FDB [INTERFACE={ type:id-range id-range ifname-list ALL }] [ADDRESS=macaddress] [VLAN={ vlanname vid }] </pre>	Clear the switch FDB for VLAN

12.1.5.2 FDB and the Ethernet Interface for the ONU

For an overview of the EPON/ONU configuration and definitions of its components, refer to [Section 8](#).

To show the associations between the VLAN and the MAC addresses for an interface, start with:

```
officer SEC>> show switch fdb
```

```
--- Switch Forwarding Database -----
```

SLOT	VLAN	MAC ADDRESS	INTERFACE ID	STATUS
4	1	00:0C:25:D5:12:04	ETH:4.1.7	Dynamic
4	2	00:0C:25:D5:16:EE	ETH:4.1.7	Dynamic
4	201	00:0C:25:D4:02:1F	ETH:4.1.7	Dynamic

To clear (aspects of) the forwarding database, the following commands are used:

- Clear all OLTs and ONUs (sicne clears all entries for the system):

```
>Clear switch fdb
```
- Completely clear all ONUs for an OLT (. * for all the ONUs for the OLT)

```
>Clear switch fdb int 4.1.*
```
- Clear an ONU FDB, and OLT FDB entries for the logical links feeding an ONU

```
>Clear switch fdb int 4.1.7
```
- Clears any OLT and ONUs if the specified MAC is in its FDB

```
>Clear switch fdb address 00:0c:00:12:ef:82
```
- Clears any OLT and ONUs if it is configured with specified VID

```
>Clear switch fdb VLAN 2
```

Note: The OLT reports MAC addresses against all of the VLANs configured for the ONU, so the FDB may show MACs learned on VLANs in addition to the one(s) for which there was a packet.

12.1.6 Egress Rules

Once the Forwarding Process has determined which ports the frame is to be forward to, the Egress Rules for each port determine whether or not the outgoing frame is VLAN-tagged with its numerical VLAN Identifier (VID).

When a port is added to a VLAN, it is configured to transmit either untagged or VLAN tagged packets, using the commands to add and change the VLAN settings. Refer to [12.2.2](#) on VLAN tagging.

12.1.7 Command Summary for Switch Learning

Table 12-2 lists the commands used for the Layer 2 Switching.

TABLE 12-2 Commands for Layer 2 Switching

Object	Verb	Syntax	Description
SWITCH	SHOW	SHOW SWITCH	Shows general fMAP product settings, and includes the settings for switch learning and switch ageing.
INTERFACE	SHOW	<pre> SHOW INTERFACE [={ type: type:id-range id-range ifname-list ALL }] [CARD=slot-list] [STATE={ UP DOWN ALL }] [DIRECTION={ NETWORK CUSTOMER INTERNAL }] [FULL] </pre>	Shows the general settings for a port or interface.
SWITCH FDB	SHOW	<pre> SHOW SWITCH FDB [INTERFACE={ type:id-range id-range ifname-list ALL }] [ADDRESS=macaddress] [HVLAN={ hvlanname vid }] SHOW SWITCH FDB [INTERFACE={ type:id-range id-range ifname-list ALL }] [ADDRESS=macaddress] [VLAN={ vlanname vid }] </pre>	<p>Displays the contents of the Forwarding Database for either a HVLAN or VLAN.</p> <p>Use wildcard (*) to filter for a range of MAC addresses.</p>

TABLE 12-2 Commands for Layer 2 Switching

Object	Verb	Syntax	Description
SWITCH LEARNING	ENABLE	ENABLE SWITCH LEARNING	Enabled is the default, so this would only be used if the switch learning had been disabled.
	DISABLE	DISABLE SWITCH LEARNING	Disables switch learning. This is not recommended (see 12.1.3).
SWITCH AGEING-TIMER	ENABLE	ENABLE SWITCH AGEINGTIMER	Enabled is the default, so this would only be used if the ageing timer had been disabled.
	SET	SET SWITCHAGEINGTIMER=10..1000000	Changes the switchageing timer (time that a source address and VLAN ID is kept in the forwarding database). The default is 300 (5 minutes). Changing the ageing time will only affect entries added to the Forwarding Database after the timer has been changed.
	DISABLE	DISABLE SWITCH AGEINGTIMER	Disables switch ageingtimer. This is not recommended (see 12.1.3).

12.2 Virtual LAN (VLAN)

12.2.1 Overview

A VLAN is a virtual subnetwork that allows devices to be grouped into one logical broadcast domain. This allows broadcasts from one VLAN to be sent only to members on the same VLAN, improving network performance.

12.2.2 VLAN Tagging

An Ethernet packet can contain a *VLAN tag*, with fields that specify VLAN membership (the VLAN ID or VID) and user priority. The VLAN tag is described in IEEE Standard 802.3ac, and is four octets that can be inserted between the Source Address and the Type/Length fields in the Ethernet packet. To accommodate the tag, Standard 802.3ac also increased the maximum allowable length for an Ethernet frame to 1522 octets (the minimum size is 64 octets). IEEE Standard 802.1q specifies how the data in the VLAN tag is used to switch frames. VLAN-aware devices are able to add the VLAN tag to the packet header. VLAN-unaware devices cannot set or read the VLAN tag.

- Ethernet packets which contain a VLAN tag are referred to as *tagged* frames.

- Switch ports that transmit tagged frames are referred to as *tagged ports*.
- Ethernet packets which do not contain the VLAN tag are referred to as *untagged frames*.
- Switch ports that transmit untagged frames are referred to as *untagged ports*.

A VLAN can therefore consist of:

- A set of untagged ports, in which the ports receive and transmit untagged packets.
- A set of tagged ports, in which all ports for the VLAN transmit tagged frames
- A mixture of tagged and untagged ports, where on some ports the VLAN receives and transmits tagged frames and on other ports the VLAN receives and transmits untagged frames.

fMAP products accept VLAN tagged frames, and support the VLAN switching required by these tags. A network can contain a mixture of VLAN aware devices and VLAN unaware devices, for example, workstations and legacy switches that do not support VLAN tagging. The fMAP product can be configured to send VLAN-tagged or untagged frames on each port, depending on whether or not the devices connected to the port are VLAN aware. By assigning a port to two different VLANs, to one as an untagged port and to another as a tagged port, it is possible for the port to transmit both VLAN-tagged and untagged frames.

When VLAN membership is determined using VLAN tagging, switch ports and network resources can be used more efficiently, since a port can belong to several VLANs. Moreover, one port can be used to uplink (trunk) all VLAN traffic between the fMAP product and another VLAN-aware switch, since this port can be configured to include all VLANs on the fMAP product.

When devices cannot include VLAN tagging, the VLAN membership is determined by which port its packets arrive on; all untagged traffic arriving on a certain port belongs to **that** VLAN.

12.2.3 Standard VLAN Configuration

Figure 12-1 shows a sample configuration for setting up a VLAN in STD mode, with the commands used. The following explanation is based on this figure.

When a standard VLAN is configured, the Forwarding Database and VLAN/port mappings are set as follows:

FDB		Port Mapping
VID=5	MAC=00:50:94:31:33:00	8.4
VID=5	MAC=00:50:94:31:60:3D	9.8

When the Control Module receives the Source Address and L2VID, it performs two steps:

1. Learning - The Source Address-VLAN ID pair are checked against the FDB, and if it is not there the values are added.
2. Forwarding - The Destination Address is checked against the port mapping, and if the port mapping exists, it forwards the data onto that port. (Otherwise, it floods all ports for that VLAN.)

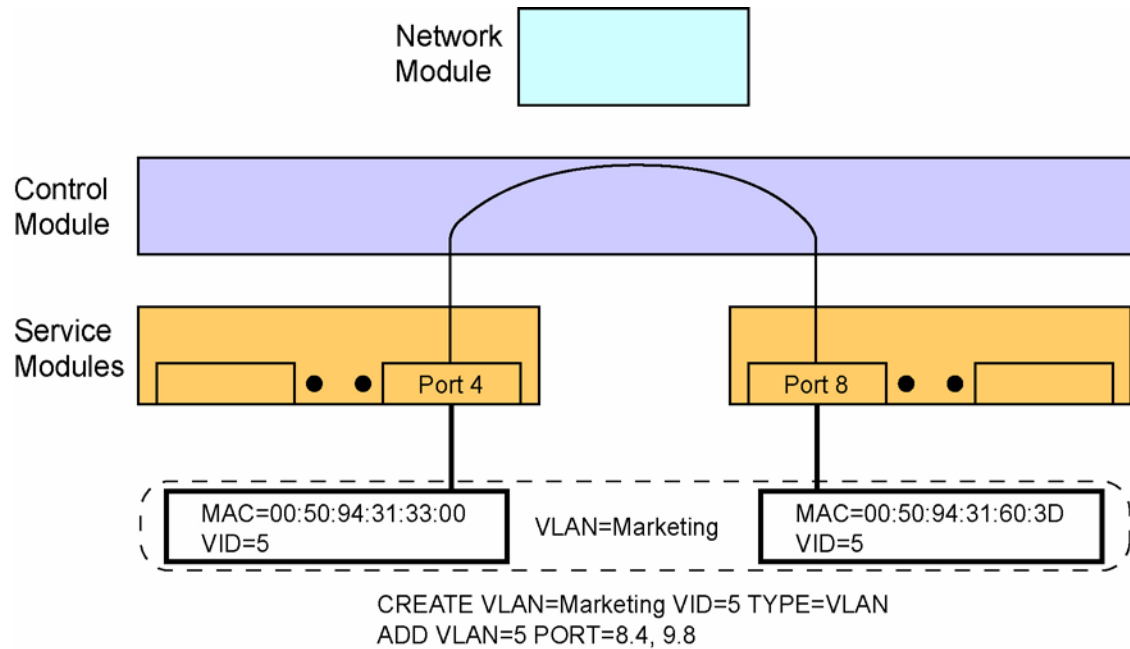


FIGURE 12-1 Standard VLAN Configuration in a fMAP Product

12.2.4 MAC Address Limiting for an Interface

In setting the VLAN attributes for an interface or an interface list, the user can specify the maximum number of MAC addresses that can be learned for an interface, or set the limit to OFF. This is useful in controlling how many MAC addresses can be learned and the customer configuration.

12.2.5 Command Summary for 802.1q VLAN

Table 12-3 lists the commands used to configure VLANs and includes notes on key parameters and use.

Note: So that future types of VLANs can be used, the user creates a VLAN (Layer 2 Virtual Network) and gives it a type called VLAN. In this release only the type VLAN can be used.

TABLE 12-3 Commands for 802.1q VLAN Provisioning

Object	Verb	Syntax	Description
VLAN	ADD	<pre> ADD VLAN={ vlaname vid } INTERFACE={ type:id-range id-range ifname-list ALL } [FRAME={ UNTAGGED TAGGED }] [TRANSLATE={ 1..4094 }] [FORWARDING={ PRIMARYUPSTREAM SECONDARYUPSTREAM DOWNSTREAM STP UCP }] </pre>	<p>Adds or associates a VLAN with either tagged or untagged interfaces. If FRAME is not included, UNTAGGED is used.</p> <p>FORWARDING is an important parameter when configuring a topology feature. Refer to Section 13.</p>
VLAN VID	CREATE	<pre> CREATE VLAN=vlaname VID=2..4094 [FORWARDINGMODE={ STD UPSTREAMONLY }] </pre>	<p>For the default configuration, VID 1 is always configured and is set at FORWARDINGMODE=UPSTREAMONLY (called UFO).</p> <p>Any VLAN created without the FORWARDINGMODE option is set to STD by default.</p> <p>The user can configure up to 24 VLANs that use UFO, and these can be anywhere in the 1-4094 range. Refer to 12.3 for details for each product.</p>
VLAN	DELETE	<pre> DELETE VLAN={ vlaname vid } INTERFACE={ type:id-range id-range ifname-list ALL } </pre>	<p>Deletes the association between the VLAN and a interface or interfaces. If the VLAN is associated with no interfaces, it can then be destroyed.</p> <p>If an interface has no VLANs associated with it, it will go to the default VLAN (1)</p>

TABLE 12-3 Commands for 802.1q VLAN Provisioning (Continued)

Object	Verb	Syntax	Description
VLAN	DESTROY	<pre>DESTROY VLAN={ vlanname vid ALL }</pre>	Deletes a VLAN.
VLAN	SET	<pre>SET VLAN={ vlanname vid } FORWARDINGMODE={ STD UPSTREAMONLY }</pre>	Changes the forwarding mode of a VLAN.
VLAN	SET	<pre>SET VLAN={ vlanname vid } INTERFACE={ type:id-range id-range ifname-list ALL } [FRAME={ UNTAGGED TAGGED }] [TRANSLATE={ 1..4094 NONE }] [FORWARDING={ PRIMARYUPSTREAM SECONDARYUPSTREAM DOWNSTREAM STP UCP }]</pre>	<p>Changes the association of a VLAN with a interface or interfaces, from tagged to untagged or vice-versa.</p> <p>FORWARDING is an important parameter when configuring a topology feature. Refer to Section 13.</p>
VLAN	SHOW	<pre>SHOW VLAN [={ vlanname vid ALL }] [FORWARDINGMODE={ STD UPSTREAMONLY ALL }] [FULL]</pre>	Displays the attributes of a specific VLAN (using the VLAN name or ID) or all VLANs.

TABLE 12-3 Commands for 802.1q VLAN Provisioning (Continued)

Object	Verb	Syntax	Description
INTERFACE	SET	<pre> SET INTERFACE={ type:id-range id-range ifname-list ALL } [ACCEPTABLE={ ALL VLAN HVLAN }] [INFILTERING={ OFF ON }] [TAGALL={ ON OFF }] [TPID=tpidvalue] [LEARNLIMIT={ 1..64 OFF }] </pre>	<p>Sets the VLAN attributes for an interface or an interface list.</p> <p>ACCEPTABLE - Sets the acceptable frame types as all (tagged and untagged) or VLAN-tagged only.</p> <p>INFILTERING - Sets the ingress filtering settings ON or OFF. Infiltering is the validation of VLANs on an interface. When ON, if a received frame's VLAN does not match the interface's VLAN membership, it is dropped.</p> <p>The default setting is ON.</p> <p>TAGALL - Controls whether all the frames are to be tagged or not. Double Tag.</p> <p>TPID - Used to identify the frame as a tagged frame. The value of the TPID for an 802.1q ethernet tagged frame is 0x8100.</p> <p>LEARNLIMIT - Specifies the maximum number of MAC addresses that can be learned for an interface.</p>
INTERFACE	SET	<pre> SET INTERFACE={ type: type:id-range id-range ifname-list ALL } DESCRIPTION={ description NONE } </pre>	<p>Sets the interface description.</p>

12.3 Upstream Forwarding Only (UFO) Mode

12.3.1 Overview

For the fMAP product, a VLAN can be created where all data from ports associated with that VLAN must be forwarded only to the upstream port, which is why it is called UFO mode. This segregation of traffic is done when:

- Certain types of services require only connections between the port and an upstream device.

- Security must be maintained (a malicious subscriber on one port cannot access a MAC or IP address on another port).

Note: Understanding UFO mode is important so that the user can understand when UFO VLANs are used for topology control. Refer to Section 13.

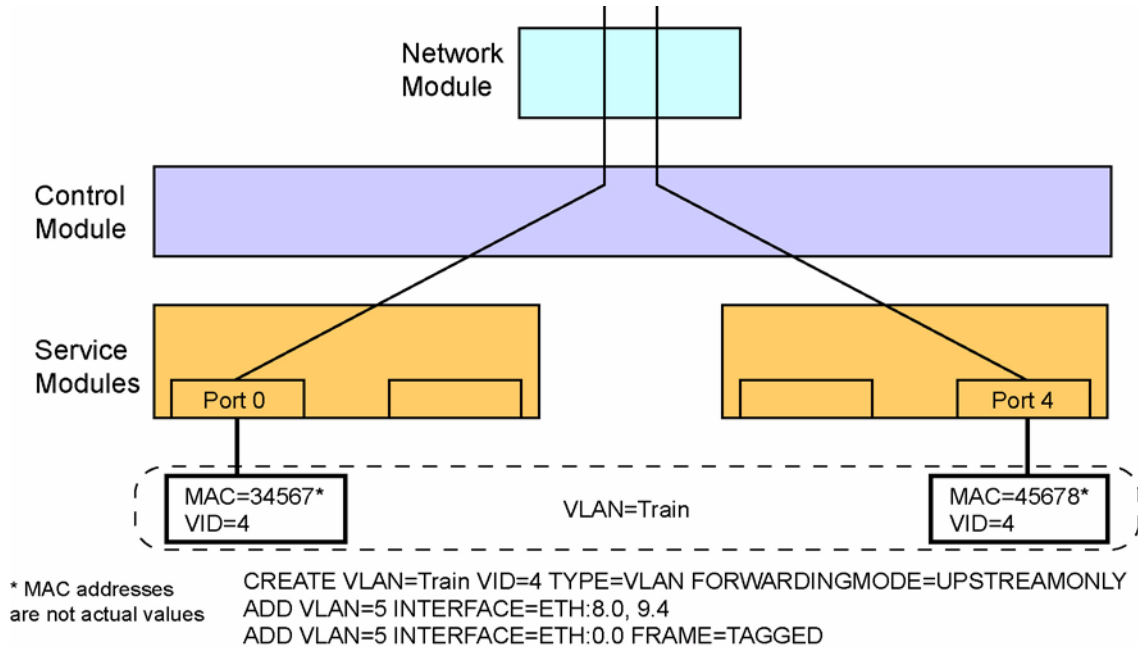


FIGURE 12-2 UFO Mode VLAN Configuration

12.3.2 Configuration Rules for VLANs (UFO and Standard) in 6.0

In Release 6.0 the user can configure an FE/FX as a UFO upstream interface on the 9x00. In addition, with the introduction of an interface being configured with the `DIRECTION` of either `NETWORK` or `CUSTOMER`; the UFO VLANs rule change as follows from release 5.0 to release 6.0:

- The UFO mode is controlled on a VLAN basis. (same)
- For all fMAP products, up to 24 VLANs can be configured in UFO mode, and they can use any VID in the 1-4094 range. (same)
- As a default, there is one VLAN (vid 1), which cannot be created or destroyed. The user has the option to change the default VLAN to STANDARD. (same)
- When all VLANs on a port are deleted, the port would revert to the default VLAN (vid 1), which would be in either UFO or Standard mode. (same)
- An interface may be set as the upstream interface (either statically or dynamically) for a UFO VLAN regardless of its `DIRECTION` setting. (new)

Note: Refer to *IGMP* for an overview of setting the *DIRECTION* of an interface, in [Section 14](#).

- Daisy chaining is supported with GE ports (and/or the FE2 on the 7100), where one fMAP is connected to another over a GE port. (same)
- For GE ports, the upstream interface can be dynamically determined by either UCP or STP/RSTP. For FE/FX ports, the upstream interface can be dynamically determined by STP/RSTP or UCP. (changed)

Note: Refer to *Topology* for more on dynamically determining upstream ports, [Section 13](#).

Once the ADD VLAN command for UFO has been invoked, the system may generate a warning message at the user's CLI session stating that classifier capacity or capabilities have been exceeded on the slot(s) impacted by the provisioning change. The user should investigate classifier-related provisioning, such as IGMP, DHCPRELAY, VLAN (for per-VLAN UFO and HVLAN), EPSR, INTERFACE (TAGALL option for HVLAN), ACCESSLIST, and CLASSIFIER to determine the reason for the message.

12.3.3 9000

On the fMAP 9000, VID 1, by default, is set to STANDARD mode. A maximum of 24 VLANs (VIDs) of the 1-4094 range can be provisioned for UFO. Only one GE/FE/FX port may be set to upstream per VLAN. LAG is not supported with UFO VLANs. LAG is not supported on any FE/FX port regardless of its *DIRECTION* setting.

12.4 VLAN - Virtual Channel (VC) Mapping

12.4.1 Overview

Note: In release 8.0, this functionality is supported on the ADSL24A cards. (It is not supported on the ADSL24 cards available up to release 5.0).

12.4.2 Overview

The VLAN to Virtual Channel (VC) mapping feature allows users to create and delete VCs, other than the default VC 0 that has always existed, and assign a Peak-Cell-Rate (PCR) on the downstream portion of the channel to VCs for voice, data, and video applications.

VC 0, the default VC, is always present in the fMAP system. It is created internally, automatically, and cannot be deleted.

VCs are provisioned for each ATM interface. These VCs are then associated with VLANs in the system.

Note: QoS per VC is not supported in this release

All VCs conform to unspecified bit rate (UBR). In past releases, there was only one VC per ATM interface and one or more VLANs were mapped to it. This functionality allows any combination of VLAN to VC mapping, such as one-to-one, many-to-one, or many-to-many. See [Figure 12-3](#).

Up to 4 VCs per ADSL Port provisioned per subscriber requirements

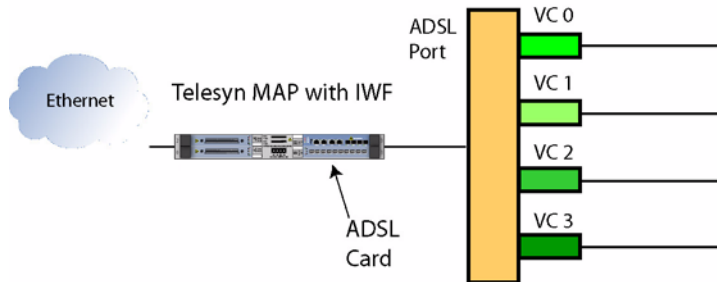


FIGURE 12-3 VLAN to VC Mapping

The fMAP system user can provision VCs as per customer requirements. For example, multiple subscribers could be served by a many-to-one VLAN to VC mapping with, for instance, Internet service provided by:

VC 0 mapped to VLAN 123 - Internet - Subscriber 1

VC 1 mapped to VLAN 123 - Internet - Subscriber 2

Another configuration for Internet and video could be as follows:

VC 0 mapped to VLAN 345 - Internet - Subscriber 1

VC 1 mapped to VLAN 678 - Video - Subscriber 1

VC 2 mapped to VLAN 345 - Internet - Subscriber 2

VC 3 mapped to VLAN 678 - Video - Subscriber 2

In this configuration, certain subscribers could be guaranteed a specified bandwidth.

The figure below is a graphical representation of the interface stack for multiple VCs. Each VC provisioned will add an AAL5 subinterface and an ETH subinterface to the ATM interface. VLANs are associated to the ETH subinterface.

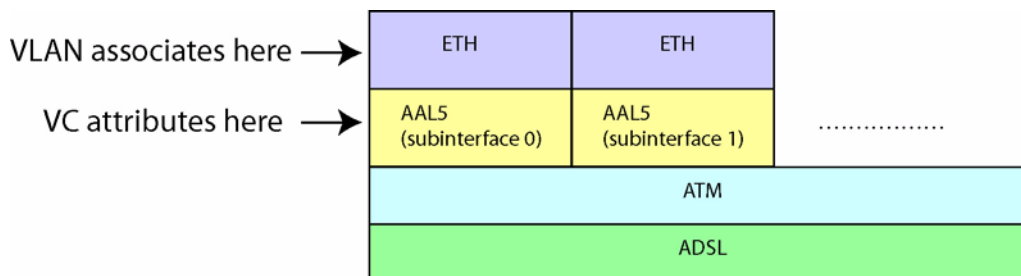


FIGURE 12-4 VLAN to VC association

12.4.3 Example VC provisioning

Here are three recommendations for how to provision for prioritization on VCs. Three models are shown.

12.4.3.1 Residential - Service-per-VC

In this model, each VC is carrying traffic with, potentially, a different p-bit priority. However, since all the VCs are UBR, all cell queues are serviced fairly, regardless of the priority of the traffic. Because ATM cell scheduling/prioritization does not consider p-bits, cells for low-priority packets on VC 1 may egress before cells for high-priority packets on VC 2 if it is VC 1's turn to transmit.

The only way that the user can provide a different treatment on a VC is to cap its PCR so that it is not always "ready to transmit". A PCR causes a VC to be "ready to transmit" only at its PCR rate. Therefore, a VC with PCR=1000 would have to wait 1ms between cells, giving other VCs an opportunity to transmit their cells in-between. However, there is a performance overhead to running the ATM QoS algorithm on a port. Adding a PCR value to a VC on a port incurs system overhead. Once the port incurs that overhead, additional PCRs on VCs on that same port do not incur additional overhead. This overhead must be considered when using PCR on a VC to ensure QoS on other VCs.

In the Triple Play scenario, if each service requires its own VC, the user should set a PCR, corresponding to the desired maximum per-service bandwidth utilization, on each VC. Setting bandwidth limits on services ensures that the remaining bandwidth is available for other services. For example:

- Video VC (VLAN 100): PCR=MAX
- Voice VC (VLAN 200): PCR=250 (~90Kbps)
- Data VC (VLAN 300): PCR=4000 (~1.5Mbps)

This will ensure that the Data VC uses no more than ~1.5Mbps (no more than one ATM cell every 250us), and the Voice VC uses no more than ~90Kbps (no more than one ATM cell every 4ms). For the Video VC, there is no required time between cells. It will still be competing against other VCs during cell times when they are allowed to transmit, but the other VCs have limited opportunity to transmit.

Note that this still does not ensure that voice and video are prioritized correctly. If correct p-bit prioritization is important, for example, to reduce voice packet latency, the user could map both voice and video, on different VLANs, to the same VC (with PCR=MAX) and the system would correctly prioritize between them within the VC. For example:

- Video/Voice VC (VLAN 100 & 200): PCR=MAX
- Data VC (VLAN 300): PCR=4000 (~1.5Mbps)

This eliminates the need for a PCR on the video or voice VC.




FIGURE 12-5 Residential subscriber with service per VC

12.4.3.2 Residential - Multiple Tenant Unit (MTU)

In this case, each VC carries packets for various services and each service’s packets may have a different p-bit priority. Since the system correctly prioritizes based on p-bits within a VC, this scenario functions well. It is expected that each “per-customer” VC would be treated fairly, unless the user has set a PCR on the VC.

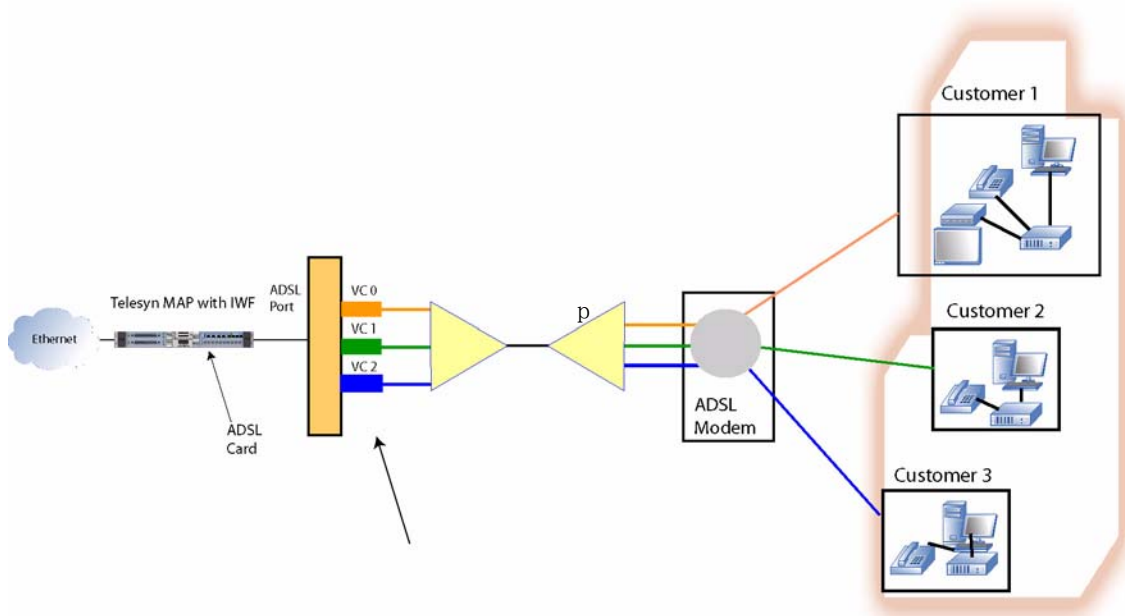


FIGURE 12-6 Residential MTU

12.4.3.3 Residential – Multiple Tenant Unit (MTU) + Service-per VC

This is the hybrid model with all the characteristics of the contributing models above. It would look like the illustration in [Figure 12-6](#), however, with a service on each VC. For instance, voice could be provisioned on VC0, video on VC1, and data on VC2.

12.4.4 Usage Notes

- This functionality is supported as follows:
 - Multiple (4) VC to xDSL interface mapping:
 - ADSL24A
- VLANs and PCR can be applied on a VC-basis. Applying other attributes to the ethernet interface of VC0 applies them to all the configured VCs on that interface.
- All VCs are UBRs, with PCR specified for each.
- VCCs can be added-to or deleted-from the ATM interface.
- VCs do not have an ADMIN state.
- The default VC with VCID=0 can never be deleted.
- A Default VC will exist at system start up and card creation with VCID=0. It will be associated with the default VLAN (VID =1). The default VPI/VCI will come from the ADSL port profile at their default values. The user can change these values.
- 1 untagged VLAN per interface also applies to VC as it does per port. Therefore, VC0 can have only 1 untagged VLAN association, plus 0-to-n tagged VLAN associations.
- The PCR limiting does not follow token/leaky bucket algorithms. It is a strict limit, there is no accommodation for transient bursts.
- Individual VCs cannot be disabled.
- The PCR value on ADSL VCs is ≥ 150 for all cards at the CLI, although this will not affect any pre-existing values that were less than 150.
- To allow over-subscribing VCs on a port.
 - The fMAP does not limit the VC's PCR to the line rate (i.e. the user can provision an individual VC's PCR to a value greater than the line rate).
 - The fMAP does not limit the sum of the PCRs of all VC's to the line rate (i.e. the user is allowed to over-subscribe the line).

Note: All ADSL have subinterfaces for their VCs.

12.4.5 Example configuration

Following is an example of configuring a VLAN to VC Association:

Add a VC to an ATM interface:

```
officer SEC>> ADD VC=2 INTERFACE=ATM: 9.0 VPI=10 VCI=40
officer SEC>> Info (010017): Operation Successful
```

Display all VCs:

```
officer SEC>> SHOW VC ALL
```

```
--- ATM - VCs ---
Interfaces VC VPI VCI Traffic-params
-----
```

Interfaces	VC	VPI	VCI	Traffic-params
ATM: 8.0	0	0	35	TXPEAKCELLRATE : MAX
	1	0	40	TXPEAKCELLRATE : MAX
	2	0	42	TXPEAKCELLRATE : MAX
	3	0	44	TXPEAKCELLRATE : MAX
ATM: 8.1	0	0	35	TXPEAKCELLRATE : MAX
ATM: 8.2	0	0	35	TXPEAKCELLRATE : MAX
ATM: 8.3	0	0	35	TXPEAKCELLRATE : MAX
ATM: 8.4	0	0	35	TXPEAKCELLRATE : MAX
ATM: 8.5	0	0	35	TXPEAKCELLRATE : MAX
ATM: 8.6	0	0	35	TXPEAKCELLRATE : MAX
ATM: 8.7	0	0	35	TXPEAKCELLRATE : MAX
ATM: 9.0	0	0	35	TXPEAKCELLRATE : MAX
	2	10	40	TXPEAKCELLRATE : MAX
ATM: 9.1	0	0	35	TXPEAKCELLRATE : MAX
ATM: 9.2	0	0	35	TXPEAKCELLRATE : MAX
ATM: 9.3	0	0	35	TXPEAKCELLRATE : MAX
ATM: 9.4	0	0	35	TXPEAKCELLRATE : MAX
ATM: 9.5	0	0	35	TXPEAKCELLRATE : MAX

(Some text omitted)

Display all VCs for an interface:

```
officer SEC>> SHOW VC INTERFACE 9.0
```

```
--- ATM - VCs ---
Interfaces VC VPI VCI Traffic-params
-----
```

Interfaces	VC	VPI	VCI	Traffic-params
ATM: 9.0	0	0	35	TXPEAKCELLRATE : MAX
	2	10	40	TXPEAKCELLRATE : MAX

The user can change attributes for one of the VCs. For example, in the following system responses, VPI and VCI were changed for interface 9.0 VC2 and VC0:

```
officer SEC>> SET VC 2 INTERFACE 9.0 VPI=15 VCI=32 TXPEAKCELLRATE=1200
officer SEC>> Info (010017): Operation Successful
officer SEC>> SHOW VC INTERFACE 9.0
```

```
--- ATM - VCs ---
Interfaces VC VPI VCI Traffic-params
-----
```

Interfaces	VC	VPI	VCI	Traffic-params
ATM: 9.0	0	0	35	TXPEAKCELLRATE : MAX
	2	15	32	TXPEAKCELLRATE : 1200

```
Info (010017): Operation Successful
officer SEC>> SET VC 2 INTERFACE 9.0 VPI=10 VCI=40 TXPEAKCELLRATE=MAX
officer SEC>> Info (010017): Operation Successful
officer SEC>> SHOW VC INTERFACE 9.0
```

```
--- ATM - VCs ---
Interfaces VC VPI VCI Traffic-params
-----
```

Interfaces	VC	VPI	VCI	Traffic-params
ATM: 9.0	0	0	35	TXPEAKCELLRATE : MAX
	2	10	40	TXPEAKCELLRATE : MAX

```
Info (010017): Operation Successful
officer SEC>> SET VC 0 INTERFACE 9.0 VPI=2 VCI=42
officer SEC>> Info (010017): Operation Successful
officer SEC>> SHOW VC INTERFACE 9.0
```

```
--- ATM - VCs ---
Interfaces VC VPI VCI Traffic-params
-----
```

Interfaces	VC	VPI	VCI	Traffic-params
ATM: 9.0	0	2	42	TXPEAKCELLRATE : MAX
	2	10	40	TXPEAKCELLRATE : MAX

Create a new VLAN to use with interface 9.0 VC2. For example, VLAN 20 could be used for Internet data:

```
officer SEC>> CREATE VLAN VLANNEW VID 20 FORWARDINGMODE STD
Info (010017): Operation Successful
```

Add interface 9.0 VC2 to the new VLAN 20:

```
officer SEC>> ADD VLAN=20 INTERFACE 9.0.2
Info (040544): One or more interfaces have been deleted from the default VLAN
Info (010017): Operation Successful
```

Display the VLAN:

```
officer SEC>> SHOW VLAN 20
--- VLAN Information -----
Type..... VLAN
Name..... VLANNEW
Identifier..... 20
Status..... static
Forwarding Mode..... Standard
IP module attached..... None
Untagged interfaces..... ETH: [9.0.2]
Tagged interfaces..... None
VLAN Translation interfaces..... None
```

The user could then create the same provisioning for video on this same interface, but use a different VC.

TABLE 12-4 VLAN to VC Mapping commands

Object	Verb	Syntax	Description
VC INTERFACE	ADD	<pre>ADD VC=vcid INTERFACE={ type:id-range id-range ifname-list } VPI=0..255 VCI=32..65535 [TXPEAKCELLRATE={ 150..65535 MAX }]</pre>	<p>Creates one or more VC(s) on one or more existing ATM interface(s). The default VC is VC-ID 0 and is created, internally, by the system, therefore it cannot be added.</p> <p>In release 6.0 TXPEAKCELL-RATE has a range.</p>
VC INTERFACE	DELETE	<pre>DELETE VC={ vcid-range ALL } INTERFACE={ type:id-range id-range ifname-list ALL } [FORCE]</pre>	<p>Delete one or more VC(s) from an ATM-interface or a list of ATM interfaces. The user cannot delete the default VC with VC-ID 0 since it is created by the system.</p>

TABLE 12-4 VLAN to VC Mapping commands

Object	Verb	Syntax	Description
VC INTERFACE	SET	<pre> SET VC=vcid INTERFACE={ type:id-range id-range ifname-list } [VPI=0..255] [VCI=32..65535] [TXPEAKCELLRATE={ 150..65535 MAX }] </pre>	<p>Allows the user to change VC interface attributes.</p> <p>In release 6.0 TXPEAKCELLRATE has a range.</p>
VC	SHOW	<pre> SHOW VC [={ vcid-range ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	Displays information about VCs.

12.5 VLAN Distribution

lists the VLAN distribution for the fMAP products.

TABLE 12-5 VLAN Distribution

255 VLANS	512 VLANs	4094 VLANs	Other
ADSL24	ADSL24A	ADSL8S	EPON2 ^a
FE10	ADSL24B	ADSL16	
FX10	SHDSL24	GE1	
FE2	NTE8	GE2RJ	
		GE3	
		GE4	
		GE8	
		9100 (with CFC12)	
		9400/9700 (with CFC24)	
		9400/9700 (with CFC56)	

a. 24 VLANs per EPON port, and 6 per ONU

13. Topology Configuration and Control

13.1 Overview

A sequence of layer-2 switches and physical LANs may be connected together in a network topology that results in more than one path between any two switches. This condition is called a network **loop**. If a loop exists, frames transmitted onto the extended LAN would circulate around the loop indefinitely, decreasing the performance of the extended LAN.

A variety of topology features can be provisioned to engineer the network properly. The following table lists the features that are available in release 6.0 and the interaction between them, since with release 6.0 there are many of these features available and users must know how they work together to properly engineer the network.

Moreover, users should understand that some features may be mutually exclusive and so should **not** be used together (usually on the same port or device).

Finally, users should understand the relationship of the physical and virtual (VLAN) topology, since in some cases users wish for a physical link (or links) to be blocked, and in other cases the virtual links (VLANs) to be blocked.

TABLE 13-1 Topology Control Features

Feature	Function	Interaction Notes
(R)STP	Controls a system wide spanning instance to avoid loops. Includes STP, which has the basic elements needed for topology control, and Rapid STP, which requires some modification to these elements. By Default, a system has RSTP disabled on all interfaces	Refer to 13.2
LAG	Allows multiple physical links to be joined into a LAG, which creates one virtual link. If one link fails, traffic is distributed to the remaining inservice links	Refer to 13.3

TABLE 13-1 Topology Control Features

Feature	Function	Interaction Notes
EPSR	Ethernet Protection Switched Ring (EPSR) is a protection scheme for ring-based network topologies. EPSR assists the multicast streams in being redirected around a faulty link in a ring network fast enough to result in an uninterrupted multicast service. EPSR is provided using the Allied Telesis Automatic Protection Switching (TAPS) protocol.	Refer to 13.4.4 . In release 6.0, EPSR and (R)STP topologies can co-exist, as long as the connection of an (R)TP sub-network is to two adjacent nodes of the EPSR ring. Refer to 13.6
UCP	A generic protocol that informs other devices in the network that it is the “upstream node” for a UFO VLAN. Moreover, using UCP protocol messages, the non-upstream nodes for the UFO VLAN can dynamically determine their upstream interfaces..	UCP actions occur independently of the topology feature being used; therefore, UCP can be used by itself as well as with STP and EPSR Refer to 13.5
MSTP	Separate spanning tree instances are created and associated with VLANs (or groups of VLANs). Each of the separate instances elect root bridges, root ports, and designated bridges independently.	Refer to 13.7
DHCP	Allows IP hosts to obtain protocol configuration parameters automatically through the network.	Refer to 13.8 .
BPDU Cop	A feature for STP instances (RSTP or MSTP), an interface that has this feature activated is placed out of service when it receives a BPDU	Refer to .13.9
LLDP	An IEEE-defined protocol to facilitate a network management system to derive the physical topology of the Network Elements	Refer to 13.10 .

13.2 Spanning Tree Protocol (STP)

13.2.1 Overview

The Spanning Tree Protocol (STP) makes it possible to automatically disable redundant paths in a network to avoid network loops, and to re-enable them when it is necessary to maintain connectivity in the event of a fault in the network (such as the failure of a link or a switch).

The spanning tree algorithm prunes redundant paths from the topology (i.e. marking paths as unavailable so frames are not transmitted over those paths). The resulting loop-free topology set of switches and active paths is called the *logical spanning tree*.

A logical spanning tree has the following elements:

- Each switch in the extended LAN has a unique **bridge ID**. This is a combination of the a switch's priority component (a value assigned by default or via manual configuration) and its the switch's MAC address.
- The switch with the *numerically lowest* bridge ID is considered the **root bridge** of the logical spanning tree.
- Each port on a switch has a unique **port ID**. This is a combination of the port's priority component (a value assigned by default or via manual configuration) and an internally assigned, unique numeric location identifier local to the bridge.
- Each port connecting a switch to a LAN has an associated **path cost**. This is a value assigned by switch software as a default based on port speed, or via manual user configuration, that provides an indication of the latency or resource consumption that would be encountered if a frame were to be transmitted on that port.
- The **root path cost** for a particular path from a port, a LAN, or switch to the root bridge is the sum of the port path costs incurred if a frame were to be transmitted on that path to the root bridge.
- The **root port** of a switch is the port on the switch with the *lowest* root path cost. If two or more ports on a switch have the same root path cost, the root port is the port with the numerically lowest port ID.
- The **designated bridge** of a LAN is the switch on the LAN with the lowest root path cost. If two or more switches on the LAN have the root path cost, the designated bridge is the switch with the lowest bridge ID.
- A **designated port** of a switch is a port that connects a LAN to its designated bridge.
- A **Spanning Tree instance** is a (named) logical representation of the underlying data structures and control mechanisms that provide a simple, fully-connected active network topology for a set of bridges and the LANs that connect them in a network.

Note: For the fMAP, RSTP is the default STP setting.

13.2.2 Protocol Concepts

13.2.2.1 Protocol Communication

To ensure that the switches in the extended LAN agree about root bridge, root port, and designated bridge elections, they must communicate information about bridge IDs and root path costs to other switches. This communication is accomplished via the exchange of messages known as **Configuration Bridge Protocol Data Units (BPDUs)**, also known as **hello messages**.

There is also the need to communicate when changes occur in the network topology (e.g. link failure or a new bridge). This type of communication is accomplished via the exchange of Topology Change Notification (TCN) BPDUs.

13.2.2.2 Spanning Tree Port States

A fMAP switch port that is participating in spanning tree operations can be in one of five states. A summary of the states is provided in [Table 13-2](#).

TABLE 13-2 Spanning Tree Port States

State	Meaning	Transition
Blocking	The port is disabled for receiving and transmitting normal traffic frames. It may receive BPDU frames, but does not transmit them. It does not add information about any MAC address from either Received BPDUs to its forwarding database	This is the initial state for each port. The switch also places the ports into his state to eliminate network loops, or if its perception if the network topology changes (new root port or root bridge).
Listening	The port does not receive or transmit traffic data frames. It may receive and transmit BPDUs. It does not add information about source MAC addresses from received BPDUs to the forwarding database.	The switch places ports into this state if it is a candidate for participating in the spanning tree topology.
Learning	The port does not receive or transmit traffic data frames. It may receive and transmit BPDUs. It adds source MAC address information from the BPDU to the forwarding database.	The switch places ports into this state upon expiration of a forwarding delay timer while in the listening state, unless something has caused the port to be placed in the blocking state.
Forwarding	The normal state for a port. The port is enabled and receiving and transmitting traffic data frames as well as BPDUs, and is adding source MAC address information for all frames to the forwarding database.	The switch places ports into this state upon expiration of a forwarding delay timer while in the learning state, unless something has caused the port to be placed in the blocking state.
Disable	No BPDUs are received or transmitted on the port	The switch places ports into this state based on manual actions.

Note: For Rapid STP (explained in 13.2.2.4), the “blocking” and “listening” states shown above are combined into a single “discarding” state

13.2.2.3 Convergence

The process by which the switches in the extended LAN come to agreement about the logical spanning tree topology is known as convergence. This process includes several key steps:

- The switches set their ports to the listening state. They elect a root bridge by exchanging hello messages to determine which switch has the lowest bridge ID.

- The root bridge initiates calculation of root path costs. Each switch uses information received from other switches, along with its own port cost information, to compute its own root path cost. It forwards this cost information along to other switches; eventually, the correct root path cost for every path in the extended LAN will be computed.
- Each switch elects a root port for that switch.
- The switches elect a designated bridge for each physical LAN, based on the root path cost for the switches
- Any port that is determined not to be a root port or a designated port is set to the blocking state.
- After the expiration of forwarding delay timers, every root port and designated port is set to the forwarding state. Once this is done, traffic may flow over the extended LAN, without any network loops being present.

If a link or switch fails, or the network topology otherwise changes, the network starts the convergence process again to reach a new spanning tree topology.

13.2.2.4 Rapid Spanning Tree (RSTP)

In the 802.1d Spanning Tree Algorithm and Protocol, timer driven processing controls how each port goes through the STP state transitions before being placed into a “forwarding” mode where normal traffic flow is supported. In the Rapid Spanning Tree Algorithm and Protocol (RSTP), significant time savings are accomplished using *rapid* STP port state transitions in many of the expected network topology change scenarios. The time savings is accomplished through additional information exchange and new “hand shake” processing between the ports of LAN connected bridges. The concept of a *point-to-point* connection is introduced to identify when a port is connected to exactly one other bridge. This condition must exist for some of the above mentioned rapid state transitions to take place. The concept of *edge ports* is also introduced to completely bypass the state transition process when a port is known to be connected to a single host.

The parameters that are associated with RSTP are included below.

13.2.3 Spanning Tree Parameters

13.2.3.1 Bridge Priority

Bridge IDs are used in root bridge elections. The root bridge is the switch in the extended LAN with the numerically lowest bridge ID value. This is guaranteed to identify a single bridge due to the unique MAC address component. The user is allowed to change the bridge priority component to override the arbitrary root selection that will result from only comparing MAC addresses when the default bridge priorities are in use.

Bridge IDs are also used in designated bridge elections. Normally the switch with the lowest root path cost is the designated bridge for a physical LAN. If more than one switch ties has the same lowest root path cost, then the designated bridge is the switch with the numerically lowest numbered bridge priority ID value.

The value of the PRIORITY parameter is used to set the writable portion of the bridge ID. The **default** bridge priority is 32768. To change the STP priority value, use the `SET STP PRIORITY` command

Note: The range is from 0 to 61,440 (a limitation of RSTP) in increments of 4096.

13.2.3.2 Port Priority

Port IDs are used in root port elections. Normally, the port with the lowest root path cost is the root port for the switch. If more than one port ties for the lowest root path cost, then the root port is the port with the lowest numerical port ID (as assigned by the system).

The default port priority value is 128. The Allied Telesis initial product release supported port priority values that could be configured on a per-port basis, as a value from zero to 255, in accordance with IEEE Std 802.1d, 1998 Edition. However, the storage space (number of bits) allocated to the priority component of the port ID had to be reduced to support bridges with larger numbers of ports, since this only left room for port numbers from 1-255.

Note: To maintain compatibility for comparison with previous versions of STP, the port priority is now considered to be a value between 0-240 that can only be provisioned in increments of 16.

13.2.3.3 Interface Path Costs

Interface path costs are used in root path cost calculations, which are a factor in root interface and designated bridge elections. By default, interface path costs are related to the bandwidth capacity of the interfaces; however, the default values may be changed by the user to reflect other factors (e.g. propagation delay, link quality, desired traffic level, etc.)

The default values and recommended ranges for path cost are as follows:

- Interface Speed: 10 Mbps
 - Default Path Cost: 100
 - Recommended Range: 50-600
- Interface Speed: 100 Mbps
 - Default Path Cost: 10
 - Recommended Range: 10-60
- Interface Speed: 1 Gbps
 - Default Path Cost: 4
 - Recommended Range: 3-10

The path cost values identified above reflect what was implemented in the initial fMAP product release as identified in IEEE Std.802.1d, 1998 Edition. The corresponding default values and recommended ranges for path cost as specified in IEEE Std. 802.1w-2001 to support RSTP and MSTP are shown in the table below.

- Interface Speed: 10 Mbps
 - Default Path Cost: 2,000,000
 - Recommended Range: 200,000-20,000,000
- Interface Speed: 100 Mbps
 - Default Path Cost: 200,000

- Recommended Range: 20,000-2,000,000
- Interface Speed: 1 Gbps
 - Default Path Cost: 20,000
 - Recommended Range: 2,000-200,000

A calculation is shown below that can be used to determine the recommended path cost value to use for intermediate link speeds:

$$20,000,000,000 / (\text{link speed in kb/s})$$

In LAN environments where bridges are in use that are operating different revision levels of STP, all the bridges must be configured to use compatible path cost value ranges. This will either require the older STP revision level bridges to be reconfigured to use the ranges specified in the newer standard, or the bridges with newer STP revisions will need to be configured to utilize the ranges from the older standard. The range of path cost values available from the older STP standard may be insufficient to support the data rates available in newer bridges.

The default PATHCOST values and the range of recommended PATHCOST values depend on the interface speed (as indicated above). If the path cost for an interface is not explicitly set, it will vary as the speed of the interface varies. Setting the path cost to a larger value on a particular interface is likely to reduce the traffic over the LAN connected to it. This may be appropriate if the LAN has lower bandwidth, or if there are reasons for limiting the traffic across it. To modify the STP interface path cost, use the command:

```
SET STP INTERFACE
```

If the path cost of an interface has been explicitly set to a particular value, it can be returned to its self-adjusting default path cost and priority, using the following command:

```
SET STP INTERFACE={type: i d-range|i d-range|i fname-l i s t|ALL} DEFAULT
```

Each interface also has a path cost, which is used if the interface is the root interface for the STP on the switch. The path cost is added to the root path cost field in configuration messages received on the interface to determine the total cost of the path to the root bridge. To modify the STP interface path cost, use the command:

```
SET STP INTERFACE={type: i d-range|i d-range|i fname-l i s t|ALL} PATHCOST=path-cost
```

Note: The range of the path-cost value for STP mode is 1..1000000. For RSTP mode, it is 1..200000000.

To display STP interface information, use the command:

```
SHOW STP INTERFACE[={type: i d-range|i d-range|i fname-l i s t|ALL}]
```

13.2.3.4 STP Timer Control Parameters

The Spanning Tree Protocol uses three configurable parameters for the time intervals that control the flow of STP information on which the dynamic STP topology depends:

- **HELLOTIME** (default 2 seconds) - This value determines how often the switch sends hello messages if it is the root bridge, or if it is trying to determine the root bridge identity in the network. Setting a shorter value

makes the network more robust, in that network changes can be detected more rapidly. Setting a longer value reduces network traffic and processing overhead.

- **MAXAGE** (default 20 seconds) - This value determines the maximum “age” of dynamic spanning tree configuration information (e.g. the root bridge ID, designated ports, and root ports). If this information has not been refreshed by hello messages before the timer expires, the information is discarded and the spanning tree must reconverge. If this timer is too short, the spanning tree will undergo reconvergence unnecessarily, resulting in network outages. If the timer is too long, the spanning tree may be slow to react to changes in network topology.
- **FORWARDDELAY** (default 15 seconds) - This value is used in the convergence process to allow for propagation of hello messages through the network. The timer represents how long ports are in the *listening* and *learning* states. By using this delay, the network has time for all the switches to agree on the spanning tree configuration. If the timer is too short, ports may reach the forwarding state before a stable topology has been reached. This may result in network loops that seriously degrade overall network performance. If the timer is too long, it will cause unnecessary delays in enabling the ports for passing bearer traffic. (At the default timer, the network will require at least 30 seconds for ports to transition from “blocking” to “forwarding”, since each port will spend 15 seconds in the “listening” state and 15 seconds in the “learning” state. All switches in the same spanning tree topology must use the same values for these parameters. The parameter values actually used by each switch are those sent by the root bridge, and forwarded to all other switches by the designated bridges.

Each switch that participates in the spanning tree (i.e. each switch in the extended LAN) must use the same values for these timers; otherwise, the convergence process would be unpredictable and unstable. To ensure that the timer values are consistent throughout the network, the timers for all the switches are set to values configured for the root bridge, once the identity of the root bridge has been determined.

The recommended relationship between the timer values can be expressed using the following formulae:

```
MAXAGE >= HELLOTIME x (number of network “hops” in longest path through network)
```

```
MAXAGE >= 2 x (HELLOTIME + 1 second)
```

```
MAXAGE <= 2 x (FORWARDDELAY - 1 second)
```

Note: A timer value, Migrate Time, is added for RSTP. It defaults to a constant value of 3 seconds.

To modify the parameters controlling these time intervals, use the command `SET STP` and the appropriate parameter.

13.2.3.5 The Priority Parameter

The value of the **PRIORITY** parameter is used to set the writable portion of the bridge ID, for example, the first two octets of the (8-octet long) Bridge Identifier. The remaining 6 octets of the bridge IDs are given by the MAC address of the switch. The Bridge Identifier parameter is used in all Spanning Tree Protocol packets transmitted by the switch. The first two octets, specified by the **PRIORITY** parameter, determine the switch’s priority for becoming the root bridge or a designated bridge in the network, with a lower number indicating a higher priority. In fairly simple networks, for instance those with a small number of switches in a meshed topology, it may make little difference which switch is selected to be the root bridge, and no modifications may be needed to the default

PRIORITY parameter, which has a default value of 32768. In more complex networks, one or more switches are likely to be more suitable candidates for the root bridge role, by virtue of being more centrally located in the physical topology of the network. In these cases, the STP PRIORITY parameters for at least one of the switches should be modified. To change the STP priority value, use the command `SET STP PRIORITY=bridge-priority`, where `bridge-priority` is 0..65535 for STP mode and 0..61440 in steps of 4096 in RSTP mode. To restore STP timer and priority defaults, use the command `SET STP DEFAULT`.

Changing the STP PRIORITY, using either of the previous commands, restarts the STP algorithm, so that elections for the root bridge and designated bridges begin anew, without resetting STP counters. To display general information about STPs on the switch, use the command `SHOW STP`.

13.2.3.6 Force Version

This parameter is used for RSTP. This parameter allows the user to specify that the bridge should operate in the STP_ORIGINAL mode, RSTP, or STP_COMPATIBLE_RSTP mode. If the STP_COMPATIBLE_RSTP mode is chosen, the RSTP will be compatible with other switches in the network that may not use RSTP and therefore use older parameter values and ranges.

13.2.3.7 Edge Port

This parameter allows the user to specify a port as an “Edge Port” when it is expected that a port will be directly connected to a host (i.e. a port at the “edge” of the Bridged LAN). Additional processing is associated with the use of this parameter to verify that a port identified as an “Edge Port” by the user is not actually connected to another bridge. This parameter and its associated processing can facilitate a port state transition directly to the forwarding state as part of the RSTP processing.

In the `SET STP INTERFACE` command, set `EDGEPORT=TRUE` to enable this for RSTP.

13.2.3.8 Point-to-Point Port

This parameter allows the user to specify a port as a *Point-to-Point Port* when it is expected that it will be connected to exactly one other bridge. Additional processing is associated with this parameter to automatically determine whether or not the port should be considered a point-to-point connection, when so indicated by the user via (auto),parameter setting

The Point-to-Point Port parameter, and its associated processing, is utilized by the RSTP to facilitate the rapid transition of a port into the forwarding state under certain conditions specific to Point-to-Point ports only.

In the `SET STP INTERFACE` command, set `POINT2POINT=TRUE` or `AUTO` to enable this for RSTP.

Note: In most cases, select `AUTO` so that the system can determine the port connection.

13.2.3.9 Transmit Hold Count

This parameter allows the user to specify the maximum BPDU transmission rate for any port on the bridge, which therefore determines how much STP control traffic is going into the network. The **default** value for this

parameter is 6, indicating that at most 3 BPDUs can be transmitted from any port in a given Hello Time period (2 seconds by default).

In the SET STP command, the parameter is TXMAX; the range is 1 to 10 (with the default of 6).

13.2.3.10 Enable/Disable STP

The default Spanning Tree instance is disabled by default at switch start up, and Spanning Tree instances created by a user are disabled by default when they are created. To enable or disable Spanning Tree instances, use the commands `ENABLE/DISABLE STP`.

13.2.3.11 Enable/Disable Port or Interface

When an STP is enabled in a looped or meshed network, it dynamically enables and disables particular ports belonging to it, to eliminate redundant links. All ports in a VLAN belong to the same STP, and their participation in STP configuration are enabled by default when STP is enabled, and hence the possibility of them being elected to the STP's active topology. To enable or disable particular ports for participation or exclusion from STP operations, use the commands `ENABLE` and `DISABLE STP INTERFACE`.

This command also supports the `TOPOLOGYCHANGE` parameter to control the detection of topology changes on the associated port. This allows the disabling of topology change detection on ports that are known to be connected to single end stations that could cause the Topology Change Notification mechanism to be triggered for the entire network when the end station is power cycled.

13.2.3.12 Display Counters

To display STP counters, use the following command, with the results shown below.

```

officer SEC> SHOW STP COUNTER

Type..... STP
STP Name..... default
State..... DISABLED
Topology Change..... FALSE
Topology Change Time(seconds)..... 35
Hold Time(seconds)..... 1
Designated Root Priority..... 32768
Designated Root Address..... 00:0C:25:00:13:78
Root Path Cost..... 0
Root Port (slot.port)..... 0
Max Age(seconds)..... 20
Hello Time(seconds)..... 2
Forward Delay(seconds)..... 15
Bridge ID Priority..... 32768
Bridge ID Address..... 00:0C:25:00:13:78
Bridge Max Age(seconds)..... 20
Bridge Hello Time(seconds)..... 2
Bridge Forward Delay(seconds)..... 15
Force Protocol Version..... Original STP
Tx Hold Count..... 6

```

13.2.3.13 Reset STP

The spanning tree algorithm can be recalculated at any time, and all timers and counters be initialized, using the command `RESET STP`.

13.2.3.14 Default STP Configuration for the fMAP Series Product

By default the fMAP series product has the following STP set-up:

- There is one STP instance that cannot be destroyed. Its name is “default” and its initial state is **disabled**.
- By default all ports will belong to the default STP.

13.2.4 An STP Network with Multiple VLANs

Since STP is a port-based topology and VLAN is a logic-based topology (over a physical port), the user needs to understand how these two work together so that the blocked links that are part of (R)STP convergence do not have unintended consequences for the VLANs that are carried over these ports.

Note: In release 6.0, with the MSTP feature, there can be an additional (R)STP instance based on a (set of) VLANs. However, the configuration rules listed here should be understood first, since they apply to understanding the MSTP instances.

Refer to the following figure, which shows an (R)STP topology in which two physical links are blocked and two VLANs are configured.

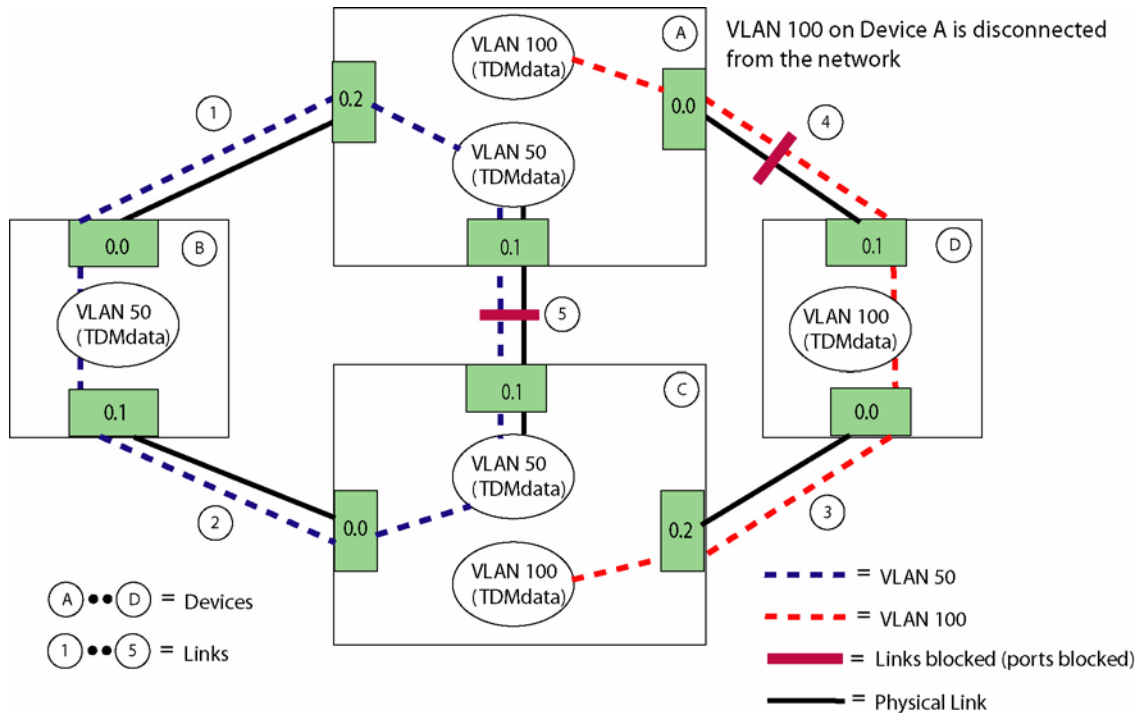


FIGURE 13-1 STP Network with Multiple VLANs - STP Blocks Two Ports and VLAN is Isolated

To prevent loops, STP convergence has blocked links 4 and 5. For VLAN 50, this is not a problem; VLAN 50 follows a physical loop and so actually mimics the loop and needs to be blocked. VLAN 100, however, is a non-looped VLAN, and so with physical link 4 being blocked, the VLAN on Device A is disconnected from the network and cannot send or receive data.

From this figure two rules follow:

- There should be no non-looped VLANs in the STP network.
- Looped VLANs should follow physical loops of the STP network.

13.2.5 STP and LAG Interaction

As explained in 13.3, LAG is another feature that allows traffic to survive component failure; with the creation of a LAG group, if a link within that group fails, the traffic will be distributed among the remaining members of the group.

Following are the feature interactions:

- When LAG and STP are configured on the same system, a LAG group is a single set of ports; if one link within the LAG group fails, the LAG group will recover the traffic and STP is not affected. However, if all

the links in the LAG group fail, STP will see this as a failure of the LAG interface and will reconverge the network.

- Any port in a LAG group can be displayed as the root port



If a system has only LAG configured, and a port in the LAG group is deleted before being disabled, there will be loops. If the system has STP and a port in the LAG group is deleted before being disabled, the STP will prevent loops. In general, though, the user should always disable a port before deleting it.

13.2.6 Command Summary for STP

In release 6.0, with the MSTP feature, the commands to create one or multiple STP instances are similar (except for the parameters and values used), so the STP commands are covered in the MSTP subsection, [13.7.8](#).

13.3 Link Aggregation (LAG)

13.3.1 Overview

The Link Aggregation Group (LAG) feature, defined in 802.3ad, allows multiple physical links to be joined into a LAG, which creates one virtual link. If one link fails, traffic is distributed to the remaining in-service links. Links can be added or deleted from the LAG, depending on traffic requirements. This release provides the following for the LAG feature:

- The LAG group ID (lag-ID) can be defined using an interface as well as a set of ports.
- Other features, such as VLAN and STP, can be configured against the lag-id or LAG interface ID.
- All ports in the LAG group must share the same untagged and tagged VLAN configuration.
- LAG commands no longer have the keyword SWITCH as part of their command syntax.

Note: For interoperability of switches from different vendors, the control of data traffic involves the Link Aggregation Control Protocol (LACP, defined in IEEE 802.3ad), which contains rules for configuring the ports on each side of the links in the group. This is not supported in the current release, so the LAG feature operates in static mode only. For a list of the MIB variables that are supported, refer to the Specification section.

The LAG feature works on both the fMAP products, but there are key differences in how they can be deployed. Refer to Section 16. for 9000 features.

13.3.2 Provisioning Rules

Following are the provisioning rules to follow when creating LAG groups:

1. All ports in a LAG group must have the same VLAN configuration (tagged and untagged ports).
2. All ports in a LAG group must share the same provisioning attributes:
 - Same speed
 - Same auto negotiate settings (AUTONEGOTIATE=ON)
 - Full duplex

The “LAG speed” is determined by the speed of the first port added to the LAG. If the user tries to add a second interface with a different speed, there will be a “port speed does not match LAG speed” error message.

Note: For the 9100 GE2RJ, there are checks made to the CLI to ensure the user sets these parameters manually. Refer to Section 16.8.3

3. All ports cannot have egress rate limiting configured.
4. Ports cannot have a traffic descriptor configured.
5. Once a port belongs to a LAG group, changing of individual port attributes is allowed using the LAG interface.

6. All ports must have the same Classifier configuration. (Refer to 15.3.)
7. The LAG group has an Operational State. An Operational State of UP means the LAG group has been provisioned and one or more ports are in an Operational State of UP.
8. LAG and (R)STP are compatible; a LAG can be created (or interfaces added to a LAG) regardless of the STP state. Also, STP can be disabled/enabled on an interface regardless of whether the interface is in a LAG or not.

Note: A LAG group does not have an Administrative state; its Operational state is determined by the Operational State of its associated ports. However, the MODE parameter does determine the traffic carrying capability of the LAG group, since the MODE must be ON before the LAG group can carry traffic.



A LAG group can be pre-provisioned with the MODE set to OFF, and then set to ON to apply the LAG group to the hardware. However, if the ports have been physically provisioned the mode can be OFF and loops can result.

13.3.3 Destroying a LAG

To destroy a LAG, the user must first perform these steps (the order of the first two steps is not important).

1. All member ports of the LAG must be deleted (i.e. no longer associated) with the LAG.

Since the LAG no longer has any associated member ports, it will change its Operational State to DOWN.

Note: Although the LAG group no longer has associated ports, the ports can remain enabled and with an Operational state of UP, since what is being deleted is an association.

2. The LAG mode must be turned to OFF.
3. The LAG can then be destroyed.

With these steps, the aggregation is taken-down in hardware, the member ports are removed from LAG provisioning, and the LAG is destroyed (removed from provisioning).



A LAG can be destroyed, but since the ports are still physically connected and enabled, traffic can run over these ports and loops may result. To avoid this, the user should either have STP configured and enabled for both the system and related ports, or all ports that were in the LAG group should be disabled except for one port.

13.3.4 Command Objects for Link Aggregation

TABLE 13-3 Commands for Link Aggregation

Object	Verb	Syntax	Description
LAG	SHOW	<pre> SHOW LAG [= { lag-list ALL }] [{ INFO STATE LACPSTATS MACSTATS }] </pre>	<p>Show information for the one lag-id, a set of lag-ids, or all lag-ids that have been created.</p> <p>INFO includes general information such as port list, speed, select criteria, and the ADMINKEY.</p> <p>MACSTATS are the MAC statistics.</p> <p>STATE and LACPSTATs are part of LACP and are not used.</p>
	SET	<pre> SET LAG=lag-id [MODE={ ON OFF PASSIVE ACTIVE }] [SELECT={ MACSRC MACDEST MACBOTH IPSRC IPDEST IPBOTH PORTSRC PORTDEST }] [ADMINKEY=1..1024] </pre>	<p>Modify the lagID attributes, such as turning the MODE to ON to activate the LAG.</p> <p>The lag-id itself cannot be changed.</p>
	DESTROY	<pre> DESTROY LAG=lag-id </pre>	<p>Before a LAG can be destroyed, the following must be true:</p> <ul style="list-style-type: none"> - the MODE must be set to OFF. - the member ports must be deleted from the LAG using the DELETE LAG=<lag-id> PORT=<port-list> command.

TABLE 13-3 Commands for Link Aggregation

Object	Verb	Syntax	Description
LAG PORT/ INTERFACE	CREATE	<pre> CREATE LAG=lac-id [INTERFACE={ type:id-range id-range ifname-list }] [MODE={ ON OFF PASSIVE ACTIVE }] [SELECT={ MACSRC MACDEST MACBOTH IPSRC IPDEST IPBOTH PORTSRC PORTDEST }] [ADMINKEY=1..1024] </pre>	<p>Create a unique LAG name for an interface or interface list.</p> <p>When the LAG is created, the default mode is OFF.</p> <p>PASSIVE and ACTIVE are part of LACP and are therefore not used.</p> <p>SELECT is the criteria to determine frame distribution among the interfaces in the LAG. The default is MACBOTH.</p> <p>The ADMINKEY is set at 0 by the system. These numbers are part of LACP and are therefore not used.</p>
	ADD	<pre> ADD LAG=lac-id INTERFACE={ type:id-range id-range ifname-list } </pre>	<p>Add to the lag-id a interface or set of interfaces. If the CREATE command was used to associate all the GE1 interfaces, this command would not be needed.</p>
	DELETE	<pre> DELETE LAG=lac-id INTERFACE={ type:id-range id-range ifname-list ALL } </pre>	<p>Deletes all interface associations with the LAG. The virtual link is stopped and traffic will revert to STP if it is configured.</p>

13.4 Networking Topologies in Release 5.0

13.4.1 Overview

*Note: The following summary of topology features for release 5.0 is included here because in release 6.0, with the introduction of Upstream Control Protocol (UCP), existing topologies will have the FORWARDING values of VLANs change automatically during upgrade from 5.0 to match new values in 6.0. (The main change is a status or RING will change to UCP or STP, depending on the topology feature used.) Certain values will not be changed; however the user has the option to change the values before upgrade so that the changes **will** be automatic. Knowledge of these features is therefore important so the user can understand and control the 6.0 topology features. Section [13.5.5](#) has a summary of these changes, but the user should read this section first to have a better understanding of the changes that are made.*

In release 5.0, the subtending configuration topologies that could be implemented for fMAP products were:

- Linear
- Daisy-chain
- Ring Network

In 5.0, subtending was when several network elements are viewed and managed by a **single** management platform and where multiple traffic streams were concentrated onto a **single** network uplink.

In each of these implementations, multiple fMAP products connect to one system that aggregates the transmission of all the attached systems and provides one network uplink. The aggregating system is referred to as the **subtending system** and each system(s) connected to it is/are referred to as the **subtended system**.

Note: It is assumed that the user has an understanding of the operation of Upstream Forwarding Only (UFO) service and IGMP for fMAP products. This understanding is important, since for many of the features available the type of VLANs used (UFO or standard) affect the functionality of the feature.

13.4.1.1 Single Shelf Network Configuration

[Figure 13-2](#) shows the deployment of the fMAP product in a single shelf network configuration. The service module ports are the downstream ports towards the subscriber and the network module ports are the upstream ports towards the network.




FIGURE 13-2 Single shelf network configuration

13.4.2 Linear Daisy Chain Network

The daisy chain is a serial link (or cascaded link) of two or more systems with one network uplink connection and one downlink connection. The subtending system aggregates the traffic from the subtended system, switches all the traffic (its own and that of the subtended systems), and provides the network uplink.

In the two system daisy chain network configuration illustrated in [Figure 13-3](#) shown below, the switch on the top is the Subtending System and the switch on the bottom is the Subtended System.

Each of the switches, whether subtending or subtended can connect to subscriber traffic via their SM cards. The relationship of subtending the system in a daisy chain configuration is as follows. Two NM cards per system are needed to configure a daisy chain. One NM connects upstream towards the network and the second NM connects downstream to a subtended switch.




FIGURE 13-3 Linear daisy-chain configuration with new ports and links indicated

In order to implement a linear daisy-chain network, a NM interface in the subtending system must be configured as a downstream interface or port. This is accomplished with the SET VLAN command.

```
SET VLAN={ vlnname | vid } INTERFACE={ type:id-range | id-range | ifname-list | ALL } [ FRAME={
UNTAGGED | TAGGED } ] [ TRANSLATE={ 1..4094 | NONE } ] [ FORWARDING={ UPSTREAM |
DOWNSTREAM | RING | PROTECTIONLINK } ]
```

In this example, the user creates a vlan called V_60, adds an interface to it called ETH:0.0, then configures it as a downstream interface.

Note: In a daisy chain configuration, the VLAN must be configured as UFO. Any UFO data received on the downstream interface is switched only to the upstream interface. IGMP reports and leaves received on the downstream port are filtered for duplicate reports on NM ports before being sent to the MC router. This prevents overload and reduces processing at the MC router. Also, when a multicast stream is received over the upstream port it is sent over the downstream port to the subtended system below.

```
officer SEC>> CREATE VLAN V_60 VID=60 FORWARDING UPSTREAM
Info (010017): Operation Successful
Warning (040526): No uplink port(s) in this UFO type of VLAN
officer SEC>> SHOW VLAN ALL
```

```
--- VLAN Information ---
Name          VID  Forwarding Tagged Interfaces      Untagged Interfaces
Mode
-----
```

Name	VID	Forwarding Mode	Tagged Interfaces	Untagged Interfaces
default	1	Upstream	None	ETH: 0.0/ETH: 8.0.0, 8.1.0, 8.2.0, 8.3.0, 8.4.0, 8.5.0, 8.6.0, 8.7.0, 9.0.0, 9.1.0, 9.2.0, 9.3.0, 9.4.0, 9.5.0, 9.6.0, 9.7.0, 9.8.0, 9.9.0, 9.10.0, 9.11.0, 9.12.0, 9.13.0, 9.14.0, 9.15.0
DATA V_60	312 512	Upstream	ETH: 1.0 None	None

```
officer SEC>> ADD VLAN V_60 INTERFACE=ETH: 0.0
Info (040544): One or more interfaces have been deleted from the default VLAN
Warning (040526): No uplink port(s) in this UFO type of VLAN
Info (010017): Operation Successful
```

```
officer SEC>> SHOW VLAN ALL
--- VLAN Information ---
Name          VID  Forwarding Tagged Interfaces      Untagged Interfaces
Mode
-----
```

Name	VID	Forwarding Mode	Tagged Interfaces	Untagged Interfaces
default	1	Upstream	None	ETH: 8.0.0, 8.1.0, 8.2.0, 8.3.0, 8.4.0, 8.5.0, 8.6.0, 8.7.0, 9.0.0, 9.1.0, 9.2.0, 9.3.0, 9.4.0, 9.5.0, 9.6.0, 9.7.0, 9.8.0, 9.9.0, 9.10.0, 9.11.0, 9.12.0, 9.13.0, 9.14.0, 9.15.0
DATA V_60	312 512	Upstream	ETH: 1.0 None	None ETH: 0.0

```
officer SEC>> SET VLAN TEST1 INTERFACE=ETH: 11.0 FORWARDING DOWNSTREAM
Warning (040526): No uplink port(s) in this UFO type of VLAN
Info (010017): Operation Successful
```

```
officer SEC>> SHOW VLAN TEST1
--- VLAN Information ---
Type..... VLAN
Name..... V_60
Identifier..... 60
Status..... static
Forwarding Mode..... Upstream Forwarding
IP module attached..... None
```

```

Untagged Upstream ports..... None
Tagged Upstream ports..... None
Untagged Downstream ports..... None
Tagged Downstream ports..... ETH: 0. 0
Untagged Ring ports..... None
Tagged Ring ports..... None
Untagged ProtectionLink ports..... None
Tagged ProtectionLink ports..... None
VLAN Translation Ports..... None

```

Figure 13-3 is an example of configuring two systems for a linear daisy-chain operation. For this example, it is assumed that both systems have been installed and provisioned and are operating normally.

The user should have previously installed and configured NMs on each system, one to be used for the downstream interface on the subtending system and the other to be used for the upstream interface on the subtended system. For this example, assume that a new NM has been installed in slot 11 on the subtending system and slot 10 on the subtended system. The cable from the NM in slot 11 of the subtending system is connected to the NM in slot 10 of the subtended system. The provisioning would be similar to the example detailed above. Once provisioning is complete, this subtending system is now sending traffic to the subtended system.

Note that when the user configures the NM in slot 10 of the subtended system, it is provisioned as upstream by default.

13.4.3 Ring-based Topology with STP

A ring network configuration is similar to the linear daisy chain configuration. However, in the case of a linear daisy chain network configuration, the aggregated link is a cascading of a number of serial links going from the subtended system to a subtending system, and so on, eventually exiting the access network at the network edge. As a result, if connectivity of any of the concatenated serial links is lost for whatever reason it results in a loss of service for all those systems situated below the point of link failure. When designing a daisy chain network configuration, the service provider should always plan an alternate path to prevent the loss of subscriber connectivity to the network for any kind of failure.

One way to achieve this capability is to chain a number of switches around a ring as shown below in [Figure 13-4](#) with the ends of the ring directly connected to an aggregating layer 2 switch. The upstream link from the aggregating layer 2 switch is the real network uplink for such a configuration.

For a ring network with fMAP switches such as shown in [Figure 13-4](#), Spanning Tree Protocol (STP) software will block a link. As a result, the topology can be considered to be two independent linear daisy chained network configurations. One of the network module ports then becomes the root port and the other network module port becomes the designated port. The root port is then considered to be the **upstream** port while the designated port is considered to be the **downstream** port.

In the case of a ring network topology, the network stabilizes itself into one or more daisy chain configurations after which traffic transmission characteristics for each of the resulting daisy chain configuration is independently handled as discussed above.

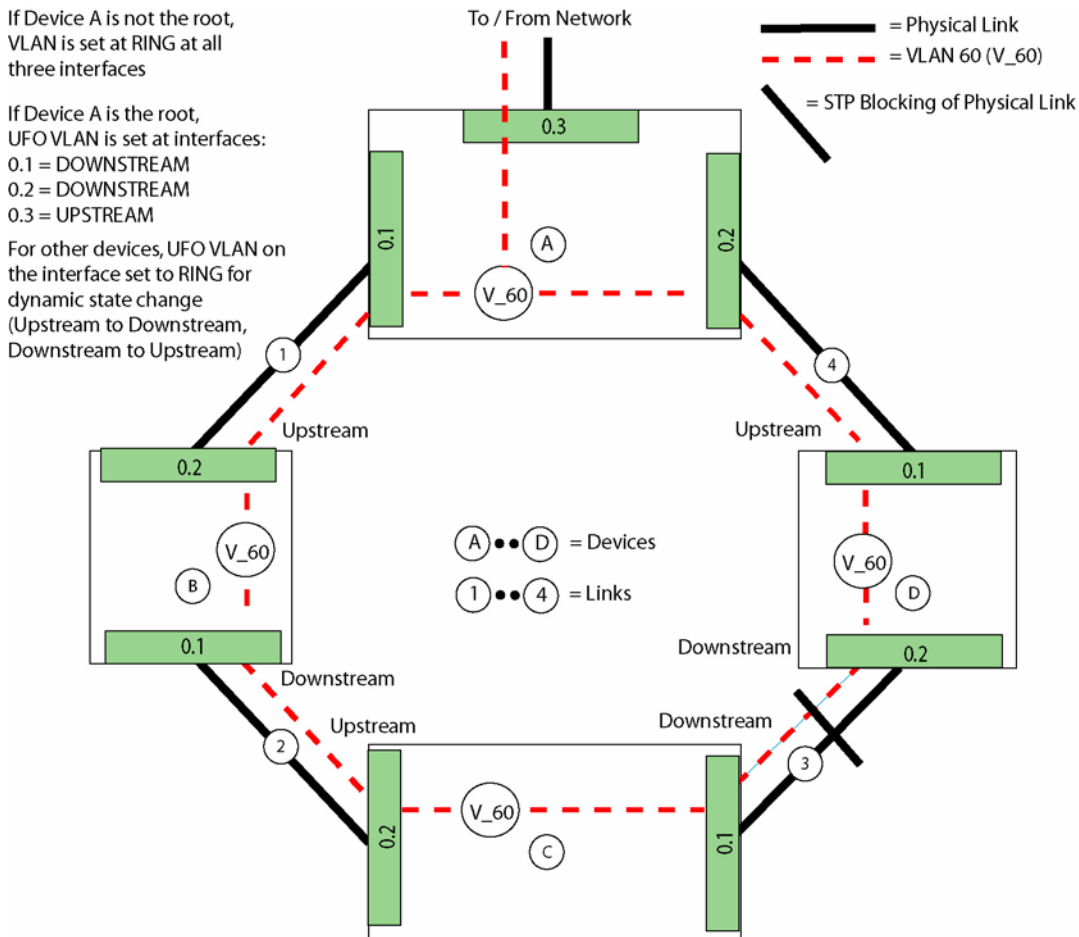


FIGURE 13-4 Ring Topology Using STP

Once the topology stabilizes, in each system one of the NM ports will become the *root* port (the one closest to the root bridge as determined by STP) and the other NM port(s) become the *designated* port. The NM port which is the root port is considered to be the upstream port and the NM port which is not the root port (designated port) is considered to be the downstream port.

For a ring-based daisy chained topology at configuration time, the user provisions the NM ports in each shelf with **FORWARDING=RING** indicating that the topology is a ring and allowing the STP protocol to determine the exact upstream NM port. Once the upstream port determination is made via the STP protocol, then the ring topology basically transforms itself into one or more linear daisy chain topologies as discussed earlier.

In order to make use of the root port and designated port designations to determine whether a NM port is upstream or downstream in a ring based daisy chain network configuration, the following engineering rule is enforced.



To prevent one of the fMAP systems from becoming the Spanning Tree root bridge, the network design must ensure that the appropriate STP parameters are set such that the root bridge is always located above the ring configuration made up of the fMAP systems.

13.4.3.1 Network topology and classifier interaction

When a ring topology is implemented in the user's network, the user may not be aware of which ring ports are designated *upstream* and which are designated *downstream*. Therefore, if the user wants to configure classifiers on these ring ports, they should apply them to both ring ports. This ensures that the traffic classification required by the user is actually applied.

This restriction strictly applies only to the switches which actually make up the ring. The aggregating layer 2 switch or any switch above it could be the root bridge for a ring network. These are configured as follows:

- If the aggregating layer 2 switch is an fMAP and does become the root bridge then the `FORWARDING` parameter for the NM ports as shown above **cannot** be set to **RING**. The upstream port from the aggregating layer 2 switch, which is the real network uplink for such a configuration, should be set to **UPSTREAM** and the other two links which make up the ring must be set to **DOWNSTREAM**.
- If the aggregating layer 2 switch is an fMAP and **not** the root bridge but one of the switches above it is, then the `FORWARDING` parameter can be set to **RING** and the status of the ports will be determined with STP's assistance similar to the other switches in the ring.

To implement a ring topology, NM interfaces must be configured as ring interfaces. This allows the STP protocol to determine the exact upstream and downstream ports. This provisioning is accomplished using the `SET VLAN` command.

```
SET VLAN={ vlnname | vid } INTERFACE={ type:id-range | id-range | ifname-list | ALL } [ FRAME={
UNTAGGED | TAGGED } ] [ TRANSLATE={ 1..4094 | NONE } ] [ FORWARDING={ UPSTREAM |
DOWNSTREAM | RING | PROTECTIONLINK } ]
```

Here is an example of provisioning a ring-based network topology. Note that the VLAN must be set to UFO.

```
officer SEC>> CREATE VLAN=V_60 VID=60
Info (010017): Operation Successful
officer SEC>> SHOW VLAN ALL
--- VLAN Information -----
Name          VID  Forwarding Tagged Interfaces      Untagged Interfaces
              Mode
-----
default t     1    Upstream   None                    ETH: 0. 0, 1. 0/ETH: 8. 0. 0,
8. 1. 0, 8. 2. 0, 8. 3. 0,
8. 4. 0, 8. 5. 0, 8. 6. 0,
8. 7. 0, 9. 0. 0, 9. 1. 0,
9. 2. 0, 9. 3. 0, 9. 4. 0,
9. 5. 0, 9. 6. 0, 9. 7. 0,
```

```

V_30          60 Standard  None          9. 8. 0, 9. 9. 0, 9. 10. 0,
                                         9. 11. 0, 9. 12. 0, 9. 13. 0,
                                         9. 14. 0, 9. 15. 0
None
    
```

```

-----
officer SEC>> SET VLAN=V_60 FORWARDING=UPSTREAM
Info (010017): Operation Successful
Warning (040526): No uplink port(s) in this UFO type of VLAN
    
```

```

officer SEC>> SHOW VLAN ALL
--- VLAN Information
Name          VID  Forwarding Tagged Interfaces      Untagged Interfaces
Mode
-----
default t    1    Upstream  None          ETH: 0. 0, 1. 0/ETH: 8. 0. 0,
                                         8. 1. 0, 8. 2. 0, 8. 3. 0,
                                         8. 4. 0, 8. 5. 0, 8. 6. 0,
                                         8. 7. 0, 9. 0. 0, 9. 1. 0,
                                         9. 2. 0, 9. 3. 0, 9. 4. 0,
                                         9. 5. 0, 9. 6. 0, 9. 7. 0,
                                         9. 8. 0, 9. 9. 0, 9. 10. 0,
                                         9. 11. 0, 9. 12. 0, 9. 13. 0,
                                         9. 14. 0, 9. 15. 0
V_30         312  Upstream  None          None
    
```

```

officer SEC>> ADD VLAN=V_30 INTERFACE=ETH: 1. 0
Info (040544): One or more interfaces have been deleted from the default VLAN
Warning (040526): No uplink port(s) in this UFO type of VLAN
Info (010017): Operation Successful
    
```

```

officer SEC>> SHOW VLAN ALL
--- VLAN Information
Name          VID  Forwarding Tagged Interfaces      Untagged Interfaces
Mode
-----
default t    1    Upstream  None          ETH: 0. 0/ETH: 8. 0. 0,
                                         8. 1. 0, 8. 2. 0, 8. 3. 0,
                                         8. 4. 0, 8. 5. 0, 8. 6. 0,
                                         8. 7. 0, 9. 0. 0, 9. 1. 0,
                                         9. 2. 0, 9. 3. 0, 9. 4. 0,
                                         9. 5. 0, 9. 6. 0, 9. 7. 0,
                                         9. 8. 0, 9. 9. 0, 9. 10. 0,
                                         9. 11. 0, 9. 12. 0, 9. 13. 0,
                                         9. 14. 0, 9. 15. 0
V_30         312  Upstream  None          ETH: 1. 0
    
```

```

officer SEC>> SET VLAN=V_30 INTERFACE=ETH: 1. 0 FRAME=TAGGED FORWARDING=RING
Info (010017): Operation Successful
    
```

```

officer SEC>> SHOW VLAN ALL
--- VLAN Information
Name          VID  Forwarding Tagged Interfaces      Untagged Interfaces
Mode
-----
    
```

default t	1	Upstream	None	ETH: 0. 0/ETH: 8. 0. 0, 8. 1. 0, 8. 2. 0, 8. 3. 0, 8. 4. 0, 8. 5. 0, 8. 6. 0, 8. 7. 0, 9. 0. 0, 9. 1. 0, 9. 2. 0, 9. 3. 0, 9. 4. 0, 9. 5. 0, 9. 6. 0, 9. 7. 0, 9. 8. 0, 9. 9. 0, 9. 10. 0, 9. 11. 0, 9. 12. 0, 9. 13. 0, 9. 14. 0, 9. 15. 0
V_30	312	Upstream	ETH: 1. 0	None

```
officer SEC>> SHOW VLAN=DATA
```

```
--- VLAN Information -----
Type..... VLAN
Name..... V_30
Identifier..... 312
Status..... static
Forwarding Mode..... Upstream Forwarding
IP module attached..... None
Untagged Upstream ports..... None
Tagged Upstream ports..... RING
Untagged Downstream ports..... RING
Tagged Downstream ports..... None
Untagged Ring ports..... None
Tagged Ring ports..... ETH: 1. 0
Untagged ProtectionLink ports..... None
Tagged ProtectionLink ports..... None
VLAN Translation Ports..... None
```

13.4.4 Protection Schemes - EPSR

Ethernet Protection Switched Ring (EPSR) is a protection scheme for Ethernet networks, specifically for ring-based network topologies. EPSR provides a 50 milliseconds switching time for an Ethernet-based ring network, similar to that provided by the Synchronous Optical Network (SONET) protocol, to maintain layer 2 redundancy in the network.

EPSR assists the multicast streams in being redirected around a faulty link in a ring network fast enough to result in an uninterrupted multicast service. EPSR is provided using the Allied Telesis Automatic Protection Switching (TAPS) protocol. This protocol provides fast protection switching to layer 2 switches which are interconnected in an Ethernet ring topology.

Note: EPSR is only supported on ring topology networks

EPSR is available only on ring topology networks comprised of nodes that are physically connected to form a ring. Each node on the ring will have two Ethernet ports connected to the ring. EPSR operates over these Ethernet ports.

13.4.4.1 Protection Scheme

The protection scheme for an Ethernet ring network basically operates by configuring an EPSR domain on the ring. The vlans that require fault protection are configured on all the ring ports and are assigned to the above mentioned EPSR domain. All such vlans are referred to as the *protected vlans*. Additionally, a *control vlan* is assigned to the EPSR domain and is used to send and receive the TAPS protocol control messages over the ring

network that are then used accordingly by all the nodes to prevent loops in the network and ensure that none of the nodes are isolated from the network.

Note: There can only be one control vlan per EPSR domain and is configured to use tagged frames. This control vlan is unique to this domain and cannot be re-used for another domain. Also, the control vlan must be provisioned to have the highest priority p-bit setting (as per IEEE 802.1p) and to be mapped to the highest priority queue in the system.

One of the nodes in the ring is designated as the **MASTER node** while all the other nodes are designated as **TRANSIT nodes**. For example, in the figures accompanying this text, **fMAP System 4** is designated the master node while all the other systems are designated transit nodes. One ring port on the master node is designated to be the **Primary Port** (PP) and the other ring port is designated to be the **Secondary Port** (SP).

Note: Like the master node, each of the transit nodes is also configured with a PP port and SP port on the ring. However, the primary/secondary port distinction is ignored since it is a transit node.

Note: The initial description of the EPSR configuration uses a standard (rather than UFO) VLAN so that the Master Node can be any node. When UFO VLANs are configured, the Master Node must be the upstream node.

When the ring is operating normally, the master node blocks its SP port for all non-control traffic (data carried over the protected vlans) belonging to the EPSR domain, preventing a loop on the ring. The layer 2 Ethernet switching and learning mechanisms operate normally on each of the nodes in the ring. However, the control vlan traffic is not blocked at the SP port and is allowed to flow through. This does not pose a problem, because the control messages originate either at a master node or transit node but always terminate at the master node.

The ring network when stabilized collapses to what is shown in [Figure 13-5](#) below:




FIGURE 13-5 An EPSR Ring Topology (Standard VLAN)

When the master node detects a physical link break in the ring, it unblocks its SP port and allows the flow of non-control traffic through EPSR domain. This mode continues until the master node determines that the break in the ring has been restored; at which point, it goes back to its normal operating procedure.

13.4.4.2 EPSR Configuration Data

To implement EPSR, the user is required to configure the TAPS protocol to support the fault detection and recovery in the network. Configuration data is as follows:

- **HelloTime** – The rate at which the TAPS protocol Health control message is sent by the master node for this EPSR domain.
- **FailOverTime** – Time for which the master node waits before declaring that it has detected a break in the ring for this EPSR domain.

- **Ringports** – The two ports that are members of the EPSR domain.
- **Protected Vlan(s)** – The vlan VIDs which require protection on the EPSR domain.
- **Control Vlan** – The unique vlan VID which will be used as the control vlan for the EPSR domain.

Configuration requirements and commands will be covered in detail later in this section.

Note: In an EPSR configuration, if both fibers are cut, and one of the fibers is restored, the fMAP device will recover and begin processing traffic. This applies to fMAP devices whether the device is a Master Node or other node on the ring. Note that this functionality only works if the device is an fMAP.

13.4.4.3 Allied Telesis Automatic Protection Switching (TAPS) Protocol

The TAPS protocol is used to provide the EPSR functionality in fMAP layer 2 systems. TAPS protocol control messages are transported around the ring network for an EPSR domain via its control vlan. The messages can originate at the master node or at the transit node; however, they will **always** terminate at the master node. These messages are used to provide fast protection switching, for a given EPSR domain, in the layer 2 systems interconnected to form the Ethernet ring.

A fMAP system with EPSR implemented can be part of more than one ring network. As a result, there can be more than one EPSR domain on such a system, one for each of the EPSR protected rings of which it is a member. Note that there can also be more than one EPSR domain running in a system when it is part of only one ring network with each domain assigned its own set of protected vlans and a control vlan managing the bandwidth available in the ring.

13.4.4.3.1 TAPS Protocol Procedures

The master node and the transit nodes that make up the Ethernet ring use the Fault Detection and Fault Restoration procedures provided by the TAPS protocol to maintain the continuous flow of the non-control traffic in the ring.

1. Fault Detection Procedure

- Master node Polling Fault Detection Procedure.
- Transit Node Unsolicited Fault Message Fault Detection Procedure.

Two kinds of fault detection procedures are defined to detect a fault. The polling procedure is the fail-safe mechanism executing in the master node in case the unsolicited fault message procedure in a transit node fails.

2. Master Node and Transit Node Fault Detected Correction Procedure.

3. Fault Restoration Procedure

- Master Node Restoration Procedure.
- Transit Node Restoration Procedure.

Each of the above procedures will be discussed next as it applies to the master and/or transit nodes.

13.4.4.3.2 Master Node Polling Fault Detection Procedure

The master node uses the polling procedure as a fail-safe mechanism to detect a fault in the ring. It does this by sending an EPSR **health** control message via its PP port (only) every HELLOTIME seconds (which is a configured value as seen earlier). Under normal conditions, when there is no fault in the ring, this health message will make it across the network and will be received by the master node over its SP port. However, if there is a fault anywhere in the network, the health message will not be received by the master node over its SP port. To detect this condition, the master node starts a failover timer (using the configured FAILOVERTIME) every time it sends the health message. If the health message is not received by the master node over its SP port before the failover timer expires, then it declares a fault in the ring and takes appropriate measure.

Note: Because of the fact that messages could get lost in the network, Allied Telesis recommends that the FAILOVERTIME configured value be at least twice the value of the HELLOTIME configured value.

13.4.4.3.3 Transit Node Unsolicited Fault Message Fault Detection Procedure

Unlike the polling procedure described above, where the burden is upon the master node to eventually detect a fault in the ring in a fail-safe manner, this procedure is used by the transit node to detect a fault on its attached ring port and immediately notify the master node of the fault. This is accomplished by sending an EPSR Links-Down control message over a functioning link. A fault link spans two nodes; therefore, both of the transit nodes that detect the fault send the EPSR Links-Down control message to the master node. When this occurs, the transit node(s) in question alter the state of the EPSR domain from Links-Up state to Links-Down and maintains this state until the transit node Fault Restoration procedures are executed. Also, the state of the faulty port is set to Blocked. However, the state of the functioning ring port is maintained at Forwarding.

13.4.4.3.4 Master Node and Transit Node Fault Detected Correction Procedure

When the master node detects a fault in the ring using either of the above described procedures, it takes the following actions:

- Declares the EPSR domain to be in a failed state (from the complete state the EPSR domain was in before the fault was detected).
- Unblocks its SP port for the non-control traffic for this EPSR domain.
- Flushes its own forwarding database (FDB) for just the two ring ports.
- Sends an EPSR Ring-Down-Flush-FDB control message to all the transit nodes via both its PP port and SP port.

As the EPSR domain non-control traffic starts flowing, each of the nodes (both master and transit) then re-learn the layer 2 addresses on the flushed ring ports again to reflect the newly collapsed network topology. The master node continues to follow the Master Node Polling Fault Detection Procedure and as long as the fault is still present in the ring, the EPSR domain will continue to remain in the failed state. This newly constructed network

topology exists until the fault in the ring is corrected; then the fault restoration procedures take over and restore the ring to its original normally operating state.

The EPSR stabilized topology under normal operating conditions is shown in Figure 13-5. For a link fault detected between fMAP System A and fMAP System B, Figure 13-6 shows the new EPSR stabilized topology after the fault detection and corrections procedures have been executed.




FIGURE 13-6 An EPSR stabilized Network after Ring Fault (Standard VLAN)

13.4.4.4 Master Node Fault Restoration Procedure

When the fault in the ring between fMAP System A and fMAP System B is fixed, the polling EPSR Health control message that was being sent by the master node over its PP port (sent even when the fault is present) is now

received over its SP port. The master node then takes the following actions to restore the ring back from that shown in [Figure 13-6](#) to its original normally operating state.

- Declares the EPSR domain to be in a **complete** state from the **failed** state it was in before the fault was corrected.
- Blocks its SP port for non-control VLAN traffic for this EPSR domain.
- Flushes its own forwarding database (FDB) for the two ring ports.
- Sends an EPSR Ring-Up-Flush-FDB control message to all the transit nodes via both its PP port and SP port.

As the EPSR domain non-control traffic starts flowing again, all nodes (both master and transit) then re-learn the layer 2 addresses again to reflect the newly collapsed network topology. The master node continues to follow the Polling Fault Detection procedure and, since the fault is no longer present, the EPSR domain continues to remain in the **complete** state.

The network topology, restored to its normally operating state, continues to operate until a fault is detected, when, again, the above mentioned procedures are re-executed. EPSR maintains a continuous, uninterrupted operation of the user's network.

13.4.4.5 Transit Node Fault Restoration Procedure

The transit node(s) that span the faulty link will delay the starting of the flow of non-control traffic over the link once the fault has been fixed and the link restored. The reason for this delay is to prevent the master node from viewing the fixed link as a **loop** in the network. The loop is caused because the transit node has corrected the fault for the domain before the master node detects that the fault is restored and blocks its SP port for the domain non-control traffic. In order to avoid this situation, the transit node(s), after detecting that the broken link has been restored, follow these fault restoration procedures:

- Ensure that the protected vlans are still in a blocked state for the repaired port. The state of the restored port was set as **blocked** earlier when the state of the domain went from **Links-Up** to **Links-Down**.
- Change the state of the EPSR domain from **Links-Down** to **Pre-forwarding**.
- Wait for the EPSR Ring-Up-Flush-FDB control message from the master node. This is the trigger that ensures that the master node has detected the restoration of the fault in the ring, flushed its FDB, and blocked its SP port for the domain non-control traffic.
- Flush its FDB, for both the ring ports, upon receiving the above trigger message from the master node.
- Change the state of the EPSR domain from **Pre-Forwarding** to **Links-Up** when the flow of the domain non-control traffic can start to flow again ensuring that there is no loop present in the ring. At this point, the state of the port is set to **Forwarding**.

Note: The user should be aware that this restoration procedure can take several seconds, since the systems must check to ensure that the above conditions are met. This ensures that traffic is not affected.

13.4.4.6 Static vs. Dynamic Determination of Upstream Port and Downstream Port

As seen from above, the TAPS protocol supports the implementation of EPSR functionality in fMAP systems that are part of a ring daisy chain network topology. EPSR provides the mechanism to ensure that there is no loop in the ring network similar to what the STP protocol does in the fMAP system for a daisy chain network topology. In addition, the STP protocol determines the root port in the system (an **upstream port**) and the non-root (designated) port (**downstream port**).

With the introduction of EPSR, instead of STP providing layer 2 redundancy for a ring network, a different mechanism has to be used to determine which ring port is an upstream port and which ring port is a downstream port at any given time.

*Note: To use EPSR functionality for the determination of the upstream and downstream port in the nodes that make up the ring, the node which has the link to the network **must** be the master node.*

The example network shown below in [Figure 13-7](#) has EPSR functionality implemented with fMAP system A configured to be the master node with a PP and SP port. The remainder of the systems are configured as transit nodes.




FIGURE 13-7 Initial forwarding configuration of ring ports in an EPSR network

The **master** node, fMAP System 1, has its port pointing towards the network configured as upstream with the two ring ports configured as downstream. This is provisioned using the SET VLAN command with the FORWARDING parameter. This is similar to the current implementation of configuring the layer 2 aggregating switch the same way when it is the root bridge according to the STP protocol. In the EPSR ring using the TAPS protocol, the master node configuration is conceptually the equivalent of it being a root bridge for this network.

In each of the transit nodes, both the ring ports are configured as ring. This is also provisioned using the SET VLAN command with the FORWARDING=RING parameter. Again, this is very similar to configuring the system ports in a daisy chain ring. However, unlike an STP implementation where a port change event is used to configure ports, for EPSR, the receipt of a TAPS protocol message is used. The ring port uses a Health message with the state of Complete received by the transit switch to configure the upstream port with the other ring port configured to be the downstream port.

13.4.4.7 Interconnected EPSR Ring Networks

The discussion above had an underlying assumption that there is an Ethernet ring access network consisting of fMAP systems that are physically connected to form a ring using EPSR functionality to provide redundancy at the layer 2 level. The master node in this EPSR-enabled ring network is the one which is considered to be the layer 2 aggregating switch with an uplink to the core. However, there could be a case where the access network consists of multiple rings, interconnected to form a more complex access network with uplinks to the core network. An example of such a network is shown below in [Figure 13-8](#).

One of the systems, fMAP System C, is common to Ethernet access ring Networks 1 and 2. As indicated, this system is a **master node** for ring network 1 and a **transit node** for the ring network 2. System 1 in ring network 2 is the **master node** for that ring. The remainder of the systems in both networks are considered to be **transit nodes**.

Also shown in [Figure 13-8](#), is the initial configuration of the **FORWARDING** parameter of the SET VLAN command for the ports in all the systems in both networks. Note that the master node ports are manually configured as **upstream** or **downstream** and remain so, whereas the transit node ports are configured as ring- eventually stabilizing as **upstream** or **downstream** as the topology dictates.




FIGURE 13-8 Configuration of Ring Ports in Interconnected EPSR Ring Networks

13.4.4.8 EPSR Engineering considerations

A maximum of 10 domains, with 50 protected VLANs per domain, can be provisioned for EPSR.

13.4.4.9 EPSR Provisioning Guidelines

When provisioning EPSR on a ring, the user should avoid creating any network loops. There are many ways to configure a network in a ring topology without producing a loop within the network. Some of approaches can involve disabling one of the ring's links while configuring the EPSR domains and the protected VLANs; however, disabling one of the ring's links may not be an acceptable approach because doing so removes the ring's redundancy while provisioning.

Note: When provisioning a system, the user can fill in the attributes "name" and "location" to identify the system. These are for administrative purposes and do not affect the working of EPSR.

13.4.5 EPSR - Interoperability

fMAP EPSR complies with RFC3619 - Extreme Networks™ Ethernet Automatic Protection Switching (EAPS)¹ Version 1 whether as the Master or Transit node and can interoperate with Extreme Networks' switches that also supports EAPS, when the Extreme Network's switch has been provisioned to be compliant with RFC3619.

When Extreme is the master, ensure that the following configuration is completed on the Extreme EAPS master switch:

```
config eaps name failtime expiry-action actionvalue
```

name is the domain name the Extreme is the master for and the actionvalue specifies the action taken by the master when the failover timer expires. actionvalue can be either **open-secondary-port** or **send-alert**.

In order to ensure that the Extreme Networks' switch interoperate strictly with RFC 3619, use the following command for configuring domains where Extreme is the master:

```
config eaps name failtime expiry-action open-secondary-port
```



DO NOT set the **expiry-action** in the above command to **send-alert**.

If **expiry-action** is set to **send-alert**, the Extreme master will not be fully compliant with RFC3619, which can cause a segment of the ring to be isolated under certain failure conditions.

13.4.6 EPSR Engineering Rules

For each system, up to ten domains can be defined.

1. Extreme Networks is a registered trademark of Extreme Networks, Inc. All Rights Reserved

For the vlan tunneling feature to work, each customer VLAN must be configured in EPSR to prevent a loop. This is only applicable to the local node, but is necessary because the GE3 strips the outer tag so EPSR sees only the inner tag.

This also allows customers who do not use multiple domains to have the limit for VLANs per domain increased to 512.

For the VLAN tunneling implementation, assigning VLANs one at a time could be too cumbersome, so the VLANs can be added as a list.

Commands used to configure EPSR are included in this table:.

TABLE 13-4 EPSR Configuration commands

Object	Verb	Syntax	Description
VLAN	ADD	<pre> ADD VLAN={ vlaname vid } INTERFACE={ type:id-range id-range ifname-list ALL } [FRAME={ UNTAGGED TAGGED }] [TRANSLATE={ 1..4094 }] [FORWARDING={ UPSTREAM DOWNSTREAM RING PROTECTIONLINK }] </pre>	The ADD VLAN command adds interfaces to the specified layer-2 virtual network. Parameters are provided for the provisioning of interfaces as either UPSTREAM, DOWNSTREAM, RING, or PROTECTIONLINK.
VLAN	SET	<pre> SET VLAN={ vlaname vid } INTERFACE={ type:id-range id-range ifname-list ALL } [FRAME={ UNTAGGED TAGGED }] [TRANSLATE={ 1..4094 NONE }] [FORWARDING={ UPSTREAM DOWNSTREAM RING PROTECTIONLINK }] </pre>	The SET VLAN command toggles the status of interfaces in a Virtual LAN (VLAN) between tagged and untagged. Parameters are provided for the provisioning of interfaces as either UPSTREAM, DOWNSTREAM, RING, or PROTECTIONLINK.

TABLE 13-4 EPSR Configuration commands

Object	Verb	Syntax	Description
EPSR INTERFACE	ADD	<pre> ADD EPSR=epsrdomain INTERFACE={ type:id-range id-range ifname-list } [TYPE={ PRIMARY SECONDARY }] </pre>	Adds an EPSR interface.
EPSR VLAN	ADD	<pre> ADD EPSR=epsrdomain VLAN={ vlanname vid } [TYPE={ CONTROL DATA }] </pre>	Adds an EPSR VLAN
EPSR	CREATE	<pre> CREATE EPSR=epsrdomain { TRANSIT MASTER [HELLOTIME=value] [FAILOVERTIME=value] [RINGFLAPTIME=value] } </pre>	Create an EPSR domain.
EPSR INTERFACE	DELETE	<pre> DELETE EPSR={ epsrdomain-list ALL } INTERFACE={ type:id-range id-range ifname-list ALL } </pre>	Delete an EPSR domain or interface.
EPSR VLAN	DELETE	<pre> DELETE EPSR={ epsrdomain-list ALL } VLAN={ vlanname vid ALL } </pre>	Delete an EPSR domain or vlan.

TABLE 13-4 EPSR Configuration commands

Object	Verb	Syntax	Description
EPSR	DESTROY	DESTROY EPSR={ epsrdomain-list ALL }	Destroy an EPSR domain.
EPSR	DISABLE	DISABLE EPSR={ epsrdomain-list ALL }	Disable an EPSR domain.
EPSR	ENABLE	ENABLE EPSR={ epsrdomain-list ALL }	Enable an EPSR domain.
EPSR	SETDEFAULTS	SETDEFAULTS EPSR={ epsrdomain-list ALL } MASTER [HELLOTIME] [FAILOVERTIME] [RINGFLAPTIME]	Set the defaults for an EPSR domain.
EPSR INTERFACE	SET	SET EPSR=epsrdomain INTERFACE={ type:id-range id-range ifname-list ALL } [TYPE={ PRIMARY SECONDARY }]	Sets an interface as an EPSR domain.
EPSR MASTER	SET	SET EPSR={ epsrdomain-list ALL } MASTER [HELLOTIME=value] [FAILOVERTIME=value] [RINGFLAPTIME=value]	Set parameters for an EPSR domain.
EPSR	SHOW	SHOW EPSR [={ epsrdomain-list ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }]	Displays an EPSR domain.

13.5 Upstream Control Protocol (UCP)

13.5.1 Overview

Section 13.4 explains how network topology protocols used in 5.0 ensure that in these topologies loops are not created, and when there is a failure there is a reconvergence so that multicast streams continue to have a pathway through the network without loops. It also explained the importance of Upstream Forwarding Only (UFO) VLANs, since these protocols rely on UFO VLANs being configured, and since UFO VLANs ensure that traffic from subscriber (downstream) interfaces is only permitted to be forwarded out of a designated upstream port into the network (usually to a multicast router).

In release 5.0, therefore, the topology features rely on UFO VLANs being configured so that convergence can occur correctly. However, this reliance causes restrictions on network topology and protocol implementation.

An example of this restriction is using the master node of an EPSR ring to determine upstream. This requires a separate EPSR domain and master node for upstream port (exit) out of the network. Moreover the master node must also be at the top of the network. If the master node fails, or if both downstream links into the network fail, there is no other exit from the network.

A protocol therefore is needed that is independent of network topology and other protocols and can therefore operate in all network environments. The protocol developed is called **Upstream Control Protocol (UCP)**.

UCP is a generic protocol used by fMAP devices so that it can inform other devices in the network that it is the “upstream node” for a UFO VLAN. Moreover, using UCP protocol messages, the non-upstream nodes for the UFO VLAN can dynamically determine their upstream interfaces. This occurs independently of the topology feature being used; therefore, UCP can be used by itself as well as with STP and EPSR.

To summarize, the user must therefore understand first how UCP works and then how it can be used in the network:

- By itself
- With STP (and MSTP, explained in 13.7)
- With EPSR

Moreover, for networks that use STP and EPSR in 5.0, the UFO VLANs had a Forwarding Mode that could be static or dynamic, meaning the interfaces for the UFO VLAN could be data filled so that the upstream or downstream direction of the interface was set manually or by the topology protocol.

With the introduction of the UCP protocol, there is a new value that can be set for the Forwarding Mode of a VLAN interface, **UCP**; with **UCP**, the direction of the interface is dynamically determined by the UCP protocol. Moreover, for existing 5.0 UFO VLANs, the Forwarding Mode may be change.

Note: For these existing 5.0 topologies, UFO VLANs on some interfaces will have the FORWARDING values for UFO VLANs change automatically during the upgrade from 5.0. Moreover, the user can reconfigure these VLANs prior to upgrade, which impact how they are

changed after upgrade. Users must therefore read this section to understand the effects of a system upgrade and whether to make any pre-upgrade changes.

13.5.2 UCP Protocol / Operation

To understand UPC operation, the user should first consider UCP enabled on a set of connected devices that does not use a topology feature; a set of interfaces are configured with a UFO VLAN, and one of the interfaces is on the upstream device and faces into the larger network.

When configuring the UFO VLAN over the interfaces to these devices, the user has the following command in release 6.0:

```
ADD (or SET)
VLAN={ vl aname
      | vi d
      }
INTERFACE={ type: i d-range
            | i d-range
            | i fname-l i st
            | ALL
            }
[ FRAME={ UNTAGGED
          | TAGGED
          } ]
[ TRANSLATE={ 1 . . 4094
             } ]
[ FORWARDI NG={ PRI MARYUPSTREAM
                | SECONDA RYUPSTREAM
                | DOWNSTREAM
                | STP
                | UCP
                | EPSR
                } ]
```

When the user sets the FORWARDING parameter to PRIMARYUPSTREAM, then for this UFO VLAN, this device is the upstream device and this interface is the upstream interface.

Moreover, the user can have a second device that also is an upstream device with an upstream interface. The user can therefore set this second interface as SECONDARYUPSTREAM, which will act as a backup if there is a failure of the primary interface device. This will be explained in detail later.

As the other interfaces are associated with the UFO VLAN, they are by default set to UCP, meaning they dynamically determine their direction during topology changes.

Once all of the interfaces are configured, they can exchange the two UCP messages:

- The Upstream Port Notification Message sent by the upstream node
- The Upstream Port Topology Change Message sent by the devices spanning a failed link (or adjacent to a failed device) when the failure occurs and when it subsequently recovers.

These are explained below.

13.5.2.1 Upstream Port Notification Message

The upstream node periodically formats this message for each of the UFO VLANs that qualifies and sends them over the UCP-enabled interfaces. This message is received by the other devices that make up this UFO VLAN configuration

Table 13-5 shows the logic of how this message is processed:.

TABLE 13-5 Processing of Upstream Port Notification Message Sent by Upstream Device

Process	If Check Condition, Outcome
A classifier rule intercepts this message based upon the layer 2 destination address value and sends it to the cpu only.	
The UCP protocol task upon receiving this message will check to see if the tagged vlan in the message is configured to be a ufo vlan.	If not, the message is discarded. If the received tagged vlan in the message is configured to be a ufo vlan, continue
Check is made to see if the port on which it was received is a tagged member of the ufo vlan	If not, the message is discarded If tagged member, continue.
Check if message is received on UCP enabled interface	If interface not configured to be UCP enabled, message discarded. If interface UCP enabled, continue.
The layer 2 source address value of the active upstream node which originally sent the message is stored against the received ufo vlan vid value. This is done so that this non-upstream node knows which node is the upstream node for this ufo vlan in the network topology in case it has to send the upstream port topology change message during a network link fail or link recovery condition.	
Check if other UCP enabled ports for this ufo vlan are configured	If no other UCP-enabled interfaces for this UFO VLAN configure, message is discarded. If other UCP-enabled ports for this UFO VLAN are configured, continue.
Received message sent over all other UCP-enabled network module interfaces that have been configured for this UFO VLAN.	If no other UCP enabled ports for this ufo vlan are configured or if any are configured to be UCP enabled but are not tagged members of the ufo vlan then the message is just discarded.

13.5.2.2 Topology Change Message

This message is used for fault and recovery scenarios:

- **Link Fault** - When a link fault occurs, each of the nodes spanning the faulty link send the upstream port topology change message for each of the ufo vlans towards its upstream node. This topology change message will indicate to the upstream node that this message is being sent as a result of a link failure in the network.
- **Link Recovery** - When a link fault gets corrected and the recovery is detected then each of the nodes spanning this recovered link send the upstream port topology change message for each of the ufo vlans towards its upstream node. This topology change message will indicate to the upstream node that this message is being sent as a result of a link recovery in the network.

The non-upstream nodes that originate this topology change message for each of the ufo vlans send the message over all the applicable UFO-enabled network ports and the ufo vlan may or may not have an upstream port (based upon where the fault is) until the node hears from the upstream node again when it receives this topology change message.

Each of the other nodes between the originating node and the upstream node receive the Topology Change Message and take the actions listed in the following table.

TABLE 13-6 Processing of Topology Change Message by Nodes Between Originating and Upstream Node

Process	If Check Condition, Outcome
A classifier rule intercepts this message based upon the layer 2 destination address value and sends it to the cpu.	
The UCP protocol task upon receiving this message will check to see if the tagged vlan in the message is configured to be a ufo vlan.	If not, the message is discarded. If the received tagged vlan in the message is configured to be a ufo vlan, continue
Check is made to see if the port on which it was received is a tagged member of the ufo vlan	If not, the message is discarded If tagged member, continue.
Check if message is received on UCP enabled interface	If interface not configured to be UCP enabled, message discarded. If interface UCP enabled, continue.
The message is sent as is over all the other applicable UCP enabled WIF ports. The ufo vlan may or may not have an upstream port (based upon where the fault is) until the node hears from the upstream node again when it receives this topology change message	
Topology Change Message received by upstream node	
A classifier rule intercepts this message based upon the layer 2 destination address value and sends it to the cpu.	
The UCP protocol task upon receiving this message will check to see if the tagged vlan in the message is configured to be a ufo vlan	If not, the message is discarded. If the received tagged vlan in the message is configured to be a ufo vlan, continue

TABLE 13-6 Processing of Topology Change Message by Nodes Between Originating and Upstream Node (Continued)

Process	If Check Condition, Outcome
Check is made to see if the port on which it was received is a tagged member of the ufo vlan	If not, the message is discarded If tagged member, continue
Check if port on which message received is configured as UFO-enabled port	If not UFO-enabled, message is discarded. If UFO-enabled, continue.
Check if this node is an upstream node for the UFO VLAN.	If this is not an upstream node, message is discarded If this is an upstream node, continue
The upstream node does not wait for its periodic timer expiry to send the upstream port notification messages over its allowed ring ports. It sends the upstream port notification message over all the allowed ring ports in rapid succession a few times after which it settles down to sending the notification message using its periodic timer	
The non-upstream nodes receive the upstream port notification message and process them as described in Table 13-5 .	

13.5.2.3 UCP Redundancy (Different Nodes)

With UCP redundancy, a standby upstream interface can be configured. In most cases these interfaces for the UFO VLAN will be on separate nodes, so redundancy is provided at the node level.

1. The active node does not actively source hello type messages. It is responsible for responding to messages received from standby nodes. This is done to reduce chatter and would be redundant for networks which do not have a standby node.
2. The active node response contains the state of the upstream port - Up or Down.
3. The active node response is to flood out all UCP enabled network ports. This is to provide information to all nodes in the network so that each can maintain an active and secondary “topology”.
4. The active node must source an unsolicited response if its configured upstream port changes state. This allows the active node to “monitor” its upstream port and to provide rapid failover and recovery characteristics.
5. The standby node is responsible for sourcing hello messages on a rapid periodic basis. These hello messages must be sent for each VLAN for which it is a standby for. These hello messages are flooded out each UCP enabled network port.
6. The standby node must assume that the active node is no longer in service if it fails to receive 2 or more hello responses.

13.5.2.4 UCP Redundancy (Same Node)

If the active and standby interfaces are on the same node, all messages are still flooded over the UCP-enabled network ports. However, since both active and standby are on the same node, the failover and message can be solely determined by port state. This operates as “protection link” did prior to 6.0. The only exception is that UCP messages (hellos plus upstream notification) will continue to be sourced.

13.5.3 UCP with EPSR Topology Feature in Release 6.0

13.5.3.1 Initial Topology

Although UCP can act as a standalone protocol when the topology control is further up in the network, it can interact with the STP and EPSR features; for example, in a ring network, the EPSR feature ensures there is no loop created over the protected domain, while the UCP is used in the non-upstream nodes to determine the upstream interface for the (protected) UFO VLANs.

[Figure 13-9](#) shows the resulting topology. Switch A is the upstream node for the UFO VLAN (V_80) in the domain, and so sends out the Upstream Port Notification message (see [13.5.2.1](#)) for each of the UFO VLANs over its two ring ports. This message is received by nodes B and C on one side of the ring and node D on the other. Note that switch C does not receive this message from Node D because the messaging is over the protected VLAN and this is blocked by EPSR.

The message when received at each node is intercepted by the classifier and sent to the CPU. If all ingress checks pass (see [13.5.2.1](#)), each node stores the VLAN ID (80) along with the MAC address of the upstream node (Node A). The message is then forwarded over the other ring port towards the next node in the ring network. Finally, the message is discarded at node C because the UFO VLAN is logically blocked.

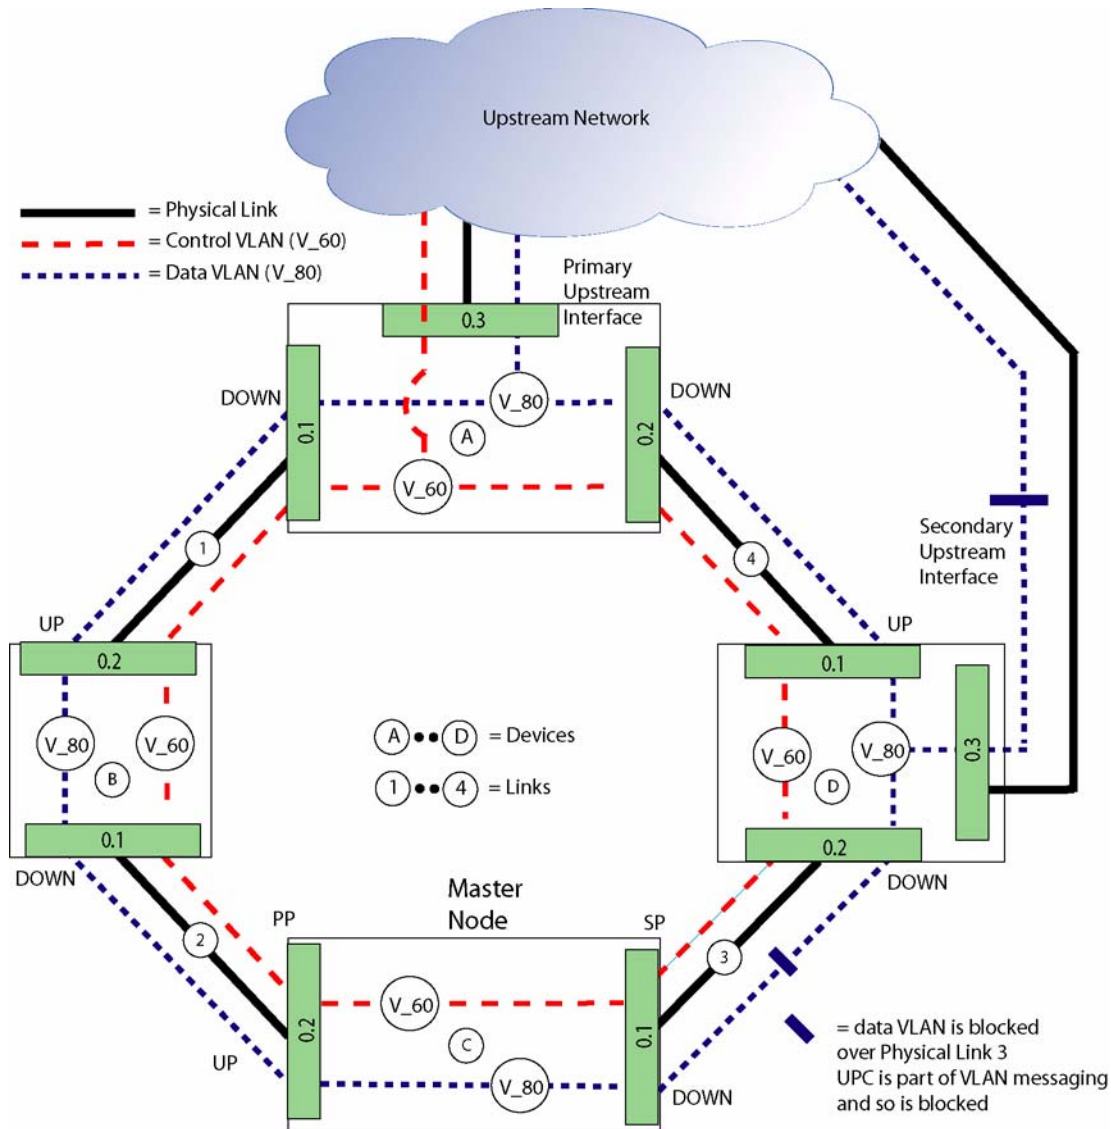


FIGURE 13-9 EPSR Topology with UCP

13.5.3.2 Fault Message and Recovery - Physical Link

Assume a fault occurs on link 2. The EPSR protocol reacts and takes steps to change the topology so that no node is isolated and no new loops are formed. The UCP protocol makes sure that the direction of the ports (upstream, downstream) are set correctly.

On node C, the UCP protocol sends the Upstream Port Topology Change (13.5.2.2) message for the UFO VLAN. This message is received and forwarded to the next node until the node that receives the message is the upstream node (A). Therefore, nodes D and A would receive the message.

On Node B, the UCP protocol would also send the Upstream Port Topology Change message for the UFO VLAN. This message is also received and forwarded to the next node until the node that receives the message is the upstream node (A). Therefore, nodes B and A would receive the message.

Node A will then send an unsolicited (non-timer) “Upstream Port Notify” message for the UFO VLAN(s) over both its ring ports a few times before settling back to its normal (timer) sending procedure. The other nodes receive this message, process it as described earlier, and the result is a reconverged topology in which the upstream/downstream direction of the interfaces are configured correctly.

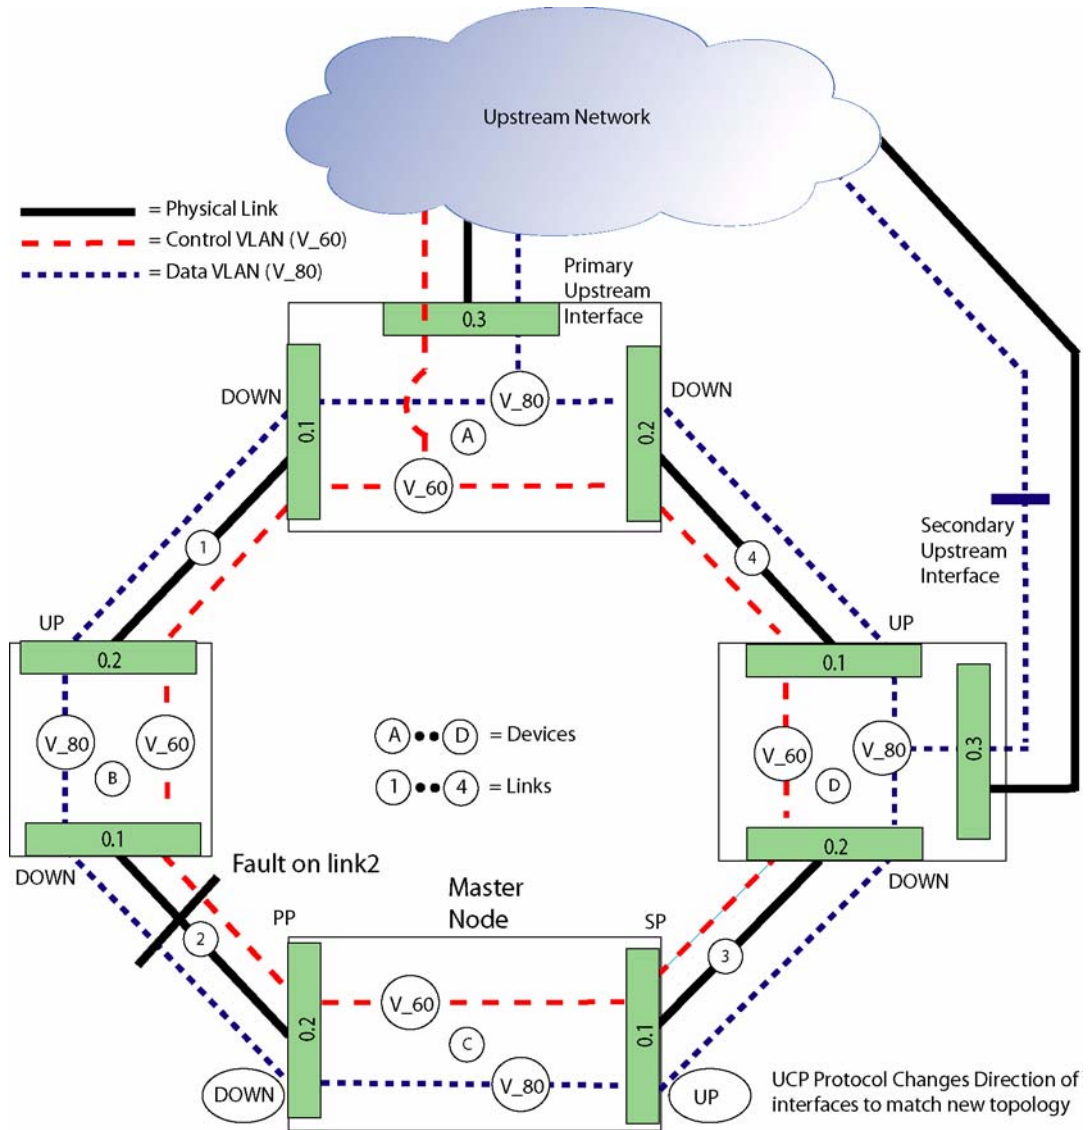


FIGURE 13-10 UCP and EPSR - Fault Recovery

13.5.4 UCP with STP Topology Feature in Release 6.0

Many of the concepts about the interaction of UCP with EPSR also apply to UCP with STP; the STP protocol ensures there are no loops in the converged (reconverged) topology, while UCP ensures that the UFO VLAN interfaces are set correctly for that topology. Since STP is a port based topology (as opposed to EPSR which is a

VLAN based topology), the STP will perform blocking on the port and therefore block **all** the VLANs on that port. The user should therefore ensure that no VLANs are isolated when STP changes the topology for the relevant nodes. (Refer to [13.2.4](#) for details and an example.)

13.5.5 Datafill for UFO VLANs when Upgrade to Release 6.0

Section [13.4](#) explained the topology features in release 5.0 and how the UFO VLANs configured on the interfaces can be data filled so that their direction (Upstream or Downstream) can be set dynamically. The values that could be used depended on the topology feature used and the place of the interface in the topology. These values in 5.0 were as follows:

- **UPSTREAM** - This was a static setting; the interface would always be upstream. (If the upstream node was the Root bridge, then the UFO VLAN over the interface or port that faced into the upstream network would be set as UPSTREAM.)
- **DOWNSTREAM** - This was a static setting; the interface would always be downstream. (If the upstream node was the Root bridge, then the UFO VLAN over this interface or port that faced into the other nodes of the topology would be set as DOWNSTREAM.)
- **RING** - This was a dynamic setting for both STP and EPSR; the direction could change depending on the current topology.
- **PROTECTIONLINK** - This was a dynamic setting for determining the upstream interface within a set of interfaces in a Layer 1 Protection Group, based on which interface was operationally “Up.”

With UCP there is now a separate protocol that can be used to dynamically change the direction of the UFO VLAN interface; moreover, there can be more than one provisioned upstream interface into the upstream network. As a result, the FORWARDING values in 6.0 can now be set as:

- **PRIMARYUPSTREAM** - This is a static setting; the interface is always upstream. (If there is only one upstream node, the UFO VLAN(s) associated with the upstream interface on that node are data filled as PRIMARYUPSTREAM.)
- **SECONDARYUPSTREAM** - This is a static setting; the interface becomes the upstream interface if the Primary Upstream fails. (If there is a second upstream node, the UFO VLAN(s) associated with the upstream interface on that node are data filled as SECONDARYUPSTREAM.)

Note: PRIMARYUPSTREAM and SECONDARYUPSTREAM are also used with connections that used PROTECTIONLINK in 5.0; this is explained in [13.11](#).

- **DOWNSTREAM** - This is a static setting; the interface is always downstream.
- **EPSR** - This is a dynamic setting for EPSR, and will change the direction depending on the current topology.
- **STP** - This is a dynamic setting for STP, and will change the direction depending on the current topology.
- **UCP** - This is a dynamic setting for UCP to determine the upstream interface. This setting is valid in any network topology (EPSR Ring, STP Network, Subtended Network) provided that the node that has the upstream connection supports UCP.

To incorporate the UCP protocol into topologies that may exist in release 5.0, the FORWARDING values will change during the upgrade from release 5.0 to release 6.0 as shown in [Table 13-7](#).

TABLE 13-7 Changes to FORWARDING Parameter when Upgrading form 5.0 to 6.0

STP	Upstream Node Interfaces	5.0	6.0
		UPSTREAM	PRIMARYUPSTREAM
		DOWNSTREAM	DOWNSTREAM
	RING	STP	
	Other Node Interfaces	RING ^a	STP
EPSR	Upstream Node Interfaces	UPSTREAM	PRIMARYUPSTREAM
		DOWNSTREAM	DOWNSTREAM
		RING	EPSR
	Other Node Interfaces	RING ^b	EPSR
Voice Process- ing Node	Layer 1 Protection Group	PROTECTIONLINK	PRIMARYUPSTREAM
		PROTECTIONLINK	SECONDARYUPSTREAM

a. If other node interfaces were not set to RING, the topology was static (linear).

b. If other node interfaces were not set to RING, the topology was static (linear).

*Note: The user can, if desired, change the FORWARDING parameter value from DOWNSTREAM to RING on the upstream node **before** the upgrade. This would make the 6.0 value for the interface automatically change, rather than remain DOWNSTREAM.*

13.5.6 Summary of UCP and Topology Engineering

Following are the general rules for FORWARDINGMODE values and their interaction with topologies:

- UCP
 - The upstream node must be an fMAP product.
 - With EPSR, any node can be the Master Node, so data can be engineered to have shorter paths (number of “hops.”)
- STP
 - The STP root node must be the Upstream Node
 - Values for all VLAN interfaces are STP except for Primary and Secondary Upstream.
- EPSR
 - Always used with the EPSR topology
 - The Master Node **must** be the upstream node.
 - Values for all VLAN interfaces are EPSR except for Primary and Secondary Upstream.

Note: Using EPSR usually has the fastest switchover (50 msec)

13.6 EPSR and (R)STP Interaction

13.6.1 Overview

In release 5.0, EPSR and (R)STP were mutually exclusive; a port that was used for an EPSR ring could not be used as part of an (R)STP subtending ring.

As explained in the previous subsections, the EPSR and (R)STP topologies conceptually do the same thing: to provide a protection scheme for the network while blocking certain links to prevent loops. The key difference between the two features, however, is that:

- EPSR requires the user to explicitly create the ring configuration and to decide where blocking will occur for the data VLAN(s).
- (R)STP configures where links are to be broken based on user provisioned values which are calculated to determine the lowest cost paths for data traffic. This is used to determine which paths allow data traffic and where links should be blocked to prevent loops.

In release 6.0, it is possible to coordinate these features (through the provisioning of key parameters) so that certain devices can take part in both EPSR and (R)STP. By data filling these parameters correctly, the blocking of links to remove loops is coordinated.

Note: One key aspect of providing this coordination is that provisioning must ensure that with ports that are part of both EPSR and (R)STP, their spanning tree states must be controlled by EPSR. This is explained in more detail later.

In release 6.0, the following configurations are supported:

- Connection of an (R)STP subnetwork to a **single** node in the EPSR ring (possible in 5.0).
- Connection of an (R)TP subnetwork to two **adjacent** nodes of the EPSR ring (new for 6.0)

Figure 13-11 shows these two configurations.

Note: As explained in 13.5, UCP is used with EPSR and STP to determine the upstream direction of the UFO VLAN interface configured on the port. In setting these interfaces as UCP, any node in the EPSR ring can be the Master Node.

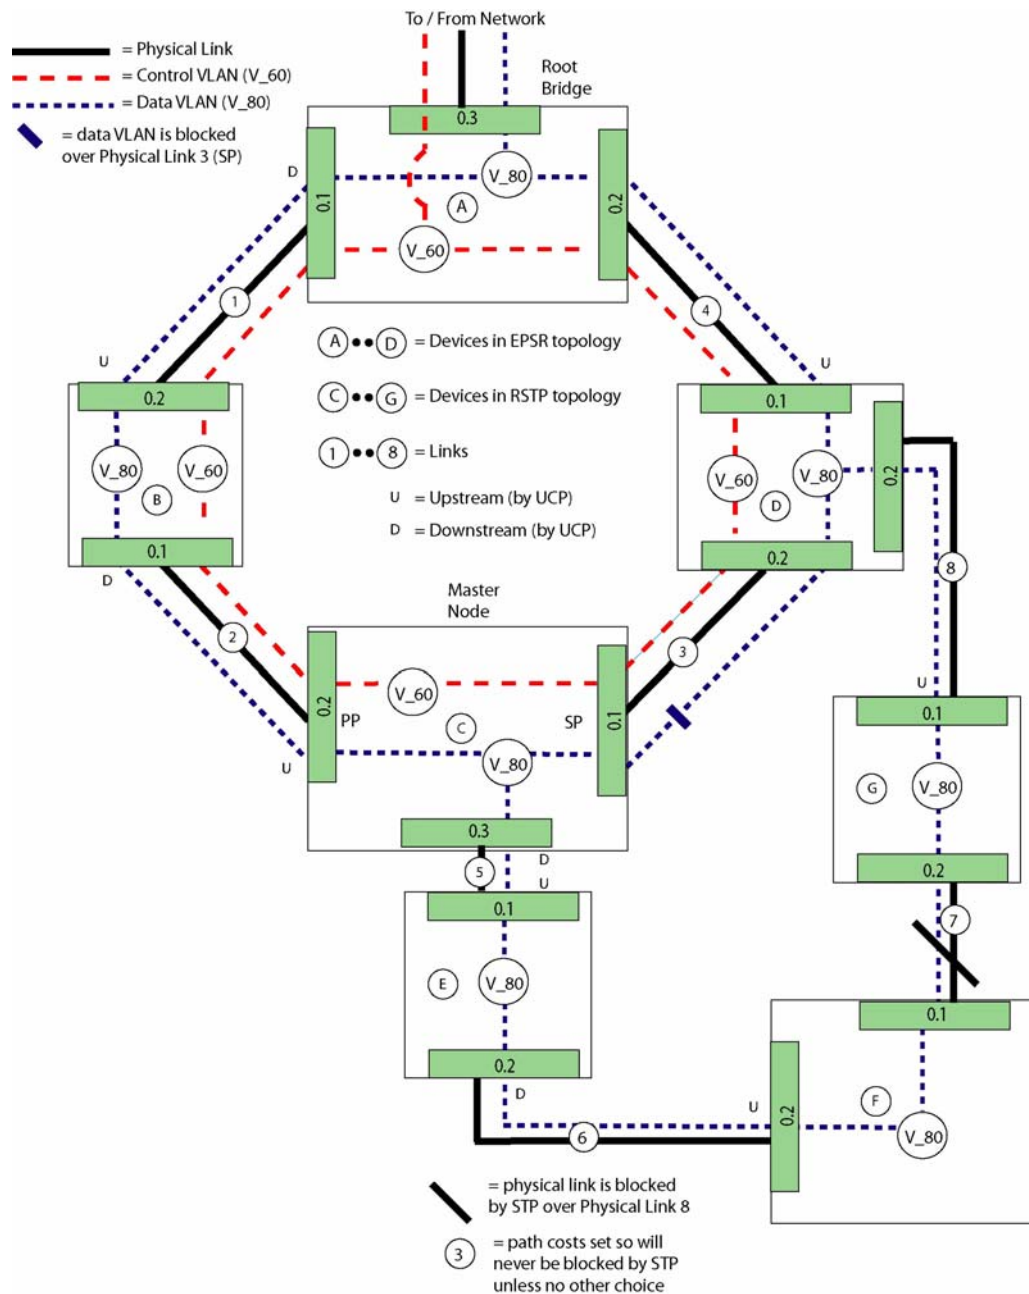


FIGURE 13-11 Possible EPSR/RSTP Configurations for 6.0

Note: In release 6.0, the Multiple Spanning Tree Protocol (MSTP) is also available (refer to [13.7](#)). In this release, however, implementing MSTP and EPSR features on the same system is prevented, and so are explained separately.

13.6.2 Summary of (R)STP

To understand the changes that have been made to make this shared configuration possible, there is first a summary of key aspects of the (R)STP feature. These are highlighted so that the user can then understand how these are modified to allow the (R)STP/EPSR interaction.

Note: For a complete description of the (R)STP feature, refer to [13.2](#).

13.6.2.1 Protocol Communication

The communication of STP/RSTP information to other bridges via the exchange of messages known as Configuration Bridge Protocol Data Units (BPDUs), Topology Change Notification (TCN) BPDUs, and RST BPDUs.

For the feature interaction, these continue to take place as part of the STP/RSTP processing on the fMAP nodes which are configured to run EPSR and STP/RSTP together, as well as the nodes which are only running STP/RSTP.

13.6.2.2 Convergence

The process by which the bridges in the extended LAN come to agreement about the logical spanning tree topology is known as convergence. This process includes several key steps:

- The bridges elect a root bridge by exchanging hello messages to determine which switch has the lowest bridge ID.
- Each bridge uses information received from other bridges, along with its own port cost information, to compute its own root path cost. It forwards this root path cost information along to other bridges; eventually, the correct root path cost for every path in the extended LAN will be computed.
- Each bridge selects a root port for that bridge.
- The bridges elect a designated bridge for each physical LAN based on root path cost, bridge ID, and/ or port ID for the bridges connected to the LAN.
- Any port that is determined not to be a root port or a designated port is set to the blocking state by STP/RSTP.
- After STP/RSTP has determined that the topology is “loop free” (e.g., all handshake processing completed or expiration of forwarding delay timers), every root port and designated port will have been set to the forwarding state. Once this is done, traffic may flow over the extended LAN.

If a link or bridge fails, or the network topology otherwise changes due to user provisioning, the network starts the convergence process again to reach a new spanning tree topology.

13.6.2.3 Configuring Ports (Path Costs)

When configuring ports, the user must datafill calculated values for:

- Port Path Costs (the value for `PATHCOST`)
- Port participation

Port path costs are used in root path cost calculations, which are a factor in root port and designated bridge elections. The **default** port path costs are related to the bandwidth capacity of the ports; however, the default values may be changed by the user to reflect other factors (e.g. propagation delay, link quality, desired traffic level, etc.)

Note: Refer to 13.2.3.3 for more details on how these values are calculated.

If ports on a switch are members of an extended LAN or VLAN that does not require use of the spanning tree protocol (i.e., if the VLAN is administered such that no network loops could exist), then spanning tree protocol operations can be disabled for those ports.

If spanning tree protocol operations are disabled for a port, it may still pass bearer traffic to and from other ports, regardless of whether or not the spanning tree protocol is used for those other ports.

The key point in configuring path costs is that the path cost for a port is added to the value of the root path cost field from a configuration message that is received on the port to determine the **total cost** of the path to the root bridge through that port. As just explained, the default path cost values and the range of recommended path cost values depend on the port bandwidth, and will vary as the speed of the port varies.

However, setting the path cost to a larger value on a particular port is likely to **reduce** the traffic over the LAN connected to it. This may be appropriate if the LAN has lower bandwidth, or if there are reasons for limiting the traffic across it. This concept is critical in configuring the feature interaction, as explained next.

13.6.3 Configuring (R)STP for Interaction with EPSR

Based on the explanation above, following are the key concepts/parameters that must be understood for the feature interaction to function correctly:

- Protocol Communication (BPDUs)

For ports that are participating in both EPSR & STP/RSTP, when the STP/RSTP processing that is “attempting” to control the spanning tree states of alternate or backup ports indicates that those ports are in a “blocked” or “discarding” state, BPDUs will not be transmitted on those ports, even though in fact they are actually “forwarding” due to EPSR control.

There is also communication added between the EPSR protocol and STP/RSTP within the fMAP node to signal when a port has been unblocked as a result of the EPSR ring being restored to full service following recovery of a failed link. This event will be processed rather than port enable event by the STP/RSTP feature.

- Convergence (selection of root bridge)

The root bridge for the overall Spanning Tree for the network in this type of configuration must either be one of the EPSR ring nodes, or a bridge which is at a “higher level” in the network and connects directly via one of the nodes on the EPSR ring. In other words the root bridge can NOT be a node from one of the STP/RSTP sub-networks, nor can it be a “higher level” network node that only connects via a link to one of the STP/RSTP sub-networks.

- Configuring Ports - MODE value for combined EPSR and (R)STP

For ports that will have combined control by the EPSR and the (R)STP feature there is a new parameter value:

- Port Protocol Mode { REGULAR_XSTP , **COMBINED_EPSR_XSTP** }

If the port needs to be returned to “regular” mode by the STP/RSTP processing then mode should be reset to the default value (MODE=REGULAR_XSTP).

13.6.4 Port Costs

When a network is being setup to utilize an EPSR ring in conjunction with STP or RSTP sub-networks, the port paths costs of **all** the links involved will need to be reviewed and potentially modified by the user. At a minimum, the port path costs for all the “shared” links from the EPSR ring will need to be set artificially low (e.g., to a value of 1,2, or 3) to keep the STP/RSTP algorithm processing from attempting to block those links.

In addition, whenever STP sub-networks are in use, it may be necessary to raise the path costs of the links in each STP sub-network such that the combined cost for a traffic path through any one of those sub-networks can **not** be lower than the cost to traverse the EPSR ring.

Note: This will only be an issue in a scenario where a link (or bridge) on the EPSR ring has failed.

This restriction is a side effect of the low magnitude and limited range of path cost values used for STP. When RSTP is in use, the same general principal applies (i.e., RSTP sub-network path costs must be greater than path cost for EPSR ring), but due to the greater magnitude of path cost value utilized for RSTP by default, this becomes much easier to accomplish through provisioning.

This concept is built into the CLI; when the user changes the mode of the STP port to use the combined mode with the command:

```
SET STP INTERFACE = { interface -list|ALL} MODE = { REGULAR_XSTP |  
COMBINED_EPSR_XSTP }
```

and selects COMBINED_EPSR_XSTP , there is a warning that the associated port path costs for the link being shared by EPSR and (R)STP need to be modified to be a very low cost value. This is done to keep STP/RSTP from attempting to “Block” any of those shared links, except when there is no other choice to avoid a loop. When the user chooses (reverts to) the REGULAR_XSTP option, another warning message will be displayed to indicate that the port path costs of the associated link should be reverted back to a nominal value.

13.6.5 Example Configuration

Figure 13-12 shows a simple example of RSTP and EPSR interaction.

13.6.6 Command Set for EPSR/(R)STP

Commands to control the interaction are included in the command set for (R)STP, and so are included in Table 13-8

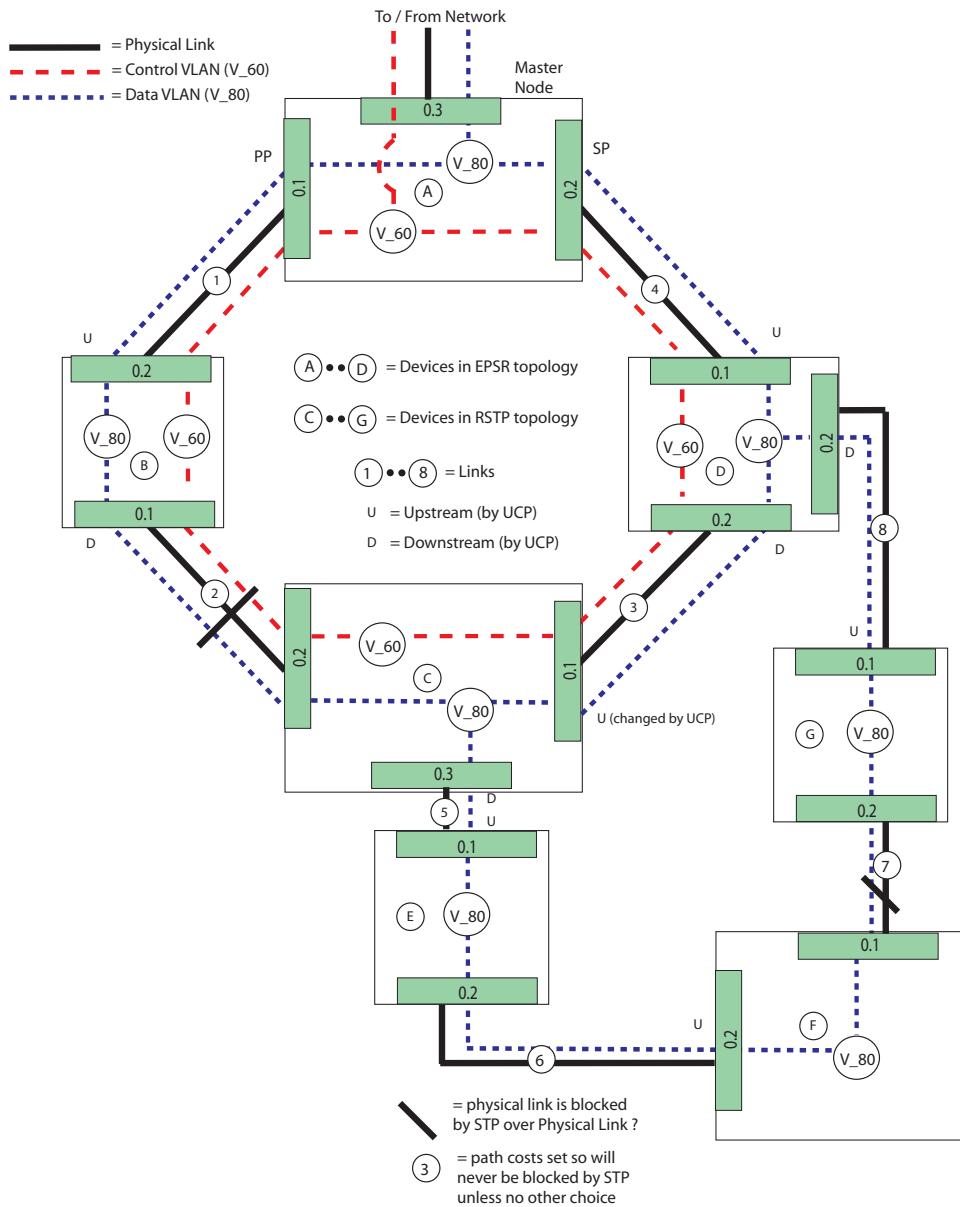


FIGURE 13-12 EPSR Ring Recovery with RSTP Interaction

13.7 MSTP

13.7.1 Overview

The previous subsections have described the STP and RSTP features and how they are configured. In release 6.0, it is possible to configure multiple (R)STP instances on a VLAN basis, so this is called the **Multiple Spanning Tree Protocol (MSTP)**.

With MSTP, separate spanning tree instances are created for VLANs (or groups of VLANs). Each of the separate instances elect root bridges, root ports, and designated bridges independently.

When an (R)STP network is configured and no VLANs are configured (only the default VLAN), each device and each port are considered part of the same extended LAN, and so all participate in the same convergence process. Therefore all devices and ports are part of a **single (R)STP instance**. As VLANs are added, they are still part of the single spanning tree instance.

Note: As VLANs are configured, the user must be careful to ensure that the physical (link) and virtual (VLAN) connections work together and do not lead to any disconnected VLANs (refer to 13.2.4).

This association of multiple VLANs with the one spanning tree is called a Common and Internal Spanning Tree, or **CIST**. Bridges configured within a CIST behave as a single spanning tree system automatically.

With MSTP, additional spanning tree instances can be created and associated with the VLANs defined on the device. These additional spanning tree instances are called Multiple Spanning Tree Instances (**MSTI**).

Note: Each VLAN can be associated with only one instance.

Once one or more MSTI (VLANs to spanning tree instance) are defined, it is recommended that all VLANs on the system be associated with a spanning tree instance other than the CIST (which removes the VLAN to CIST associations); in other words, once any MSTI is defined, others should be defined so that all VLANs for the system are associated with an MSTI, and the CIST no longer has any associated VLANs.

Bridges that share a common set of MSTIs (each with their associated set of VLANs) make up an **MST region**, with each MSTI forming a logical network topology; this is explained below.

Figure 13-13 shows an example of a network using MSTP. Note that the CIST has been omitted for simplicity.

Note: Since MSTP is a set of RSTP instances, the user should be familiar with the concepts of the single (R)STP instance, explained in previous subsections. This includes understanding of Upstream Ports, Root Bridges, static versus dynamic setting of ports as upstream or downstream, etc.

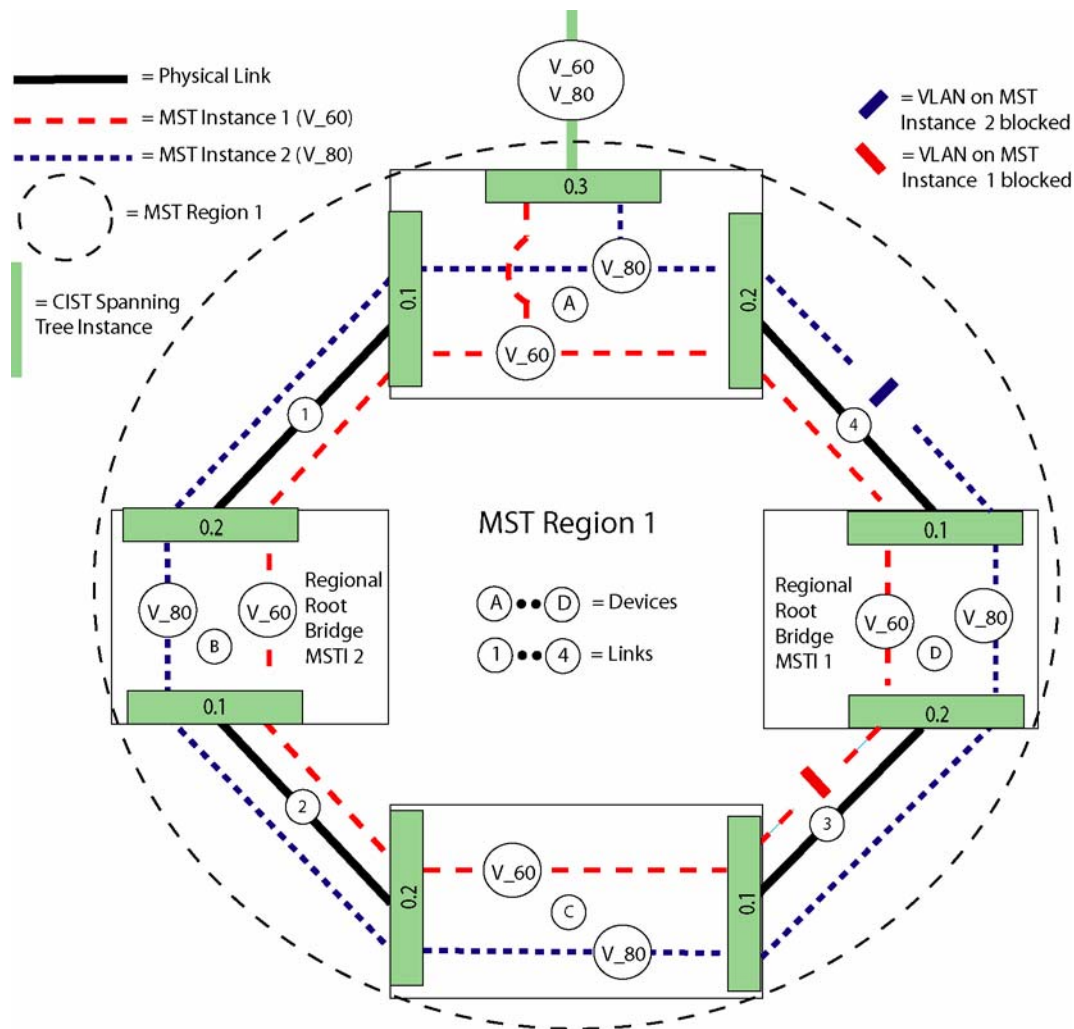


FIGURE 13-13 Concept of an MSTP Network

In [Figure 13-13](#), There are two MST instances, Instance 1 with VLAN 60 and Instance 2 which includes VLAN 80. Only one VLAN is associated with each instance; more than one VLAN can be associated with an MST instance, but this simple example helps to demonstrate key concepts.

For MST Instance 1, the VLAN is blocked on physical link 3, so that no traffic over VLAN 60 can traverse between bridges C and D. For MST Instance 2, the VLAN is blocked on physical link 4, so that no traffic over VLAN 80 can traverse between bridges A and D. With this topology, no loops are formed for each RSTP instance.

With this topology, if link 2 is now physically blocked, there will also be a block over MST Instances 1 and 2 over Physical Link 2. As a result, Bridge C is blocked from the network for MST Instance 2 (VLAN 80) and Bridges C and D are blocked from the network for MST Instance 1 (VLAN 60).

To correct this, MST Instance 1 will unblock its VLAN (V_60) over physical link 3, and Instance 2 will unblock its VLAN (V_80) over Physical Link 4. The resulting topology will now allow for no loops and no bridge is isolated. Refer to [Figure 13-14](#).

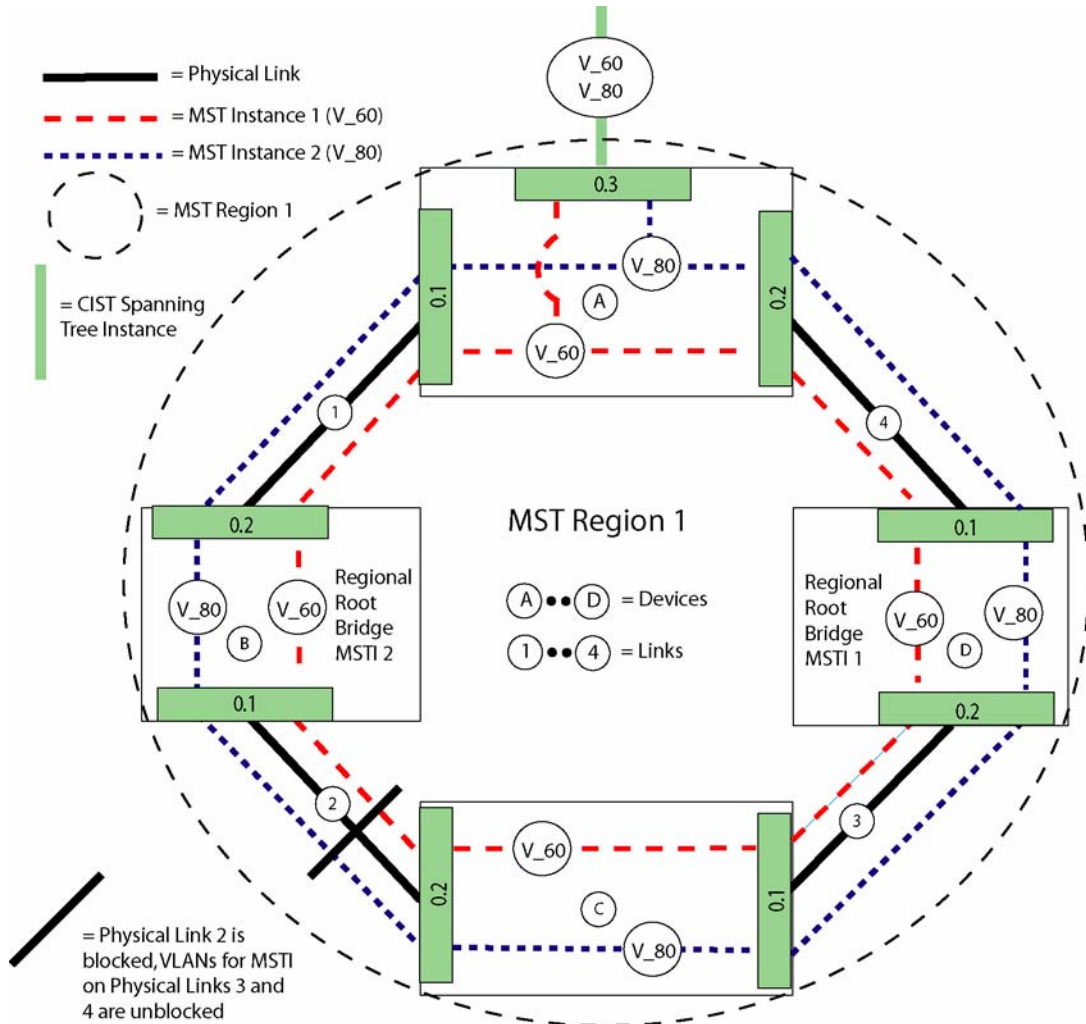


FIGURE 13-14 MSTP Recovery when Physical Link Blocked

13.7.2 MSTP with a Primary and Secondary Upstream Port

In release 6.0, it is possible to designate the VLAN interface (using the `FORWARDING` parameter) as `PRIMARYUPSTREAM` and `SECONDARYUPSTREAM`. This can be done for the VLANs for each MST instance, as shown in [Figure 13-15](#). In this example, if Node A failed, the VLANs on interfaces designated as `SECONDARYUPSTREAM` could carry traffic into the network.

Note: In this example, each MST Instance has only one VLAN and its upstream interface is configured as `PRIMARY` or `SECONDARYUPSTREAM`. If there are multiple VLANs for an MST Instance, each VLAN should be set as `PRIMARY` or `SECONDARYUPSTREAM` over the same upstream port. All non-upstream ports should be set to `STP`.

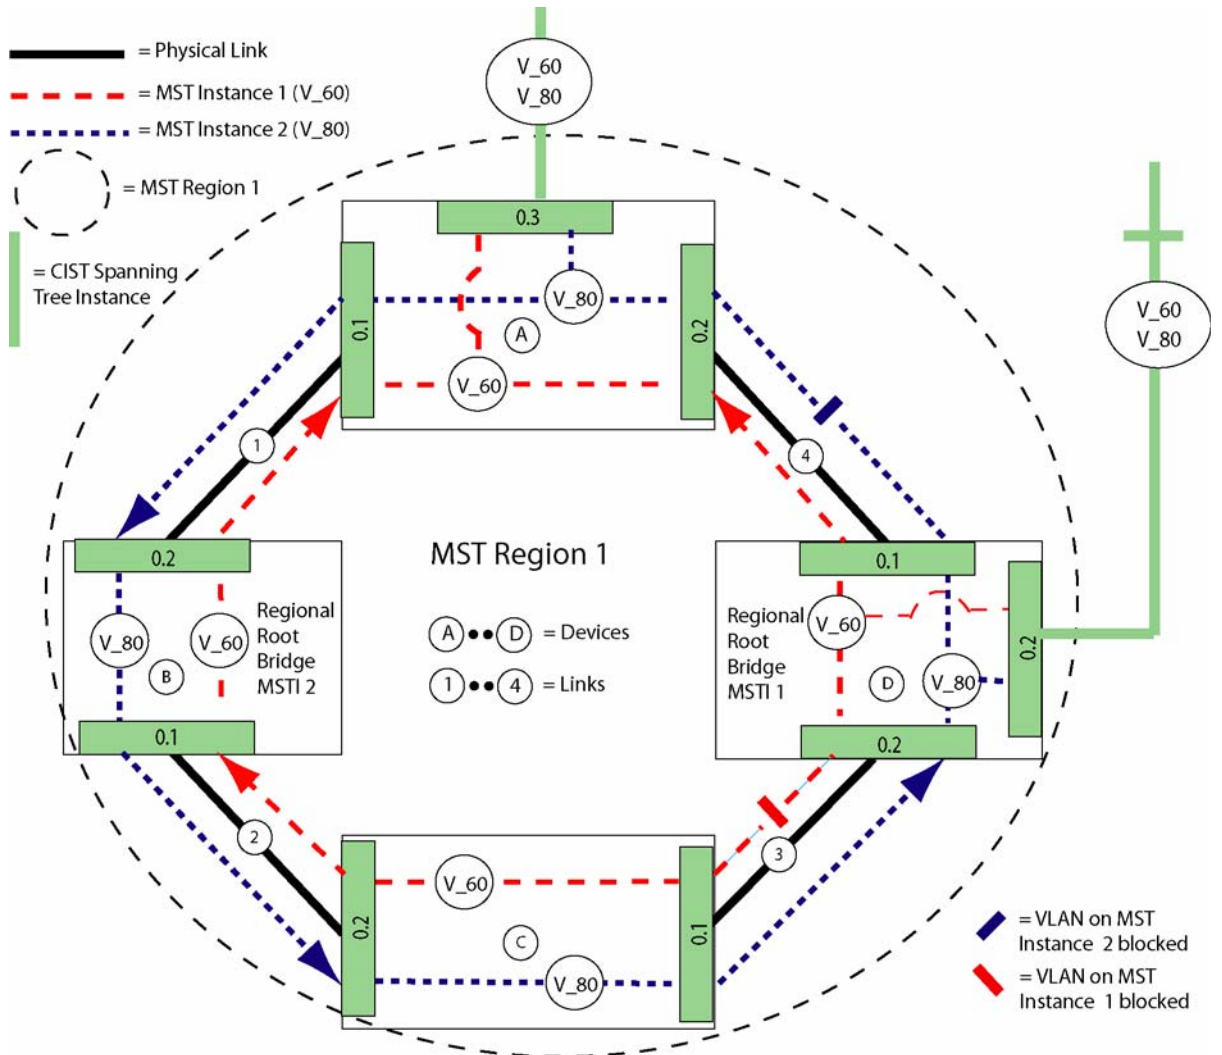


FIGURE 13-15 MSTP with Primary and Secondary Upstream Ports

13.7.3 MTSP Region

When a set of switches have the same MSTI configuration (meaning the set of switches have the same MSTIs and their VLAN associations), these switches can make an **MSTP region**. This allows the group of switches to be placed under a common administration; the region appears as one large bridge to the rest of the network spanning tree (i.e. the CIST). Since there is one overall network instance, which connects all the regions, blocking on boundary ports would occur so that there would be no loops into and out of the MST Region. Refer to the following figure.

Note: One feature available in 6.0, Cisco Compatible STP Mode, allows the fMAP to participate in the same MSTP region with one or more adjacent Cisco bridges. Refer to [13.7.7](#)

To form an MSTP Region, all bridges that make up the region must share these attributes:

- MSTP Instances
- VLANs associated with these instances
- MSTP Region Name
- MSTP Region Revision Level

Refer to [Figure 13-16](#), which shows the MST Region as part of the larger CIST. The CIST represents a spanning tree outside the MST region, but also has a spanning tree inside the region (the IST), and can carry all VLAN traffic outside the MST region.

Note that it is not required that VLANs are configured on all the ports (interfaces), although it is necessary if the user wishes traffic for a specific VLAN (which is part of an Instance) to be carried over that port. Not configuring VLANs on the port can be useful in the following scenarios:

- The user wishes to block VLAN traffic without changing the existing spanning tree
- As the MST Region is created, no loops are created which could result in packet storms.




FIGURE 13-16 Concept of an MSTP Region

13.7.4 Provisioning Parameters

Many of the commands and parameters for MSTP are similar to (R)STP, since the user is still creating an (R)STP instance that must go through a convergence process. However, some parameters are unique for MSTP, or some value for a common parameter is different, and these are highlighted.

Following are the key parameters that are data filled for (R)STP; for each there is a summary for the parameter (or a reference to an earlier subsection, especially [13.2.3](#)), and how MSTP uses the parameter.

13.7.4.1 Bridge ID

Bridge IDs are used in root bridge elections. The root bridge is the switch in the extended LAN with the numerically lowest bridge ID value. This is guaranteed to identify a single bridge due to the unique MAC address component. The user is allowed to change the bridge priority component to override the arbitrary root selection that will result from only comparing MAC addresses when the default bridge priorities are in use.

Bridge IDs are also used in designated bridge elections. Normally the switch with the lowest root path cost is the designated bridge for a physical LAN. If more than one switch has the same lowest root path cost, then the designated bridge is the switch with the numerically lowest bridge ID value.

The default bridge priority value is 32768. The fMAP previously supported a bridge priority that could be configured as a value from zero to 65536, in accordance with IEEE Std 802.1D, 1998 Edition. For MSTP, however, the priority component of the bridge ID is reduced to support MSTP operations, to allow for the unique identification of each MSTI in a bridge as part of a “system ID” that represents a (12-bit) numerical extension to the MAC address. This avoids the potential need to allocate up to 4094 additional MAC addresses per bridge to uniquely identify each MSTI. The reallocation of (bits in) the bridge ID contents was done in a manner that supports backwards compatibility with IEEE Std. 802.1D, 1998 Edition.

As a result, the bridge priority component has been modified to be a (4-bit) value between 0-61,440 that can only be provisioned in increments of 4096. This was done to allow for direct comparison with values from earlier versions of STP.

For Bridges that are running MSTP, there will be MSTI definitions to support the different VLANs defined for the bridge. Each of the MSTIs will have its own Bridge Identifier with the composition described above, except that each will include the Bridge MAC address as a component of the Bridge ID. Each will have a priority component, as described above, which can be independently provisioned from the other spanning tree instances defined for the same bridge.

The final component is an identifier called the “system ID extension” that is used to uniquely identify each of the MSTIs for a bridge. The CIST for each bridge will use the system ID extension value of **zero**. Any other MSTI defined for the bridge will utilize a value called the **MSTID** that identifies the MSTI. The MSTID parameter is described in a later section.

13.7.4.2 Port ID

Port IDs are used in root port elections. Normally, the port with the lowest root path cost is the root port for the switch. If more than one port ties for the lowest root path cost, then the root port is the port with the lowest numerical port ID

The default port priority value is 128. The IEEE Std 802.1D, 1998 Edition includes priority values on a per-port basis from zero to 255. For the fMAP, the storage space (number of bits) allocated to the priority component of the port ID had to be reduced to support bridges with larger numbers of ports, since this only left room for port numbers from 1-255. To maintain compatibility for comparison with previous versions of STP, the port priority is a value between 0-240 that can only be provisioned in increments of 16.

For Bridges that are running MSTP, the priority component of the Port ID is repeated for the CIST, and each MSTI defined for the bridge. **This allows the user complete independent control over the port configurations for each Spanning Tree instance.**

13.7.4.3 Port Path Costs

Port path costs are used in root path cost calculations, which are a factor in selecting root ports and designated bridges. By default, port path costs are related to the bandwidth capacity of the ports; however, the default values may be changed by the user to reflect other factors (e.g. propagation delay, link quality, desired traffic level, etc.)

The values for port path costs are listed in [13.2.3.3](#). For MSTP, the internal port path cost and the external port path cost are represented by one port path cost parameter described there.

13.7.4.4 Timer Durations

The spanning tree convergence process is dependent on certain timers, explained in [13.2.3.4](#). As explained there, RSTP has a parameter, Migrate Time, with a set default of 3 seconds. With MSTP, this is a variable.

13.7.4.5 Port participation

If ports on a switch are members of an extended LAN or VLAN that does not require use of the spanning tree protocol (i.e., if the VLAN is administered such that no network loops could exist), then spanning tree protocol operations can be disabled for those ports. However, if a port is a member of multiple VLANs, then the spanning tree protocol must be enabled on that port for all those VLANs or none of them; a mixed configuration is not supported.

If spanning tree protocol operations are disabled for a port, it may still pass bearer traffic to and from other ports, regardless of whether or not the spanning tree protocol is used for those other ports.

For bridges that run MSTP, port participation in the spanning tree may be disabled on a per MSTI basis. This means that VLAN traffic associated with the disabled MSTI may flow freely through those ports

Note: Spanning tree instances (MSTIs) themselves may not be disabled individually.

13.7.4.6 Force Version

Refer to [13.2.3.6](#). The value specifies whether STP or RSTP is to be used on a bridge. For MSTP the value 2 (RSTP) is used.

13.7.4.7 Edge Port

Refer to [13.2.3.7](#), and is part of RSTP processing in that it identifies a port is directly connected to the host. For MSTP, this value is also used.

13.7.4.8 Point-to-Point Port

Refer to [13.2.3.8](#), and is used by MSTP.

Note: In most cases, select AUTO so that the system can determine the port connection.

13.7.4.9 Transmit Hold Count

Refer to [13.2.3.9](#), and is used by MSTP processing.

13.7.4.10 Max Hops (Unique for MSTP)

For MSTP, an additional mechanism is added to control the circulation of old information within a spanning tree instance (CIST and MSTIs). Each BPDU sent for MSTP will contain a “remaining hop count” field. The value is initially set by the root bridge for the spanning tree instance (i.e., the Regional Root) where the BPDU originates. The field is then decremented at each bridge that the information passes through. Once the field reaches zero, the information stops circulating.

The Max Hops parameter allows the user to specify the value that this field will be initialized to should the bridge for which it is provisioned become the root of a spanning tree instance. By default, the max hops parameter is set to the same value as max age.

13.7.4.11 Multiple Spanning Tree Instance ID (Unique for MSTP)

When the user creates a new Spanning Tree instance for a bridge, a number between 1 and 4096 **must** be specified to uniquely identify this Multiple Spanning Tree Instance (MSTI) to other bridges connected to this bridge using LAN segments. The user also provides a name when the spanning tree instance is created which may be used for all commands on the local bridge which require a spanning tree instance to be identified. When the name “MAIN” is provided, the Common and Internal Spanning Tree instance (CIST) is utilized by the command processing. The MSTID for the CIST is zero.

13.7.4.12 MST Configuration Table (Unique for MSTP)

To associate one or more VLANs with a spanning tree instance, The `ADD STP` command is used with the name or MST ID of a spanning tree instance and one or more VLANs (by name or VID). The MST Configuration Table contains the VLAN to Spanning Tree mapping for a particular Bridge that is running MSTP. By default, all VLANs defined for a Bridge running MSTP will be mapped to CIST. As additional MSTIs are defined for the Bridge, this table will be filled in as the user provisions the desired mapping of defined VLANs for the bridge to the new MSTI.

Note: Once additional MSTIs are defined for the bridge, all created VLANs should be mapped out of the CIST and into the MSTI(s). (The CIST cannot be destroyed, and it is used when an MSTP region is configured.)

13.7.5 Network Engineering and Balancing VLAN/Port Configurations

To properly engineer MSTP in a network, the user should consider the following two rules:

- To meet the redundancy and load balance needs of the network, the minimum number of MSTIs necessary should be created. This implies putting as many bridges as possible into an MST Region. With this minimum number of MSTIs, multiple VLANs can be assigned to each one.
- When multiple VLANs are mapped to the MSTI, port blocking will occur for **all** the VLANs on the ports that are blocked by the MSTI.

13.7.6 Command Overview

The following commands highlight where MSTP is to be provisioned. For other commands and parameters, the same concepts apply as for (R)STP, with the user including the INSTANCE parameter to choose a specific MSTI.

By default, the switch has one CIST for (R)STP, which cannot be destroyed, and the (R) STP is disabled.

To change the protocol used for the switch, part of the SET STP command is used:

```
SET STP
PROTOCOL={ STP_ORIGINAL
            RSTP
            STP_COMPATIBLE_RSTP
            MSTP
            CISCO_COMPATIBLE_MSTP
            }
```

Note: The CISCO_COMPATIBLE_MODE is explained in the next section.

By default, the switch has one CIST which cannot be destroyed. To create an MSTI, the user must set to an MSTP mode and then create an instance with the command:

```
CREATE STP INSTANCE=stpname
MSTID=1..4094 [ PRIORITY=0..65535 ]
```

(Refer to [13.7.4.1](#) on numbering for the PRIORITY.)

The MSTI can also be renamed.

By default, all VLANs (and therefore all ports), belong to the CIST. Once created, VLANs can be associated with the MSTI using the command:

```
ADD STP
INSTANCE={ stpname
           mstid
           }
VLAN={ vlanname
       vid-range
       }
```

(VLANs can also be dis-associated with the MSTI as well.)

The user can continue to associate VLANs with MSTIs until there are no VLANs associated with the CIST.

By default, an fMAP device has one system wide STP instance that is disabled for all interfaces. The sequence to create the MSTIs and associate these with VLANs over interfaces depends on the installation scenario.

One scenario might be as follows:

1. Exclude interfaces that will not take part in the STP instance(s), such as ADSL interfaces.
2. Create the MSTI(s) and give them the appropriate attributes.
3. Associate the VLANs with the MSTI(s). (When done, there should be no VLAN to CIST association.)
4. Enable STP. This will ensure that no physical (port) loops are created.
5. Associate the VLANs with the interfaces. As this is done, the topology will move from physical (port-based) to logical (VLAN-based).

13.7.7 Cisco-Compatible MSTP Mode

13.7.7.1 Overview

The Cisco compatible MSTP mode allows the fMAP to interoperate with Cisco's proprietary implementation of MSTP, allowing the fMAP to participate in the same MSTP region with one or more adjacent Cisco bridges, and supporting up to 16 MST instances and up to 4094 associated VLANs. While in Cisco compatible MSTP mode, the fMAP can also simultaneously interoperate with other adjacent bridges that comply with the IEEE 802.1s-2002 standard. Refer to [Figure 13-16](#).

To turn on Cisco compatibility for MSTP, the user must first set the STP protocol version to `CISCO_COMPATIBLE_MSTP`, rather than `MSTP`. (Refer to [13.7.6](#).)

When in Cisco compatible MSTP mode, the fMAP will try to detect connections with Cisco bridges by looking for a non-zero Format Selector field in a received MSTP BPDU. An explicit Cisco compatible MSTP mode is provided, because if another bridge vendor also uses a non-zero Format Selector, the user may not want the fMAP to treat it like a Cisco bridge, and can then set the STP protocol version to `MSTP`.

Because the Cisco's calculation of the MST Configuration Digest does not follow the IEEE 802.1s-2002 standard and its algorithm cannot be obtained, the fMAP will simulate Cisco's MST Configuration Digest by learning and storing the value in flash. This allows the user to configure the fMAP to be in the same MST region as a Cisco bridge.

Note: Because of this simulation, user intervention is required to determine that all MSTI/VLAN associations configured on a fMAP match that of a Cisco bridge in the same MST region.

To be compatible with Cisco's MSTP, the fMAP first determines that the adjacent bridge is a Cisco bridge by detecting that the Format Selector field of the MST Configuration Id in a received BPDU is non-zero. (This happens to Cisco BPDUs because their Version 3 Length information is encoded one byte further than the standard.) If the STP protocol version of the fMAP is `CISCO_COMPATIBLE_MSTP`, then the fMAP will set the `rcvdCisco` flag for the port that the BPDU was received on, and all subsequent transmitted BPDUs on that port will be encoded using Cisco's format.

For the fMAP to be in the same MST region as a Cisco bridge, the fMAP must send the same MST Configuration Identifier information in its BPDUs as in the Cisco BPDUs. As discussed above, Cisco's MST Configura-

tion Digest is not calculated according to the IEEE 802.1s-2002 standard, and therefore cannot be recreated by the fMAP based on the MSTI/VLAN associations. Instead, the fMAP saves the MST Configuration Digest from the first Cisco BPDU that has the same MST region name and revision level as the fMAP. The fMAP then encodes Cisco's MST Configuration Digest back into BPDUs destined for Cisco bridges (BPDUs over ports with `rcvdCisco` flag set).

As a result, the fMAPs only initial protection from being placed in an incompatible MST region with a Cisco bridge is its provisioned MST region name and revision level; if both the fMAP and Cisco bridge have the same MST region name and revision level (and the fMAP does not already have a stored Cisco Configuration Digest), then both bridges will communicate as if in the same MST region even if the MSTI/VLAN associations differ between them.

13.7.7.2 Changing from Cisco-Compatible MSTP to MSTP Mode

Changing a VLAN / MST instance association when the fMAP is in Cisco Compatible MSTP Mode will cause the fMAP's learned Cisco configuration digest to change, and the fMAP will not participate in the MST region. This is the expected result, according to IEEE 802.1s.

However, because of the fMAPs limitations with respect to Cisco compatibility, the Cisco configuration digest must be re-learned for the fMAP to interoperate with a Cisco bridge in another MST region. To ensure that no loops are formed when changing a VLAN / MST instance association and re-learning the Cisco configuration digest, the following procedure should be followed.

1. Change the region name or revision level on the fMAP that the VLAN / MST instance association change will be made on. Note that this change may be made before or after the VLAN / MST instance association change. The new region name or revision level adds an extra level of security to prevent the fMAP from learning the wrong Cisco configuration digest from a different MST region, and thus causing a network loop.
2. Make all required VLAN / MST instance association changes.
3. Verify all adjacent Cisco bridges that are intended to be in the same MST region as the fMAP, have the same MST configuration as the fMAP (i.e., region name, revision level, and all VLAN / MST instance associations are the same).
4. Make sure the region name or revision level was changed (step 1), then perform a `RESET STP LEARN-CISCODIGEST` command.

13.7.7.3 Changing MSTP Configuration on an Adjacent Cisco Bridge when in Cisco Compatible MSTP Mode

The following procedure should be followed when a VLAN / MST instance association on an adjacent Cisco bridge is participating in the same MST region as the fMAP.

1. Change the region name or revision level on the Cisco bridge adjacent to the fMAP. Note that this change may be made before or after the VLAN / MST instance association change. The new region name or revision level adds an extra level of security to prevent the fMAP from learning the wrong Cisco configuration digest from a different MST region, and thus causing a network loop.
2. Make all required VLAN / MST instance association changes.

13.7.7.4 Limitations

- User Intervention Required for Proper MSTP Region Boundary Detection

Because the fMAP cannot reproduce a Cisco MST Configuration Identifier without copying the MST Configuration Digest from a Cisco BPDU, the user must ensure the MST configuration information of the fMAP matches that of other Cisco bridges in the same MST region. In particular, the MSTI/VLAN associations must be the same on all bridges in the region. The consequence of the fMAP and a Cisco bridge in the same MST region with a differing MSTI/VLAN association is one or more VLAN traffic loops.

To resolve a traffic loop caused by the last scenario, the user should change the MST region name or revision level on the fMAP, correct the MSTI/VLAN associations on the fMAP or the Cisco bridge, perform a `RESET STP LEARNCISCODIGEST`, then change back the MST region name or revision level.

If the MSTI/VLAN associations are changed on the fMAP by the user, the stored Cisco MST Configuration Digest on the fMAP will change as expected, but will not match any Cisco MST Configuration Digest. Once the MST configuration information is provisioned to be the same between the fMAP and a Cisco bridge, the `RESET STP LEARNCISCODIGEST` command should be issued to get the fMAP to relearn the Cisco MST Configuration Digest.

- Misleading Interface Summary Information

In Release 6.0, the `SHOW STP` command displays a summary of the interfaces that are participating in a particular MST instance. As well as indicating other things, the type field in the display will also indicate whether an interface is connected to a Cisco bridge or not. Because the fMAP can only detect a Cisco bridge by its BPDU format, the interface summary from the `SHOW STP` command may not indicate an interface is connected to a Cisco bridge when it should. This will occur if the RSTP port role of the Cisco port is alternate port role for all MSTIs associated with that port., since alternate ports do not transmit BPDUs. A `RESET STP` command will rectify this.

- Cisco Compatibility

Although the fMAP is compatible with Cisco's implementation of Spanning Tree Protocol (STP) and Multiple Spanning Tree Protocol (MSTP 802.1s), the fMAP is not compatible (in general) with Cisco's Per Vlan Spanning Tree (PVST) or Per Vlan Spanning Tree Plus (PVST+).

13.7.7.5 Command Overview

- `SET STP PROTOCOL=CISCO_COMPATIBLE_MSTP`

This will set the fMAP into Cisco compatible MSTP mode, and cause it to transmit and receive Cisco MSTP BPDUs to and from Cisco bridges.

- `RESET STP LEARNCISCODIGEST`

This will cause the fMAP to copy and store the MST Configuration Digest from the first received Cisco MSTP BPDU, provided that the MST region name and revision level from the BPDU matches that on the fMAP. Note, this command only applies to Cisco compatible MSTP mode.

- SET STP CISCODIGEST

This will set the Cisco configuration digest. This command is not intended to be used by the user. Instead, it is provided to allow the Cisco configuration digest to be backed up or restored from a text config file.

- SET STP CISCOLEARNEDIIF

This will specify the interface that the Cisco configuration digest was learned on. This command is **not** intended to be used by the user. Instead, it is provided to allow the interface that the Cisco configuration digest was learned on to be backed up or restored from a text config file.

13.7.8 Command Summary

Table 13-8 lists the commands used for STP and therefore includes parameters specific for RSTP and MSTP.

TABLE 13-8 Commands for Spanning Tree Protocol

Object	Verb	Syntax	Description
STP INSTANCE MSTID	CREATE	CREATE STP INSTANCE=stpname MSTID=1..4094 [PRIORITY=0..65535]	For MSTP, creates the STP instance that is not the initial instance (CIST).
VLAN VID	CREATE	CREATE VLAN=vlanname VID=2..4094 [FORWARDINGMODE={ STD UPSTREAMONLY }]	Creates a VLAN other than the default VLAN (1). This is where the VLAN is set to UFO (UPSTREAMONLY).
VLAN INTERFACE		ADD VLAN={ vlanname vid } INTERFACE={ type:id-range id-range ifname-list ALL } [FRAME={ UNTAGGED TAGGED }] [TRANSLATE={ 1..4094 }] [FORWARDING={ PRIMARYUPSTREAM SECONDARYUPSTREAM DOWNSTREAM STP UCP }]	Associates the VLAN with the specified interfaces. The FORWARDING parameter value depends on how the interface fits into the network.

TABLE 13-8 Commands for Spanning Tree Protocol (Continued)

Object	Verb	Syntax	Description
STP	SET	<pre> SET STP { INSTANCE={ stpname mstid MAIN ALL } { DEFAULT PRIORITY=0..65535 INTERFACE={ type:id-range id-range ifname-list ALL } { DEFAULT [PATHCOST=path-cost] [PORTPRIORITY=port-priority] [EDGEPORT={ TRUE FALSE }] [POINT2POINT={ TRUE FALSE AUTO }] } } DEFAULT [PRIORITY=0..65535] [FORWARDDELAY=4..30 [HELLOTIME=1..10] [MAXAGE=6..40] [TXMAX=1..10] [MAXHOPS=6..40] [MSTREGION=regionname] [REVISIONLEVEL=0..65535] [CISCODIGEST=hexstring] [CISCOLEARNEDINTERFACE={ type:id id }] PROTOCOL={ STP_ORIGINAL RSTP STP_COMPATIBLE_RSTP MSTP CISCO_COMPATIBLE_MSTP } [FORCE] INTERFACE={ type:id-range id-range ifname-list ALL } { DEFAULT [PATHCOST=path-cost] [PORTPRIORITY=port-priority] [EDGEPORT={ TRUE FALSE }] [POINT2POINT={ TRUE FALSE AUTO }] } } </pre>	<p>Sets the STP parameters for the specific instance; the parameters that are used depend on the type of STP (PROTOCOL) being configured.</p> <p>Refer to 13.2.3.4.</p>

TABLE 13-8 Commands for Spanning Tree Protocol (Continued)

Object	Verb	Syntax	Description
STP INSTANCE VLAN	ADD	<pre> ADD STP INSTANCE={ stpname mstid } VLAN={ vlanname vid-range } </pre>	Associates the STP Instances with the (set of) VLANs.
	ENABLE	<pre> ENABLE STP [{ [INSTANCE={ stpname mstid MAIN ALL } INTERFACE={ type:id-range id-range ifname-list ALL } [{ TOPOLOGYCHANGE RSTPCHECK }]] [INTERFACE={ type:id-range id-range ifname-list ALL } [{ TOPOLOGYCHANGE RSTPCHECK }]] }]] </pre>	<p>Enables an interface or set of interfaces for the STP instance. This ensures no loop capabilities can take place on that interface or interfaces.</p> <p>Including TOPOLOGYCHANGE is only modifying whether a Topology Change will be declared when a failure is detected on the set of interfaces.</p> <p>RSTPCHECK is a single command to initiate the checking of adjacent bridges (i.e. no disable). Refer to 13.2.3.11.</p>
	SHOW	<pre> SHOW STP [{ [INSTANCE={ stpname mstid MAIN ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] [FULL] COUNTER }]] </pre>	Shows the STP information for the specified STP instance.

TABLE 13-8 Commands for Spanning Tree Protocol (Continued)

Object	Verb	Syntax	Description
STP	DISABLE	<pre> DISABLE STP [{ [INSTANCE={ stpname mstid MAIN ALL } INTERFACE={ type:id-range id-range ifname-list ALL } [TOPOLOGYCHANGE]] [INTERFACE={ type:id-range id-range ifname-list ALL } [TOPOLOGYCHANGE]] }] </pre>	<p>Disables the STP instance over an interface or set of interfaces.</p> <p>TOPOLOGYCHANGE controls the detection of topology changes on the associated interface. This allows the disabling of topology change detection on ports that are known to be connected to single end stations that could cause the Topology Change Notification mechanism to be triggered for the entire network when the end station is power cycled</p>
STP INSTANCE VLAN	DELETE	<pre> DELETE STP INSTANCE={ stpname mstid ALL } VLAN={ vlanname vid-range ALL } </pre>	<p>Disassociates the STP Instances with the (set of) VLANs. There is a warning if the STP instance is still enabled.</p>
STP INSTANCE	DESTROY	<pre> DESTROY STP INSTANCE={ stpname mstid ALL } </pre>	<p>Once the STP instance and VLANs are disassociated, the STP instance can be destroyed. There is a warning is there are still any VLANs associated with the STP Instance.</p>

13.8 DHCP Relay

13.8.1 Overview

Dynamic Host Configuration Protocol (DHCP) provides a mechanism through which IP hosts can obtain protocol configuration parameters automatically through the network. Such configuration parameters may include, among other things, an IP address, the subnet mask, and the default gateway.

DHCP is based on a client-server model in which a client contacts a server to obtain its configuration parameters. The server is normally centrally located in the network and is maintained by the administrator of that network. Since each client requires an IP address to communicate with other clients in an IP network, DHCP eases the administrative burden of manually configuring each client with an IP address. In addition, if a client moves, DHCP manages its ability to obtain another IP address in the subnet it moved to.

Note that a client must be allocated an IP address in the subnet it is part of. This would imply that a DHCP server would be needed on each subnet since DHCP messages are sent as broadcast messages (both at layer 2 and 3) and routers normally would not broadcast such messages. However, with the DHCP relay agent, the fMAP system intercepts the broadcast DHCP messages and can forward the messages as unicast to the appropriate server(s). (The DHCP messages can also be forwarded as broadcast, a feature for 6.1 as explained in [13.8.5](#).) Messages from the server (unicast) are intercepted and sent to the client over the correct port.




FIGURE 13-17 Network with DHCP server - Physical View

The network illustrated in [Figure 13-17](#) includes two subnets with fMAP systems configured as DHCP relay agents, a router (configured as a DHCP relay agent), and a single DHCP server (instead of individual DHCP servers in each subnet). Since the server is allocating IP addresses to clients on many subnets without actually being connected to each of these subnets, it must have knowledge of which subnet to allocate an IP address to a client that requested it. This is accomplished by the DHCP protocol.

13.8.2 DHCP Relay Agent Functionality in 6.0

13.8.2.1 Overview

With the DHCP Relay Agent function, the fMAP system emulates a DHCP server, in essence acting on behalf of the DHCP server and providing the IP address information to the client.

Figure 13-18 shows how the DHCP packets are sent and received between the server and client with the DHCP Relay Agent. Refer to this figure while reading this subsection.




FIGURE 13-18 DHCP Relay - Release 6.0

13.8.2.2 Client Requesting an IP Address (Dynamically Allocated)

The DHCP relay agent functionality allows the router to forward the request to a DHCP server on behalf of a DHCP client on the same subnet as the router interface. When the DHCP Relay Agent receives a DHCP packet from a client (configured as a CUSTOMER interface), the following sequence occurs:

1. The packet is discarded for the following conditions; otherwise it is forwarded.
 - Option 82 has already been added to the packet
 - The gateway IP address is not 0.0.0.0
2. The relay agent adds the DHCP relay agent information option consisting of the following:
 - **agent circuit id option** - Identifies the exact interface the request was received on. (Also, when a reply to the request is received by the agent, the agent circuit id option is used to determine the interface to send the reply to.) The agent circuit id is generated **automatically** and is known based upon the slot/port from where the request was received.
 - **agent remote id option** - Identifies the subnet address space from which the IP address for the client should be allocated from. This is **manually** configured using commands.

Note: The tag value for the DHCP relay agent information option is 82. The length field includes all the bytes of the sub-options. The value field is the sub-option bytes. Each sub-option also follows the tag-length-value format. The tag for the agent circuit id sub-option is 1 and the tag for the agent remote id option is 2.

3. The DHCP Relay Agent inserts its own IP interface address, over which the DHCP request was received on, in the **IP source address** field of the packet (i.e. the source IP address of 0.0.0.0 is replaced by the router interface address which incidentally is part of the same subnet as the client). This allows the server to respond via unicast to the relay agent.
4. The DHCP Relay Agent also replaces the “all networks broadcast address” in the **IP destination address** field of the packet with the IP address of the server. If the relay agent is aware of more than one server, it does this for each of the server addresses so that the same request is sent to each of these servers.
5. The DHCP Relay Agent populates a field called the “gateway IP address” in the request with the IP address of the interface over which the DHCP request was received from the client. This informs the server that it has to allocate the client an IP address from this subnet space, since it uses this IP address value to determine the subnet from where certain DHCP messages originate.
6. The DHCP snooping agent sends as a unicast message the DHCP packet over the **inband management VLAN**.

13.8.2.3 DHCP Server Sending IP Address

When the snooping agent receives a DHCP packet from the server over the inband management VLAN, the following occurs:

1. The DHCP Snooping agent checks if the packet has option 82. If option 82 is not present, it is discarded.
2. The DHCP snooping agent checks the agent remote id, which identifies a specific fMAP system.

- If it is not this system, the packet is flooded over all interfaces configured as NETWORK for the inband management VLAN.
- If it is this system, the option 82 information (circuit id) is used to determine the interface to send the reply to. The option 82 is removed from the packet and the packet is sent to the destination (CUSTOMER) interface.

13.8.2.4 Static IP Allocation

The client may also request a static IP address; if the DHCP server can allocate the address, the client will have the address without a time limit.

13.8.2.5 Classification Rules

The DHCP Relay feature utilizes classifiers to extract the client's DHCP packets from the customer interfaces to be processed by the DHCP Relay function. These classifiers are installed by the system on all interfaces that are provisioned with a DIRECTION of CUSTOMER and are not installed on interfaces that are provisioned with a DIRECTION of NETWORK. The classifiers are only installed when DHCP Relay is enabled for a given interface. DHCP Relay utilizes the IP stack for its communication with the DHCP server(s); this occurs through Inband Management.

Once DHCPRELAY has been enabled, the system may generate a warning message at the user's CLI session stating that classifier capacity or capabilities have been exceeded on the slot(s) impacted by the provisioning change. The user should investigate classifier-related provisioning, such as IGMP, DHCPRELAY, VLAN (for per-VLAN UFO and HVLAN), EPSR, INTERFACE (TAGALL option for HVLAN), ACCESSLIST, and CLASSIFIER to determine the reason for the message.

13.8.2.6 IP Filtering

DHCP Relay has an option for automatically filtering on the Customer interface, based on the source IP address. The IP address used in the filter is based on the DHCP Discover/Offer/Request/Ack exchange between the DHCP client and the DHCP Server. When an IP address is handed out, the DHCP Relay software will create an IP "pass" filter to allow packets with the source IP address to be received from the port that was given the IP address; all other IP packets are dropped.

Note: In release 6.0, the maximum number of IP filters is 5. In release 6.1, the maximum is eight.

13.8.2.7 DIRECTION of Interfaces

DHCP Relay can only be enabled on interfaces that are provisioned with a DIRECTION of CUSTOMER (the interface is at the port/Ethernet level). When the first Customer interface is enabled, all Network interfaces are enabled for DHCP Relay. When the last Customer port is disabled for DHCP Relay, then all of the Network interfaces are disabled for DHCP Relay.

When an interface's DIRECTION is changed from CUSTOMER to NETWORK, all learned IP addresses and all IP filters will be removed. If there are any other interfaces that have a DIRECTION of CUSTOMER, then the changed interface will be enabled as a Network Interface for DHCP Relay. If there are no other Customer interfaces that are enabled for DHCP Relay, then all Network Interfaces will be disabled for DHCP Relay.

When an interface's `DIRECTION` is changed from `NETWORK` to `CUSTOMER`, DHCP Relay will change the interface's role from being a possible interface to the DHCP server to that of a DHCP client, but DHCP Relay will be `DISABLED` on that interface (and no IP filters).

13.8.3 DHCP Relay Agent Functionality in 6.1

13.8.3.1 Overview

In release 6.1, the following enhancements are made to the DHCP Agent feature. These are then described in more detail.

- **DHCP Relay Mode** - A fMAP system can support multiple instances of DHCP Relay (RFC 3046). Moreover, each instance has mapping to a VLAN. The functionality of the DHCP relay is the same as in release 6.0 except for the features described here.

Note: Counters are displayed and reset on an Interface basis

- **DHCP Snooping Mode** - DHCP Snooping is a configurable option for each DHCP Relay instance. The key difference between DHCP Relay and DHCP snooping is that DHCP Snooping allows normal (non-MGMT) VLANs to be configured to have DHCP functionality. DHCP packets have the 82 option added and are forwarded as a broadcast (rather than unicast) message; as a result, the DHCP Snooping Agent floods the DHCP packet out of all interfaces that are configured as `NETWORK` for the VLAN.
- **Auto Ageing and IP Filter Removal** - There is the option to age the DHCP IP address and the optional IP Filter for any or all customer interfaces based on the lease time in the DHCP Ack packet. The Auto Ageing feature is configurable on a per interface (`CUSTOMER`) basis.

13.8.4 DHCP Relay Mode

13.8.4.1 Overview

Refer to the following figure that shows how multiple DHCP Relay Instances could be used. The DHCP Relay instance provides VLAN-level control but also allows users that don't require VLAN level of control simpler provisioning:

1. For users that require a different set of configuration data (DHCP Server list, Agent Remote ID, CID format) for each VLAN, they could create separate DHCP Relay Instances and map one VLAN to each Instance.
2. For users that want to have a set of VLANs use the same DHCP configuration, they could create the required number of instances and then map one or more VLANs to each DHCP Instance. Each DHCP Instance could match a different ISP.
3. Users that only require one DHCP Instance (as supported in release 6.0) only need to provision one instance and have all VLANs map to that one instance.

13.8.4.2 Configuration Rules

- The MGMT VLAN can be non-UFO.

- Various DHCP Instance/VLAN mapping are allowed; however, one VLAN **cannot** map to multiple DHCP instances.
- If IP Filtering is set to ON for an interface, there is a limit of eight VLANs that can participate in DHCP Relay mode on that interface.
- Each DHCP Relay instance can have up to 10 servers.

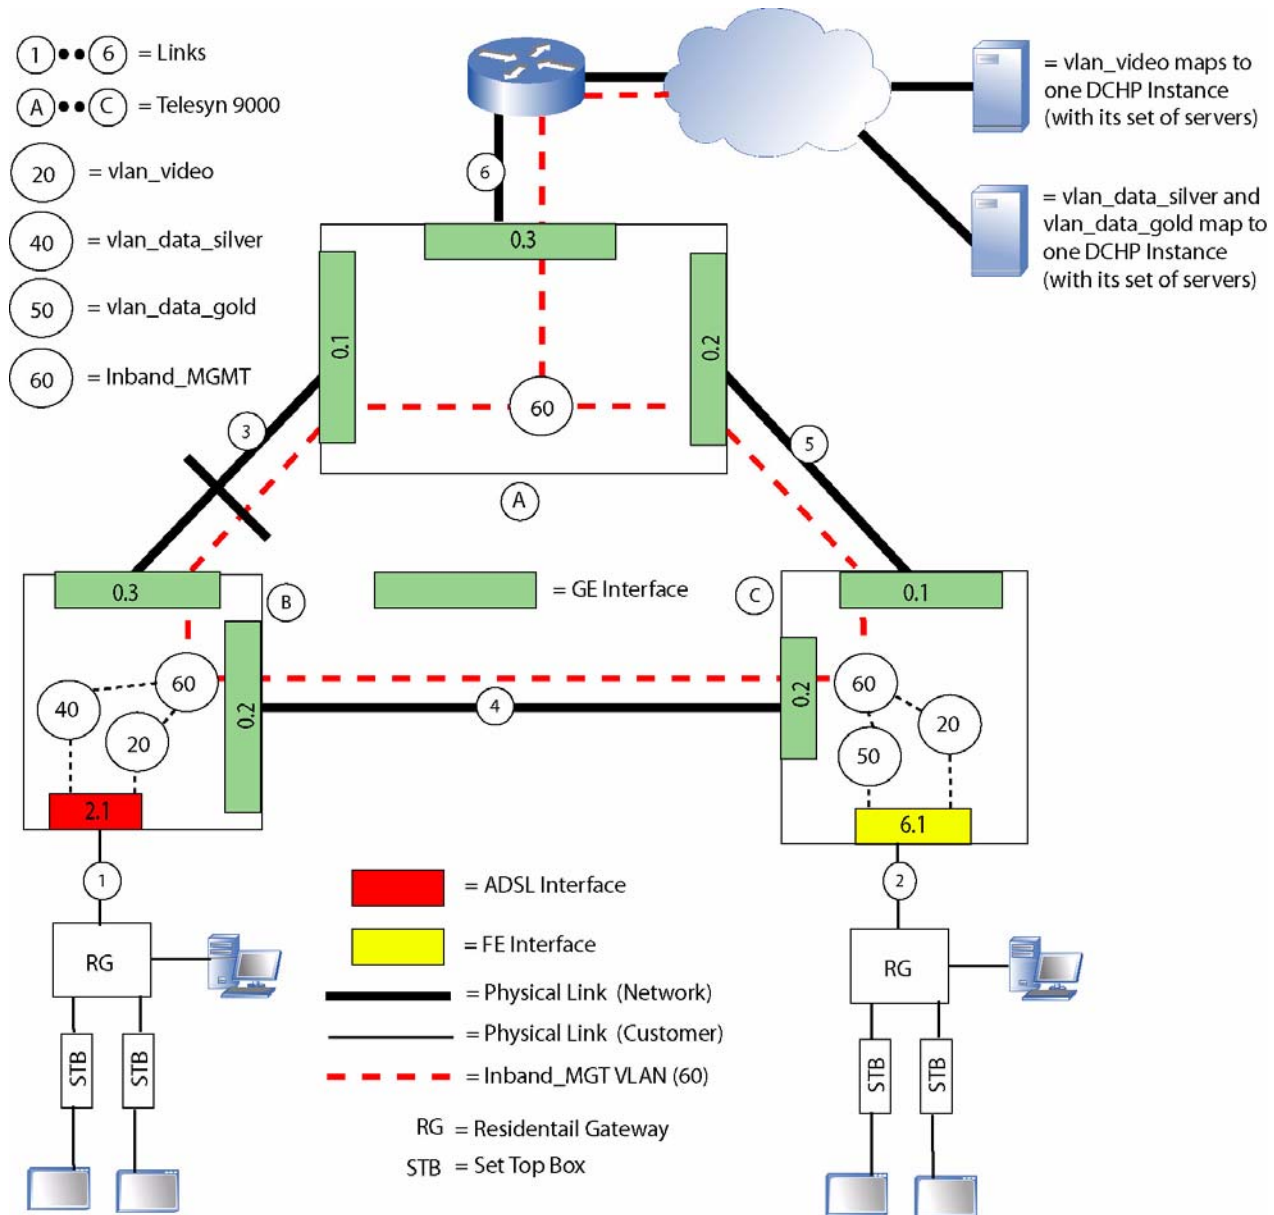


FIGURE 13-19 Example DHCP Relay in 6.1

13.8.5 DHCP Snooping Mode

DHCP Snooping Mode allows **any** VLAN to be configured for DHCP. Because of this, different configuration rules apply, such as the router must have DHCP relay agent configured. These are explained in more detail below.

Refer to [Figure 13-20](#) for an example configuration.

13.8.5.1 Client Requesting an IP Address (Dynamically Allocated)

- Steps 1 and 2 are the same as in [13.8.2.2](#).
- Subsequent steps are:
 1. DHCP Snooping leaves the “all networks broadcast address” in the IP destination field of the packet with the IP address of the server.
 2. The DHCP Snooping agent leaves the “gateway IP address” in the request with the 0.0.0.0 IP address.
 3. The DHCP Snooping agent floods the DHCP packet out all `NETWORK` interfaces that are members of the VLAN.

In essence, DHCP Snooping allows a VLAN in which option 82 can be inserted.

13.8.5.2 Receiving Requests from the DHCP Server

The steps are the same as in [13.8.2.2](#).

13.8.5.3 Configuration Rules

- The VLAN type used (UFO or non-UFO) is independent of including DHCP, since the DHCP packets are flooded only on `NETWORK` interfaces
- The upstream router must be configured with a DHCP relay agent since it must route the broadcast DHCP message to the correct set of servers. If IP Filtering is set to ON for an interface, there is a limit of eight VLANs that can participate in DHCP Relay mode on that interface.

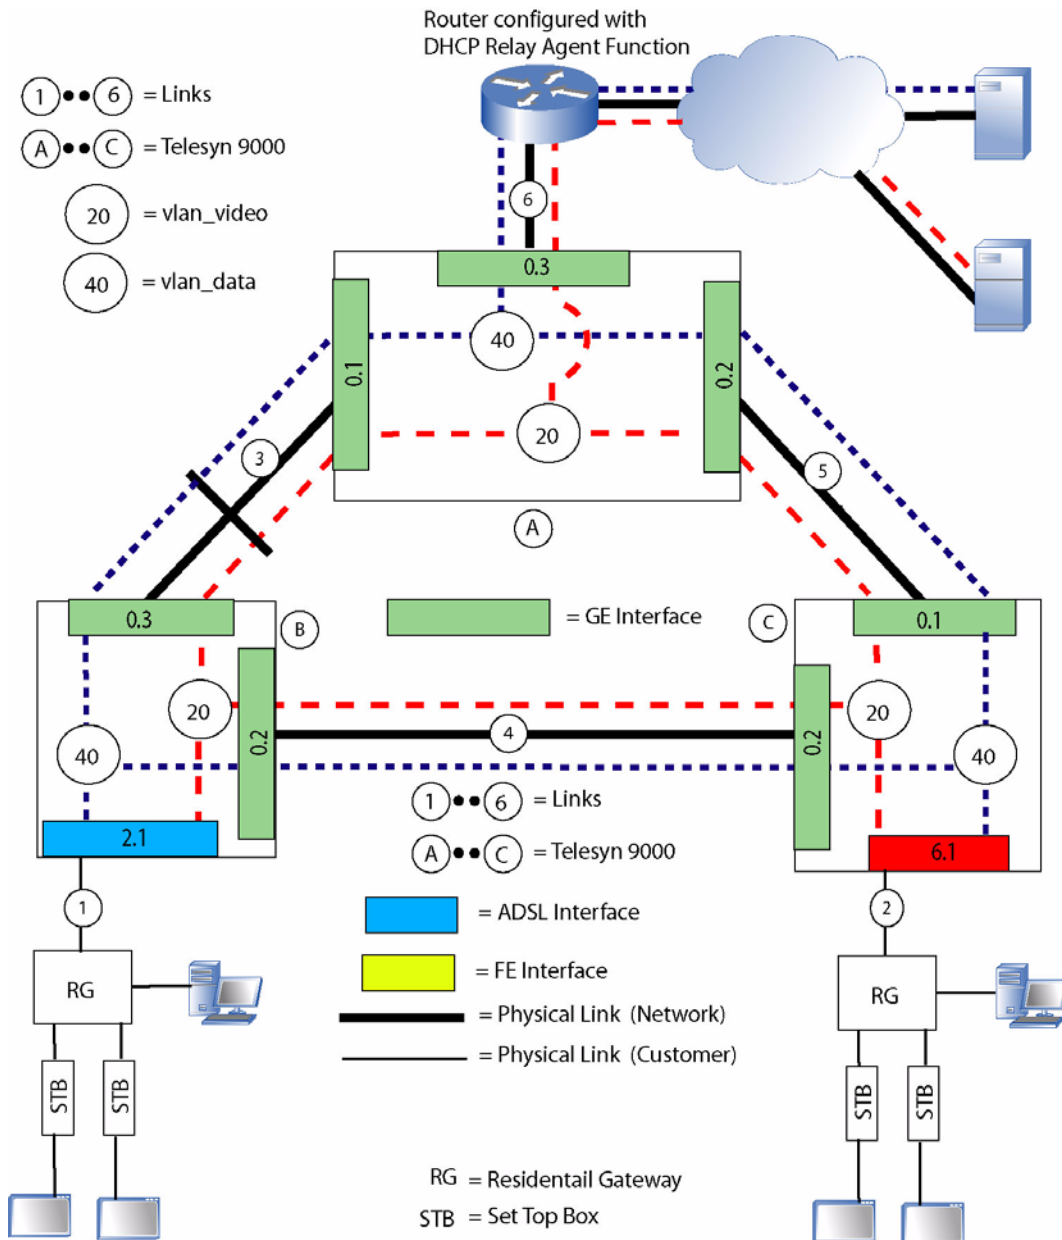


FIGURE 13-20 DHCP Relay Snooping

13.8.6 Auto Ageing and IP Filter Removal

In release 6.1, the fMAP supports the option to age the DHCP IP address and the optional IP Filter for any or all customer interfaces based on the lease time in the DHCP Ack packet. The Auto Ageing feature is configurable on a per interface (CUSTOMER) basis.

When DHCP Relay is enabled (in either Relay or Snooping mode) for a given customer interface, it can be configured to automatically age the allocated IP address and then to automatically remove the IP address from the DHCP Relay tables, thereby making room for another entry in its tables. When DHCP Relay has been configured to automatically create an IP Filter based on the allocated IP address (FILTER=ON) and auto age is also enabled (AGEING=ON), then DHCP Relay will automatically remove the IP address when the IP address ages out. The ageing will be based upon the lease-time specified in the DHCP Ack packet that is received from the DHCP server.

The Auto Age feature will save the Lease time and the DHCP server's IP ("server identifier" option field) address upon receiving a DHCP Ack packet. If the DHCP Ack packet has a lease time of 0, there is no ageing (infinite), and this will replace any previous lease value.

The DHCP Relay entry and the IP filter will be removed when one of the following occurs:

- Lease time expires (unless the lease is infinite)
- Receive a DHCP NAK from the DHCP server that the current IP address was allocated
- Receive a DHCP Decline packet from the DHCP client with the ciaddr of the currently allocated IP address
- Receive a DHCP Release packet from the DHCP client with the ciaddr of the currently allocated IP address (this would occur if the client moved to another port).

13.8.7 Default and Upgrade Configuration

When a fMAP is **initially booted up** (or from a "Purge database"), DHCP Relay will be configured as follows:

- There is one DHCP Relay instance called MAIN. All interfaces are enabled for MAIN, and no VLANs are associated with MAIN.
- MAIN cannot be renamed or deleted. (This is true even if the user subsequently disassociates all VLANs with MAIN.)
- The per-interface IP Filters will be disabled (off)
- The per-interface Auto Age status will be disabled (off)
- The per-interface counters will be cleared (all 0).

Caution: Ensure that DHCP is enabled on the fMAP network (upstream) interfaces. Otherwise, DHCP will not work.

During **software upgrade to release from R6.0 to R6.1** the configuration is affected as follows:

- If DHCP Relay was configured in R6.0 (single instance), then a single instance will be created (MAIN) with all VLANs as a member of the MAIN instance and the "mode" set to "Relay".

- The Interface status (enable/disable) will be reflected in the single DHCP Relay instance (MAIN).
- The Automatic Ageing option for an interface will be set to disabled.
- If DHCP was configured on a customer-side interface, DHCP RELAY is enabled for the Network side interfaces automatically.

13.8.8 Counters

Counters for DHCP can be viewed per interface (both NETWORK and CUSTOMER), not per Instance. There are also global counters that give totals for a system.

13.8.9 DHCP Relay / Snooping on Network Interfaces (6.1.3)

Prior to release 6.1.3, on a Network interface, DHCP Relay/Snooping would only forward previously Relayed/Snooped DHCP packets to the DHCP server. The network devices would have static IP addresses and would not be performing DHCP to get an IP address. Therefore, DHCP Relay/Snooping discarded DHCP packets that it received from a Network Interface that was not already Relayed or Snooped (did not already have option82).

In release 6.1.3 forward, DHCP Relay has been enhanced as follows:

- Upon receiving a (client side) DHCP packet without option82, DHCP Relay will add option82 (with remote id + slot/port/vid of the network interface) and forward the DHCP packet to the list of servers.
- Upon receiving a (client side) DHCP packet without option82, DHCP Snooping will add option82 (with remote id + slot/port/vid of the network interface) and flood the DHCP packet out the other Network interfaces that are a member of the vlan.
- DHCP Relay/Snooping tracks up to 20 concurrent DHCP client/server exchanges at any one time per Network interface. If there are more than 20 concurrent DHCP client/server exchanges on any one interface, then DHCP Relay/Snooping will overwrite one of the current client/server entries only when the client has already received an IP address, with the new one. This means that if there are 20 DHCP requests in progress from the same (network) interface and none of the 20 have completed (have not received an ACK/NACK), and the 21st DHCP request comes in, the 21st DHCP request will be dropped.
- DHCP Relay/Snooping will not install filters based on IP address received from the DHCP server.
- DHCP Relay/Snooping will not auto age out the IP filters based on the lease-time.

13.8.10 DHCP Relay Commands

TABLE 13-9 DHCP Relay commands

Syntax	Samples for DHCP Relay Agent Configuration	Description
<pre>CREATE DHCPRELAY=dhcpname [AGENT REMOTEID={ remote-id DEFAULT }] [MODE={ RELAY SNOOPING }] [SERVER={ ipaddress-list NONE }] [VLAN={ vlanname-list vid-range ALL }]</pre>		<p>Creates a DHCP Relay instance and its attributes.</p> <p>DHCP Relay or DHCP Snooping is controlled with the MODE parameter.</p> <p>The AGENT, MODE, SERVER, and VLAN attributes can be defined with the SET DHCPRELAY INSTANCE command as well.</p> <p>REMOTEID is used by DHCP Servers to identify this Relay Agent. Setting this parameter is optional. The default REMOTEID is the MAC address of the system the Relay Agent is running on. The user can specify the REMOTEID by entering a string containing 1..31 ASCII characters.</p>
<pre>ADD DHCPRELAY={ dhcpname- list MAIN ALL } SERVER=ipaddress-list</pre>		<p>Specifies the list of up to ten (10) DHCP Server IP addresses for the DHCPRELAY instance. For DHCP Relay, the relay agent will forward all DHCP broadcast messages to each configured IP address. For messages containing Server Identifier Option (54), the message will be unicast to the Server IP address in the message</p> <p>MAIN is the default DHCP instance.</p>
<pre>ADD DHCPRELAY={ dhcpname MAIN } VLAN={ vlanname-list vid-range ALL }</pre>		<p>Associates or maps a VLAN (or set of VLANs) to the DHCPRELAY Instance.</p> <p>MAIN is used if only one DHCP Relay Instance (DEFAULT) has been configured.</p>

TABLE 13-9 DHCP Relay commands (Continued)

Syntax	Samples for DHCP Relay Agent Configuration	Description
<pre>SET DHCPRelay={ dhcpname- list MAIN ALL } [AGENT REMOTEID={ remote-id DEFAULT }] [MODE={ RELAY SNOOPING }] [FORCE]</pre>		<p>Configures the DHCP Relay parameters for the instance(s).</p> <p>The parameters are defined in the CREATE DHCPRelay command.</p>
<pre>RENAME DHCPRelay=dhcpname TO=dhcpname</pre>		<p>Renames an existing DHCP Relay instance. (All of the existing attributes are unchanged.) Note that no components need to be disabled to perform this.</p>
<pre>ENABLE DHCPRelay={ dhcpname- list MAIN ALL } INTERFACE={ type:id- range id-range ifname-list ALL }</pre>		<p>Enables the DHCPRelay instance(s) on the specified interface(s).</p> <p>The default instance is ALL.</p> <p><i>Note:</i> Ensure that DHCPRelay is enabled on the Network-side interfaces.</p>
<pre>SET DHCPRelay INTERFACE={ type:id- range id-range ifname-list ALL } [FILTER={ ON OFF }] [AGEING={ ON OFF }]</pre>		<p>Enables/disables the IP Filtering and Ageing function on the specified interface.</p> <p>When FILTER=ON, a maximum of eight (8) IP Filters will be applied to the interface, based on the number of learned MAC addresses that have been assigned IP addresses from DHCP Servers. IP Filtering is off by default for interfaces with DHCP Relay Agent enabled. Additionally, FILTER=ON has meaning only when DHCP Relay is enabled on the interface.</p> <p>The AGEING option applies only when FILTER=ON.</p> <p>See section 13.8.6 for more information on this IP filtering and ageing functionality.</p>

TABLE 13-9 DHCP Relay commands (Continued)

Syntax	Samples for DHCP Relay Agent Configuration	Description
<pre> CLEAR DHCPRELAY INTERFACE={ type:id- range id-range ifname-list ALL } { [IPADDRESS={ ipaddress-list ALL }] [MACADDRESS={ macaddress ALL }] } </pre>		<p>Allows the manual deletion of IP filters applied to subscriber interface. Note that there is no means of manually adding an IP filter to an interface (this is performed by DHCP Relay Agent).</p>
<pre> RESET DHCPRELAY COUNTER INTERFACE={ type:id- range id-range ifname-list ALL } </pre>	<pre> RESET DHCPRELAY COUNTER INTERFACE=ALL </pre>	<p>Resets the DHCP Relay packet counts on the specified interface. If an interface is not specified, the cumulative DHCP Relay counter will be reset, as well as all individual interface DHCP Relay counters.</p>
<pre> SHOW DHCPRELAY [INTERFACE={ type:id- range id-range ifname-list ALL }] </pre>	<pre> SHOW DHCPRELAY INTERFACE=ALL </pre>	<p>Displays information about the DHCP Relay Agent for one or more interfaces. If more than one interface is specified (or if no interface is specified) the following information will be displayed, which includes the list of DHCP Servers, cumulative DHCP Relay Statistics, and list of interfaces that are configured to run DHCP Relay.</p>
<pre> SHOW DHCPRELAY COUNTER [INTERFACE={ type:id- range id-range ifname-list ALL }] [FULL] </pre>		<p>Displays counter information about the DHCP Relay Agent for one or more interfaces.</p>

TABLE 13-9 DHCP Relay commands (Continued)

Syntax	Samples for DHCP Relay Agent Configuration	Description
<pre>DISABLE DHCPRELAY={ dhcpname- list MAIN ALL } INTERFACE={ type:id- range id-range ifname-list ALL }</pre>		Disables the DHCP Relay function on the specified interface
<pre>DELETE DHCPRELAY={ dhcpname- list MAIN ALL } SERVER={ ipaddress-list ALL } [FORCE]</pre>		Deletes for the DHCP Relay instance the Server IP addresses from the list. Note that the FORCE option can be used to override CLI command confirmation behavior if implemented.
<pre>DELETE DHCPRELAY={ dhcpname- list MAIN ALL } VLAN={ vlanname-list vid-range ALL }</pre>		Disassociates a DHCP Relay instance (or set of instances) with a VLAN (or set of VLANs).
<pre>DESTROY DHCPRELAY={ dhcpname- list ALL } [FORCE]</pre>		Destroys the DHCP Relay instance (or set of instances). FORCE will override a warning that the instance still has associations.

13.8.11 DHCP Relay Example

13.8.11.1 Overview

In an example configuration, there is the following for an EPSR topology:

- EPSR Master node : 172.16.66.158

- Interface 11.1 is Primary
- Interface 11.2 is Secondary
- VLAN 10 is the control VLAN
- 4 dhcp instances (main, voice, video, data)
- All 4 instances are in Relay mode
- EPSR Transit node: 172.16.66.167
 - 4 dhcp instances (main, voice, video, data)
 - All 4 instances are in Relay mode
 - 3 customer ports (with one STB each) – ports 16.4,16.5,19.6
- EPRS Transit node: 172.16.66.153
 - 4 dhcp instances (main, voice, video, data)
 - All 4 instances are in Snooping mode
 - 1 customer port with a PC – port 15.6

There are four VLANs that are used as part of DHCP Relay:

- 105 - data
- 310 - voice
- 512 - video
- 402 - inband mgmt

13.8.11.2 System 172.16.66.158 (EPSR Master Node)

Once the control VLAN (10) is created (with the Primary and Secondary interfaces), and the VLANs created, the DHCP instances are created and associated with the VLANs, as follows:

```
officer SEC>> create dhcp video
add dhcp video vlan 512
add dhcp video serv 192.168.201.64
set dhcp video agent remoteid D1
set dhcp video mode relay
```

```
create dhcp data
add dhcp data vlan 105
add dhcp data serv 172.16.65.38
set dhcp data agent remoteid H1
set dhcp data mode relay
```

```
create dhcp voice
add dhcp voice vlan voice
add dhcp voice serv 172.16.18.45
set dhcp voice agent remoteid MAP_P1
set dhcp voice mode relay
```

The inband management VLAN interface is also enabled, as follows:

```

////////////////////////////////////
//      Inband
////////////////////////////////////

add ip int vlan: 402.0 ipadd 172.16.66.158 sub 255.255.255.0 gate 172.16.66.1
en ip int vlan: 402.0

```

13.8.11.3 System 172.16.66.167 (EPSR Transit Node)

Once the control VLAN (10) is created (with the Primary and Secondary interfaces), and the VLANs created (with associations as type data and interfaces), the DHCP instances are created and associated with the VLANs, as follows:

```

officer SEC>> create dhcp video
add dhcp video vlan 512
add dhcp video serv 192.168.201.64
set dhcp video agent remoteid D1
set dhcp video mode relay

create dhcp data
add dhcp data vlan 105
add dhcp data serv 172.16.65.38
set dhcp data agent remoteid H1
set dhcp data mode relay

create dhcp voice
add dhcp voice vlan voice
add dhcp voice serv 172.16.18.45
set dhcp voice agent remoteid MAP_P1
set dhcp voice mode relay

```

The inband management VLAN interface is also enabled, as follows:

```

////////////////////////////////////
//      Inband
////////////////////////////////////

add ip int vlan: 402.0 ipadd 172.16.66.167 sub 255.255.255.0 gate 172.16.66.1
en ip int vlan: 402.0

```

13.8.11.4 System 172.16.66.153 (EPSR Transit Node in Snooping DHCP Mode)

Once the control VLAN (10) is created (with the Primary and Secondary interfaces), and the VLANs created (with associations as type data and interfaces), the DHCP instances are created and associated with the VLANs, as follows:

```

officer SEC>> create dhcp video
add dhcp video vlan 512
add dhcp video serv 192.168.201.64
set dhcp video agent remoteid D1
set dhcp video mode snoop

create dhcp data
add dhcp data vlan 105
add dhcp data serv 172.16.65.38
set dhcp data agent remoteid H1
set dhcp data mode snoop

```

```

create dhcp voice
add dhcp voice vlan voice
add dhcp voice serv 172.16.18.45
set dhcp voice agent remoteid MAP_P1
set dhcp voice mode snoop

```

The inband management VLAN interface is also enabled, as follows:

```

////////////////////////////////////
//      Inband
////////////////////////////////////

add ip int vlan:402.0 ipadd 172.16.66.153 sub 255.255.255.0 gate 172.16.66.1
en ip int vlan:402.0

```

13.8.11.5 System output - System 172.16.66.158

Once the system is up and running, the user can query the DHCP configuration as follows:

```
officer SEC>> SHOW DHCPRELAY ALL
```

```
--- DHCP Agent Summary -----
```

```

DHCP Cumulative
Statistics
-----
Discover..... 10
Offer..... 10
Request..... 21
Decline..... 0
Ack..... 20
Nack..... 1
Release..... 0
Inform..... 0
-----

```

```
DHCP Instance Information
-----
```

DHCP Instance Name	Mode	Remote ID	VLAN Vid List
MAIN	RELAY	00:0C:25:00:05:A6	None
video	RELAY	D1	512
data	RELAY	H1	105
voice	RELAY	MAP_P1	310

```
officer SEC>>SHOW DHCPRELAY INTERFACE=11.0
```

```
--- DHCP Interface Status -----
```

```
DHCP Interface Status (system wide)
-----
```

```

Interface..... 11.0
Circuit ID..... AUTOMATIC
Filtering..... OFF

```


Auto Age..... OFF

IP Address Allocations			Statistics	
Mac Address	Vi d	IP Address	DHCP Packet	Count
			Di scover.....	0
			Offer.....	10
			Request.....	0
			Decl i ne.....	0
			Ack.....	20
			Nack.....	1
			Rel ease.....	0
			Inform.....	0

DHCP Interface Control (per instance)

Instance	Status
MAI N	Enabl ed
vi deo	Enabl ed
data	Enabl ed
voi ce	Enabl ed

Info (010017): Operation Successful

offi cer SEC>>SHOW DHCPRELAY INTERFACE=11.1

--- DHCP Interface Status ---

DHCP Interface Status (system wide)

Interface..... 11.1
 Circuit ID..... AUTOMATIC
 Filtering..... OFF
 Auto Age..... OFF

IP Address Allocations			Statistics	
Mac Address	Vi d	IP Address	DHCP Packet	Count
			Di scover.....	10
			Offer.....	0
			Request.....	21
			Decl i ne.....	0
			Ack.....	0
			Nack.....	0
			Rel ease.....	0
			Inform.....	0

DHCP Interface Control (per instance)

Instance	Status
MAI N	Enabl ed
vi deo	Enabl ed
data	Enabl ed

```
voice                               Enabled
```

```
-----
Info (010017): Operation Successful
```

```
officer SEC>> SHOW DHCPRELAY INTERFACE=11.2
```

```
--- DHCP Interface Status -----
```

```
-----
DHCP Interface Status (system wide)
```

```
-----
Interface..... 11.2
Circuit ID..... AUTOMATIC
Filtering..... OFF
Auto Age..... OFF
```

```
-----
IP Address Allocations
```

```
-----
Statistics
```

Mac Address	Vid	IP Address	DHCP Packet	Count
			Discover.....	0
			Offer.....	0
			Request.....	0
			Decline.....	0
			Ack.....	0
			Nack.....	0
			Release.....	0
			Inform.....	0

```
-----
DHCP Interface Control (per instance)
```

Instance	Status
-----	-----
MAIN	Enabled
video	Enabled
data	Enabled
voice	Enabled

Note that the interfaces have Filtering and Auto Ageing OFF (the default), and that interface 11.2 has no statistics because the (data) VLANs are blocked as part of the EPSR topology.

13.8.11.6 System output - System 172.16.66.167

Once the system is up and running, the user can query the DHCP configuration as follows:

```
officer SEC>> SHOW DHCPRELAY
```

```
--- DHCP Agent Summary -----
```

```
-----
DHCP Cumulative
Statistics
-----
Discover..... 10
Offer..... 10
Request..... 21
Decline..... 0
Ack..... 20
```

```
Nack..... 1
Release..... 0
Inform..... 0
-----
```

DHCP Instance Information

DHCP Instance Name	Mode	Remote ID	VLAN Vid list
MAI N	RELAY	00:0C:25:00:FB:84	None
vi deo	RELAY	D2	512
data	RELAY	H2	105
voi ce	RELAY	MAP_P2	310

officer SEC>> SHOW DHCPRELAY INTERFACE 11.0

--- DHCP Interface Status -----

DHCP Interface Status (system wide)

```
Interface..... 11.0
Circuit ID..... AUTOMATIC
Filtering..... OFF
Auto Age..... OFF
```

IP Address Allocations			Statistics	
Mac Address	Vid	IP Address	DHCP Packet	Count
			Discover.....	4
			Offer.....	0
			Request.....	18
			Decline.....	0
			Ack.....	0
			Nack.....	0
			Release.....	0
			Inform.....	0

DHCP Interface Control (per instance)

Instance	Status
MAI N	Enabl ed
vi deo	Enabl ed
data	Enabl ed
voi ce	Enabl ed

Info (010017): Operation Successful

officer SEC>> SHOW DHCPRELAY INTERFACE 11.1

--- DHCP Interface Status -----

DHCP Interface Status (system wide)

```
Interface..... 11.1
Circuit ID..... AUTOMATIC
Filtering..... OFF
Auto Age..... OFF
```

IP Address Allocations

```
-----
Mac Address      Vid      IP Address
-----
```

Statistics

```
-----
DHCP Packet Count
-----
Discover..... 0
Offer..... 10
Request..... 0
Decline..... 0
Ack..... 20
Nack..... 1
Release..... 0
Inform..... 0
```

DHCP Interface Control (per instance)

```
-----
Instance          Status
-----
MAIN              Enabled
video             Enabled
data              Enabled
voice             Enabled
```

Info (010017): Operation Successful**officer SEC>> SHOW DHCPRELAY INTERFACE 16.4**--- DHCP Interface Status ---

DHCP Interface Status (system wide)

```
Interface..... 16.4
Circuit ID..... AUTOMATIC
Filtering..... OFF
Auto Age..... ON
```

IP Address Allocations

```
-----
Mac Address      Vid      IP Address
-----
00:02:02:00:AB:9C  512     192.168.69.237
```

Statistics

```
-----
DHCP Packet Count
-----
Discover..... 2
Offer..... 2
Request..... 1
Decline..... 0
Ack..... 1
Nack..... 0
Release..... 0
Inform..... 0
```

DHCP Interface Control (per instance)

Instance	Status
MAI N	Enabl ed
vi deo	Enabl ed
data	Enabl ed
voi ce	Enabl ed

Info (010017): Operation Successful

offi cer SEC>>SHOW DHCPRELAY INTERFACE 16.5

--- DHCP Interface Status ---

DHCP Interface Status (system wide)

Interface..... 16.5
 Circuit ID..... AUTOMATIC
 Filtering..... OFF
 Auto Age..... ON

IP Address Allocations

Mac Address	Vi d	IP Address
00: 02: 02: 00: AB: 84	512	192. 168. 69. 234

Statistics

DHCP Packet	Count
Di scover.....	2
Offer.....	2
Request.....	1
Decl ine.....	0
Ack.....	1
Nack.....	0
Rel ease.....	0
Inform.....	0

DHCP Interface Control (per instance)

Instance	Status
MAI N	Enabl ed
vi deo	Enabl ed
data	Enabl ed
voi ce	Enabl ed

offi cer SEC>>SHOW DHCPRELAY INTERFACE 19.6

--- DHCP Interface Status ---

DHCP Interface Status (system wide)

Interface..... 19.6
 Circuit ID..... AUTOMATIC

```

Filtering..... OFF
Auto Age..... ON

```

IP Address Allocations			Statistics	
Mac Address	Vid	IP Address	DHCP Packet	Count
00:02:02:00:AB:16	512	192.168.69.251	Discover.....	2
			Offer.....	2
			Request.....	1
			Decline.....	0
			Ack.....	1
			Nack.....	0
			Release.....	0
			Inform.....	0

DHCP Interface Control (per instance)

Instance	Status
MAI N	Enabl ed
vi deo	Enabl ed
data	Enabl ed
voi ce	Enabl ed

13.8.11.7 System output - System 172.16.66.153

Once the system is up and running, the user can query the DHCP configuration as follows:

```
officer SEC>> SHOW DHCPRELAY
```

```
<cr>
dhcpname-list|MAIN|ALL - Show DHCP Relay
INTERFACE -
```

```
officer SEC>> SHOW DHCPRELAY
```

```
--- DHCP Agent Summary ---
DHCP Cumulative
Statistics
-----
Discover..... 2
Offer..... 1
Request..... 8
Decline..... 0
Ack..... 3
Nack..... 1
Release..... 0
Inform..... 0
-----
```

DHCP Instance Information

DHCP Instance Name	Mode	Remote ID	VLAN Vid list
MAI N	RELAY	00:0C:25:00:FA:B8	None
vi deo	SNOOP	D3	512
data	SNOOP	H3	105

voice SNOOP MAP_P3 310

officer SEC>> SHOW DHCPRELAY INTERFACE=10.0

--- DHCP Interface Status -----

DHCP Interface Status (system wide)

```
Interface..... 10.0
Circuit ID..... AUTOMATIC
Filtering..... OFF
Auto Age..... OFF
```

IP Address Allocations

Statistics

Mac Address	Vid	IP Address	DHCP Packet	Count
			Discover.....	0
			Offer.....	0
			Request.....	0
			Decline.....	0
			Ack.....	0
			Nack.....	0
			Release.....	0
			Inform.....	0

DHCP Interface Control (per instance)

Instance	Status
MAIN	Enabled
video	Enabled
data	Enabled
voice	Enabled

Info (010017): Operation Successful

officer SEC>> SHOW DHCPRELAY INTERFACE=11.0

--- DHCP Interface Status -----

DHCP Interface Status (system wide)

```
Interface..... 11.0
Circuit ID..... AUTOMATIC
Filtering..... OFF
Auto Age..... OFF
```

IP Address Allocations

Statistics

Mac Address	Vid	IP Address	DHCP Packet	Count
			Discover.....	1
			Offer.....	1
			Request.....	4

```
Decline..... 0
Ack..... 3
Nack..... 1
Release..... 0
Inform..... 0
```

DHCP Interface Control (per instance)

```
-----
Instance                               Status
-----
MAIN                                    Enabled
video                                   Enabled
data                                     Enabled
voice                                    Enabled
```

Info (010017): Operation Successful

officer SEC>>SHOW DHCPRELAY INTERFACE=15.6

--- DHCP Interface Status ---

DHCP Interface Status (system wide)

```
-----
Interface..... 15.6
Circuit ID..... AUTOMATIC
Filtering..... OFF
Auto Age..... ON
```

IP Address Allocations

Statistics

Mac Address	Vid	IP Address	DHCP Packet	Count
00:06:5B:D8:72:FF	512	192.168.69.19	Discover.....	1
			Offer.....	1
			Request.....	4
			Decline.....	0
			Ack.....	3
			Nack.....	1
			Release.....	0
			Inform.....	0

DHCP Interface Control (per instance)

```
-----
Instance                               Status
-----
MAIN                                    Enabled
video                                   Enabled
data                                     Enabled
voice                                    Enabled
```


13.9 BPDU Cop

Some devices (i.e. switches), when connected to an interface, could have a loop implemented on them which could leave the fMAP connected to a source of a broadcast storm. These devices could forward BPDUs, and when a loop was implemented on that device, it could forward these BPDUs back to the fMAP over the same interface.

This has been fixed with a new feature called BPDU Cop. The parameter (BPDUCOP={ ON | OFF }) is part of SET STP and is set against the specific interface(s).

When a BPDU is received and the feature is activated, the operational state goes to DOWN and the state changes to AUTO-DISABLED, as shown in the following logs:

```
PORT007 2006-01-06 16:04:08 6594 INFO
```

```
Location: Slot: 4 Port: 0
Description: Port state change
From: UP-UP-Online To: UP-DOWN-Failed
```

```
* PORT003 2006-01-06 16:04:08 6593 FAULT
Location: Slot: 4 Port: 0
Description: Port Fault Set
Reason Code: Unexpected BPDU Received
```

The SHOW INTERFACE command gives details of the fault, as follows:

```
officer SEC>> sh int 4.0

--- FE Interfaces ---

Interface..... 4.0
Type..... FE
State..... UP-DN-AutoDisabled
Description..... video

Interface Faults

Unexpected BPDU Received..... Minor

Provisioning

Provisioning Profile..... AutoProv
Direction..... Customer
Auto Negotiation..... On
Speed..... Auto
Duplex..... Auto
Flow Control..... Auto
Remote Monitoring..... On

Actual

Direction..... Customer
Physical Address..... 00:0C:25:00:05:BF
```

VLAN Information

```

Acceptable Frame Types..... All
Ingress Filtering..... On
TPID..... 0x8100
TAGALL..... Off
Dynamic MAC Learning Limit..... Off
Untagged VLAN..... 512
    
```

Packet Statistics

	Input	Output
Octets.....	8448	229303
Unicast Packets.....	0	0
Discarded Packets.....	0	0
Errored Packets.....	0	0
Unknown Proto Packets.....	0	N/A

Note that once the interface is set to AutoDisabled by the system, the user must disable and enable the interface to clear the alarm and bring the interface back into service. In other words, **once an interface is placed out of service by the BPDU Cop feature, it will stay disabled unless it is explicitly brought back up by the user.**

The following shows the sequence when the user disables and then re-enabled the interface, bringing the interface back into service:

```

officer SEC>> disable int 4.0 force

  CLI002 2006-01-06 16:07:54 6605 INFO

User: "officer" on system console entered CLI command:

DISABLE INTERFACE=4.0

* PORT004 2006-01-06 16:07:54 6606 FAULT
Location: Slot: 4 Port: 0
Description: Port Fault Cleared
Reason Code: Unexpected BPDU Received

  PORT007 2006-01-06 16:07:54 6607 INFO
Location: Slot: 4 Port: 0
Description: Port state change
From: UP-DOWN-Failed To: DOWN-DOWN-Terminating

  PORT007 2006-01-06 16:07:54 6608 INFO
Location: Slot: 4 Port: 0
Description: Port state change
From: DOWN-DOWN-Terminating To: DOWN-DOWN-Offline

Info (039512): Operation Successful (FE10 Slot 4 Port 0)

officer SEC>>
officer SEC>> enable int 4.0

enable int 4.0

  CLI002 2006-01-06 16:08:01 6609 INFO
    
```

User: "officer" on system console entered CLI command:

```
ENABLE INTERFACE=4.0
```

```
PORT007 2006-01-06 16:08:01 6610 INFO
Location: Slot: 4 Port: 0
Description: Port state change
From: DOWN-DOWN-Offline To: UP-DOWN-Configuring
```

```
Info (039512): Operation Successful (FE10 Slot 4 Port 0)
```

```
PORT003 2006-01-06 16:08:01 6611 FAULT
Location: Slot: 4 Port: 0
Description: Port Fault Set
Reason Code: Loss of Link
```

```
PORT007 2006-01-06 16:08:01 6612 INFO
Location: Slot: 4 Port: 0
Description: Port state change
From: UP-DOWN-Configuring To: UP-DOWN-Failed
```

```
officer SEC>> PORT004 2006-01-06 16:08:03 6613 FAULT
Location: Slot: 4 Port: 0
```

```
Description: Port Fault Cleared
Reason Code: Loss of Link
```

```
PORT007 2006-01-06 16:08:03 6614 INFO
Location: Slot: 4 Port: 0
Description: Port state change
From: UP-DOWN-Failed To: UP-UP-Online
```

13.10 Link Layer Discovery Protocol (LLDP)

13.10.1 Overview

A network topology, as explained throughout this section, is a set of Network Elements (NEs) and the network links that interconnect them.

The IEEE (Institute of Electrical and Electronic Engineers) defined Link Layer Discovery Protocol (LLDP) standard 802.1ab is an application protocol that runs directly over layer 2 in network elements (NEs) to facilitate a centrally located network manager to derive the physical network topology the NEs are part of.

Note: The LLDP runs over network links, and not over links that connect the NE to the customer endpoint. This is covered by another protocol, LLDP Media Endpoint Discovery, or LLDP-MED), which is not part of this feature.

To assist in the discovery of the network topology, LLDP has two components:

- **Transmitting** - This transmits or advertises the local NE discovery-related data on a per link basis to the remote NE at the other end of the link. Since a network link is spanned by two physical ports on each of the

two NEs (local and remote), the NEs are advertising their local network port related discovery data to each other.

- **Receiving** - This collects the advertised data received over the network link from the transmitting NE, resulting in both the local NE and the remote NE having the port discovery data at each end of the network link.

Using both LLDP protocol components allows the NE to possess discovery-related data for each of the network links between itself and the neighboring NE. This data can then be used to deduce the physical topology of the network which the NEs and the network (port) links are part of.

Note that the responsibility of the LLDP protocol lies in only with providing the protocol message formats, the type of discovery data to be advertised and the procedures required to send and receive the above mentioned data on a per link basis. The transmission of the data (collected via the LLDP protocol) however to a central location like the Network Management System (NMS) and how it uses the data to deduce the network topology is outside of the scope of the LLDP protocol. The LLDP does not favor one method over another in the method of transmitting the data to the NMS, although Simple Network Management Protocol (SNMP) is the most common.

An SNMP based interaction between an NMS and the network agent in the NE is shown below.

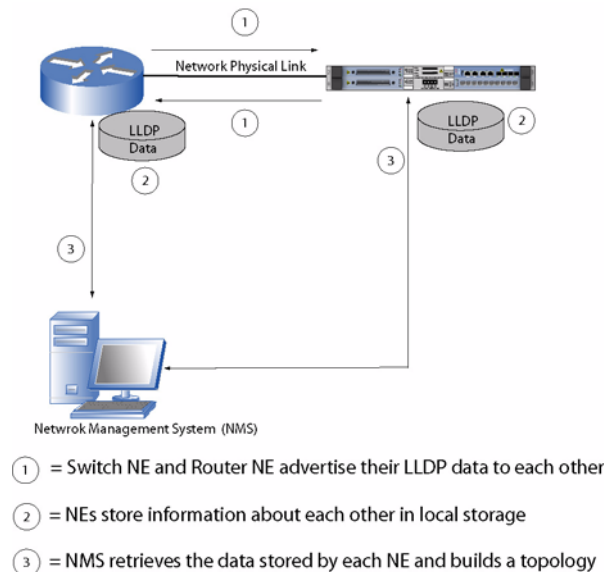


FIGURE 13-21 Main LLDP Components

The LLDP protocol is restricted to a single Local Area Network (LAN). LLDP is prohibited from relaying information from one port to another port on one NE device and therefore prohibited from relaying information through a bridged network.

Note that the main purpose of the LLDP protocol is to **advertise** data that is useful for discovering pertinent information about a network link port. It is not intended to be used as a configuration protocol in a NE. In other words, the LLDP protocol is not intended to be used to configure the receiving port in a NE based upon the

advertised information for the corresponding peer port received from the neighboring NE via the network link spanning them.

During the operation of the LLDP protocol, the advertised data could also be used by NMS to discover port data configuration inconsistencies at the local NE and remote NE ends, which results in communication failure between the two NEs. Note however that the LLDP protocol does not provide any mechanisms or procedures to correct these inconsistencies. It is left to the network provider to correct the inconsistencies discovered via the LLDP protocol with the least possible service interruption.

13.10.2 Concepts

The following terms are used when describing the LLDP protocol:

- **LLDP-Aware NE** - A network element that has implemented the LLDP protocol.
- **Non-LLDP NE** - A network element that is not capable of executing the LLDP protocol procedures.
- **Advertisement** - An LLDP protocol data unit (LLDPDU).
- **MSAP Identifier** - The identification of the local network port in an outgoing LLDPDU.
- **LLDP Neighbor** - An LLDP aware NE that is either directly connected to another LLDP aware NE or connected to an LLDP aware NE via a non-LLDP NE such as a hub.

The LLDP protocol is a **one-way protocol**. An LLDP-aware NE can advertise information associated with an MSAP identifier and also receive advertise data associated with the peer MSAP identifier from the neighboring LLDP aware NE.

It is important to note however that the LLDP protocol does not provide a means to solicit the advertisement data from the neighboring NE for an MSAP identifier.

Since the receiving of advertisements for an MSAP identifier by an NE is unsolicited, **the LLDP transmitting component and receiving component can be enabled and disabled independent from each other**. This is reflected in the command parameters as explained below.

13.10.3 Protocol Configuration Parameters

The following tables list the LLDP protocol procedure configuration options for an LLDP aware NE. They are divided into two categories:

1. System wide protocol configuration options that are applicable to the entire NE as a whole.
2. Protocol configuration options that are applicable on a per network port (interface) basis.
 - Two parameters (MODE and NOTIFY) are always defined for the interface and have defaults
 - Other parameters (defined as OPTIONS) are optional and are not required.

The following tables list and define these parameters.

- [Table 13-10](#) lists the parameters that are set system wide and have default values that can be changed or set back to their default values.

- [Table 13-11](#) lists the interface parameters that are set for interface(s), and have default values that can be changed.
- [Table 13-12](#) lists the optional parameters that define what is to be included in the outgoing LLDPDU .

TABLE 13-10 System-Wide Protocol Configuration Options for LLDP

Parameter	Range	Description
TXINTERVAL	5..32768 (30) (in seconds)	Determines the protocol transmit periodic timer value and controls the sending of the LLDP message by the transmitting component for the applicable network ports.
TXHOLD	2..10 (4) (multiplying factor)	Used as a multiplying factor for the TXINTERVAL parameter above to compute the value of the transmission component TTL TLV discussed earlier. The default value of this parameter is 4 and the valid range is defined from a minimum of 2 to a maximum of 10.
TXDELAY	1..8192 (2) (in seconds)	Determines the delay in seconds before successive LLDP messages for a MSAP identifier are transmitted due to a change in value in its advertised data. The relationship between the TXDELAY parameter and the TXINTERVAL parameter is : $TXINTERVAL \geq (4 * TXDELAY)$
REINITDELAY	1..10 (2) (in seconds)	Determines the time in seconds for when a MSAP identifier state is allowed to go from a LLDP disabled state to a transmit enabled state again. It is not applicable the first time when the MSAP identifier goes from a disabled state to a transmit enable state.
NOTIFYINTERVAL	1-3600 (5) (in seconds)	Determines the delay in seconds before successive unsolicited notifications are sent for to the NMS for a network port whose NMS notification status option (NOTIFY) is enabled. The use of this parameter results in only one trap being sent to the NMS even if multiple traps are generated within the NMS notification interval for a network port whose NMS notification status parameter is enabled.

TABLE 13-11 Port (Interface) Protocol Configuration Required Parameters for LLDPDU

Parameter	Range	Description
MODE		Controls the transmission or reception of LLDP packets. Options are below. Note: In order for these option values (except OFF) to have any affect for a network port the LLDP protocol functionality must be enabled on a system-wide basis.
	TX	Allows the transmission of LLDP messages to the neighboring NE over the local network port. Any received LLDP messages from the neighboring NE over this network port is dropped. Use of this setting prevents an NE from discovering its neighbors connected to this network port
	RX	Allows the receipt of LLDP messages from neighboring NEs connected to this network port. LLDP messages are not transmitted to neighboring NEs connected to this network port. Use of this setting prevents a neighboring NE connected to this network port from discovering this NE.
	BOTH	Allows the transmission and reception of LLDP messages to and from the neighboring NE connected to this network port. If the same setting is configured on either end of the link spanning the network ports, then it allows both the NEs to discover each other.
	OFF	Disallows the transmission and reception of LLDP messages to and from the neighboring NE connected to this network port. Use of this setting prevents this NE from discovering the neighboring NE connected to this network port and also prevents the neighboring NE from discovering this NE.
NOTIFY	ON OFF	In order for the above option values to have any affect for a network port the LLDP protocol functionality must be enabled on a system-wide basis.

TABLE 13-12 Port (Interface) Protocol Configuration Options for an Outgoing LLDPDU

Parameter	Meaning	Fields - Description
PORTDESC	Port Description	The configured value describing the port (The “Description” field in the <code>show int . . .</code> CLI command).
SYSNAME	System Name	The configured value assigned to the system (The “Host” field in the <code>show system</code> CLI command).
SYSDESC	System Description	The configured value describing the system (The “Location” field in the <code>show system</code> CLI command)
SYSCAP	System Capabilities	4 (=Bridge Capability Supported)
PORTVLAN	Port VLAN Identifier (PVID)	The default VID value associated with the network port for received untagged frames if one exists else it is set to zero (0).
VLANNAME	VLAN Name	VLAN ID (VID) - The default VID value associated with the network port for received untagged frames if one exists else it is set to zero (0). VLAN Name Length - The VLAN name length in bytes VLAN NAME - The configured VLAN Name
PROTOVLAN	Port and Protocol VLAN ID	Flag - Bit 1: - 0 (Port and protocol VLAN not supported) Flag - Bit 2: - 0 (Port and protocol VLAN not enabled) Flags - Bits 3 through 8: - 0 (Reserved for future use) Port and Protocol VLAN ID (PPVID): - 0 (Valid value when flags above indicates port and protocol VLAN is not supported and not enabled).

TABLE 13-12 Port (Interface) Protocol Configuration Options for an Outgoing LLDPDU

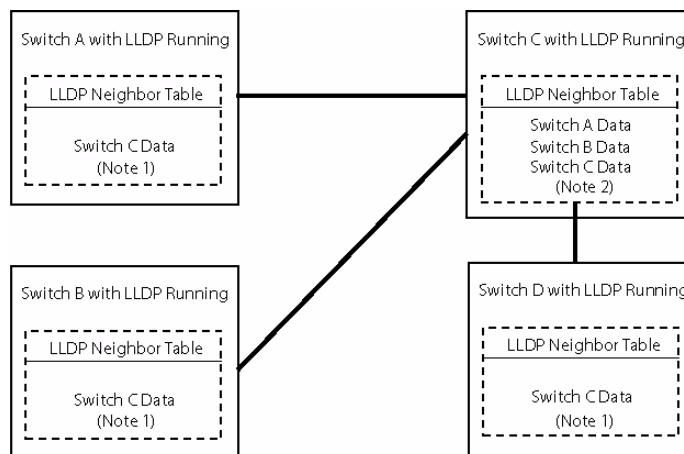
Parameter	Meaning	Fields - Description
PROTOCOL	Protocol Identity	<p>- Protocol Identity Length - The length of the Protocol Identity field below in bytes.</p> <p>- The number of bytes (indicated by the protocol identity length field above) whose values identify the protocol that is being advertised.</p> <p>For example, when advertising the STP protocol the TLV field values are:</p> <ul style="list-style-type: none"> - Protocol Identity Length: 8 bytes - Protocol Identity: <ul style="list-style-type: none"> - Bytes 0 and 1 – xSTP length field - Bytes 2, 3 and 4 – LLC header bytes - Bytes 5 and 6 – xSTP Protocol ID - Byte 7 – xSTP Protocol Version
MACPHYCONFIGSTATUS	Auto-negotiation Support/Status	<p>Bit 0: 0 (If port auto-neg not supported) 1 (If port auto-neg is supported)</p> <p>Bit 1: 0 (If port auto-neg is not enabled) 1 (If port auto-neg is enabled)</p> <p>Bits 2 to 7: 0 (Reserved for future use)</p>
POWERVIAMDI	MDI Power Support	<p>Bit 0: 1 (PSE: Powered Sourcing Equip.)</p> <p>Bit 1: 0 (Not Supported)</p> <p>Bit 2: 0 (Disabled)</p> <p>Bit 3: 0 (Pair selection cannot be controlled)</p> <p>Bits 4 through &: 0 (Reserved for future use)</p>

TABLE 13-12 Port (Interface) Protocol Configuration Options for an Outgoing LLDPDU

Parameter	Meaning	Fields - Description
LINKAGGREGATION	Aggregation Status	Bit 0: 0 (If port cannot be aggregated) 1 (If port can be aggregated) Bit 1: 0 (If currently not in aggregation) 1 (If currently in aggregation) Bits 2 through 7: 0 (Reserved for future use)
MAXFRAMESIZE	Maximum 802.3 Frame Size	1522 – Indicates the support of both tagged and untagged frames.
EPSR		
UCP		
ALL		

13.10.4 Functional Overview

The following figures show the general, outgoing, and ingoing functions of LLDP.



Note 1 - The neighbor tables in Switch A, B, and D contain information on Switch C only, since Switch C is the only neighbor for these switches.

Note 2 - The neighbor table for Switch C contains information on switches A, B, and D, since all three switches are neighbors of Switch C.

FIGURE 13-22 General Mode for LLDP Operation

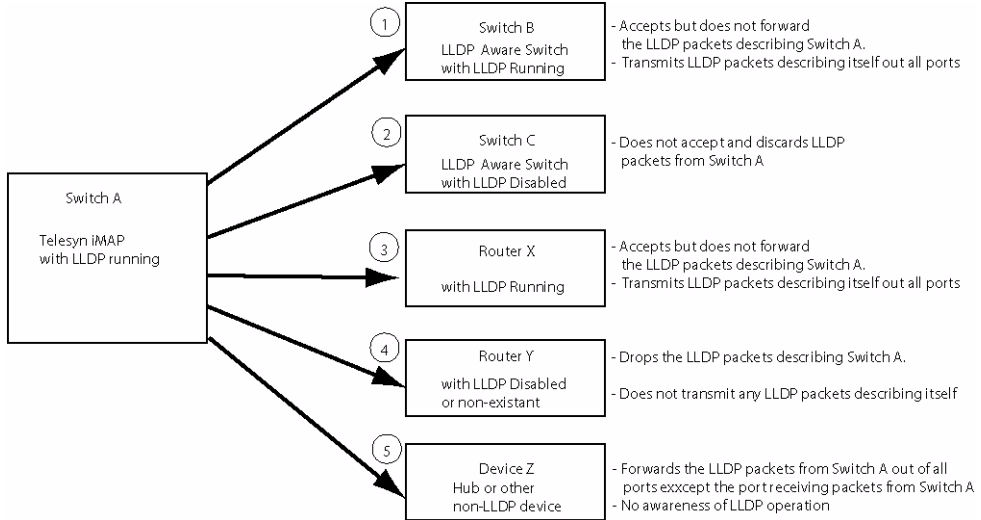


FIGURE 13-23 Processing of Outgoing LLDP Packets

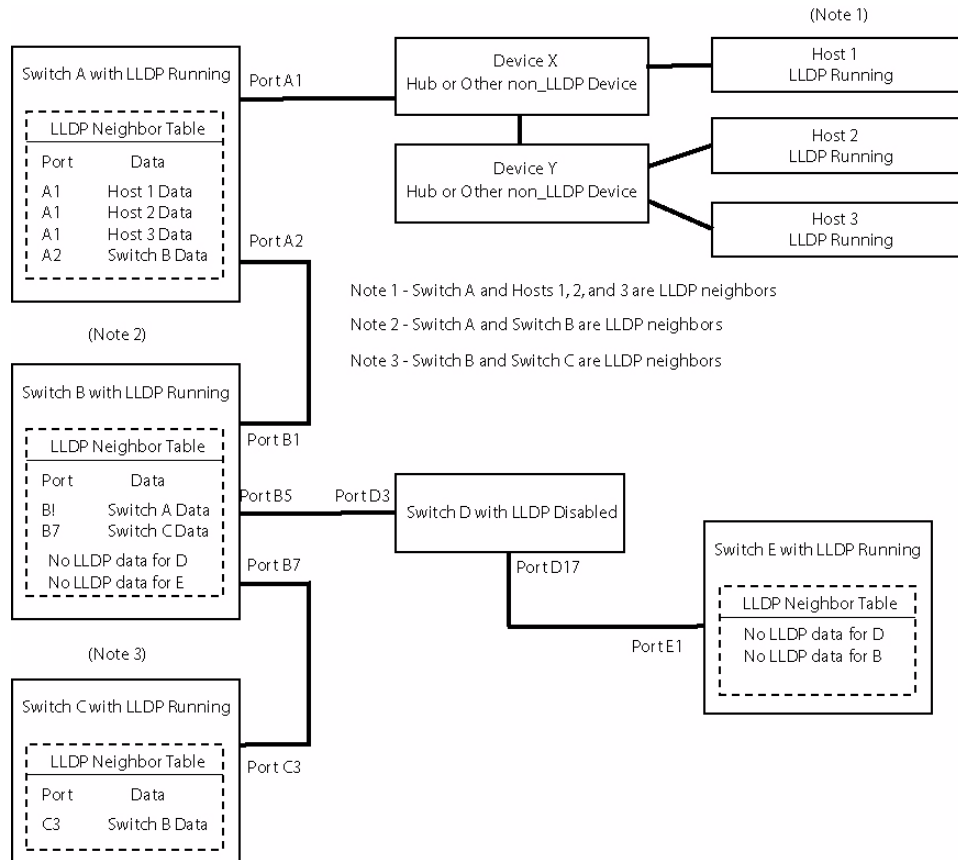


FIGURE 13-24 Processing of Incoming LLDP Packets

In [Figure 13-24](#), the following are neighbors:

- Switch A and hosts 1, 2, and 3
- Switch A and Switch B
- Switch B and Switch C

Switch D is receiving LLDP packets describing switches B and E. However since LLDP is disabled in switch D, it does not accept these LLDP packet and so has no LLDP awareness of either switch B or switch E. Switch D also does not transmit any LLDP packets describing itself to switch B and switch E. As a result, Switch B and E are not LLDP aware of D. Note that switch D does not forward the LLDP packets received from B onto E and vice-versa. So switch B does not know about E and switch E does not know about B.

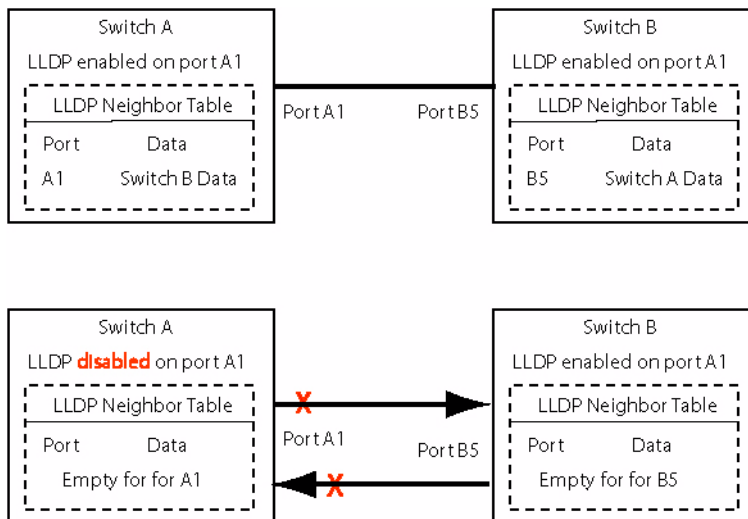


FIGURE 13-25 Effect of Disabling LLDP on a Single Interface

The top of [Figure 13-25](#) shows that with LLDP enabled and running on both switches for both directions, both switches discover each other and populate their tables with the neighbor's data for that interface. The bottom of the figure shows that with LLDP disabled on switch A, the following occurs:

- Switch A stops transmitting LLDP packets describing itself to Switch B, so the data for Switch A ages out in the Switch B table. Switch B is now not LLDP aware of Switch A anymore.
- Switch A does not accept and discards the packets that are sent by Switch B, so Switch A is not LLDP aware of Switch B.

13.10.5 Command Set for LLDP

TABLE 13-13 Commands for Link Layer Discovery Protocol (LLDP)

Object	Verb	Syntax	Description
LLDP INTERFACE	SET	<pre> SET LLDP INTERFACE={ type:id-range id-range ifname-list ALL } [MODE={ TX RX BOTH OFF }] [NOTIFY={ ON OFF }] </pre>	<p>Used to enable or disable LLDP for the interface(s).</p> <p>To enable, set the MODE command to TX, RX, or BOTH. (If the MODE command is not entered, the default is BOTH).</p> <p>To disable, set the MODE to OFF.</p> <p>NOTIFY controls whether traps are sent to an NMS if there is a change in the link set.</p>
	SETDEFAULTS	<pre> SETDEFAULTS LLDP INTERFACE={ type:id-range id-range ifname-list ALL } [MODE] [NOTIFY] </pre>	Controls the default settings for LLDP for the specified interface(s).
LLDP INTERFACE OPTIONS	ADD	<pre> ADD LLDP INTERFACE={ type:id-range id-range ifname-list ALL } OPTIONS [PORTDESC] [SYSNAME] [SYSDESC] [SYSCAP] [PORTVLAN] [VLANNAME] [PROTOVLAN] [PROTOCOL] [MACPHYCONFIGSTATUS] [POWERVIAMDI] [LINKAGGREGATION] [MAXFRAMESIZE] [EPSR] [UCP] [ALL] </pre>	<p>Adds one or more OPTIONS to the Interface(s) for LLDP. Note that this command does not enable LDP (that is the SET LLDP INTERFACE MODE command), but adds values for these optional parameters.</p> <p>Refer to Table 13-12 for a description of these OPTIONS.</p>

TABLE 13-13 Commands for Link Layer Discovery Protocol (LLDP)

Object	Verb	Syntax	Description
LLDP INTERFACE OPTIONS	DELETE	<pre>DELETE LLDP INTERFACE={ type:id-range id-range ifname-list ALL } OPTIONS [PORTDESC] [SYSNAME] [SYSDESC] [SYSCAP] [PORTVLAN] [VLANNAME] [PROTOVLAN] [PROTOCOL] [MACPHYCONFIGSTATUS] [POWERVIAMDI] [LINKAGGREGATION] [MAXFRAMESIZE] [EPSR] [UCP] [ALL]</pre>	<p>Deletes one or more OPTIONS to the Interface(s) for LLDP. Note that this command does not disable LDP (that is the SET LLDP INTERFACE MODE command), but adds values for these optional parameters.</p> <p>Refer to Table 13-12 for a description of these OPTIONS.</p>
LLDP	SET	<pre>SET LLDP [TXINTERVAL=5..32768] [TXHOLD=2..10] [TXDELAY=1..8192] [REINITDELAY=1..10] [NOTIFYINTERVAL=5..3600]</pre>	Sets the global LLDP values
	SETDE- FAULTS	<pre>SETDEFAULTS LLDP [TXINTERVAL] [TXHOLD] [TXDELAY] [REINITDELAY] [NOTIFYINTERVAL]</pre>	Sets the defaults for the global LLDP values. This command would be used to change back to the default values that had been changed by the SET LLDP command.
	SHOW	<pre>SHOW LLDP [INTERFACE { = { type:id-range id-range ifname-list ALL }] [FULL]]</pre>	Shows the attributes of how LLDP has been set on the system.

TABLE 13-13 Commands for Link Layer Discovery Protocol (LLDP)

Object	Verb	Syntax	Description
LLDP COUNTER	SHOW	<pre> SHOW LLDP COUNTER [INTERFACE={ type:id-range id-range ifname-list ALL }] [FULL] </pre>	<p>Allows the user to view the runtime data (local system LLDP counters) for each interface specified. Packets counted include frames as well as Type-Length-Value (TLV) information elements, which are variable in length.</p> <p>Counters include:</p> <p>Receive Frames:</p> <ul style="list-style-type: none"> - Total, with Errors, Discarded <p>Received TLVs:</p> <ul style="list-style-type: none"> - Discarded, Unrecognized <p>Transmitted Frames:</p> <ul style="list-style-type: none"> - Total <p>If the INTERFACE parameter is not provided, then all interfaces are displayed. The runtime information is displayed in tabular format, unless the FULL parameter is provided.</p>
	RESET	<pre> RESET LLDP COUNTER [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	<p>Allows the user to clear the LLDP interface counters. data (local system LLDP counters) for each interface specified.</p> <p>If the INTERFACE parameter is not provided, then all interfaces are reset.</p>

13.11 Layer 1 Protection Groups

13.11.1 Overview

When fMAP users configure connections to other network elements, such as General Bandwidth's G6, Layer 1 protection groups must be configured and provisioned for the links to these network elements. The following commands are used to implement layer 1 protection groups.

CREATE PROTECTIONGROUP=groupname

Creates a protections group and its name.

ADD PROTECTIONGROUP=groupname INTERFACE={type:id-range|ifname-list}

Adds an interface to a protection group.

```
DELETE PROTECTIONGROUP=groupname INTERFACE={type:id-range|ifname-list|ALL}
```

Removes an interface from a protection group.

```
DESTROY PROTECTIONGROUP={groupname-list|ALL}
```

Deletes a protection group.

```
SHOW PROTECTIONGROUP={groupname-list|ALL}
```

Displays a protection group.

Another command important to the provisioning of Layer 1 protection groups is the SET VLAN command:

```
SET VLAN={ vlnname | vid } INTERFACE={ type:id-range | id-range | ifname-list | ALL } [ FRAME={  
UNTAGGED | TAGGED } ] [ TRANSLATE={ 1..4094 | NONE } ] [ FORWARDING={PRIMARYUPSTREAM |  
SECONDARYUPSTREAM | DOWNSTREAM | STP | UCP | EPSR }
```

This command is used to set the FORWARDING mode on VLANs to PRIMARYUPSTREAM and SECONDARYUPSTREAM.

13.11.2 Usage Notes

- Only GE NM interfaces are allowed to be put into a protection group.
- Destroying a protection group will automatically delete all interfaces from it. The user will receive a warning at the CLI stating that all interfaces will be deleted when destroying a protection group.
- Deleting interfaces from a protection group or destroying the protection group will raise a warning at the CLI if the interfaces are port of a UFO-based VLAN that has FORWARDING mode set to “PRIMARYUPSTREAM”.

Note: When adding an interface to a UFO based VLAN, if the forwarding mode is designated as “PRIMARYUPSTREAM” for a connection for a member of a Protection Group, a warning is generated by the system to ensure that the interface is a member of a Layer 1 protection group.

13.11.3 State and Alarm Handling Rules for Layer 1 Protection Group

- No alarms are raised on interfaces in a protection group; instead alarms are raised against the protection group itself when failure conditions occur. Fault conditions on the interfaces (such as LOS) are masked but not alarmed, and can be seen using the SHOW PORT or SHOW ALARMS FULL commands.
- As long as the status of one interface in the protection group is UP-UP-ONLINE, all other interfaces in the protection group with status that are not UP-UP-ONLINE are marked as UP-DOWN-STANDBY.
- The normal interface statuses for the protection group is that one interface is UP-UP-ONLINE and the other interface is UP-DOWN-STANDBY. In this case there are no alarms reported against the protection group.
- If both links to the network element are down, the status of both interfaces will be UP-DOWN-FAILED and the protection group itself raises an “All links down” critical alarm.

- If a failover occurs, the former ONLINE interface goes UP-DOWN-STANDBY, the other interface goes UP-UP-ONLINE and the protection group raises a “Failover” minor alarm. This alarm is intended to prompt the user to investigate the fault condition on the failed link. Once the interface is fixed, the General Bandwidth G6 will not automatically initiate a switchover back to the normal condition. The secondary interface status (from the General Bandwidth G6 perspective) remains UP-UP-ONLINE, while the primary (from the General Bandwidth G6 perspective) status remains UP-DOWN-STANDBY and the “Failover” alarm remains active until manually cleared.
- If a failover occurs while the “Failover” alarm is present from a previous failover, the former ONLINE interface changes status to UP-DOWN-STANDBY, the other interface changes status to UP-UP-ONLINE (assuming it is capable), and the protection group retains a “Failover” major alarm. The “Failover” alarm remains active until manually cleared.
- Logs are generated each time interface configuration changes or alarms occur on the protection group.

The user can manually clear the “Failover” alarm on the protection group using the following command:

```
CLEAR ALARMS PROTECTIONGROUP={groupname-list|ALL}
```

The user would manually use this command to clear alarms on a protection group *after* clearing the fault that caused the failover.

An example follows:

Add the protection group:

```
officer SEC>> CREATE PROTECTIONGROUP=potsprotgrp1
Info (010017): Operation Successful
```

Add Ethernet interface 10.0 to the protection group:

```
officer SEC>> ADD PROTECTIONGROUP=potsprotgrp1 INTERFACE=ETH: 10.0
Info (010017): Operation Successful
```

Note: When adding an interface to a UFO based VLAN, if the forwarding mode is designated as "PROTECTIONLINK", a warning is generated by the system to ensure that the interface is a member of a Layer 1 protection group.

Display the protection group:

```
officer SEC>> SHOW PROTECTIONGROUP
```

```
-----
Group Name                               Interface
-----
potsprotgrp1                             10.0
-----
```

Delete the interface from the protection group:

```
officer SEC>> DELETE PROTECTIONGROUP=potsprotgrp1 INTERFACE=ETH: 10.0
The interface has been removed.
Info (010017): Operation Successful
```

Display the protection group:

```
officer SEC>> SHOW PROTECTIONGROUP
```

```
-----
```

Group Name	Interface
-----	-----
potsprotgrp1	
-----	-----

Destroy the protection group:

```
officer SEC>> DESTROY PROTECTI ONGROUP=potsprotgrp1
Info (010017): Operati on Successful
```

Verify that the group was destroyed:

```
officer SEC>> SHOW PROTECTI ONGROUP
No protecti on groups provi si oned.
```


14. IGMP

14.1 Internet Group Management Protocol (IGMP) Snooping

14.1.1 Overview of IGMP

The Internet Group Management Protocol (IGMP) is used to dynamically register (join or leave) individual hosts in a multicast group on a particular LAN with a **multicast router**. Hosts join group memberships by sending IGMP report messages to the multicast group and leave group memberships by sending leave messages to their local IP multicast router. The IP multicast router listens to the IGMP messages in the group to determine which multicast groups are active (meaning there is at least one or more host members). Also, the router periodically sends out IGMP queries to discover which of the multicast groups are still active on a particular subnet.

In the fMAP product, IGMP snooping is a feature that allows the product to conserve network bandwidth by limiting the layer 2 forwarding of IP multicast packets only to the LAN segments that have expressed interest in receiving packets addressed to a multicast group. This function is performed by monitoring the Internet Group Management Protocol messages between IP hosts and a multicast router.

The fMAP product, through its learning process, knows which unicast MAC address can be reached via which of its ports and programs its forwarding database (MAC forwarding table) accordingly, and the forwarding process uses this to forward received unicast frames. Broadcast and multicast frames are forwarded by the 7000/9000 out of all the ports other than the port the frame was received on. This flooding approach is acceptable for a broadcast frame since that is exactly what broadcast implies. However the same approach used for forwarding multicast frames (containing the IP multicast packet) is less network bandwidth efficient if only a subset of hosts have joined the IP multicast group and hence are interested in receiving the multicast traffic.

This deficiency can be overcome and the IP multicasting in a layer 2 switching (LAN switching) networking environment can be handled more efficiently by implementing the **IGMP snooping** capability in the 7000/9000 product.

14.1.2 IGMP Snooping

IGMP snooping allows the fMAP product to conserve the local area network bandwidth by not flooding (broadcasting) the multicast frames but rather forwarding the multicast frames only to those ports that have expressed an interest in receiving such frames. The product must examine (or **snoop**) some layer 3 information (**join** and **leave**) in the IGMP host membership report message and the IGMP host leave group messages sent by the host to

a multicast router. The snooping of these messages is used to learn (or forget) which ports are interested (or not interested) in receiving multicast packets.

In simple terms, upon the receipt of an IGMP host membership report message for a particular multicast group, the *IGMP learning process* adds the port to the MAC address table against the multicast MAC address if it is not already present. Upon the receipt of an IGMP host leave group message for a multicast group, the IGMP learning process deletes the port from the MAC address table if it is present.

The forwarding process then utilizes the MAC address table populated by the learning process above to do efficient forwarding of the received multicast frame.

14.1.2.1 Default settings

- The system-wide default setting for IGMP snooping is enabled.
- The per-port default setting for IGMP snooping is enabled.
- Flooding of unknowns default setting for IGMP snooping is disabled.

14.1.3 IGMP Snooping Disabled

When IGMP snooping is disabled, the treatment of multicast frames by the fMAP product is the same as any other layer 2 switch.

- Each time a frame is received, the learning process reads the source MAC address and updates the address tables if required. The forward process then uses these address tables to do an address lookup on the destination MAC address to determine where to forward the frame.
- Initially, the fMAP product starts out by broadcasting/flooding (default forwarding) the received unicast frames on all its ports other than the port it was received on. This continues until the learning process learns and populates the MAC address table (consisting of MAC address - port entries) after which the received unicast frames are forwarded only to the intended destination. The exact port of the intended destination is obtained by using the destination MAC address in the received frame as a key to locate the address - port entry (inserted by the learning process earlier on) in the MAC address table. This is called the address lookup phase as part of the forwarding process to exactly forward the unicast frames. Note that there is one entry per unicast MAC address in the MAC address table since the unicast addresses are unique.
- For any broadcast frames (with a destination MAC address of all 1's), the frame is forwarded on all the LAN switch ports (flooding) by the forwarding process obviously not including the port the broadcast frame was received on.
- For any multicast frames the lookup fails to determine the ports to send this frame on, and so floods them to all ports in the VLAN. There is no Source Address with the Multicast Address since it has not been learned.

Note: Creating a VLAN of type VLAN is actually the same, except the frame may be flooded on only member ports of the VLAN.

14.1.4 IGMP Snooping Enabled

As mentioned above, IGMP snooping allows the 7000/9000 to conserve local area network bandwidth by not broadcasting a received multicast frame but rather forwarding the multicast frame only to those ports that have expressed an interest in receiving such frames. (The default forwarding behavior of a LAN switch for unicast and broadcast frames is not affected.) The snooping will configure the hardware to allow multicast streams for that group to be forwarded only to ports that have requested that stream.

Once IGMP has been enabled, the system may generate a warning message at the user's CLI session stating that classifier capacity or capabilities have been exceeded on the slot(s) impacted by the provisioning change. The user should investigate classifier-related provisioning, such as IGMP, DHCPRELAY, VLAN (for per-VLAN UFO and HVLAN), EPSR, INTERFACE (TAGALL option for HVLAN), ACCESSLIST, and CLASSIFIER to determine the reason for the message.

14.1.5 Interaction Between System and Ports/Interfaces

Since the IGMP Snooping feature can be enabled and disabled at both the system and port level, the following interactions apply:

- If IGMP Snooping is disabled system wide, all multicast packets will be flooded (within the VLAN) including IGMP control packets (Reports, General Queries, Groups Specific Queries). In this case, no IGMP control messages are forwarded to the CPU for processing; they are just be switched (flooded).
- If IGMP Snooping is enabled system wide and the port level control is enabled, then the port is snooped; IGMP Snooping software in the CFC will receive Reports and Leaves and process them as follows:
 - Unrequested (no Report processed by IGMP Snooping) multicast packets are dropped at the CFC switch.
 - Requested multicast packets are only sent to the ports where an IGMP Report is received.
 - If IGMP Snooping is enabled system wide and the port level control is disabled, then any IGMP Control Packet (Reports and Leaves) that are received from that port will be discarded (not processed and not flooded). This port will not be able to participate in IGMP.

14.1.6 Reserved Multicast Range Behavior

For the fMAP series, packets within a **subset** (see [16.2.1](#)) of the reserved multicast address range of 224.0.0.x (x = 0..255) will be allowed to be flooded within the VLAN. The rest are dropped.

14.2 IGMP Changes to Support FE/FX Upstream Interface

14.2.1 Overview

With the 6.0 feature of an FE/FX interface being both a customer interface and an interface to a multicast router, there are changes to the setting of an interface as well as the settings for IGMP that must be understood if the FE/FX configuration is to be deployed correctly. The changes are as follows:

- IGMP snooping will need to know whether interface's `DIRECTION` is a network or customer interface.
- IGMP Snooping needs to know whether to perform `INTERNAL`, `EXTERNAL`, or `MCPASSTHROUGH` (explained below) on each of the interfaces.

The defaults for the `DIRECTION` attribute for an interface are as follows:

- GE2 interfaces will default to `NETWORK`.
- FE interfaces on FE2 cards will default to `NETWORK`
- All other interfaces will default to `CUSTOMER`.

Note: This `DIRECTION` attribute can be set whether the interface is administratively UP or DOWN, and will take effect immediately. The attribute is persisted in the system database and mirrored to the inactive CFC.

The default IGMP Snooping settings (when IGMP snooping settings are enabled on the interface) are as follows:

- GE2 interfaces will default to `MCPASSTHROUGH`.
- FE interfaces on FE2 cards will default to `MCPASSTHROUGH`.
- All other interfaces will default to `INTERNAL`.

14.2.2 Router Interface

A multicast router is dynamically detected on any Network interface by detecting an IGMP Membership Query packet. When an IGMP Membership Query packet is received on a GE port, that port is designated as a multicast router port for the VLAN the packet was received on. A classifier is used to detect the IGMP Membership Query packets; these classifiers are only installed on the Network interfaces.

14.2.3 Customer Interface

IGMP Membership Report and Leave Packets will be extracted from all customer interfaces and are processed by IGMP Snooping. The interface's `DIRECTION` will be determined by querying the interface `DIRECTION` value. When IGMP Snooping is in `INTERNAL` mode, it will reconfigure the hardware to limit the forwarding of multicast packet only to the ports that have expressed interest in the multicast group. When IGMP Snooping is in `MCPASSTHROUGH` mode, it will filter IGMP packets and will flood all multicast traffic that is received from the multicast router to the Network Interfaces that are set to `MCPASSTHROUGH` and that are a member of the VLAN.

A classifier is used to extract the IGMP Membership Report and Leave packets; these classifiers are installed on all of the customer interfaces.

Note: The third option, EXTERNAL, is used when a device further towards the customer performs IGMP snooping.

14.2.4 Network Interface

IGMP Membership Report and Leave Packets are extracted from all network interfaces and are processed by IGMP Snooping. The interface's DIRECTION will be determined by querying the interface DIRECTION value. When IGMP Snooping is in INTERNAL mode, it will reconfigure the hardware to limit the forwarding of multicast packet only to the ports that have expressed interest in the multicast group. When IGMP Snooping is in MCPASSTHROUGH mode, it will filter IGMP packets and will flood all multicast traffic that is received from the multicast router to the Network Interfaces that are set to MCPASSTHROUGH and that are a member of the VLAN.

The extraction of the IGMP packets from the GE ports is performed to filter the Reports and Leaves that are being sent from a subtended system to the multicast router. A classifier is used to extract the IGMP Membership Report and Leave packets; these classifiers are installed on the network interfaces.

Note: The third option, EXTERNAL, is used when behind the port is a device (usually another fMAP) that performs IGMP snooping. This is especially important in a daisy chain configuration when the upstream port is an FE/FX port.

14.2.5 Multicast Stream Counts

IGMP Snooping will keep track of the number of multicast streams that are associated with the non NETWORK (GE and FE2) cards. When an IGMP Host requests to join a multicast group, the card's stream count is checked to ensure that the card's provisioned stream count is not exceeded. If the Report would exceed the card's provisioned stream count, then the Report is not allowed. The stream count processing is not performed on the GE interfaces because they receive all multicast traffic.

14.2.6 Key IGMP Snooping Syntax

Setting the interface DIRECTION is part of setting the overall attributes of an interface, as follows:

```
SET
INTERFACE={ type:
             | type:id-range
             | id-range
             | ifname-list
             | ALL
            }
interface_type or port_type
(other attributes and values)
[ DIRECTION={ NETWORK
              | CUSTOMER
            }
  [ FORCE ] ]
[ DESCRIPTION=description ]
```

Controlling the IGMP snooping on an interface is controlled as part of setting the overall attributes for IGMP snooping on an interface, as follows:

```

SET
IGMPSNOOPING
{ CARD={ slot-list
        | ALL
      }
  MCASTGROUPLIMIT=1..512
  | INTERFACE={ type:id-range
               | id-range
               | ifname-list
               | ALL
             }
  SNOOPINGMODE={ INTERNAL
                 | EXTERNAL
                 | MCPASSTHROUGH
               }
  [ FLOODUNKNOWNNS={ ON
                    | OFF
                  } ]
  [ ROUTERAGEINGTIMER=10..1200 ]
  [ GENQUERYTIMER=5..120 ]
  [ DUPREPORTTIMER=5..120 ]
}

```

14.3 IGMP Interactions

14.3.1 IGMP Group Limit (MCASTGROUPLIMIT) for Each ADSL24A Card

The number of IGMP groups for each ADSL24A can be set from 1 up to 512, depending on bandwidth requirements (usually for the number of video channels). The default is 25. Note that when the number of IGMP groups reaches 80% of the configured number, a management log is produced, and at 100% an alarm is produced.

14.3.2 IGMP Multicast Handling for FLOODUNKNOWNS (EPON2 Card)

With `FLOODUNKNOWNS=ON`, when MCAST packets come into the system from any port (SM or NM) with an unknown destination, they will be flooded to all ports (normal bridging).

With `FLOODUNKNOWNS=OFF`, when MCAST packets come into the system from any port (SM or NM) with an unknown destination, they will be dropped. This is the **default** for this command if the `FLOODUNKNOWNS` option is not specified

Support for IPV6 traffic is available. When IGMP is enabled, all unrequested (snooped) multicast traffic is dropped. When IGMP `FLOODUNKNOWNS` is turned ON, IPv6 groups will flood.

The EPON2 card is different than the rest of the system in that:

- For IPV6, multicast packets are flooded regardless of the `FLOODUNKNOWNS` setting.
- The EPON2 card will **not** flood multicast other than IPV6.
- When `IGMPSNOOPING` is enabled system wide, the EPON2 does not support `IGMPSNOOPING FLOODING`. You can set the EPON2 to support flooding by disabling IGMP system wide **and** disabling IGMP on the ONU interfaces.
- When the user sets `FLOODUNKNOWNS=ON`, there is the following message:

```
officer SEC>> set igmp floodunknowns=on
Warning (040292): One or more Service Modules (SMs) were found that do NOT
                  support the flooding of unknown multicast packets. Please
                  check system alarms to determine effected SMs.
```

When the user inputs the command to check the alarms, the output shows the slots where cards do not flood unknown multicast; the EPON2 card.

```
officer SEC>> show alarms
```

```
--- Card Alarms ---
```

Slot	Fault	Severity
16	Inconsistent Flooding	Major

14.3.3 MAC Limiting

MAC limiting restricts the ability to learn MAC addresses on a port. When the MAC learning limit is reached, all frames are dropped, including Broadcast and Multicast frames. This could be part of a subscriber's Service Level Agreement.

14.3.4 ICommand Summary for IGMP

TABLE 14-1 Commands for IGMP

Object	Verb	Syntax	Description
IGMPSNOOPING FLOODING	ADD	<pre> ADD IGMPSNOOPING FLOODING { ALLSTANDARD DVMRP OSPFALL OSPFDESIGNATED RIP2 IGRP DHCPRELAY PIM RSVP CBT VRRP DXCLUSTER CISCONHAP HSRP MDNS CUSTOM=groupname GROUPADDRESS=ipaddress } { VID=vid-list } </pre>	
IGMPSNOOPING INTERFACE	ADD	<pre> ADD IGMPSNOOPING INTERFACE={ type:id-range id-range ifname-list ALL } MACADDRESS={ macaddress-list partial-macaddress-list } </pre>	<p>The ADD IGMPSNOOPING command is used to configure the set-top box (STB) MAC address connected to an interface on the switch. Up to five (5) STB MAC addresses can be configured for a given interface. The purpose of this command is prevent STB mobility and prevent theft of broadcast video service.</p> <p>The user can also specify a partial MAC address. This provides a filtering function so a specific vendor can participate. There must be at least one hex pair. (Vendors use three or four pairs.)</p>

TABLE 14-1 Commands for IGMP (Continued)

Object	Verb	Syntax	Description
IGMPSNOOPING FLOODING	DELETE	<pre> DELETE IGMPSNOOPING FLOODING { ALL ALLSTANDARD DVMRP OSPFALL OSPFDESIGNATED RIP2 IGRP DHCPRELAY PIM RSVP CBT VRRP DXCLUSTER CISCONHAP HSRP MDNS CUSTOM=groupname } [VID={ vid-list ALL }] </pre>	
IGMPSNOOPING INTERFACE	DELETE	<pre> DELETE IGMPSNOOPING INTERFACE={ type:id-range id-range ifname-list ALL } MACADDRESS={ macaddress-list partial-macaddress-list ALL } </pre>	<p>The DELETE IGMPSNOOPING command is used to delete some or all of the set-top box (STB) MAC addresses that have been configured for a given interface.</p>
IGMPSNOOPING INTERFACE	DISABLE	<pre> DISABLE IGMPSNOOPING [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	<p>The DISABLE IGMPSNOOPING command is used to disable the snooping feature and stop the interception and monitoring of IGMP protocol messages. This command is used to disable the IGMP on a system-wide basis or on a single interface.</p>

TABLE 14-1 Commands for IGMP (Continued)

Object	Verb	Syntax	Description
IGMPSNOOPING INTERFACE	ENABLE	<pre> ENABLE IGMPSNOOPING [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	The ENABLE IGMPSNOOPING command is used to enable snooping for the intercept and monitoring of IGMP protocol messages that set-up the layer 2 multicast group information in the forwarding database. It is used to enable the IGMP on a system-wide basis or on a single interface.
IGMPSNOOPING COUNTER	RESET	<pre> RESET IGMPSNOOPING COUNTER [{ STANDARD MESSAGERESPONSE INTERFACE={ type:id-range id-range ifname-list ALL } CARD={ slot-list ALL } }] </pre>	The RESET IGMPSNOOPING command is used to reset IGMP counters/statistics.

TABLE 14-1 Commands for IGMP (Continued)

Object	Verb	Syntax	Description
IGMPSNOOPING	SET	<pre> SET IGMPSNOOPING { CARD={ slot-list ALL } MCASTGROUPLIMIT=1..512 INTERFACE={ type:id-range id-range ifname-list ALL } SNOOPINGMODE={ INTERNAL EXTERNAL MCPASSTHROUGH } [FLOODUNKNOWN= { ON OFF }] [ROUTERAGEINGTIMER=10..1200] [GENQUERYTIMER=5..120] [DUPREPORTTIMER=5..120] } </pre>	<p>The SET IGMPSNOOPING command is used to set various configurable IGMP settings in the switch, including setting the multicast stream count per slot, the flooding of unknown multicast packets, and various timers.</p> <p>In release 6.0, IGMPSNOOPING can be set by INTERFACE, and the SNOOPINGMODE can be selected.</p>

TABLE 14-1 Commands for IGMP (Continued)

Object	Verb	Syntax	Description
IGMPSNOOPING	SHOW	<pre> SHOW IGMPSNOOPING [{ STATUS MCASTGROUPS [FULL] COUNTER [{ STANDARD MESSAGERESPONSE INTERFACE={ type:id-range id-range ifname-list ALL } CARD={ slot-list ALL } }] INTERFACE={ type:id-range id-range ifname-list ALL } [FULL] CARD={ slot-list ALL } [FULL] }] </pre>	<p>The SHOW IGMPSNOOPING command is used to display the current status of the IGMP snooping feature, the currently configured value of the IP multicast stream count value, or the set-top box (STB) MAC address(es) that have been configured for a particular interface in the switch or for all of the interfaces in the switch.</p>

15. Traffic Management

15.1 Overview

15.1.1 Summary Table and Notes

Traffic management involves a set of features to ensure the following:

- Traffic types are given a certain level of priority (usually involving queues or classes of service)
- Traffic is filtered by some criteria so that it can or cannot pass through some point.

This table summarizes components and their traffic management feature availability. (“X” means supports, with qualifiers having footnotes, blank means not supported.).

TABLE 15-1 Traffic Management Summary Table - DSL- and TDM-based Cards

Classifier Match Fields	ADSL 24A^a
ETHFORMAT	X ^b
INNERVID	
INNERVPRIORITY	
IPTOS	X
IPDSCP	X
IPPROTOCOL	X
IPSOURCE	X ^c
IPDEST	X (c)
LSAP	X ^d
MACDEST	X
MACSOURCE	X
PROTOCOL	X
TCPFLAGS	
TCPPORTDEST	X

TABLE 15-1 Traffic Management Summary Table - DSL- and TDM-based Cards (Continued)

Classifier Match Fields	ADSL 24A ^a
TCPPORTSOURCE	X
UDPPORTDEST	X
UDPPORTSOURCE	X
VID	X
VPRIORITY	X
Classifier Actions	
DROP	X
FORWARD	X
COUNT	X
SETVPRIORITY	X
SETIPTOS	
SETIPDSCP	
MOVEPRIOTOTOS	
MOVETOSTOPRIO	
ARP Filtering	X
Ingress Metering Increments	64 Kbps
Egress Port Rate Limit Increments ^e	
VC Mapping to DSL Interface	4
Peak Cell Rate Limiting (on a VC basis)	X
Number of Egress Queues	8
Number of Classifiers	(i)

a. The PAC24A card has the same ADSL features as the ADSL24A card.

- b. Acceptable match values for the ETHFORMAT field are ANY (which matches any Ethernet format), ETHII (along with any other match fields), and 802.3 (w/ LSAP=0xAAAA (SNAP) or LSAP=0x0F0F (NETBIOS) match rule).
- c. Only whole addresses are matched (no masked values).
- d. Only NETBIOS and SNAP values are supported.
- e. For all ADSL cards, configuration of egress rate limiting is not supported directly, but can be set using the ADSL maximum downstream rate.

TABLE 15-2 Traffic Management Summary Table - Ethernet-based Cards

Classifier Match Fields	FE10/ FX10	GE4/ GE2RJ	GE8	EPON2 ^a
ETHFORMAT	X	X	X	
INNERVID	X		X	
INNERVPRIORITY	X		X	
IPTOS	X	X	X	
IPDSCP	X	X	X	
IPPROTOCOL	X	X	X	
IPSOURCE	X	X	X	X
IPDEST	X	X	X	
LSAP	X			
MACDEST	X	X	X	
MACSOURCE	X	X	X	
PROTOCOL	X	X	X	
TCPFLAGS	X	X	X	
TCPPORTDEST	X	X	X	
TCPPORTSOURCE	X	X	X	
UDPPORTDEST	X	X	X	
UDPPORTSOURCE	X	X	X	
VID	X	X	X	X

TABLE 15-2 Traffic Management Summary Table - Ethernet-based Cards (Continued)

Classifier Match Fields	FE10/ FX10	GE4/ GE2RJ	GE8	EPON2 ^a
VRIORITY	X	X	X	
Classifier Actions				
DROP	X	X	X	X
FORWARD	X	X	X	X
COUNT	X	X	X	
SETVPRIORITY	X	X	X	
SETIPTOS	X	X	X	
SETIPDSCP	X	X	X	
MOVEPRIOTOTOS	X	X	X	
MOVETOSTOPRIO	X	X	X	
ARP Filtering	X			
Ingress Metering Incre- ments	64 Kbps	8 Mbps	8 Mbps	b
Egress Port Rate Limit Increments	1 Mbps	8 Mbps	1 Mbps	(c)
Number of Egress Queues	4	4	8	
Number of Classifiers ^c	(d)	128	128	d

- a. Not the interface for traffic management; this is done by the ONU, which is the UNI interface port. The EPON2 models the ONU as an fMAP extension.
- b. QoS policies (SLAs) allow limiting of services at specific rates.
- c. Classifier capacity is an approximation. The features enabled on a port or interface, such as such as IGMP, DHCPRELAY, VLAN (for per-VLAN UFO and HVLAN), EPSR, INTERFACE (TAGALL option for HVLAN), ACCESSLIST, and CLASSIFIER, classifiers on surrounding ports, content of the user classifiers, and many other factors influence the number of classifiers available on a given port or interface.e.
- d. Refer to [15.1.3](#).

15.1.2 Usage Notes:

- Classifiers on LAGs only support filtering and remarking; not metering.

- LAGs do not support egress port rate limiting.
- Ingress Meters (TRAFFIC DESCRIPTORS) are limited to 1 per classifier. Their actions on out-of-profile packets are NCFORWARD (the default), NCDROP, and NCCOUNT.
- All interfaces support p-bit mapping to queues.
- All interfaces support only Strict Priority, Tail-drop queuing discipline.
- FE10/FX10/ADSL24A - Not supported with double tagging of 802.3 packets. Rejects classifiers with ETHFORMAT=802.3x if the classifier contains an INNERVID or INNERVPRIORITY rule. Rejects classifiers with ETHFORMAT=802.3x if the classifier contains is an HVLAN configured on that port. Any generic classifier that does not specify an EthFormat of 802.3 will NOT match 802.3 packets. In other words, classifiers with no ETHFORMAT rule or a rule of ETHFORMAT=ANY will not match 802.3 packets.
- FE10, FX10, ADSL24A
 - If a classifier contains a rule that specifies one of the 802.3 EthFormats, a filter will be created that will match 802.3 packets. The fMAP system assumes that if the user specifies 802.3x and IP-related fields, that this is an indicator that IP uses LLC-SNAP and as a result the system allows for an 8-byte SNAP header. The system does NOT support 802.3 headers of any size other than 8 bytes when matching IP-related fields.
 - Rejects classifiers containing INNERVID or INNERVIDPRIORITY rule(s) if there is NOT a HVLAN configured on that interface/port.
 - If a classifier fails to be installed due to any reason related to the configuration of other features on that port (such as, an HVLAN being configured on that port), the system will keep the data related to the specified classifier. When the condition(s) preventing the classifier's installation is corrected (such as, HVLAN is removed from the port) the classifier will be automatically installed and its status will be updated at the user's CLI session.
 - Supports ETHFORMATs of 802.3 and EthII. Note that each of these formats consumes twice as many hardware resources as ETHFORMAT=802.3tagged, 802.3untagged, EthIItagged, or EthIIuntagged.
- Interfaces/ports configured with client-side IGMP classifiers are configured to drop IGMP router packets.
- FE10, FX10, ADSL24A - MACs for dropped or filtered packets will appear in the Forwarding Database. They will be displayed when the SHOW SWITCH FDB command is entered.

15.1.3 EPON2 Traffic Management

As shown in [Table 15-2](#), the EPON2 has the following traffic management support:

- ARP Filtering is not supported.
- For classifiers, only the IPSOURCE and VLAN match fields are used.
- For classifier actions, only FORWARD and DROP are supported.
- There is a limit of approximately 20 classification rules that each ONU can support.

15.2 Quality of Service (QoS) Model

15.2.1 Overview

Packet-based networks provide primarily three types of services:

- Data
- Voice
- Video

Providers must deliver these services at a level of quality acceptable to the customer. Service quality or service level can be defined by controlling:

- Availability
- Delay
- Delay variation (jitter)
- Lost packet ratio (bandwidth)

Additionally, various applications (e-mail, file transfer, teleconferencing, video conferences) can be considered as real-time versus non real-time applications.

- Real-time applications (such as voice) have a lower tolerance to delay or delay variation, but can handle some packet loss.
- Non-real-time applications are not as adversely affected by delay or delay variation, but are highly affected by packet loss.

A Service Level Agreement (SLA) details the *level of service* the service provider and customer negotiate. Providers use Quality of Service (QoS) functions to segregate traffic and then manage the service quality through the network to meet the customer's needs.

Figure 15-1 shows the general flow for QoS; refer to this figure while reading the rest of this subsection.

Note: The rest of this subsection describes in general what traffic management provides. For the capabilities of the fMAP products refer to the rest of this section. Refer to the next two sections for details (capabilities, restrictions) on a specific product.

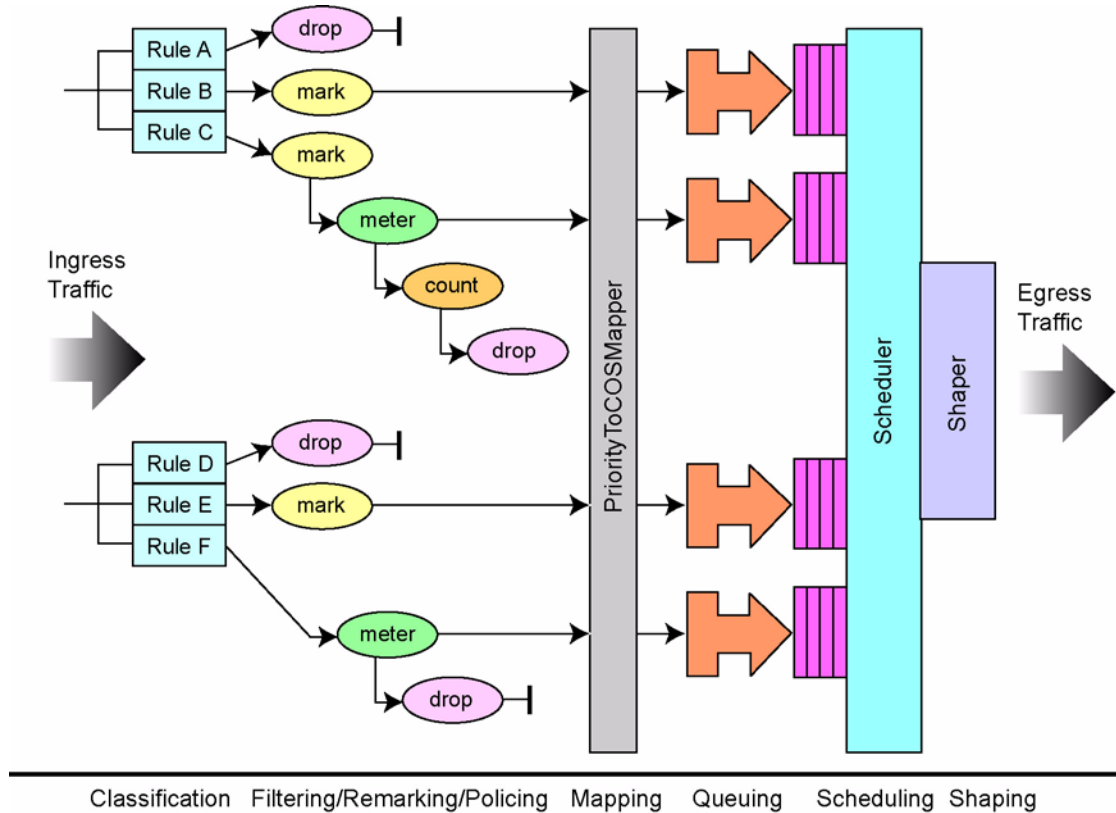


FIGURE 15-1 Model of Traffic Flow for a QoS-capable Device

15.2.2 Ingress Traffic

The main strategy in providing QoS is to first **classify** and **segregate** traffic into separate flows. These flows can then be managed separately through the provider's network with each flow getting a specified level of service. Traffic classification and segregation are performed when traffic from a customer enters the network through the network edge device. Traffic is classified and segregated according to set of criteria or **rules**. Once the traffic is classified, the packets will have certain actions performed upon them as configured by the provider. These actions are **mark**, **meter**, **count**, and **drop**.

To **mark** a packet in the traffic flow means that once a packet has been allowed to ingress the port, it will be associated with a certain flow. Marking the packet means to identify the packet with a **Class of Service (COS)** that will be applied to the packet as it moves through the device and into the network.

- For ethernet frames, these are defined as the 802.1p user priority bits or class of service bits.

- For **IP packets** there are the **DSCP** field and the **TOS** field. The COS identifier can specify both a service level priority and the precedence for dropping packets, but this is not done at the ethernet level.

Once a packet is marked, it may go immediately to a COS queue, but in many cases the traffic is metered. To **meter** the packet flow is to monitor or **police** the rate of traffic flow and to see if incoming traffic exceeds the bandwidth specified in the SLA. When packet flow exceeds the bandwidth allocated, they are labeled as Out Of Profile (**OOP**). This policing function is done using the **leaky bucket** algorithm. The bucket has a capacity and an output rate and packets enter and leave the bucket.

If packets arrive at a rate faster than contracted for in the SLA for a continuous period of time, the bucket will overflow. These overflow packets are classified as *out of profile* and another action can be applied to them, such as **drop** (throw away), or **remark** them in such a way that these OOP packets have a higher probability of being thrown away when congestion points are encountered through the device or further on in the network. Packets that exceed this bandwidth are labeled as out of profile with the SLA.

By metering the rate at which packets arrive, the provider can control bandwidth, since the SLA may include a minimum bandwidth availability and as well as a maximum (for short periods). These are defined as follows:

- Committed Information Rate (**CIR**) is the minimum guaranteed rate the provider network will provide under normal conditions, and is measured in bits per second.

Note: Any service that provides a non-zero bandwidth guarantee must have a CIR. A CIR of zero indicates the service will provide no minimum guarantee for frame delivery.

- Committed Burst Size (**CBS**) is the maximum number of bytes that can be sent at the CIR and is measured in kilobytes (KB) or megabytes (MB)
- Peak Information Rate (**PIR**) is the maximum rate at which frames/packets are allowed to burst above the CIR and is measured in bytes per second.
- Peak Burst Size (**PBS**) is the maximum number of bytes that can be sent at the PIR and is measured in kilobytes (KB) or megabytes (MB).

For a service that requires only a CIR and CBS, a single meter is used; the capacity of a single bucket is equal to the CBS and the leak rate is the CIR. For services that require all four parameters, two meters are required.

15.2.3 Egress Traffic

Once the traffic flows have passed through the policing function and are allowed to continue, the COS identifier (obtained when the traffic was marked) is used to map the traffic with a set of queues and to assign a priority. Each queue is associated with a level of service of low versus high. If a QoS network provides four levels of service, there will typically be four queues. Traffic flows will be associated with a priority (using the COS identifier bit) and therefore a queue.

As the packets are placed in the queues, there may still be conditions where packets may need to be dropped. One method of handling overflow is **tail-drop**; when a queue is in an overflow state, all newly arriving packets are dropped. If the potential for queue overflow was too high, the queue size(s) were increased. However, dynamic managing of queue depth can result in better network performance, and one method is **Random Early Discard (RED)**, which improves network throughput and lowers the probability of packet discard.

As the traffic passes through the queues, they are then scheduled for output. The common type of **scheduler** is the Strict Priority (SP) scheduler, which selects a packet at the head of the highest priority queue (usually allows no delay) and continues to select packets in that queue until it is empty; only then are packets chosen from other queues. When a network has little or no congestion, all queues are scheduled equally. However, in a heavily congested network, the highest priority queue may always have packets, and so the lower priority queues are never scheduled and are therefore blocked.

The **Weighted Round Robin (WRR) scheduler** associates an additional weight to each queue, so that the scheduler spends at least some time (although limited because of its lower weighting) with the other queues.

All of these functions together help ensure that traffic is classified/filtered and then metered to ensure that its bandwidth does not exceed the SLA. However, this does not guarantee that all network resources are available so that the bandwidth agreed to can be propagated through the network. **Call Admission Control (CAC)** is an accounting algorithm that qualifies the available network resources against the SLA. Within a network device, the CFC function balances downstream congestion and full utilization of available bandwidth. CAC takes into account all factors and calculates an equivalent maximum bandwidth.

Finally, the provider should use the performance monitoring tools provided by the system to verify the levels of service that have been negotiated. These tools measure packet loss, delay, jitter, availability, and failure recovery time.

Note: Weighted round robin (WRR) scheduling is not supported.

15.2.4 Traffic Management Throughout the Network

Since the traffic flows occur throughout the network, each device must give these flows the same treatment as they pass through the device. For Rapiere devices, there are QoS features that can be configured so that traffic prioritization is similar to what is configured for the fMAP products.

15.3 Classifier Management

15.3.1 Overview

Classifiers provide the remarking, metering, counting, etc., actions for interfaces. Classifiers define packet matching rules that classify packets into data flows so that they may be processed in a similar manner. For example, all packets with the same destination TCP/IP port may be defined to form a flow (such as Telnet or HTTP traffic).

Once packets are defined by the use of classifiers, the QoS functions associate the classifier rule and subsequent actions on the packets.

Classifiers perform the following key functions:

- Traffic filtering - Filters traffic so that only required traffic goes on to the VLAN. Sample filters would be by protocol, IP address and applications such as HTTP or SMTP.

- Traffic quality of service - Prioritizes frames based on their classification. For instance, voice over IP traffic could be given a higher priority than web traffic.

15.3.2 Classifier Match Rules

A set of packet matching rules can be created by the user. The classifiers can identify any single packet based upon the following criteria:

- Layer 2 protocols - Ethernet encapsulation type - Packets are classified depending on the specific protocol type of each frame. Different values indicate how the packet is formatted. For more details on values see the ETHFORMAT parameter in the CREATE CLASSIFIER command in [Table 15-3](#).
- Layer 3 protocols - Frames are classified based on any value for Layer 3 protocols. The system can match based on any Layer 3 field regardless of the Layer 2 frame type (as long as it is supported).
- Source/destination IP address - Frames are classified based on an IP mask so that frames can be allowed on a partial match.
- Layer 4 protocol (TCP/UDP, etc.) - Frames are classified based on specific Layer 4 TCP or UDP destination and source port numbers contained within the header of an IP frame.

[Table 15-3](#) lists the available classifiers match rules.

TABLE 15-3 Classifiers Match Rules

Parameter	Description
[VID={1..4095 ANY}]	The source VLAN the packet is associated with when received by the switch. - ANY - match all packets with any VLANID value.
[VPRIORITY={0..7 ANY}]	This matches the VLAN ID specified with the User Priority frame - ANY - match all packets with any VPRIORITY value. This match rule is used to set up the class of service queues. Refer to 15.8 .
[INNERVID={1..4095 ANY}]	When the HVLAN feature is used, this is the VLAN ID of the customer-based VLAN. - ANY - match all packets with any VLAN ID value.
[INNERVPRIORITY={0..7 ANY}]	When the HVLAN feature is used, this is the value of the User Priority frame of the customer-based VLAN. - ANY - match all customer VLAN packets with any VPRIORITY value.

TABLE 15-3 Classifiers (Continued) Match Rules

Parameter	Description
[ETHFORMAT={802.3TAGGED 802.3UNTAGGED ETHIITAGGED ETHIUNTAGGED ETHII 802.3 ANY}]	<p>The ethernet encapsulation type for the packet:</p> <ul style="list-style-type: none"> - 802.3TAGGED: matches IEEE 802.3 format with a VLAN tag. - 802.3UNTAGGED: matches IEEE 802.3 format without a VLAN tag. - 802.3: matches IEEE 802.3 format regardless of tags. - ETHIITAGGED: matches Ethernet II format with a VLAN tag. - ETHIUNTAGGED: matches Ethernet II format without a VLAN tag. - ETHII: matches Ethernet II format regardless of tags. - ANY: matches any Ethernet encapsulation. <p>If an ETHFORMAT is specified, the PROTOCOL parameter must be used, specifying the protocol-type.</p> <p>Note: For the CFC4, FE10, and ADSL24, ETHFORMAT=ANY or no ETHFORMAT rule will result in a match for EthII formats only. For the CFC24 and GE3, ETHFORMAT=ANY or no ETHFORMAT rule will result in a match for either EthII or 802.3 formats.</p>
[LSAP={NETBIOS lsap-value ANY}]	<p>The LSAP match rule field matches on any packet with the specified LSAP value.</p> <p>LSAP refers to the combination of the DSAP (Destination Service Access Point) and SSAP (Source Service Access Point) octets in an 802.3 Ethernet frame. The value may be entered in decimal or in hex but must be less than or equal to 4095.</p> <p>The value "NETBIOS" can be used to specify the LSAP value for that protocol (0xF0F0). The value "ANY" matches any LSAP value.</p>
[IPDEST={ipaddress-mask MULTICAST ANY}]	<p>The destination IP address (either host or subnet) of the IP packet.</p> <p>IP address ranges are specified using a valid IP address or valid subnet and mask. A range is specified using a '/' character (such as 1.0.0.0/8).</p> <p>MULTICAST means all packets with a multicast address.</p> <p>ANY - match all IP packets with any IPDEST value.</p>
[IPSOURCE={ipaddress-mask ANY}]	<p>The source IP address (either host or subnet) of the IP packet.</p> <p>IP address ranges are specified using a valid IP address or valid subnet and mask. A range is specified using a '/' character (such as 1.0.0.0/8).</p> <p>ANY - match all IP packets with any IPSOURCE value.</p>
[IPDSCP={0..63 ANY}]	<p>The code point field with the DiffServ byte of an IP packet. This parameter cannot be specified with the IPTOS parameter.</p> <p>ANY - match all IP packets with any IPDSCP value.</p>

TABLE 15-3 Classifiers (Continued) Match Rules

Parameter	Description
[IPPROTOCOL={TCP UDP ICMP IGMP ipprotocol-number ANY}]	The layer 4 IP protocol of the IP packet. If the command specifies a TCP/IP packet matching rule, (e.g. TCP-PORT-DESC is specified), then the value for this parameter will be TCP. If the command specifies a UDP/IP packet matching rule, (i.e. the UDP-PORT-DESC parameter is specified), then the value for this parameter will be UDP. ipprotocol-number is expressed as a one byte hexadecimal or decimal number. ANY - match any layer 4 IP protocol.
[IPTOS={0..7 ANY}]	The value of the precedence field within the TOS byte of an IP packet. This parameter cannot be used with the IPDSCP parameter. ANY means to accept any IPv4 packets but not other types, such as ARP. If no IPTOS (including ANY) is specified, there are no checks for the packet being an IPv4 packet, and so other packet types such as ARP are allowed. ANY - match all IP packets with any IPTOS value.
[MACDEST={macaddress MULTI-CAST ANY}]	The destination MAC address for the packet. MULTICAST is for multicast packets. ANY - match all packets with any MACDEST value.
[MACSOURCE={macaddress ANY}]	The source MAC address. ANY - match all packets with any MACSOURCE value.
[PROTOCOL= {IPV4 IPV6 protocol-type ANY}]	The type field of the payload. ANY - match all IP packets with any PROTOCOL value.
[TCP-PORT-DEST={tcp-port ANY}]	The TCP destination port of a TCP/IP packet. ANY - match all IP packets with any TCP-PORT-DEST value.
[TCP-PORT-SOURCE={tcp-port ANY}]	The TCP source port of a TCP/IP packet. ANY - match all IP packets with any TCP-PORT-SOURCE value.
[TCP-FLAG= {URG ACK RST SYN FIN PSH}[...] ANY]	The control bytes used in the TCP header. (Refer to RFC793.) - URG: Urgent Pointer field significant - ACK: Acknowledgment field significant - RST: Reset the connection - SYN: Synchronize sequence numbers - FIN: No more data from sender - PSH: Push Function ANY - match all IP packets with any TCP-FLAG value.

TABLE 15-3 Classifiers (Continued) Match Rules

Parameter	Description
[UDPPORTDEST={udp-port-list ANY}]	The UDP destination port of a UDP packet. ANY - match all IP packets with any UDPPORTDEST value. Note: In order to filter (block) a subscriber's port and prevent it from acting as a DHCP client, add a filter of UDPPORTDEST=67, dropping any packets destined for a DHCP server. To filter packets from an upstream DHCP server to the subscriber port, add a filter of UDPPORTDEST=68.
[UDPPORTSOURCE={udp-port ANY}]	The UDP source port of a UDP packet. ANY - match all IP packets with any UDPPORTSOURCE value.

15.3.3 Classifier Actions and COUNTS

A classifier or set of classifiers then have actions associated with them:

- **DROP** - discard the packet at the card. This action excludes the packet.
- **FORWARD** - allow traffic to be forwarded. This action includes the packet.
- **COUNT** - count the number of packets that have been forwarded or dropped. These are displayed with the SHOW CLASSIFIER command.

Note: The outputs associated with the COUNT setting is as follows: If the COUNT is combined with a DROP action, then the **Filter** Count is incremented in the output. If combined with a TRAFFICDESCRIPTOR (for policing), then the **Policed** Count is incremented in the output. If neither is associated, then the **Match** Count is incremented in the output. To view these outputs, use the command SHOW CLASSIFIER COUNTER <interface>. Refer to the Command Handbook for more details.

- **Remark the 802.1q priority field** - The priority bits can be set (remarked) on ingress, and that priority is used throughout the network devices at each egress queue. This can be set as follows:
 - The value is set directly (using the SETVPRIORITY action).
 - The value is set with the incoming value from the IP TOS field (using the MOVETOSTOPRIO command).

Note: To correlate the p-bit value with a queue, the SET QOS VLANQUEUEMAP command is used, as explained in 15.8.

- Set the IP TOS and IP DSCP fields - These can be set as follows:
 - The values can be set directly (using the SETIPTOS and SETIPDSCP action).
 - The value is set using the 802.1q priority field (using the MOVETOSTOPRIO command).

15.3.4 Classifier Association with a Port or Interface (Precedence)

When a classifier is associated with a port or interface, it is given a precedence, with the lowest number receiving the highest precedence. Classifiers on the same port cannot share the same precedence number.

If the user wishes to further qualify a traffic flow, metering can be applied to the ingress interface before the classifier is associated with that interface. Refer to [15.6](#).

Note: The precedence setting for classifiers should be 51 to 68 for classifiers that perform a filtering action, with 69 used for dropping packets that do not match any of the filtering criteria.

Note: The precedence setting for classifiers that remark packets for the QoS function should be 146 to 199.

Note: DHCP Relay uses precedence setting 69, therefore use precedence setting 69 only if DHCP Relay functionality is not implemented on the system.

15.3.5 Default Classifier `telesyn_default_video` Behavior in Release 5.0

From release 2.1 to release 5.0, there is a classifier `telesyn_default_video` that exists by default, as shown below:

```
officer SEC> SHOW CLASSIFIER=ALL
--- Classifier Configuration Data -----
Name                Field Match(es)                Action(s)
-----
telesyn_default_video  IPDEST=MULTICAST                SETVRIORITY=1
```

This classifier setting takes all IP data with a multicast destination and remarks the VLAN priority (p-bit) to a setting of 1.

Note: As mentioned in [15.3.4](#), the classifiers for remarking packets should be in the range 146-199. The `telesyn_default_video` will be set to a precedence of 145, as shown below.

The `telesyn_default_video` classifier can be configured as part of the overall filtering/remarking that is part of traffic control, as shown in below..

```
officer SEC> SHOW CLASSIFIER=ALL PORT=ALL
--- Classifier Configuration Data -----
Port Rank Name                Field Match(es)                Action(s)
-----
4.0 51  ipfilt1                    IPSOURCE=1.0.0.0/8              FORWARD
                                         COUNT
52  ipfilt2                    IPSOURCE=2.2.0.0/16            FORWARD
                                         IPDEST=7.7.0.0/16             COUNT
                                         PROTOCOL=IPV4
53  ipfilt3                    TCPPODEST=8080                  FORWARD
                                         IPSOURCE=3.3.0.0/16           COUNT
```

	145	telesyn_default_video	IPDEST=MULTICAST	SETVPRIORITY=1
4.1	51	ipfl1	IPSOURCE=1.0.0.0/8	FORWARD COUNT
	52	ipfl2	IPSOURCE=2.2.0.0/16 IPDEST=7.7.0.0/16 PROTOCOL=IPV4	FORWARD COUNT
	53	ipfl3	TCPPORTDEST=8080 IPSOURCE=3.3.0.0/16	FORWARD COUNT
	145	telesyn_default_video	IPDEST=MULTICAST	SETVPRIORITY=1

The user can set the p-bit to a different value and place it in a different queue. For example, multicast video traffic could be placed at the highest priority by remarking the p-bit to 7 which, by default, places it in queue 7 on interfaces supporting 8 queues and on queue 3 on interfaces supporting 4 queues.

Note: If the user wishes, the telesyn_default_video can be destroyed and have no remarking of p-bits; the p-bit settings should then be set by the upstream device.

When configuring classifiers, the FULL display for the classifiers for an interface or port will also include information that has been derived from the classifiers. This will be shown with a **(D)** next to the classifier attribute.

As an example, a user has configured a classifier set as shown below.

```
officer SEC> SHOW CLASSIFIER=ALL PORT 5.0
--- Classifier Configuration Data -----
Port Rank Name          Field Match(es)          Action(s)
-----
5.0  51  ip1                    IPSOURCE=1.1.1.1/32
      52  ip2                    IPSOURCE=1.1.1.2/32
      53  ip3                    IPSOURCE=1.1.1.3/32
      54  ip4                    IPSOURCE=1.1.1.4/32
      55  ip5                    IPSOURCE=1.1.1.5/32
      56  ip6                    IPSOURCE=1.1.1.6/32
      57  ip7                    IPSOURCE=1.1.1.7/32
      69  ipde                   IPSOURCE=ANY              DROP
```

When the user inputs the SHOW command with the FULL option, the details for the port, including Allied Telesis-assigned internal classifiers, are shown.

15.3.6 Classifier telesyn_default_video Behavior in Release 6.0

In release 6.0, the default telesyn_default_video classifier is no longer included with the system at initial start up. However, the classifier is automatically retained during an upgrade. Therefore, to have an equivalent classifier to handle the QoS for video, the user should either:

- Create a classifier that has the QoS attributes (for a new system in 6.0)
- Replace the default classifier if they no longer wish to use it (system upgraded from 5.0)

In a newly installed 6.0 system, the user can **recreate** the `telesyn_video_default` with the following commands, which configure the following:

- The ports 11.0-11.3 are all upstream
- The precedence setting is 146 (within the remarking range, `telesyn_default_video` was 145)
- The name of the classifier is `remark_multicast`

```
officer SEC>> create class remark_multicast IPDEST=MULTICAST
officer SEC>> add action remark_multicast SETVPRIORITY=4
officer SEC>> add class remark_multicast interface 11.0-11.3 precedence 146
```

To **replace** an existing `telesyn_video_default` is more complex, and involves the following major steps:

1. Note the attributes of the current `telesyn_default_video` classifier

```
officer SEC>> show class all interface all
--- Classifier Configuration Data ---
Interface Rank Name          Field Match(es)                Action(s)
-----
ETH:11.0  145  telesyn_default_video      PROTOCOL= IPV4 (D) SETVPRIORITY=4
                                     IPVERSION= 4 (D)
                                     IPDEST= MULTICAST
ETH:11.1  145  telesyn_default_video      PROTOCOL= IPV4 (D) SETVPRIORITY=4
                                     IPVERSION= 4 (D)
                                     IPDEST= MULTICAST
ETH:11.2  145  telesyn_default_video      PROTOCOL= IPV4 (D) SETVPRIORITY=4
                                     IPVERSION= 4 (D)
                                     IPDEST= MULTICAST
```

2. Create a new classifier (`remark_multicast`). The two classifiers now exist concurrently.

```
officer SEC>> create class remark_multicast IPDEST=MULTICAST
Info (010017): Operation Successful
officer SEC>> show class all
--- Classifier Configuration Data ---
Name          Field Match(es)                Action(s)
-----
telesyn_default_video  IPDEST= MULTICAST SETVPRIORITY=4
remark_multicast      IPDEST= MULTICAST
```

3. Associate this new classifier with similar attributes

```
officer SEC>> add action remark_multicast SETVPRIORITY=4
Info (010017): Operation Successful
officer SEC>> show class all
--- Classifier Configuration Data ---
Name          Field Match(es)                Action(s)
-----
telesyn_default_video  IPDEST= MULTICAST SETVPRIORITY=4
remark_multicast      IPDEST= MULTICAST SETVPRIORITY=4
```



```
-----
officer SEC>> add class remark_multicast int 11.0-11.3 precedence 146
Info (010017): Operation Successful
```

4. Compare the attributes of the two classifiers

```
officer SEC>> show class all interface all
```

```
--- Classifier Configuration Data ---
-----
```

Interface	Rank	Name	Field Match(es)	Action(s)
ETH:11.0	145	telesyn_default_video	PROTOCOL= IPV4 IPVERSION= 4 IPDEST= MULTICAST	(D) SETVPRIORITY=4 (D)
	146	remark_multicast	PROTOCOL= IPV4 IPVERSION= 4 IPDEST= MULTICAST	(D) SETVPRIORITY=4 (D)
ETH:11.1	145	telesyn_default_video	PROTOCOL= IPV4 IPVERSION= 4 IPDEST= MULTICAST	(D) SETVPRIORITY=4 (D)
	146	remark_multicast	PROTOCOL= IPV4 IPVERSION= 4 IPDEST= MULTICAST	(D) SETVPRIORITY=4 (D)
ETH:11.2	145	telesyn_default_video	PROTOCOL= IPV4 IPVERSION= 4 IPDEST= MULTICAST	(D) SETVPRIORITY=4 (D)
	146	remark_multicast	PROTOCOL= IPV4 IPVERSION= 4 IPDEST= MULTICAST	(D) SETVPRIORITY=4 (D)

```
-----
```

5. Destroy the telesyn_default_video (first from the interfaces, then the classifier itself)

```
officer SEC>> delete class telesyn_default_video int 11.0-11.3
Delete classifier(s) from interface(s) (Y/N)? y
Info (010017): Operation Successful
officer SEC>> destroy class telesyn_default_video
Destroy classifier(s) (Y/N)? y
Info (010017): Operation Successful
```

6. Ensure the new classifier has the correct attributes

```
officer SEC>> show class all
```

```

--- Classifier Configuration Data -----
Name                               Field Match(es)                               Action(s)
-----
remark_multicast                   IPDEST= MULTICAST                             SETVPRIORITY=4
-----

```

officer SEC>> show class all int all full

```

--- Classifier Configuration Data -----
Interface Rank Name                 Field Match(es)                               Action(s)
-----
ETH:11.0  146  remark_multicast                   PROTOCOL= IPV4                               (D) SETVPRIORITY=4
                                           IPVERSION= 4                               (D)
                                           IPDEST= MULTICAST
ETH:11.1  146  remark_multicast                   PROTOCOL= IPV4                               (D) SETVPRIORITY=4
                                           IPVERSION= 4                               (D)
                                           IPDEST= MULTICAST
ETH:11.2  146  remark_multicast                   PROTOCOL= IPV4                               (D) SETVPRIORITY=4
                                           IPVERSION= 4                               (D)
                                           IPDEST= MULTICAST
-----

```

15.3.7 Derived Classifiers (D)

The following shows that classifiers are added by the system when they can be derived. For the IPSOURCE classifiers, the PROTOCOL (IPV4) and IPVERSION (4) are derived from the IPSOURCE and are added with the (D) added, telling the user these were added by the system.

officer SEC>> SHOW CLASSIFIER=ALL PORT=18.0 FULL

```

--- Classifier Configuration Data -----
Port Rank Name                       Field Match(es)                               Action(s)
-----
18.0  51  ip1                                  PROTOCOL= IPV4                               (D) DROP
                                           IPVERSION= 4                               (D)
                                           IPSOURCE= 1.1.1.1/1
      68  ipde                                PROTOCOL= IPV4                               (D) DROP
                                           IPVERSION= 4                               (D)
                                           IPSOURCE= ANY
-----

```

15.3.8 Set Match Rule Defaults (SETDEFAULTS)

Classifier match rule defaults can be reset using the SETDEFAULTS command. This command is useful if the user wishes to change a match rule setting without having to delete the classifier. An example follows.

Here a classifier, ipfilt1, is created with IPSOURCE set.

officer SEC>> CREATE CLASSIFIER=ipfilt1 IPSOURCE=172.16.5.0/28

```

Info (010017): Operation Successful
officer SEC>> SHOW CLASSIFIER=IPFLT1
--- Classifier Configuration Data ---
Name                Field Match(es)                Action(s)
-----
ipflt1              IPSOURCE= 172.16.5.0/
                    28

```

An action is added to the classifier to drop the IPSOURCE ipaddress.

```

officer SEC>> ADD ACTION CLASSIFIER=ipflt1 DROP
Info (010017): Operation Successful
officer SEC>> SHOW CLASSIFIER=IPFLT1
--- Classifier Configuration Data ---
Name                Field Match(es)                Action(s)
-----
ipflt1              IPSOURCE= 172.16.5.0/
                    28                DROP

```

Now, set a PROTOCOL filter on the classifier.

```

officer SEC>> SET CLASSIFIER=ipflt1 PROTOCOL=IPV4
Info (010017): Operation Successful
officer SEC>> SHOW CLASSIFIER=IPFLT1
--- Classifier Configuration Data ---
Name                Field Match(es)                Action(s)
-----
ipflt1              PROTOCOL= IPV4                DROP
                    IPSOURCE= 172.16.5.0/
                    28

```

Set a IPDEST filter.

```

officer SEC>> SET CLASSIFIER=ipflt1 IPDEST=10.0.0.0/8
Info (010017): Operation Successful
officer SEC>> SHOW CLASSIFIER=IPFLT1
--- Classifier Configuration Data ---
Name                Field Match(es)                Action(s)
-----
ipflt1              PROTOCOL= IPV4                DROP
                    IPSOURCE= 172.16.5.0/
                    28
                    IPDEST= 10.0.0.0/8

```

Using the SETDEFAULTS command, set the IPDEST back to it's default value.

```

officer SEC>> SETDEFAULTS CLASSIFIER=IPFILTER1 IPDEST
Info (010017): Operation Successful
officer SEC>> SHOW CLASSIFIER=IPFILTER1
--- Classifier Configuration Data -----
Name                Field Match(es)                Action(s)
-----
ipfilter1           PROTOCOL= IPV4                 DROP
                    IPSOURCE= 172.16.5.0/
                    2

```

15.3.9 System Monitoring for Errors (NORES, ERR, NOSPT)

When creating classifiers, the user should consider all configuration guidelines, restrictions and limitations, some of which are described in previous sections. The CLI provides outputs that help the user understand a classifier configuration and why a certain command was accepted or rejected. These are explained below.

15.3.9.1 Classifier Resources Exceeded (NORES)

This could occur when all registers are full after the command is invoked.

Note: The system will generate a warning message informing the user if or when classifier resource capacity or capabilities have been exceeded on the slot(s) impacted by the provisioning change. The user should investigate classifier-related provisioning, such as IGMP, DHCPRELAY, VLAN (for per-VLAN UFO and HVLAN), EPSR, INTERFACE (TAGALL option for HVLAN), ACCESSLIST, and CLASSIFIER to determine the reason for the message.

Exceeding classifier resources raises a NORES alarm. An example of setting the NORES alarm is illustrated below.

```

officer SEC> CREATE CLASSIFIER IPS1 IPSOURCE=10.10.10.1
Info (010017): Operation Successful

// 12 more classifiers were created, IPS2 through IPS13.

officer SEC> CREATE CLASSIFIER IPDROP protocol=ipv4
Info (010017): Operation Successful
officer SEC> ADD ACTION CLASSIFIER IPDROP DROP
Info (010017): Operation Successful
officer SEC> ADD ACTION CLASSIFIER IPDROP COUNT
Info (010017): Operation Successful
officer SEC> ADD ACTION CLASSIFIER IPS1,IPS2,IPS3,IPS4,IPS5,IPS6,IPS7,IPS8,IPS9,IPS10,IPS11,
IPS12,IPS13 FORWARD
Info (010017): Operation Successful
officer SEC> ADD CLASS IPS1 INTERFACE 4.4 PRECEDENCE 51
Info (010017): Operation Successful

// 13 more classifiers were added to interface 4.4, IPS2 through IPS13.

```

```
officer SEC> ADD CLASSIFIER IPDROP INTERFACE 4.4 PRECEDENCE 68
```

```
Info (010017): Operation Successful
```

```
officer SEC> SHOW CLASS ALL INTERFACE 4.4
```

```
--- Classifier Configuration Data ---
```

Interface	Rank	Name	Field Match(es)	Action(s)
ETH: 4.4	51	ips1	IPSOURCE= 10.10.10.1/32	FORWARD
	52	ips2	IPSOURCE= 10.10.10.2/32	FORWARD
	53	ips3	IPSOURCE= 10.10.10.3/32	FORWARD
	54	ips4	IPSOURCE= 10.10.10.4/32	FORWARD
	55	ips5	IPSOURCE= 10.10.10.5/32	FORWARD
	56	ips6	IPSOURCE= 10.10.10.6/32	FORWARD
	57	ips7	IPSOURCE= 10.10.10.7/32	FORWARD
	58	ips8	IPSOURCE= 10.10.10.8/32	FORWARD
	59	ips9	IPSOURCE= 10.10.10.9/32	FORWARD
	60	ips10	IPSOURCE= 10.10.10.10/32	FORWARD
	61	ips11	IPSOURCE= 10.10.10.11/32	FORWARD
	62	ips12	IPSOURCE= 10.10.10.12/32	FORWARD
	63	ips13 (NORES)	IPSOURCE= 10.10.10.13/32	FORWARD
68	ipdrop	PROTOCOL= IPV4	DROP COUNT	

```
officer SEC> SHOW ALARMS ALL
```

```
Slot 19, Port 00          Classifier Resources      Minor
                          Exceeded
```

15.3.9.2 Error (ERR)

This error would occur in the instance of a software error. This would be different from a NOSPT or NORES.

In the example below, the user has added three classifiers to a port and attempts to add a fourth. An error appears saying the card cannot accept the fourth classifier because the number of masks supported by the card has been exceeded. As a result, when the user displays the classifiers for the port a No Resources error appears next to the classifier. The user can delete the fourth classifier and the (ERR) is removed from the display.

```
officer SEC> ADD CLASSIFIER=ip1 PORT=11.0 PRECEDENCE=51
```

```
Info (010017): Operation Successful
```

```
P
```

```
officer SEC> ADD CLASSIFIER=ip2 PORT=11.0 PRECEDENCE= 52
```

```
Info (010017): Operation Successful
```

```
officer SEC> ADD CLASSIFIER=ip3 PORT=11.0 PRECEDENCE= 53
```

```
Info (010017): Operation Successful
```

```
officer SEC> ADD CLASSIFIER=ip4 PORT=11.0 PRECEDENCE= 54
```

```
Info (010017): Operation Successful
```

```
officer SEC> SHOW CLASSIFIER=ALL PORT=11.0
--- Classifier Configuration Data -----
Port Rank Name           Field Match(es)           Action(s)
-----
11.0 51  ip1 (ERR)             IPSOURCE=1.1.1.1/1
      52  ip2 (ERR)             IPSOURCE=1.1.1.1/2
      53  ip3 (ERR)             IPSOURCE=1.1.1.1/3
      54  ip4 (ERR)             IPSOURCE=1.1.1.1/4
      145 tel esyn_defaul t_ IPDEST=MULTI CAST           SETVPRI ORITY=1
          vi deo (ERR)
```

```
officer SEC> DELETE CLASSIFIER=ip4 PORT=11.0
Delete classifier(s) from port(s) (Y/N)? y
Info (010017): Operation Successful
officer SEC> SHOW CLASSIFIER=ALL PORT=11.0
--- Classifier Configuration Data -----
Port Rank Name           Field Match(es)           Action(s)
-----
11.0 51  ip1                   IPSOURCE=1.1.1.1/1
      52  ip2                   IPSOURCE=1.1.1.1/2
      53  ip3                   IPSOURCE=1.1.1.1/3
      145 tel esyn_defaul t_ IPDEST=MULTI CAST           SETVPRI ORITY=1
          vi deo
```

15.3.9.3 No Support (NOSPT)

An example of when this error will be raised is when classifiers are configured on a port and the software for the card where the port resides is downgraded to a release that doesn't support classifiers. The system will generate a NOSPT error.

In the example below, the user has created a classifier with an IP source and has associated this with a VRIORITY. The system allows this, but when the user tries to associate the classifier with a port, a message is output stating that pbit marking is only supported for IP mulitcast in a certain range. The command is accepted, but when the user lists the classifiers for that port, No Support (NOSPT) is displayed next to the classifier, meaning the classifier will not be used.

```
officer SEC> SHOW CLASSIFIER=ALL PORT=11.0
--- Classifier Configuration Data -----
Port Rank Name           Field Match(es)           Action(s)
-----
11.0 51  ip1                   IPSOURCE=1.1.1.1/1
      52  ip2                   IPSOURCE=1.1.1.1/2
      53  ip3                   IPSOURCE=1.1.1.1/3
      145 tel esyn_defaul t_ IPDEST=MULTI CAST           SETVPRI ORITY=1
          vi deo
```

```

officer SEC> CREATE CLASS=badclass ips=3.3.3.3
Info (010017): Operation Successful

officer SEC> ADD ACTION CLASSIFIER=badclass SETVPRIORITY=2
Info (010017): Operation Successful

officer SEC> ADD CLASSIFIER=badclass PORT=11.0 PRECEDENCE=146
Info (010017): Operation Successful

officer SEC> SHOW CLASSIFIER=ALL PORT=11.0

```

```

--- Classifier Configuration Data -----
Port Rank Name          Field Match(es)          Action(s)
-----
11.0 51  ip1                IPSOURCE=1.1.1.1/1
      52  ip2                IPSOURCE=1.1.1.1/2
      53  ip3                IPSOURCE=1.1.1.1/3
      145 tel esyn_defaul t_ IPDEST=MULTICAST      SETVPRIORITY=1
      vi deo
      146 badclass (NOSPT) IPSOURCE=3.3.3.3/32      SETVPRIORITY=2

```

Another scenario is when a combination of classifiers and specific values for match fields is not allowed. For example, with the ADSL24A card, if the user installs a classifier that tries to match the LSAP field to a value other than NETBIOS and SNAP (refer to [Table 15-1](#)), then the NOSPT error code appears.

15.3.10 Classifiers and Subinterface interactions

The user cannot provision classifiers against any subinterface other than VC0 on the DSL interface cards. The command will be rejected. For example:

```

officer SEC>> ADD CLASSIFIER TEST1 INTERFACE=17.10.1 PRECEDENCE=51
Warning (040805): Classifiers not supported on the following interfaces:
                  ETH: [17.10.1]
Error (010009): Operation Failure

```

All classifiers provisioned on VC0 are, in fact, applied to all the VCs on that port, because classifiers are port based. Refer to [12.4.4](#).

15.4 Access Control List

15.4.1 Overview

Access Control Lists (ACLs) provide traffic filtering functionality. They are shortcuts for creating classifiers. Unlike classifiers, ACLs are a more easily understood syntax and a more common method for applying filters.

ACLs give the user the ability to define traffic types by protocol (in English words) without the need to know the exact IP/TCP/UDP characteristics of the protocol specified.

An ACL is composed of a set of rules, each rule specifies a traffic stream to be permitted or denied entry into the provider's network.

Provisioning allows one access list per port or interface. It can be applied to the ingress traffic on the specified interface.

In addition to the standard WAN and LAN physical interfaces, the user can apply an access list to control traffic associated with the management interface (MGMT and inband). The management interface refers to either the physical Ethernet port on the control module faceplate or the virtual management port accessed through in-band traffic paths.

The following lists the packet attributes and protocols that can be provisioned in an ACL. These attributes may be combined to form an expression to compare against the attributes of a packet as it enters or exits an external port.

- Ethernet MAC source and/or destination address.
- Layer 2 protocol type field.
- IP source and/or destination address with a subnet mask.
- IP protocol type field.
- UDP source and/or destination port numbers.
- TCP source and/or destination port numbers.
- APPLICATION abstract rule types that provide a predefined set of rules such as a rule to permit or deny NETBIOS, DHCP and subscriber multicast traffic (FUM). These rule attributes are expanded by the traffic management system into one or more rules.
- An access list can be created and provisioned by the user as a standalone configuration.
- The access list is managed by *name*.
- Rules may be added, modified or deleted at any time. The order of rules in an access list convey an evaluation priority. Earlier rules that may overlap with rules that occur later in the list will be given priority if the actions on the two rules conflict.
- The user can apply the access list to an interface or a set of interfaces. The system will reject a users request if an attribute of the access list is not compatible with interface's capabilities.

ACLs will be qualified by the fMAP system as follows:

- Conflicting rules will be rejected.
- Internally, there is some automatic match rule derivation. For example, if the user configures an access list with a TCP source port rule, the system will automatically add match rules for the layer 2 protocol field to be IPv4, and the IP protocol field to be TCP.
- Some fMAP hardware may not support certain packet fields as specified in the access list rules. Qualification of this capability will be based on the associated card's provisioned preference-load before accept-

ing the access list for an interface. This qualification will not guarantee the accepted configuration will be successful, there can still be alarms resulting from hardware resource exhaustion or conflicts caused by certain combination of rules.

15.4.2 Usage Notes

- All manually configured filters **MUST** be removed prior to configuring ACL filters. This guarantees that there will be no conflicts.
- Dynamic DHCP IP filters are preserved.
- The user can set the default DENY or PERMIT rule for accesslists using the CREATE ACCESSLIST command. See the example that follows for details.
- When upgrading to Release 3.0, previously defined IP filtering classifier sets must be manually converted by the user, the system will not automatically convert the previous classifiers.
- Filtering can be applied to the MGMT and inband interfaces. This allows the user to block certain packets at the CFC interface preventing them from entering the user's network.
- Hardware classification resources on ingress ports are limited. In the event the system experiences contention for resources, an alarm will be raised on the port.
- The user is not allowed to add an access list to a port that currently has classifiers in the precedence range reserved for access lists. The user must remove those classifiers on the port before being allowed to add the access list.
- Mapping of a classifier configuration port alarm will not be direct. If an accesslist configuration error occurs, a system alarm or error indication will be generated. The user can observe, using the SHOW ALARMS command, an error against an ACL. From there, the user can use the **SHOW ALARMS** command on the port in combination with **SHOW ACCESSLIST <acl-name> INTERFACE <interface-name>** and **SHOW CLASSIFIER ALL** on the interface command to understand the root cause of the alarm. The cause of the error will be revealed in the **SHOW ACCESSLIST <acl-name> INTERFACE <interface-name>** output. Users can normally diagnose the error from that output. To see exactly which classifier caused the problem use **SHOW CLASSIFIER ALL INTERFACE <interface-name> FULL** (note that this is usually not required).
- The user must be careful when applying the FUM (From User Multicast) application rule. If applied to the wrong upstream port, for example a GE port, multicast video could be disabled for the whole system.
- Because accesslists use classifiers, the user may observe classifier configuration failure logs/alarms when configuring ACLs. Refer to the **fMAP Log Manual** for information about classifier configuration failure.
- Whenever IGMP snooping is enabled on a port, classifiers are configured against the ingress side of that port. There is potential for interaction between the classifiers configured as a result of an access list and that associated with IGMP snooping. For example, if the access list is set up to block all IP traffic from entering the shelf, then effectively IGMP ports will be filtered.
- Similar to classifiers applied to a LAG, an access list cannot be applied to a member port of a LAG, instead it must be applied to the LAG itself. The user also cannot apply a classifier or access list to an empty LAG (i.e. one with no port members).

- If Subtending functionality is implemented on the system, the designated downstream port will disable the user specified access list currently applied to that port. In this case, the user should be aware that the access list is currently not applied. See Network Topologies [13.4](#).

Note: The system will generate a warning message informing the user if or when resources have been exceeded. The user should investigate classifier-related provisioning, such as IGMP, DHCPRELAY, VLAN (for per-VLAN UFO and HVLAN), EPSR, INTERFACE (TAGALL option for HVLAN), ACCESSLIST, and CLASSIFIER to determine the reason for the message.

15.4.3 Examples

First, the user can set the default rule. Here the user sets the rule to PERMIT:

```
officer SEC> CREATE ACCESSLIST metro_stb_range DEFAULTRULE=PERMIT
Info (010017): Operation Successful
officer SEC> SHOW ACCESSLIST ALL
```

```
--- Access Lists -----
Name                Interfaces          Rule Action  Fields
-----
metro_stb_range     -- PERMIT
```

Referring to the CREATE CLASSIFIER command below, the user sets the rule to DENY:

```
officer SEC>> CREATE ACCESSLIST rural_stb_range DEFAULTRULE=DENY
Info (010017): Operation Successful
officer SEC>> SHOW ACCESSLIST rural_stb_range
```

```
--- Access Lists -----
Name                Interfaces          Rule Action  Fields
-----
rural_stb_range     -- DENY
```

Continuing with the example, assume the user wishes to only allow traffic originated from a range of IP addresses assigned to customers using the user's set-top boxes (172.16.5.0 – 172.16.5.15).

```
officer SEC> CREATE ACCESSLIST metro_stb_range RULE=PERMIT IPSOURCE=172.16.5.0 SOURCEMASK=255.255.255.240
Info (010017): Operation Successful
officer SEC>> SHOW ACCESSLIST metro_stb_range
```

```
--- Access Lists -----
Name                Interfaces          Rule Action  Fields
-----
metro_stb_range     1 PERMIT      IPSOURCE=172.16.5.0
                   -- DENY        SOURCEMASK=255.255.255.240
```

This command creates the configuration record. Nothing is changed with respect to an external port on the fMAP shelf. The SOURCEMASK specification qualifies the IPSOURCE address to apply to the 16 contiguous addresses (in other words, do not compare the least significant four bits). The SOURCEMASK specification must be a contiguous region starting with the most significant portion of the address.

The value is specified as a mask for the IPSOURCE field. For example, an IPSOURCE of 192.168.1.0 with a SOURCEMASK of 255.255.255.0 matches 192.168.1.0 to 192.168.1.255. If no mask is provided then 255.255.255.255 is assumed. For example, 255.255.240.00 would be valid but 255.0.255.0 would not.




FIGURE 15-2 IPSOURCE value of 192.168.1.0/24

In addition, it should be noted there are other packets fields that will be qualified in addition to the IP address field such as the IP version field (version 4) and the etherframe protocol field of IP.

```
officer SEC> ADD ACCESSLIST METRO_STB_RANGE RULE=DENY PROTOCOL=IPV4
Info (010017): Operation Successful
```

This rule will deny all other IP packets from entering the equipment.

```
officer SEC> ADD ACCESSLIST METRO_STB_RANGE RULE=PERMIT
Info (010017): Operation Successful
```

Let other implicit traffic such as ARP etc. into the equipment.

```
officer SEC> ADD ACCESSLIST METRO_STB_RANGE INTERFACE=ETH: 0. 4
Info (010017): Operation Successful
```

This command will associate the access control list “metro_stb_range” with the specified Ethernet interface (slot 0 port 4).

```
officer SEC>> SHOW ACCESSLIST=METRO_STB_RANGE
--- Access Lists -----
Name          Interfaces      Rule Action  Fields
-----
metro_stb_range  ETH: [0. 4. 0]  1  PERMIT    IPSOURCE=172. 16. 5. 0
                                     SOURCEMASK=255. 255. 255. 240
                                     2  DENY      PROTOCOL=IPV4
                                     3  PERMIT
                                     --  DENY
```

This example will set up a more complex access list to block NETBIOS and DHCPSEVER traffic, allow only traffic sourced from known MAC addresses, and deny all DHCP traffic. Apply this to port 5 on the 16 port interface card in slot 0 of a fMAP system.

```
officer SEC> CREATE ACCESSLIST mal_n_acl RULE DENY APPLICATION=NETBIOS
Info (010017): Operation Successful
```

Create a new access list configuration record and set the top priority rule to deny all IP traffic destined for NET-BIOS. This includes traffic from the permitted MAC source address in the subsequent rule.

```
officer SEC> ADD ACCESSLIST main_acl RULE=PERMIT MACSOURCE=11: 22: 33: 44: 55: 66
Info (010017): Operation Successful
```

This rule will permit all remaining traffic from the specified Ethernet MAC source address.

```
officer SEC> ADD ACCESSLIST main_acl RULE=DENY APPLICATION=DHCPSEVER
Info (010017): Operation Successful
```

Deny all DHCP traffic. This does not include any DHCP traffic originated from MAC source address 11:22:33:44:55:66.

```
officer SEC>> ADD ACCESSLIST main_acl INTERFACE=eth: 0. 5
Info (010017): Operation Successful
```

This command will associate the access control list “main_acl” with the specified Ethernet interface (ports 5 on the interface card in slot 0).

```
officer SEC>> SHOW ACCESSLIST ALL
```

```
--- Access Lists -----
Name           Interfaces          Rule Action  Fields
-----
main_acl       ETH: [0. 5. 0]      1   DENY      APPLICATION=NETBIOS
                2   PERMIT    MACSOURCE=11: 22: 33: 44: 55: 66
                3   DENY      APPLICATION=DHCPSEVER
                --  DENY
metro_stb_range ETH: [0. 4. 0]      1   PERMIT    IPSOURCE=172. 16. 5. 0
                SOURCEMASK=255. 255. 255. 240
                2   DENY      PROTOCOL=IPV4
                3   PERMIT
                --  DENY
rural_stb_range                --  DENY
```

This example will illustrate how to modify the access list, main_acl, configured and applied in the previous example. In this case, the user chooses to raise the priority of the rule to deny all DHCP traffic to higher than the rule to permit traffic from a specific MAC source address.

```
officer SEC> DELETE ACCESSLIST main_acl RULE=3
Info (010017): Operation Successful
```

Referring to the output of the SHOW ACCESSLIST command from the previous example, rule number 3 is the rule to deny DHCP traffic. However, this rule should have been configured as rule number 2. This command will remove that rule from the access list named “main_acl”. In this case, since the access list is currently associated with a set of external ports, this command will affect those ports.

```
officer SEC> ADD ACCESSLIST main_acl RULE=DENY APPLICATION=DHCPSEVER BEFORE=2
Info (010017): Operation Successful
```

This command adds the rule *before* rule number 2 as listed in the previous SHOW ACCESSLIST command. This will result in all DHCP traffic being denied from entering the fMAP equipment on the interface eth:0.5.

```
officer SEC> SHOW ACCESSLIST ALL
```

```
--- Access Lists -----
Name          Interfaces      Rule Action  Fields
-----
mai_n_acl     ETH: [0. 5. 0]  1   DENY     APPLICATI ON=NETBI OS
               2   DENY     APPLICATI ON=DHCPSERVER
               3   PERMI T   MACSOURCE=11: 22: 33: 44: 55: 66
               --   DENY
metro_stb_range ETH: [0. 4. 0]  1   PERMI T   IPSOURCE=172. 16. 5. 0
               SOURCEMASK=255. 255. 255. 240
               2   DENY     PROTOCOL=I PV4
               3   PERMI T
               --   DENY
rural_stb_range      --   DENY
```

The user can set a default rule on an accesslist using the SET ACCESSLIST <accesslist name> DEFAULTRULE command. For example, the user originally set the accesslist **rural_stb_range** to DENY. This rule can be reset to PERMIT using the SET ACCESSLIST <accesslist name> DEFAULTRULE command. An example follows:

```
officer SEC>> SHOW ACCESSLIST= rural_stb_range
```

```
--- Access Lists -----
Name          Interfaces      Rule Action  Fields
-----
rural_stb_range      --   DENY
```

```
officer SEC>> SET ACCESSLIST rural_stb_range DEFAULTRULE=PERMIT
```

```
Info (010017): Operation Successful
```

```
officer SEC>> SHOW ACCESSLIST= rural_stb_range
```

```
--- Access Lists -----
Name          Interfaces      Rule Action  Fields
-----
rural_stb_range      --   PERMI T
```

15.4.4 ACL and Manual Classifier configuration

Users can configure classifiers using ACLs or manually. The following table compares some commands from the two methods.

TABLE 15-4 ACL and Manual Classifier configuration

ACL	Manual
CREATE ACCESSLIST metro_stb_range RULE=DENY IPSOURCE=172.16.5.0 SOURCE- MASK=255.255.255.240	CREATE CLASSIFIER=ipfilt1 IPSOURCE=172.16.5.0/28 ADD ACTION CLASSIFIER=ipfilt1 DROP
ADD ACCESSLIST metro_stb_range RULE=DENY PROTOCOL=IPV4	CREATE CLASSIFIER=ipfilt1 SET CLASSIFIER=ipfilt1 PROTOCOL=IPV4 ADD ACTION CLASSIFIER=ipfilt1 DROP
ADD ACCESSLIST main_acl RULE=PERMIT MAC- SOURCE=11:22:33:44:55:66	CREATE CLASSIFIER=netbios_filter03 SET CLASSIFIER=netbios_filter03 MAC- SOURCE=11.22.33.44.55.66
ADD ACCESSLIST main_acl RULE=DENY APPLICA- TION=DHCPSEVER	CREATE CLASSIFIER=netbios_filter03 SET CLASSIFIER=netbios_filter03 UDPPORTD- EST=67 ADD ACTION CLASSIFIER NETBIOS_FILTER03=DROP
SET ACCESSLIST rural_stb_range DEFAULT- RULE=PERMIT DEFAULTRULE - Rule which defines what to do with pack- ets which don't match any other rules in access list. Can either be PERMIT all or DENY all.	CREATE CLASSIFIER=ipfilt1 IPSOURCE=172.16.5.0/28 ADD ACTION CLASSIFIER=ipfilt1 DROP SET CLASSIFIER=ipfilt1 PROTOCOL=IPV4 SET CLASSIFIER=ipfilt1 IPDEST=10.0.0.0/8 SETDEFAULTS CLASSIFIER=IPFILT1 IPDEST SETDEFAULTS CLASSIFIER - clears the speci fied user defined match rule (or rules) from the CLASSIFIER.

Table 15-5 lists the available classifiers match rules.

TABLE 15-5 Access List associated command parameters

Parameter	Description
accesslistname	An arbitrary name. The name of the ACCESSLIST being affected.
RULE	The RULE key word indicates the following arguments are to form a single rule entry on the specified access list. A RULE must be marked with either the PERMIT or DENY action. PERMIT indicates to allow packets matching the accompanying match arguments to pass into the fMAP shelf. DENY prevents matching traffic from entering the shelf.

TABLE 15-5 Access List associated command parameters

Parameter	Description
MACDEST	MAC destination address to match (e.g. 00:0C:25:00:13:8C)
MACSOURCE	MAC source address to match (e.g. 00:0C:25:00:13:8C)
PROTOCOL	The value of the Layer 2 protocol field to match. The command will accept IPV4, IPV6, or a numeric value.
IPSOURCE	IPv4 specified source address and optional address mask via the MASK keyword. Creates a IPv4 field mask entry if it doesn't already exist and creates a rule match entry.
IPDEST	IPv4 specified destination address and optional address mask via the MASK keyword. Creates a IPv4 field mask entry if it doesn't already exist and creates a rule match entry.
IPPROTOCOL	IP protocol field value to match. The command will accept TCP, UDP, ICMP, IGMP, or a protocol number.
UDPPORTSRC	The value of the UDP source port to match.
UDPPORTDEST	The value of the UDP destination port to match. Note: Filters can be added for DHCPSEVER and DHCPCLIENT. For example, ADD ACCESSLIST main_acl RULE=DENY APPLICATION=DHCPSERVER, adds a rule to the access list main_acl to filter (deny) packets to a DHCP server. The other form of the filter for DHCP is APPLICATION=DHCPCCLIENT.
TCPPORTSRC	The value of the TCP source port to match.
TCPPORTDEST	The value of the TCP destination port to match.
APPLICATION	The name of the APPLICATION to match on. APPLICATION is one of several predefined match rules. For example, the APPLICATION TELNET matches all packets with TCPDEST=23. Allowed values are NETBIOS, DHCP, FUM, TELNET, SSH, or SNMP.
INTERFACE	This option identifies one or more interfaces to associate this access list with.
type:id-range	The interface type followed by a range of identifiers (e.g. eth:1.3-1.8)
id-range	Range of interfaces IDs, such as "4.0, 4.5"
ifname-list	A list of one or more logical interface names.
BEFORE	The RULE number that the new RULE will be inserted ahead of.

TABLE 15-5 Access List associated command parameters

Parameter	Description
DESTMASK	DESTMASK works with the IPDEST match rule field to match on any IPv4 packet with the specified IP destination address. The value is specified as a mask for the IPDEST field. For example, an IPDEST of 192.168.1.0 with a DESTMASK of 255.255.255.0 matches 192.168.1.0 to 192.168.1.255. If no mask is provided then 255.255.255.255 is assumed. The DESTMASK must be a contiguous series of bits starting with the MSB. For example, 255.255.240.00 would be valid, but 255.0.255.0 would not.
SOURCEMASK	SOURCEMASK works with the IPSOURCE match rule field to match on any IPv4 packet with the specified IP source address. The value is specified as a mask for the IPSOURCE field. For example, an IPSOURCE of 192.168.1.0 with a SOURCEMASK of 255.255.255.0 matches 192.168.1.0 to 192.168.1.255. If no mask is provided then 255.255.255.255 is assumed. The SOURCEMASK must be a contiguous series of bits starting with the MSB. For example, 255.255.240.00 would be valid, but 255.0.255.0 would not.

15.5 Traffic Management Alarms

When a command to configure a classifier would allow something the system cannot support, a rejection message appears and the user is prevented from entering the command.

In cases where the user wishes to pre-provision a card and the user inputs commands that the system can support, an alarm message can still be raised when the card, when physically inserted, does not support the QoS features that have been already provisioned.

For a card, these messages are:

- QOS configuration failed
- QOS resources exceeded
- QOS configuration not supported

The resolution is to ensure the card being inserted does match the product version and that the QoS configuration has been correctly pre-provisioned.

Note: These alarms are minor if the card is enabled. If the card is disabled, the messages are for information only.

For a port, there are alarm messages as well:

- Classifier configuration failed
- Classifier resources exceeded
- Classifier configuration not supported

In most cases, the user should verify the classifier configuration is correct. For more information, refer to the Troubleshooting Section.

15.6 Ingress Metering

To provide metering (policing) of ingress traffic, classifiers are defined and then associated with a RATE (the CIR) and BURSTSIZE (the CBS). The actions that are then taken when the traffic exceeds the profile is DROP or FORWARD the packets. The NCCOUNT option allows for counting non-conforming packets.

The major steps and commands used to implement ingress metering are as follows:

- Create a set of classifiers, in this example using an IPSOURCE and VLANID values.


```
CREATE CLASSIFIER=i pMeter1 IPSOURCE=1. 0. 0. 0/8
CREATE CLASSIFIER=vlanMeter1 VLANID=1
CREATE CLASSIFIER=i pMeter2 IPSOURCE=1. 2. 3. 4
CREATE CLASSIFIER=del ete IPSOURCE=ANY
```
- Create a set of meters.


```
CREATE TRAFFICDESCRIPTOR=small_rate RATE=128K BURSTSIZE=8KB
CREATE TRAFFICDESCRIPTOR=large_rate RATE=2M BURSTSIZE=64KB
CREATE TRAFFICDESCRIPTOR=throttle RATE=1M BURSTSIZE=4KB
```
- Associate the meters with the classifiers.


```
ADD TRAFFICDESCRIPTOR=small_rate CLASSIFIER=i pMeter1 NCDROP NCCOUNT=OFF
ADD TRAFFICDESCRIPTOR=large_rate CLASSIFIER=i pMeter2 NCDROP NCCOUNT=ON
ADD TRAFFICDESCRIPTOR=throttle CLASSIFIER=del ete NCCOUNT=ON
ADD TRAFFICDESCRIPTOR=large_rate CLASSIFIER=vlanMeter1
```
- Associate a classifier (with a meter) to an interface.


```
ADD CLASSIFIER=i pMeter1 INTERFACE=ETH: 7. 0 PRECEDENCE=65
ADD CLASSIFIER=vlanMeter1 INTERFACE=ETH: 0. 1, 0. 2, 7. 1 PRECEDENCE=65
```

Note: If the user attempts to associate a classifier (metered) with an interface that cannot support the CIR rate of the meter, the command is rejected.

Once the ADD CLASSIFIER command has been invoked, the system may generate a warning message at the user's CLI session stating that classifier capacity or capabilities have been exceeded on the slot(s) impacted by the provisioning change. The user should investigate classifier-related provisioning, such as IGMP, DHCPRELAY, VLAN (for per-VLAN UFO and HVLAN), EPSR, INTERFACE (TAGALL option for HVLAN), ACCESSLIST, and CLASSIFIER to determine the reason for the message.

15.7 Egress Port Rate Limiting

The fMAP product supports egress port rate limiting on both the SM and NM sides, with some limitations.

The major steps and commands used to implement egress port rate metering are as follows:

- Create a set of egress limiter names with a RATE and BURSTSIZE.


```
CREATE EGRESSLIMITER=sub_gold RATE=1M BURSTSIZE=128KB
CREATE EGRESSLIMITER=sub_silver RATE=1M BURSTSIZE=64KB
```

```
CREATE EGRESSLIMITER=sub_bronze RATE=1M BURSTSIZELIMIT=16KB
CREATE EGRESSLIMITER=NMLimit RATE=8M BURSTSIZELIMIT=8KB
```

- Assign each limiter to a port/port range or interface/interface range

```
ADD EGRESSLIMITER=sub_gold INTERFACE=ETH: 7. 0
ADD EGRESSLIMITER=sub_silver INTERFACE=ETH: 7. 1
ADD EGRESSLIMITER=sub_bronze INTERFACE=ETH: 7. 2
ADD EGRESSLIMITER=NMLimit INTERFACE=ETH: 0. 1
```

Note: For all ADSL cards, configuration of egress rate limiting is not supported directly, but can be set using the ADSL maximum downstream rate. Refer to [Table 15-1](#).

15.8 Priority Queueing

15.8.1 Overview

A priority may be set on the p-bit ingress (by a remarking action on a classifier), and then that priority is used throughout the system at each egress queue. Tagged traffic retains its p-bit settings if it is not remarked and untagged traffic receives a p-bit setting of 0 if not remarked. The values in the `SET QoS VLANQUEUEMAP` command accept a queue number (0 = lowest priority; higher numbers = higher priority).

There are two system settings which indicate which egress queue a packet is inserted into based on the p-bit value of the packet. There is one system setting for interfaces that are capable of supporting 4 egress queues and one for interfaces that are capable of supporting 8 egress queues.

- The CFC24, CFC6, and NM interfaces hosted by these cards will support 8 queues.
- The CFC4 and all SMs support 4 queues.

Each card can be configured with the corresponding p-bit to queue mapping system setting for the maximum number of queues the card's interface supports.

System default settings are:

- 0,1,2,3,4,5,6,7 for 8 queues
- 0,0,1,1,2,2,3,3 for 4 queues.

For software upgrades to release 5.0, the system will maintain the existing p-bit to queue map setting as the 4 queue setting and copy it to the 8 queue setting. However, if the system is starting from a default database, the defaults will be set.

Note: Only strict priority (SP) will be supported. Weighted round robin (WRR) scheduling or any other scheduling algorithm is not supported.

The user can map the 4 queue and 8 queue settings. The system will validate the attempted settings and will not allow the user to provision conflicting values for the 4 queue and 8 queue map settings. Conflicting values are those that give a different *relative priority* for any particular p-bit when comparing each set of values. For example:

- If the user entered 3,3,0,0,2,2,1,1 and 6,7,1,0,4,5,3,2 for the 4 and 8 queue map respectively, this would be *valid* as both settings order the p-bit in the relative priority order of p-bit 3,2,7,6,4,5,0,1.
- However, if the user had entered 3,3,0,3,2,2,1,1 and 6,7,1,0,4,5,3,2 for the 4 and 8 queue map respectively, it would be rejected by the system as *p-bit 3 has a different relative priority* in each p-bit to queue map.

p-bit values range from 0, lowest priority, to 7, the highest priority.

The user can provision and display the p-bit to queue mappings for the system with the following commands:

```
SET QOS [VLAN4QUEUEMAP=value-map][VLAN8QUEUEMAP=value-map]
```

```
SHOW QOS
```

The value-map is a list of *eight* integers, 0-3 for the VLAN4QUEUEMAP and a list of **eight** integers, 0-7 for the VLAN8QUEUEMAP. The “*eight*” is the number of p-bit values which is independent of the number of queues. Command syntax allows both to be configured simultaneously if desired. An example is shown below.

```
officer SEC>> SHOW QOS
```

```
--- Quality of Service Queue Mapping -----
Priority Level                QoS Egress Queue
-----
p-bit value   8 Queue Capable Interface   4 Queue Capable Interface
-----
0              0                          0
1              1                          0
2              2                          1
3              3                          1
4              4                          2
5              5                          2
6              6                          3
7              7                          3
-----
```

```
officer SEC>> SET QOS VLAN4QUEUEMAP=3, 3, 0, 0, 2, 2, 1, 1 VLAN8QUEUEMAP=6, 7, 1, 0, 4, 5, 3, 2
Info (010017): Operation Successful
```

```
officer SEC>> SHOW QOS
```

```
--- Quality of Service Queue Mapping -----
Priority Level                QoS Egress Queue
-----
p-bit value   8 Queue Capable Interface   4 Queue Capable Interface
-----
0              6                          3
1              7                          3
2              1                          0
3              0                          0
4              4                          2
5              5                          2
6              3                          1
7              2                          1
-----
```

15.8.2 IGMP Multicast queue priority with telesyn_default_video

IGMP multicast traffic is remarked by the “`telesyn_default_video`” classifier to p-bits=4. This corresponds to queue 4 of 8 queues or queue 2 of 4 queues. If the user assigns lower priority traffic to a higher queue than IGMP traffic, IGMP priority has effectively been lowered. However, under most conditions, RTP, MGCP, and EPSR traffic should all be higher priority than IGMP.

Note: In release 6.0, the `telesyn_default_video` classifier will survive an upgrade but is not created for a new installation. If the user wishes to have the same classifier behavior for a new installation, the user will need to create an equivalent one. Refer to [15.3.6](#).

15.8.3 Changing Queue Mapping and Disabling/Enabling Interfaces

The user should understand that the setting of the queue mapping is a network engineering issue, and the setting or changing of this mapping is not trivial. Moreover, for the fMAP system, when the command to change the queues is entered, all affected interfaces are automatically disabled, so this is a service affecting procedure. As part of the command, the interfaces are then enabled and the new queue mapping takes effect.

Caution: For the GE4 and GE2RJ, all interfaces must be manually disabled before entering the queue mapping command. After entering the command, the user must manually enable all the interfaces. Otherwise, packets on these interfaces may not be prioritized according to the new settings.

15.9 Product Support for Rate and Burst Size

[Table 15-6](#) lists the Burst Size support when configuring the NM and SM ports.

TABLE 15-6 Burstsize Support for Ingress Limiters and Traffic Descriptors

Type	Component	Burst Size Support
Ingress Metering	FE10,FX10,ADSL24A	- Burst size starts at 4K (byte) and doubles up to 64K (byte) (4K, 8K, 16K, 32K, 64K)
Egress Port Rate Limiting	FE10,FX10,ADSL24A	Burst size starts at 4KB and doubles up to 64KB (4KB, 8KB, 16KB, 32KB, 64KB)

15.10 Example Configurations

15.10.1 Example (IP Source)

The PCs being used for internet data services have IP addresses. Incoming data to the Ethernet interface ports can be filtered by IP address or a range of IP addresses. This can be used primarily as a security feature since filtering allows the following:

- Traceability - The ability to correlate an IP address with a particular port.
- The prevention of problems due to incorrect configuration of an IP address.

The classifier for IP source addresses works as follows:

- Classifiers are created that are associated with an IP address or IP address range.
IP address ranges are specified using a valid IP address or valid subnet and mask. A range is specified using a '/' character (such as 1.0.0.0/8).
- A classifier can be created with the option ANY for all IP addresses and therefore all packets (such as raw ethernet packets).
- A classifier or set of classifiers then has an action associated with them:
 - DROP - discard the packet at the card. This action excludes the packet.
 - FORWARD - allow traffic to be forwarded. This action includes the packet.
- For the action chosen, a COUNT can be selected as well.

Note: Other types of packets (such as ARP) would be passed and would not be included in these counters. However, they would increment the PMON counter.

- Each classifier is then associated with one or more SM ports.
- Each classifier/port(s) association is given a precedence, with the lowest number receiving the highest precedence. Classifiers on the same port cannot share the same precedence number.

The following shows an example of IP filtering. (The “operation successful” message and confirmation messages have been excluded.)

Note: When listing classifier data for a number of ports, the sort order is by port, then rank.

(Create the classifiers)

```
officer SEC> CREATE CLASSIFIER=ipfilt1 IPSOURCE=192.168.1.0/24
```

```
officer SEC> CREATE CLASSIFIER=ipfilt2 IPSOURCE=10.0.0.0/24
```

officer SEC> **CREATE CLASSIFIER=ipdrop IPSOURCE=ANY**

(Associate actions)

officer SEC> **ADD ACTION CLASSIFIER=ipfilt1 FORWARD**

officer SEC> **ADD ACTION CLASSIFIER=ipfilt2 FORWARD**

officer SEC> **ADD ACTION CLASSIFIER=ipdrop DROP**

officer SEC> **ADD ACTION CLASSIFIER=ipdrop COUNT**

officer SEC> **SHOW CLASSIFIER=ALL**

--- Classifier Configuration Data -----

Name	Field Match(es)	Action(s)
ipdrop	IPSOURCE=ANY	DROP COUNT
ipfilt1	IPSOURCE=192.168.1.0/24	FORWARD
ipfilt2	IPSOURCE=10.0.0.0/24	FORWARD

(Associate ports)

officer SEC> **ADD CLASSIFIER=ipfilt1 PORT=5.0 PRECEDENCE=51**

officer SEC> **ADD CLASSIFIER=ipfilt2 PORT=5.0 PRECEDENCE=52**

officer SEC> **ADD CLASSIFIER=ipdrop PORT=5.0 PRECEDENCE=69**

officer SEC> **SHOW CLASSIFIER PORT=5.0**

--- Classifier Configuration Data -----

Port	Rank	Name	Field Match(es)	Action(s)
5.0	51	ipfilt1	IPSOURCE=192.168.1.0/24	FORWARD
	52	ipfilt2	IPSOURCE=10.0.0.0/24	FORWARD
	69	ipdrop	IPSOURCE=ANY	DROP COUNT

15.10.2 Example (Class of Service)

15.10.2.1 Overview

Figure 15-3 shows a sample QoS configuration for a network, and it has the following attributes:

- The types of traffic flows are associated with specific VLANs.
- The video Head End uses the VLAN ID (VID) range of 3xx; these are then divided up into regions so that certain VLANs are configured on upstream devices that connect with a fMAP product.
- The ISP Head End uses the VID range of 5xx.
- The voice over IP gateway uses the VID range of 7xx.
- The quality of service is defined entirely through priority queueing, so classifier filters are not used.

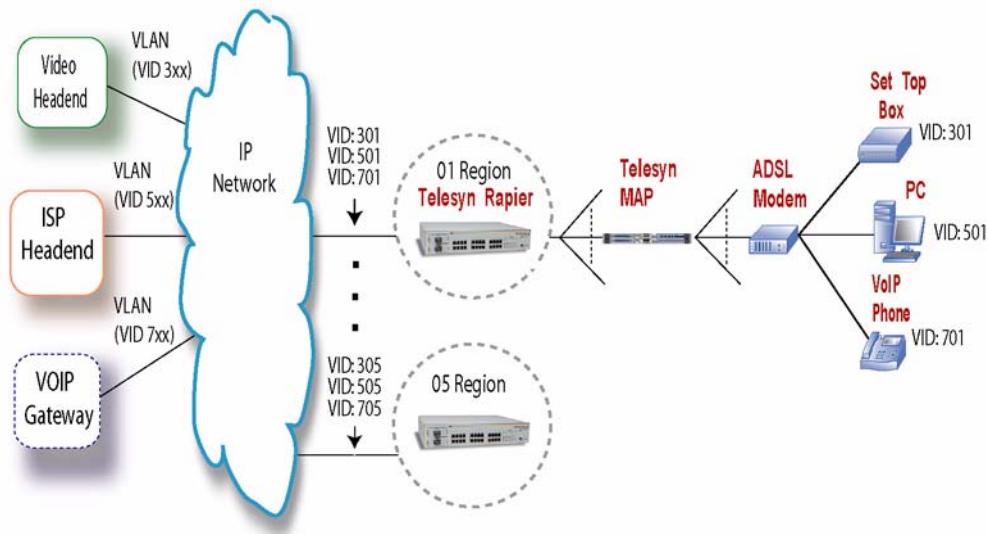


FIGURE 15-3 Sample QoS Network

Table 15-7 lists the quality attributes for these classes of service.:

TABLE 15-7 Classes of Service for a Subscriber

Class of Service	Application	Delay	Jitter	Packet Loss
qos_voice	Voice	Low	Low	Low
qos_video	Video	Low	Undefined	Low

TABLE 15-7 Classes of Service for a Subscriber

Class of Service	Application	Delay	Jitter	Packet Loss
qos_data	internet access	Undefined	Undefined	Low
default	Non-critical	Undefined	Undefined	Undefined

- For qos_voice, the subscriber is set up on a VLAN with the VID 701, and the voice traffic is the only incoming stream with a VLAN tag; all other traffic is untagged and is given a tag by its port association.
- For qos_data the subscriber is set up on a VLAN with VID 501.
- For qos_video the subscriber is set up on a VLAN with VID 301, and will connect to the subscriber’s Set Top Box (STB).

Note: In many scenarios, the qos_data is one queue above the lowest priority queue so that it will not be grouped with all other traffic. In this example, the data traffic is associated with a specific VLAN which is in turn associated with a VRIORITY.

15.10.2.2 Example Commands

The following commands are used to configure the fMAP product to support this configuration:

- Set up the VLAN to queue mapping.

The priority bits and the egress queues are correlated as follows:

TABLE 15-8

Class of Service	802.1 Priority Bits	Egress Queues	Priority
default	0,2,3,4,6	0	Lowest Highest
qos_data	1	1	
qos_video	5	2	
qos_voice	7	3	

The command to provide this would be:

```
SET QOS VLAN8QUEUEMAP=0,1,0,0,0,2,0,3 VLAN4QUEUEMAP=0,0,0,0,0,2,0,3
```

- Create the classifiers for these classes of service and associate them with a VID. The classifiers qos_voice, qos_video, and qos_data are created, and these are associated with the VIDs.

```
CREATE CLASSIFIER=qos_voice VID=701
CREATE CLASSIFIER=qos_video VID=501
CREATE CLASSIFIER=qos_data VID=301
```

- Associate an action with each classifier. In this case, the action is to set the 802.1 p-bits for each classifier.

```
ADD ACTION CLASSIFIER=qos_voice SETVRIORITY = 7
ADD ACTION CLASSIFIER=qos_video SETVRIORITY = 5
```



```
ADD ACTION CLASSIFIER=qos_data SETVPRORITY = 1
```

Note: For these classifiers, the other p-bit values for each class of service could have been used.

- Associate each classifier with an interface or port. Since this configuration will be for the system, the ALL value is used.

```
ADD CLASSIFIER=qos_voice INTERFACE=ALL PRECEDENCE=146
```

```
ADD CLASSIFIER=qos_video INTERFACE=ALL PRECEDENCE=147
```

```
ADD CLASSIFIER=qos_data INTERFACE=ALL PRECEDENCE=148
```

Note: In this example, the precedence order of the QoS classifiers is not important since the flows are mutually exclusive and no filtering is being performed.

- To limit the upstream flow for data, a TRAFFICDESCRIPTOR could be created that has rate and burst values. This is then associated with the classifier. Moreover, these packets that exceed these metering values will be dropped and counted.

```
CREATE TRAFFICDESCRIPTOR=limit_data RATE=128K BURSTSIZE=512KB
```

```
ADD TRAFFICDESCRIPTOR=limit_data CLASSIFIER=qos_data
```

```
NCDROP NCCOUNT=ON
```

15.11 Address Resolution Protocol (ARP) filtering

The ARP is a network protocol that maps a network layer (L3) protocol address to a data link layer hardware address, and is described in RFC 826.

ARP filtering is the ability to “authenticate” ARP messages to ensure that *unauthorized* ARP spoofing is not permitted. ARP spoofing is the act of sending ARP messages with phony IP addresses encoded therein thus corrupting a router's ARP tables. When ARP filtering is enabled the default action is to **drop ALL** ARP packets.

This feature is accomplished by checking the encoded IP address against IP pass filters that are configured for a given interface. If the IP address in the ARP packet matches any IP pass filter on that interface, the ARP is allowed, if not, the ARP is dropped.

ARP filtering conditions any ARP packets with a L3 source address that matches the source address of any IP source filters present on that interface. The system allows ARP packets to pass if there is an IP pass classifier (i.e., IPSOURCE match rule plus a FORWARD action) on the port allowing the IP source address that is in the ARP packet's L3 sender address field. For example:

Consider this classifier configuration with ARP filtering enabled; all ARP packets are dropped as if no classifier were configured because no matching **FORWARD** action classifier was configured.

Interface	Rank	Name	Field Match(es)	Action(s)
ETH:7.0	51	ipf1	IPSOURCE= 10.10.9.1/32	DROP

COUNT

Now, consider these classifier configurations with ARP filtering enabled; any ARP packet with address 10.10.9.x is forwarded, all others are dropped.

Interface	Rank	Name	Field Match(es)	Action(s)
ETH:7.0	51	ipf1	IPSOURCE= 10.10.9.1/32	DROP COUNT
	59	ipf2	IPSOURCE= 10.10.9.0/24	FORWARD COUNT

In the next example, a residence is provisioned with the static IP address 10.9.9.9. ARP packets are initially used to populate the system. As part of the provisioning sequence, ARP filtering is then enabled, and the following classifiers are configured:

Interface	Rank	Name	Field Match(es)	Action(s)
ETH:7.0	51	ipf1	IPSOURCE= 10.10.9.9	FORWARD COUNT
	59	ipf2	IPSOURCE= ANY	DROP

The result of this configuration is that IP packets with an IP source of 10.10.9.9 are allowed, and all ARP packets that do not have this IP source are dropped.

Note that if a count action is present on the IP source filter the associated counter is **not** incremented for matching an ARP packet. If a count action is present on the IP source filter, packets that match the IP are counted.

Also, the L3 match rule fields present in the classifier must be IP SOURCE; the relevant match rule for ARP filtering must be an IP match.

The only L2 match rule fields that may be present are PROTOCOL (= IPv4) and optionally VID/INNERVID.

The fMAP user can configure ARP filtering using the commands described below.

TABLE 15-9 ARP filtering commands

Syntax	Samples for ARP filtering	Description
ENABLE ARPFILTER INTERFACE={type:id-range id-range ifname-list ALL}	ENABLE ARPFILTER INTERFACE 4.4	Enables ARP filtering on an interface(s).

TABLE 15-9 ARP filtering commands

Syntax	Samples for ARP filtering	Description
DISABLE ARPFILTER INTERFACE={type:id-range id-range ifname-list ALL}	DISABLE ARPFILTER INTERFACE 4.4	Disables ARP filtering on an interface(s).
SHOW ARPFILTER [INTERFACE={type:id-range id-range ifname-list ALL}]	SHOW ARPFILTER SHOW ARPFILTER INTERFACE 4.4	Displays information about ARP filters.

System logs are associated with ARP filtering. They are described here.

ARP logs are generated against a system card. They are CARD005 and CARD006. Three logs will be generated with their respective messages as follows:

- CARDARPFILTERINGFAILED - “ARP Filtering Configuration Failed”
- CARDARPFILTERINGEXCEEDED - “ARP Filtering Resources Exceeded”
- CARDARPFILTERINGNOTSUPPORTED - “ARP Filtering Configuration Not Supported”

See the **fMAP Log / Troubleshooting Manual** for more information on system logs.

Following is an example of what occurs when ARP filtering is enabled on a system. Refer to [Figure 15-4](#) when reading the next section of text.




FIGURE 15-4 ARP filtering

1. No classifiers have been configured, therefore all traffic passes.
2. Add a IP filter for 'X' in [Figure 15-4](#). Deny all other IP packets. Note that ARP filtering is not enabled, so ARP and any other non-IP packets can pass.

Example:

```
officer SEC> CREATE ACCESSLIST=iparpflt RULE=DENY IPSOURCE=192.168.1.0 SOURCEMASK=255.255.255.0
officer SEC> ADD ACCESSLIST=iparpflt INTERFACE=17.4
officer SEC> SHOW ACCESSLIST ALL
```

```

--- Access Lists -----
Name          Interfaces      Rule Action  Fields
-----
i parpfilt    ETH: 17. 4      1   DENY     IPSOURCE=192. 168. 1. 0
                                     SOURCEMASK=255. 255. 255. 0
                                     --   PERMIT

```

At this point, IP addresses in the range 192.168.1.0 through 192.168.1.255 are now being specified.

The IPSOURCE value configured at the CLI was "192.168.1.0" with accompanying SOURCE-MASK=255.255.255.0. This means that the first 24 bits are being filtered. The last byte, bits numbered 25-32, imply a decimal value of 0-255, so addresses from 192.168.1.0 through 192.168.1.255 will be blocked because of the DENY rule, which states, do not allow packets matching this RULE to pass. PERMIT, which means, allow packets matching this RULE to pass, would allow the address range of 192.168.1.0 through 192.168.1.255 to pass. The user should keep in mind, however, that ARP packets are still passing through because they have not been specifically blocked.

3. ENABLE ARPFILTER. The system now filters ARP packets so that only the IP address for 'X' in Sender L3 Addr. passes.

Note: The reason the user may want to use ARP filtering and block ARP packets is because other hosts and servers may learn incorrect ARP data.

```

officer SEC> ENABLE ARPFILTER INTERFACE=17. 4
Info (010017): Operation Successful
officer SEC> SHOW ACCESSLIST i parpfilt
--- Access Lists -----
Name          Interfaces      Rule Action  Fields
-----
i parpfilt    ETH: 17. 4      1   DENY     IPSOURCE=192. 168. 1. 0
                                     SOURCEMASK=255. 255. 255. 0
                                     --   DENY
-----
officer SEC> SHOW ARPFILTER INTERFACE=17. 4
--- ARP Filtering Configuration Data -----
Interface      Status
-----
ETH: 17. 4     ENABLED
officer SEC> SHOW ARPFILTER
--- ARP Filtering Configuration Data -----
Interface      Status
-----
ETH: 11. 0     DISABLED
ETH: 17. 0     DISABLED
ETH: 17. 1     DISABLED
ETH: 17. 2     DISABLED
ETH: 17. 3     DISABLED
ETH: 17. 4     ENABLED
ETH: 17. 5     DISABLED

```

ETH: 17. 6	DI SABLED
ETH: 17. 7	DI SABLED
ETH: 17. 8	DI SABLED
ETH: 17. 9	DI SABLED
ETH: 17. 10	DI SABLED
ETH: 17. 11	DI SABLED
ETH: 17. 12	DI SABLED
ETH: 17. 13	DI SABLED
ETH: 17. 14	DI SABLED
ETH: 17. 15	DI SABLED
ETH: 19. 0	DI SABLED
ETH: 19. 1	DI SABLED
ETH: 19. 2	DI SABLED
ETH: 19. 3	DI SABLED
ETH: 19. 4	DI SABLED
ETH: 19. 5	DI SABLED
ETH: 19. 6	DI SABLED
ETH: 19. 7	DI SABLED
ETH: 0	DI SABLED

ARP packets are now blocked on classifier iparpfilt on interface 17.4.

15.12 Dynamic IP Filtering Using DHCP Relay

15.12.1 Overview

IP filters can be automatically configured on specific subscriber ports to allow only the DHCP-allocated IP address to pass traffic.

15.12.1.1 IP Filtering Using DHCP Relay

When DHCP relay agent functionality is enabled, and based upon DHCP server responses, IP filters are applied to subscriber pMACorts allowing **only** the DHCP allocated IP addresses to pass traffic. When DHCP relay agent functionality is enabled, a maximum of eight (8) IP filters will be applied to the interface, based on the number of learned MAC addresses that have been assigned IP addresses from DHCP Servers.

Note: This maximum was five (5) in 6.0.

The IP filters have the following attributes:

- They persist through reboots
- They do not age out
- They can be manually deleted
- They are updated whenever a known MAC is allocated a new IP address by the DHCP relay agent
- Up to 8 addresses will be allowed

Note: *Manually configured IP filters cannot be provisioned on a port when dynamic IP filters are being learned by DHCP on that port.*

Note: *The user may notice that when they execute the SHOW CLASSIFIER ALL FULL command, some internal DHCP classifiers are no longer displayed.*

Refer to section [13.8](#) for more information on DHCP relay agent and details on enabling and disabling IP filtering.

15.13 Treatment of DHCP Packets (MAC/VID only)

Software filtering treats DHCP protocol packets differently; instead of applying all classifiers to these packets, it applies only MAC/VID classifiers. Following are restrictions:

1. This treatment only applies to DHCP packets originating from the DHCP client, i.e. with UDP destPort = 67 and UDP srcPort = 68. If they are pulled off as exception packets, DHCP packets originating from the DHCP server (UDP srcPort = 67) will have all classifiers applied, just like any other exception packet (IGMP, etc.).
2. This functionality affects FE10/FX10/ADSL24 software filtering, where classifiers in both hardware and software are used. In addition, the software filtering change affects inband and management port filtering on the CFC (i.e. if a DHCP client packet is received on these interfaces, user classifiers are selectively applied as described above).
3. The changed functionality is applied to all ports described by restriction 2 above, regardless of whether those ports are considered to be network-facing or customer-facing ports.
4. Classifiers that will be applied to DHCP client packets are ones that contain only rules relating to the following fields:
 - MAC DA
 - MAC SA
 - VID
 - VPRIORITY

If the classifier contains rules related to any other fields, then it will not be applied.

15.14 Classifier Provisioning and Topology Control

The user must be aware of any network topology configurations implemented on the system. Network topology must be taken into consideration when classifiers are configured. If STP and/or EPSR are implemented on the system, the user must ensure that when classifiers are configured on the links used by either of these network topology features, that the same classifier configuration(s) is provisioned on the interface(s) used by the feature. This is because STP and/or EPSR can change upstream links to downstream and downstream links to upstream. See Network Topologies [13.4](#).

15.15 Summary of IP Filtering Options

There are three ways to implement IP filtering:

1. ACL (static)
2. Classifiers (static)
3. Via DHCP relay (dynamic)

ARP filtering works the same as IP filtering for all three methods. It is generic to these three methods of implementing IP filtering. The user can turn it on or off regardless of which method of IP filtering is being used.

15.16 Command Summary for Traffic Management

Table 15-10 lists the commands for traffic management.

TABLE 15-10 Commands for Classifier Management

Noun/Key Word	Verb	Syntax	Description
<i>Note: For complete information on parameters, refer to the fMAP Command Manual</i>			
CLASSIFIER	CREATE	<pre> CREATE CLASSIFIER=classifiername [VID={ 1..4095 ANY }] [VRIORITY={ 0..7 ANY }] [INNERVID={ 1..4095 ANY }] [INNERVPRIORITY={ 0..7 ANY }] [ETHFORMAT={ 802.3 802.3TAGGED 802.3UNTAGGED ETHII ETHIITAGGED ETHIUNTAGGED ANY }] [LSAP={ NETBIOS lsap-value ANY }] [IPDEST={ ipaddress-mask MULTICAST ANY }] [IPSOURCE={ ipaddress-mask ANY }] [IPDSCP={ 0..63 ANY }] [IPPROTOCOL={ TCP UDP ICMP IGMP ipprotocol-number ANY }] [IPTOS={ 0..7 ANY }] [MACDEST={ macaddress MULTICAST ANY }] [MACSOURCE={ macaddress ANY }] [PROTOCOL={ IPV4 IPV6 protocol-type ANY }] [TCPPTDEST={ tcp-port-list ANY }] [TCPPTSOURCE={ tcp-port ANY }] [TCPFLAGS={ { URG ACK RST SYN FIN PSH } [,...] ANY }] [UDPPORTDEST={ udp-port-list ANY }] [UDPPORTSOURCE={ udp-port ANY }] </pre>	Refer to Table 15-3 .

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
CLASSIFIER	SET	<pre> SET CLASSIFIER=classifiername-list [VID={1..4095 ANY}] [VPRIORITY={0..7 ANY}] [INNERVID={1..4095 ANY}] [INNERVPRIORITY={0..7 ANY}] [ETHFORMAT={802.3 802.3TAGGED 802.3UNTAGG ED ETHII ETHIUNTAGGED ANY}] [LSAP={NETBIOS lsap-value ANY}] [IPDEST={ipaddress-mask MULTICAST ANY}] [IPSOURCE={ipaddress-mask ANY}] [IPDSCP={0..63 ANY}] [IPPROTOCOL={TCP UDP ICMP IGMP ipprotocol -number ANY}] [IPTOS={0..7 ANY}] [MACDEST={macaddress MULTICAST ANY}] [MACSOURCE={macaddress ANY}] [PROTOCOL={IPV4 IPV6 protocol-type ANY}] [TCPPORTDEST={tcp-port-list ANY}] [TCPPORTSOURCE={tcp-port ANY}] [TCPPFLAGS={{URG ACK RST SYN FIN PSH} [,...] ANY}] [UDPPORTDEST={udp-port-list ANY}] [UDPPORTSOURCE={udp-port ANY}] </pre>	Change the criteria values for the classifier.
CLASSIFIER	SETDEFAULTS	<pre> SETDEFAULTS CLASSIFIER=classifiername [VID] [VPRIORITY] [INNERVID] [INNERVPRIORITY] [ETHFORMAT] [LSAP] [IPDEST] [IPSOURCE] [IPDSCP] [IPPROTO] [IPTOS] [MACDEST] [MACSOURCE] [PROTOCOL] [TCPPORTDEST] [TCPPORTSOURCE] [TCPPFLAGS] [UDPPORTDEST] [UDPPORTSOURCE] </pre>	Allows the user to reset classifier attributes back to factory defaults.

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
CLASSIFIER	SHOW	<pre> SHOW CLASSIFIER COUNTER [{ INTERFACE={ type:id-range id-range ifname-list ALL } }] SHOW CLASSIFIER COUNTER [{ PORT={ port-list ALL } }] SHOW CLASSIFIER={ classifiername-list ALL } [{ PORT={ port-list ALL } }] [{ SUMMARY FULL }] SHOW CLASSIFIER={ classifiername-list ALL } [{ INTERFACE={ type:id-range id-range ifname-list ALL } }] [{ SUMMARY FULL }] </pre>	<p>Show the ports or interfaces associated with the classifier and its attributes.</p> <p><i>Note: The user may notice that when they execute the SHOW CLASSIFIER ALL FULL command, some internal DHCP classifiers are no longer displayed.</i></p>

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
CLASSIFIER PORT/INTER- FACE PRECEDENCE	ADD	<pre> ADD CLASSIFIER=classifiername INTERFACE={ type:id-range id-range ifname-list ALL } PRECEDENCE=1..255 ADD CLASSIFIER=classifiername PORT={ port-list ALL } PRECEDENCE=1..255 </pre>	<p>Associate the classifier with a port or interface range.</p> <p>PRECEDENCE of the highest rank is 1, and then in descending order.</p>
ACTION CLASSIFIER	ADD	<pre> ADD ACTION CLASSIFIER=classifiername-list { DROP FORWARD COUNT SETVPRIORITY=0..7 SETIPTOS=0..7 SETIPDSCP=0..63 MOVEPRIOTOTOS MOVETOSTOPRIO } </pre>	<p>Associate an action with a classifier or set of classifiers. Each classifier is separated with a comma. Actions are:</p> <p>DROP - Discard the packet</p> <p>FORWARD - Allow the packet to be forwarded.</p> <p>One action or the other can be assigned, but not both.</p> <p>COUNT starts the counting of the actions for the classifier(s).</p> <p>SETVPRIORITY sets the 802.1p bits to the specified value. This value will impact selection of the egress CoS queue.</p> <p>SETIPTOS sets the IP TOS field.</p> <p>SETIPDSCP sets the IP DiffServ CodePoint field.</p> <p>MOVEPRIOTOTOS copies the 802.1q priority field to the IP TOS field</p> <p>MOVETOSTOPRIO copies the IP TOS field to the 802.1q priority field. This new value will determine selection of the egress CoS queue</p>

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
ACTION CLASSIFIER	DELETE	<pre>DELETE ACTION CLASSIFIER=classifiername-list { DROP FORWARD COUNT SETVPRRIORITY SETIPTOS SETIPDSCP MOVEPRIOTOTOS MOVETOSTOPRIO ALL }</pre>	<p>Delete an action associated with a classifier or set of classifiers. If choosing an action, all classifiers must have that same action. If choosing ALL, options may be different for classifiers, but all are dropped.</p>
CLASSIFIER	DELETE	<pre>DELETE CLASSIFIER=classifiername-list INTERFACE={ type:id-range id-range ifname-list ALL } DELETE CLASSIFIER=classifiername-list PORT={ port-list ALL }</pre>	<p>Delete an association between the classifier(s) and the port(s). If desired only some but not all associations can be deleted.</p> <p>Note that the classifier is not deleted, only the port associations.</p>
CLASSIFIER	RESET	<pre>RESET CLASSIFIER=classifiername</pre>	<p>Resets (makes null) the match rules for a classifier.</p>
CLASSIFIER INTERFACE COUNTER	RESET	<pre>RESET CLASSIFIER COUNTER INTERFACE={ type:id-range id-range ifname-list ALL } RESET CLASSIFIER COUNTER PORT={ port-list ALL }</pre>	<p>Resets all classifier counters on the interface(s).</p>
CLASSIFIER	DESTROY	<pre>DESTROY CLASSIFIER={ classifiername-list ALL }</pre>	<p>Destroys the classifier(s). Before doing this, the user must delete all port associations to the classifier. All port associations can be deleted in one command:</p> <pre>DELETE CLASSIFIER=<name> PORT=ALL</pre>

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
TRAFFICDE- SCRIPTOR RATE BURSTSIZE	CREATE	<pre>CREATE TRAFFICDESCRIPTOR=tdname RATE=bits-per-second BURSTSIZE={ 4KB 8KB 16KB 32KB 64KB 128KB 256KB 512KB }</pre>	<p>Creates the ingress metering rates that are associated with classifiers to provide metering/policing for a traffic flow.</p> <p>The RATE (CIR) must be in 1kb increments and can use K or M, such as “128K” or “3M”.</p> <p>The BURSTSIZE (CBS) is in 2ⁿ increments, with the values already provided.</p> <p>CAC is not supported in this release.</p>
TRAFFICDE- SCRIPTOR	SET	<pre>SET TRAFFICDESCRIPTOR=tdname-list [RATE=bits-per-second] [BURSTSIZE={ 4KB 8KB 16KB 32KB 64KB 128KB 256KB 512KB }]</pre>	Changes the ingress metering attributes.
TRAFFICDE- SCRIPTOR CLASSIFIER	ADD	<pre>ADD TRAFFICDESCRIPTOR=tdname CLASSIFIER=classifiername-list { NCDROP NCFORWARD } [NCCOUNT={ ON OFF }]</pre>	<p>Associate the traffic descriptor with a classifier to provide metering/policing for a traffic flow.</p> <p>If the rate is exceeded and therefore the traffic flow is Out Of Profile (OOP), the packets can be dropped or forwarded, with the option to count the dropped or forwarded packets.</p>
TRAFFICDE- SCRIPTOR CLASSIFIER	DELETE	<pre>DELETE TRAFFICDESCRIPTOR={ tdname-list ALL } CLASSIFIER={ classifiername-list ALL }</pre>	Delete the association between the traffic descriptor and the classifier
TRAFFICDE- SCRIPTOR	DESTROY	<pre>DESTROY TRAFFICDESCRIPTOR={ tdname-list ALL }</pre>	Destroy the traffic descriptor. If a classifier is still associated with the traffic descriptor, an error message is displayed.

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
TRAFFICDESCRIPTOR	SHOW	SHOW TRAFFICDESCRIPTOR [= { tdname-list ALL }]	Display the attributes for the traffic descriptor(s).
EGRESSLIMITER RATE BUFFERSIZE	CREATE	CREATE EGRESSLIMITER=limitername RATE=bits-per-second BURSTSIZE={ 4KB 8KB 16KB 32KB 64KB 128KB 256KB 512KB }	Create egress limiter names (usually a service class or limit) and give it the Rate and Buffersize attributes. The RATE must be in 1kb increments and can use K or M, such as "128K" or "3M".
EGRESSLIMITER	SET	SET EGRESSLIMITER=limitername [RATE=bits-per-second] [BURSTSIZE={ 4KB 8KB 16KB 32KB 64KB 128KB 256KB 512KB }]	Change the attributes for the egress-limiter.
EGRESSLIMITER INTERFACE	ADD	ADD EGRESSLIMITER=limitername INTERFACE={ type:id-range id-range ifname-list ALL }	Associate an egress limiter with an interface on the subscriber or network side. Only the ethernet interface can be associated with an egress limiter. (LAG interfaces cannot be associated)
EGRESSLIMITER INTERFACE	DELETE	DELETE EGRESSLIMITER=limitername INTERFACE={ type:id-range id-range ifname-list ALL }	Delete the association between the egress limiter and the interface.
EGRESSLIMITER	DESTROY	DESTROY EGRESSLIMITER={limitername-list ALL}	Destroy the egress limiter. If an interface is still associated with the egress limiter, an error message is displayed.

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
EGRESSLIMITER	SHOW	<pre>SHOW EGRESSLIMITER [={ limitername-list ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }]</pre>	Display the attributes for the egress limiter(s).
QOS VLAN4QUEUEMAP VLAN8QUEUEMAP	SET	<pre>SET QOS [VLAN4QUEUEMAP=value-map] [VLAN8QUEUEMAP=value-map]</pre>	Set the p-bit ingress (by a remarking action on a classifier) for 4 queue to 8 queue mapping and then use that priority throughout the system at each egress queue. The values accept a queue number (0 = lowest priority; higher numbers = higher priority).
QOS	SHOW	<pre>SHOW QOS</pre>	Show the mapping between priority queues and the egress queues.
ACCESSLIST	ADD	<pre>ADD ACCESSLIST=accesslistname INTERFACE={ type:id-range id-range ifname-list }</pre>	Associate an access list with one or more interfaces.
ACCESSLIST	ADD	<pre>ADD ACCESSLIST=accesslistname RULE { PERMIT DENY } [IPSOURCE={ ipaddress ANY } [SOURCEMASK=mask]] [IPDEST={ ipaddress ANY } [DESTMASK=mask]] [MACSOURCE={ macaddress ANY }] [MACDEST={ macaddress ANY }] [APPLICATION={ DHCPSEVER DHCPCLIENT NETBIOS FUM TELNET SSH SNMP FTP TFTP }] [TCPDEST={ tcp-port-list ANY }] [TCPSOURCE={ tcp-port ANY }] [UDPDEST={ udp-port-list ANY }] [UDPSOURCE={ udp-port ANY }] [PROTOCOL={ IPV4 IPV6 protocol-type ANY }] [IPPROTOCOL={ TCP UDP ICMP IGMP ipprotocol-type ANY }] [BEFORE=rulenum]</pre>	Add a new rule to an access list. The BEFORE= option allows inserting the new rule before an existing rule as defined by its number provided in the SHOW ACCESSLIST command. Otherwise, the new rule will be appended to the existing accesslist.

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
ACCESSLIST	CREATE	<pre> CREATE ACCESSLIST=accesslistname [DEFAULTRULE={ PERMIT DENY }] [RULE { PERMIT DENY } [IPSOURCE={ ipaddress ANY } [SOURCEMASK=mask]] [IPDEST={ ipaddress ANY } [DESTMASK=mask]] [MACSOURCE={ macaddress ANY }] [MACDEST={ macaddress ANY }] [APPLICATION={ DHCPSEVER DHCPCLIENT NETBIOS FUM TELNET SSH SNMP FTP TFTP }] [TCPPORTDEST={ tcp-port-list ANY }] [TCPPORTSOURCE={ tcp-port ANY }] [UDPPORTDEST={ udp-port-list ANY }] [UDPPORTSOURCE={ udp-port ANY }] [PROTOCOL={ IPV4 IPV6 protocol-type ANY }] [IPPROTO={ TCP UDP ICMP IGMP ipprotocol-type ANY }]] [INTERFACE={ type:id-range id-range ifname-list }] </pre>	Create a new access list configuration.
ACCESSLIST	DELETE	<pre> DELETE ACCESSLIST=accesslistname RULE=rulenummer </pre>	Delete a rule entry from an access list.
ACCESSLIST	DELETE	<pre> DELETE ACCESSLIST={ accesslistname-list ALL } INTERFACE={ type:id-range id-range ifname-list } </pre>	Remove the association between an access list and one or more interfaces.
ACCESSLIST	DESTROY	<pre> DESTROY ACCESSLIST={ accesslistname-list ALL } [FORCE] </pre>	Destroy an access list. The request will be rejected if an access list configuration is currently associated with one or more interfaces.
ACCESSLIST	RESET	<pre> RESET ACCESSLIST=accesslistname RULE=rulenummer [{ PERMIT DENY }] </pre>	Reset a RULE for an access list.

TABLE 15-10 Commands for Classifier Management (Continued)

Noun/Key Word	Verb	Syntax	Description
ACCESSLIST	SET	<pre> SET ACCESSLIST=accesslistname [DEFAULTRULE={ PERMIT DENY }] [RULE=rulenumbr [{ PERMIT DENY }]] [IPSOURCE={ ipaddress ANY }] [SOURCEMASK=mask]] [IPDEST={ ipaddress ANY }] [DESTMASK=mask]] [MACSOURCE={ macaddress ANY }] [MACDEST={ macaddress ANY }] [APPLICATION={ DHCPSEVER DHCPCLIENT NETBIOS FUM TELNET SSH SNMP FTP TFTP }] [TCPDEST={ tcp-port-list ANY }] [TCPSOURCE={ tcp-port ANY }] [UDPDEST={ udp-port-list ANY }] [UDPSOURCE={ udp-port ANY }] [PROTOCOL={ IPV4 IPV6 protocol-type ANY }] [IPPROTOCOL={TCP UDP ICMP IGMP ipprotocol-type ANY }] </pre>	<p>Allows the user to edit an existing access list configuration. Only currently existing rules can be modified. Removing or adding rules is done using the DELETE or ADD ACCESSLIST commands. Existing attributes of a rule can be modified or new attributes can be added.</p>
ACCESSLIST	SHOW	<pre> SHOW ACCESSLIST={ accesslistname-list ALL } [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	<p>Display the list of access list configuration records. This list can be qualified with the optional INTERFACE= parameter by limiting the display access list only associated with the specified list of interfaces. If there are any errors associated with application of the access list against the interface, they will be displayed in parenthesis underneath the interface name.</p>

16. Implementing Allied Telesis™ fMAP Features

16.1 Overview

Table 16-1 lists the features that are either specific to the fMAP series or require further explanation for how a feature is implemented.

TABLE 16-1 Features for the fMAP

Feature	Description	Reference
IGMP	Reserved multicast address range	16.2.1
HVLAN	An outer tag or service provider tag can be added to a VLAN Port Based	16.3
VLAN Translation	The user can change (translate) a customer VID into a unique VLAN ID for transport across the network	16.4
Traffic Management	The features explained in Section Traffic Management are listed as they apply to the fMAP as well as any special conditions or restrictions.	16.7
LAG	Allows multiple physical links to be joined into a virtual link.	16.8

16.2 IGMP on the fMAP

16.2.1 Channel Usage for IGMP

As mentioned in Section 14.1, for the fMAP series, packets within a **subset** (see below) of the reserved multicast address range of 224.0.0.x (x = 0..255) will be allowed to be flooded within the VLAN. The rest are dropped.

The list of 16 reserved addresses are as follows:

- 01:00:5e:00:00:**01** (The changing part of the address is in bold.)
- 01:00:5e:00:00:**02**
- 01:00:5e:00:00:**04**
- 01:00:5e:00:00:**05**
- 01:00:5e:00:00:**06**
- 01:00:5e:00:00:**09**

- 01:00:5e:00:00:0a
- 01:00:5e:00:00:0c
- 01:00:5e:00:00:0d
- 01:00:5e:00:00:0e
- 01:00:5e:00:00:0f
- 01:00:5e:00:00:12
- 01:00:5e:00:00:23
- 01:00:5e:00:00:65
- 01:00:5e:00:00:66
- 01:00:5e:00:00:fb

Note: The user cannot add, change, or delete this list.

The fMAP supports up to 512 multicast groups, and up to 255 per SM card. The reserved multicast group entries do not take away from the 512.

16.2.1.1 User provisioned MCAST addresses

The 9000 system user can provision specific MCAST addresses that are considered important and are not auto-populated. This allows the user to pass other application-specific protocols that are outside the reserved multicast addresses that the 9000 system IGMP policy would otherwise block. With IGMP snooping enabled, reserved MCAST addresses (224.0.0.1->224.0.1.255) will be dropped unless joined. For this reason, the system provides the user the ability to use CLI commands to statically configure MCAST addresses that are considered important; but, that were not auto-populated. Furthermore, the ability to pass other application specific protocols that are outside the reserved multicast addresses which the 9000 system IGMP policy would otherwise block.

The default behavior is that with IGMP Snooping disabled, the system will flood all MCAST. However, with IGMP Snooping enabled, the system will automatically enter 224.0.0.1 and 224.0.0.2 as entries. The user then has the option, using CLI commands, to add and delete these protocol forwarding addresses.

Protocol forwarding addresses can be displayed using the SHOW IGMP SNOOPING command with the MCASTGROUPS FULL option. An example follows.

Display the multicast groups:

```
officer SEC>> SHOW IGMP SNOOPING MCASTGROUPS FULL
--- System-wide IGMP snooping multicast(MC) groups -----
MC VID      MC MAC      MC IP      Card(s) receiving MC stream
-----
-           -           -           -
-----
--- Statically Provisioned Reserved Multicast(MC) Addresses -----
MC Address Name      MC MAC      MC IP
```

- - -

Provision a MCAST address:

officer SEC>> **ADD IGMP Snooping Flooding PIM**

Info (010017): Operation Successful

officer SEC>> **SHOW IGMP Snooping MCAST Groups Full**

--- System-wide IGMP snooping multicast(MC) groups -----

MC VID	MC MAC	MC IP	Card(s) receiving MC stream
--------	--------	-------	-----------------------------

- - -

--- Statically Provisioned Reserved Multicast(MC) Addresses -----

MC Address Name	MC MAC	MC IP
-----------------	--------	-------

PIM 01:00:5E:00:00:0D 224.0.0.13

Info (010017): Operation Successful

Delete the MCAST address:

officer SEC>> **DELETE IGMP Snooping Flooding PIM**

Info (010017): Operation Successful

officer SEC>> **SHOW IGMP Snooping MCAST Groups Full**

--- System-wide IGMP snooping multicast(MC) groups -----

MC VID	MC MAC	MC IP	Card(s) receiving MC stream
--------	--------	-------	-----------------------------

- - -

--- Statically Provisioned Reserved Multicast(MC) Addresses -----

MC Address Name	MC MAC	MC IP
-----------------	--------	-------

- - -

TABLE 16-2 Commands for MCAST addresses

Object	Verb	Syntax	Description
IGMPSNOOPING FLOODING	ADD	<pre> ADD IGMPSNOOPING FLOODING { ALLSTANDARD DVMRP OSPFALL OSPFDESIGNATED RIP2 IGRP DHCPRELAY PIM RSVP CBT VRRP DXCLUSTER CISCONHAP HSRP MDNS CUSTOM=groupname GROUPADDRESS=ipaddress } </pre>	Add a MCAST address.

TABLE 16-2 Commands for MCAST addresses (Continued)

Object	Verb	Syntax	Description
IGMPSNOOPING FLOODING	DELETE	<pre> DELETE IGMPSNOOPING FLOODING { ALL ALLSTANDARD DVMRP OSPFALL OSPFDESIGNATED RIP2 IGRP DHCPRELAY PIM RSVP CBT VRRP DXCLUSTER CISCONHAP HSRP MDNS CUSTOM=groupname } </pre>	Delete an MCAST address.
IGMPSNOOPING	SHOW	<pre> SHOW IGMPSNOOPING [{ STATUS MCASTGROUPS [FULL] COUNTER [{ STANDARD MESSAGERESPONSE INTERFACE={ type:id-range id-range ifname-list ALL } CARD={ slot-list ALL } }] INTERFACE={ type:id-range id-range ifname-list ALL } [FULL] CARD={ slot-list ALL } [FULL] }] </pre>	Display MCAST addresses.

16.2.1.2 Continuous Flooding of Unicast Data

In some configurations, unicast data coming from the network (flowing from upstream to downstream) can be flooded for an indefinite period of time.

When two switches are connected to form a LAG they will hash and route packets independently of each other. This may result in packets (between the switches) being transmitted upstream on one Network Module port and received downstream on another Network Module port. It is possible, in this condition, for L2 addresses to not be learned properly in the Forwarding Database and thus cause flooding.

To prevent this problem from occurring all ports in a LAG group should be connected to the same Network Module.

16.2.2 Multicast stream availability

The number of recommended multicast streams can be up to 25 per card for most fMAP products. This is the default for most service modules. However, the fMAP can provide a greater number of multicast streams. Therefore, the user may want to configure more than 25 streams per card up to the maximum of 512.

16.3 HVLAN (Port Based)

16.3.1 Overview for Release 6.0

As explained in Section 12.2, a VLAN allows broadcast traffic to flood only ports that are members of that VLAN. Moreover, ports can be tagged or untagged, with a tagged Ethernet frame including the VID field that uniquely identifies the VLAN of the frame. The number of VLANs that can be configured across the service provider network is restricted to the 12-bit VID field (1 to 4095).

To help overcome this limitation, an additional or **outer tag** can be added on top of the 802.1q tagged or untagged frame. At the SM port, incoming customer frames are wrapped with an outer tag that is used to switch the traffic across the network. At the SM port for the outgoing traffic, the outer tag is removed and the frame is delivered to the customer's VLAN.

By using this outer tag,, fMAP system users can expand service to customers in the following ways:

1. Two VLAN tags are used to identify the customer VLAN, in theory expanding the number range of customer VLAN tags to 4094 * 4094.
2. Since the inner tag is used by each customer, the VLAN ID for different customers may be the same (overlap)
3. The customer VLAN ID is preserved and unchanged as it crosses the network.

By using this outer tag, service providers can tunnel the VLANs of each customer into a single VLAN (the VLAN ID of the outer tag) and send them across the network, allowing businesses to interconnect devices from multiple locations in a service provider area. The use of the additional tag creates a hierarchical VLAN (HVLAN).

16.3.2 VLAN Frame Flow (TAGALL and TPID Values)

To understand the HVLAN feature, the 802.1q tagged ethernet frame and the fields it contains must be fully understood. These are listed in Table 16-3.

TABLE 16-3 VLAN Tag Fields

Field Name	Length	Description
Tag Protocol Identifier (TPID)	2 octets	The TPID is used to identify the frame as a tagged frame. The value of the TPID for an 802.1q ethernet tagged frame is 0x8100
User Priority	3 bits	The User Priority field can represent up to eight priority levels. (This field is explained in greater detail when discussing traffic management, in Section 16.7.

TABLE 16-3 VLAN Tag Fields

Field Name	Length	Description
Canonical Format Indicator (CFI)	1 bit	The CFI is a flag to indicate whether all MAC address information that may be present in the MAC data carried by the frame is in canonical format.
VLAN ID (VID)	12 bits	The VID identifies which VLAN the frame belongs to, with a range of 1 to 4094. It consists of the Tag Protocol Identifier (TPID) and the Tag Control Information (TCI).

The TPID, which is used to identify the frame as a tagged frame in 802.1q, has a value of 0x8100. The TPID value for the HVLAN (the outer tag), is configurable, and should be set depending on the vendor's recommendation.

Note: To obtain the TPID value that each vendor supports, consult the interconnecting vendor's documentation.

16.3.3 Provisioning Rules

Following are the rules for setting the TAGALL and TPID values for a port-based HVLAN configuration:

- When adding a Network Interface to an HVLAN, the user should include FRAME=TAGGED (example `Add hvlan Btunnel interface=10.0 frame=tagged`)
- When adding a Customer Interface to an HVLAN, the user should include FRAME=UNTAGGED (example `Add hvlan Btunnel interface=5.2 [frame=untagged]`) - Note untagged is the default.
- The Port-based HVLAN tunnel is not operational until the interface is set to `tagall=on`. (An example is `set interface 5.2 tagall=on`)
- DHCP Relay/Snooping and IGMP Snooping must be disabled for the Customer Interface prior to setting the interface to `tagall=on`. Examples of the IGMP and DHCP commands are:
 - `Disable igmp interface=5.2`
 - `Disable dhcp all interface=5.2`
- An SM interface must be set to `TAGALL=ON` if that interface is a member of an HVLAN.
- An interface with a TPID value set can be a member of more than one HVLAN.
- An HVLAN can consist of one or more interfaces with TPIDs set at different values.
- A TPID can be set on an NM interface even if the interface is not a member of the HVLAN, with the default 0x8100 value. However, this is not recommended.
- When an interface is removed from an HVLAN, it will become a member of the an 802.1q default VLAN, and the TPID value should be set at the default 0x8100 value.

16.3.4 Sample Configuration and Commands

[Figure 16-1](#) shows a possible configuration where both 802.1q and HVLANs are created. The 802.1q VLAN (10), is used for multicasting for video. The HVLAN is provisioned for a business customer who has their own network and wishes to connect this network to devices on fMAP systems.

Note: The configuration rules for Port-based HVLANs have not changed in release 7.0; however because of the 7.0 feature VLAN-based HVLANs, the command syntax has changed with the need to define an HVLAN type (PORTTUNNEL versus VLANTUNNEL).




FIGURE 16-1 Port-based HVLAN Configuration for Release 6.0

Configuring the HVLAN follows these four main steps:

1. Create the HVLAN (HVLAN).
2. Associate member ports with the HVLAN.
3. Configure the SM ports with the TAGALL=ON value.
4. Configure the NM ports with the outer tag TPID value.

For the configuration shown in [Figure 16-1](#), each 9400 is configured as follows:

Commands to provision System A.

- Create the port-based HVLAN called HVLANC.
`CREATE HVLAN=H_100 VI D=100 TYPE=PORTTUNNEL`
- Create a standard 802.1q standard VLAN.
`CREATE VLAN=STDVLAN VI D=10`
- Associated the H_100 to NM ports 0.1/0.2, and SM ports 2.0/2.1.
`ADD HVLAN=H_100 I NTERFACE=ETH: 01, 0. 2 FRAME=TAGGED`
`ADD HVLAN=H_100 I NTERFACE=2. 0, 2. 1`
- ADD the SM port 2.1 and NM ports 0.1, 0.2 into the standard VLAN.
`ADD VLAN=STDVLAN I NTERFACE=2. 2`
`ADD VLAN=STDVLAN I NTERFACE=0. 1, 0. 2 FRAME=TAGGED`
- Set the SM port 2.0, 2.1 to TAGALL.
`SET I NTERFACE=2. 0, 2. 1 TAGALL=ON`
- Set the NM ports 0.1 and 0.2 to a TPID value. (This applies to the STDVLAN.)
`SET I NTERFACE=0. 1, 0. 2 TPID=0x9100`

To deprovision any of the HVLANs:

- Delete ALL the interfaces from the VLANs
- Destroy the VLANs
- Delete ALL the interfaces from the HVLANs
- Destroy the HVLANs
- Set the TPID values back to 0x8100
- Set the TAGALL to OFF.

16.3.5 Port-based HVLAN for Release 7.0 (9100)

Note: The 9100 does not support the VLAN-based HVLAN feature, although it can pass the VLAN-based HVLANs.

The 9100 also supports the port-based HVLAN feature, as follows:

- **GE4/GE2RJ:**
 - An interface can be configured as a tagged HVLAN member only.
 - An interface can be configured as a member of multiple HVLANs.
 - An interface that has an HVLAN configured can also be configured as a tagged member of regular VLANs.
 - An interface that has an HVLAN configured can also be configured as an untagged member of a regular VLAN.
 - IGMP, DHCP, etc. can be configured on ports that are HVLAN members.
 - An interface cannot have its TPID set to anything other than 0x8100.
- **FE10/FX10:**
 - An interface can be configured as an untagged HVLAN member only (so only one HVLAN can be configured).
 - Interfaces must be configured with TAGALL=ON to make HVLAN work.
 - IGMP, DHCP, and ARP filtering must be disabled before TAGALL can be set to ON.
 - Regular VLANs can be configured on an interface that has an HVLAN on it.

16.4 VLAN Translation

16.4.1 Overview

When customer networks are connected through service provider networks, customers may want to keep their existing VLAN assignments. It is not uncommon for the VLAN IDs to be same for different customers (overlap). To allow this overlap, a service provider needs to be able to change (translate) a customer VID into a unique VLAN ID for transport across the network.

To do this, an 802.1q tagged VLAN can be configured with this translations option.

The general flow of commands to perform this translation are:

- Create a VLAN
`CREATE VLAN=VLAN100 VID=100`
- Add a Service Module port to the VLAN as a tagged port
`ADD VLAN=VLAN100 INTERFACE=1.0 FRAME=TAGGED`
- Turn the translation option on for the SM port for a customer VLAN ID
`SET VLAN=VLAN100 INTERFACE=1.0 TRANSLATE=10`

This will result in the following:

- When a tagged frame with a VLAN ID of 10 enters the SM port 1.0, the VLAN ID will be translated to VLAN ID 100.
- When the tagged frame with VLAN ID 100 leaves the SM port 1.0, the VLAN ID will be translated (back) to VLAN ID 10.

This is shown in more detail in the example in [16.4.3](#).

16.4.2 Provisioning Rules

To configure the VLAN translation option, the following rules apply:

- The maximum number of customer VLANs that can be translated on an FE10/FX10 port is 16.
- The SM card must support the translation option.

Note: Only the FE10 and FX10 SM cards for the fMAP support this feature.

- There is a one-to-one mapping of between the customer VLAN ID and the VLAN ID used for crossing the service provider network. (Each customer VLAN ID can be translated into only one VLAN ID, and vice versa.)
- The customer VLANs to be translated must be tagged.

16.4.3 Example Configuration

An example command set for two systems is below.

1. Commands for one 9400:

- Create the VLAN 100 and 200

```
CREATE VLAN=VLAN100 VID=100
```

```
CREATE VLAN=VLAN200 VID=200
```

- Add the NM port 1.1 and SM ports 7.1 and 7.2 to VLAN100 and VLAN200

```
ADD VLAN=100 INTERFACE=1.1,7.1 FRAME=TAGGED
```

```
ADD VLAN=200 INTERFACE=1.1,7.2 FRAME=TAGGED
```

- Set the translation option on SM port 7.1 to translate the Customer 1 VLAN ID 10 to 100.

```
SET VLAN=100 INTERFACE=7.1 TRANSLATE=10
```

- Set the translation option on SM port 7.2 to translate the Customer 2 VLAN ID 10 to 200

```
SET VLAN=200 INTERFACE=7.2 TRANSLATE=10
```

2. Commands for the second 9400:

- Create the VLAN 100 and 200

```
CREATE VLAN=VLAN100 VID=100
```

```
CREATE VLAN=VLAN200 VID=200
```

- Add the NM port 1.1 and SM ports 7.1 and 7.2 to VLAN100 and VLAN200

```
ADD VLAN=100 INTERFACE=1.1,7.1 FRAME=TAGGED
```

```
ADD VLAN=200 INTERFACE=1.1,7.2 FRAME=TAGGED
```

- Set the translation option on SM port 7.1 to translate the Customer 1 VLAN ID 10 to 100.

```
SET VLAN=100 INTERFACE=7.1 TRANSLATE=10
```

- Set the translation option on SM port 7.2 to translate the Customer 2 VLAN ID 10 to 100

```
SET VLAN=200 INTERFACE=7.2 TRANSLATE=10
```

16.5 Port-Based HVLAN and Translation Feature Interactions

With the port-based HVLAN and translation options, VLAN configurations interact with each other and other features as follows:

- If an SM port is a member of an HVLAN, the port cannot be added to an 802.1q VLAN; to have both types, **one** SM port can be configured with an 802.1q VLAN, and a **different** port with an HVLAN, so the two features can co-exist on the same fMAP product.

- The Port-based HVLAN and translation feature are not compatible on the same port. Once a port is configured with the HVLAN option, it cannot use the translation feature, and vice-versa.
- For traffic management, classifiers are used to filter traffic according to certain criteria, and this may be affected with the double tagging of frames. Refer to [16.7.1](#).
- Link Aggregation (LAG) can still be enabled for a port that has an HVLAN as long as all member ports of the LAG group belong to the same VLANs, both tagged and untagged.
- **IGMP Snooping and Port-based HVLAN are mutually exclusive features.** If IGMP snooping is enabled system wide and a port has IGMP snooping enabled, that port cannot participate in the HVLAN; if a port is part of an HVLAN, IGMP snooping cannot be enabled on that port.
- Spanning Tree Protocol can be enabled on an HVLAN port, as long as the following applies:

When customer traffic at multiple sites is tunneled over the service provider network, every customer VLAN will need to build a spanning tree that includes the multiple sites across the VLAN. To enable this, the Bridge Protocol Data Unit (BPDU) will need to be tunneled across the network.

Note: The VLAN-based HVLAN and Translation features can be supported on one system, but in most network engineering solutions, either one or the other is used.

16.6 Command Summary for HVLAN

TABLE 16-4 Commands for HVLAN and VLAN Translation

Object	Verb	Syntax	Description
HVLAN	ADD	<pre>ADD HVLAN={ hvlanname vid } INTERFACE={ type:id-range id-range ifname-list ALL } [FRAME={ UNTAGGED TAGGED }]</pre>	Associate the HVLAN with an interface (and therefore the associated ports) and set the frames as tagged or untagged.
HVLAN	CREATE	<pre>CREATE HVLAN=hvlanname VID=2..4094 [TYPE={ PORTTUNNEL VLANTUNNEL }]</pre>	<p>Create the HVLAN that will be used as the outer VLAN.</p> <p>PORTTUNNEL is for port-based HVLANs.</p> <p>VLANTUNNEL is for VLAN-based HVLANs.</p>
VLANTUN- NELMAP VLAN HVLAN	ADD	<pre>ADD VLANTUNNELMAP VLAN={ vlanname-list vid-range } HVLAN={ hvlanname vid }</pre>	Make the association of the VLAN to a VLAN-based HVLAN tunnel. The tunnel is defined by the HVLAN and its interface membership
VLANTUN- NELMAP VLAN HVLAN	DELETE	<pre>DELETE VLANTUNNELMAP VLAN={ vlanname-list vid-range ALL } HVLAN={ hvlanname vid }</pre>	Disassociate a VLAN from a VLAN-based HVLAN tunnel
HVLAN	DELETE	<pre>DELETE HVLAN={ hvlanname vid } INTERFACE={ type:id-range id-range ifname-list ALL }</pre>	Delete the association between the HVLAN and the interface.

TABLE 16-4 Commands for HVLAN and VLAN Translation (Continued)

Object	Verb	Syntax	Description
HVLAN	DESTROY	DESTROY HVLAN={ hvlaname vid ALL }	Destroy the HVLAN. If there is still an association with an interface, there is an error message.
HVLAN	SET	SET HVLAN={ hvlaname vid } INTERFACE={ type:id-range id-range ifname-list ALL } [FRAME={ UNTAGGED TAGGED }]	Change the HVLAN association with the interface.
HVLAN	SHOW	SHOW HVLAN [={ hvlaname vid ALL }] [FULL]	Displays HVLAN information.

16.7 Traffic Management for the fMAP

16.7.1 Overview

Table 16-5 lists the features used in a QoS traffic management system, highlights what the fMAP will support in this release, lists the key parameters used to enable the feature, and includes a reference when relevant.

TABLE 16-5 QoS Features for this release for the fMAP

Feature for QoS Model	Subscriber Side		Network Side
	FE10, FX10	ADSL8S, ADSL24A/B, SHDSL24	
Classification/Filters			
Field Classification	Fully Supported	Only IP Address and VLANID, not VPRIORITY, INNERVPRIORITY, INNERVID, IPTOS, IPDSCP, TCPFLAGS	Fully Supported

TABLE 16-5 QoS Features for this release for the fMAP (Continued)

Feature for QoS Model	Subscriber Side		Network Side
Ingress Monitoring (Classifier Counters)	Match, Filter, and Policed	FORWARD, DROP, and COUNT	FORWARD, DROP, and COUNT
ARP Filtering	Fully Supported	When IP filtering is ON for a port or interface, an ARP filter will also be added (Can be turned ON or OFF)	When IP filtering is ON for a port or interface, an ARP filter will also be added (Can be turned ON or OFF)
MAC Limiting	Fully Supported	When the learning limit is reached, all frames are dropped, including Broadcast and Multicast frames	When the learning limit is reached, all frames are dropped, including Broadcast and Multicast frames
Field Marking			
Field Marking	Y (N if DTAG/ HVLAN)	No support	COS priority bits within the fMAP
QOS LAG Support	Not Supported	QOS LAG Support	
Policing / Metering			
Policing / Metering	Supported	One per classifier	
Traffic Descriptors (Metering)	Supported	Send and Drop Counts per Queue	
Egress Rate Limiting			
Egress Rate Limiting	Supported	Not Supported ^a	Egress Port Rate Limiting
Egress Monitoring	Not Supported	Out of Profile Actions are DROP and/or COUNT	
Scheduling and Queuing			
Queue Scheduling	Not Supported	Strict Priority, Tail Drop	Tail drop only
Number of Queues	4	4	8, SP only
p-bit Mapping	Supported	Supported	Fully Supported

- a. For all ADSL cards, configuration of egress rate limiting is not supported directly, but can be set using the ADSL maximum downstream rate.

Note: The IGMP default traffic queue has some special considerations related to provisioning.

Note: The counters match, filter, and policed are analogous to FORWARD, DROP and COUNT. Match is the number of packets that meet the matching criteria. Filter is the number of packets that are dropped. Policed is the number of packets that are non-conforming

16.7.2 Possible Conflict with Classifier Combinations - fMAP

In some cases, the user may wish to apply different classifiers with the same mask on a port or interface. The following shows an example of doing this.

```
officer SEC> SHOW CLASSIFIER PORT=4, 4
```

```

--- Classifier Configuration Data ---
-----
Port Rank Name           Field Match(es)           Action(s)           Notes
-----
4.4  51  ipfilt1                IPSOURCE=1.0.0.0/8       FORWARD
                                     COUNT
                                     DROP
69   69  ipdeadend              IPSOURCE=ANY              COUNT
                                     FORWARD
201  201  ipmark                 IPSOURCE=1.0.0.0/8       FORWARD
                                     SETIPTOS=5
-----

```

Normally, packets matching the mask 1.0.0.0/8 would be forwarded and counted (precedence or rank of 51). All remaining IP packets would be dropped. The classifier ipmark would then be applied and the packets matching 1.0.0.0/8 would be forwarded with an IPTOS values set to 5.

If classifiers match packets in which an unintended action may occur because of a conflict or mismatch of actions, the packet may be dropped.

To understand this situation, the user must look at the “Field Match” column in combination with the “Action” column. This situation occurs when there is a classifier with a conflicting action in combination with other classifiers. An action conflict occurs when a packet matches more than one classifier and the actions of those classifiers cannot all be executed. In the output above, the DROP action of the ipdeadend classifier conflicts with the FORWARD action of the ipfilt1 classifier and the FORWARD action of the ipmark classifier, because the packet can either be forwarded or dropped.

Here are some other examples of conflicting actions:

- One classifier has a FORWARD action while another has a DROP action
- Two classifiers have SETVLANPRIORITY with different values
- Two classifiers have SETIPTOS with different values
- Two classifiers have SETIPDSCP with different values

For this to occur, the classifier with the conflicting action must match the same packet as the higher precedence classifier. In the output above, if a packet comes in with an IP source address of 1.2.3.4, this packet will match both the ipfilt1 classifier (IPSOURCE=1.0.0.0/8) and the ipdeadend classifier (IPSOURCE=ANY). This can occur even if the classifiers are not all examining the same field. For example, the ipfilt1 classifier is configured with another classifier vlan100 that matches VLANID=100 and has a DROP action. If a packet were sent in that had an IP source address of 1.2.3.4 and a VLAN ID of 100, then that packet would match both classifiers, and an action conflict would occur.

If a classifier has a conflicting action, the user should search the “Field Match” column for classifiers that use the same classification field(s). The unexpected behavior will occur if the classifier with the conflicting action has a precedence value between the other two classifiers (as shown in the output above).

Note: The phrase “same classification field” refers to two classifiers that attempt to match exactly the same fields in the packet header.

Also note that **field width** plays a role in this problem. Fields such as MACDEST have a fixed width while fields such as IPDEST have a variable width (depending on the subnet mask).

Refer to the two filters shown below:

Port	Rank	Name	Field Match(es)	Action(s)
4.4	100	ipfilt1	IPSOURCE=1.0.0.0/8	FORWARD COUNT
	200	ipfilt2	IPSOURCE=1.0.0.0/16	FORWARD COUNT

Since **ipfilt1** uses a subnet mask of /8 and **ipfilt2** uses a subnet mask of /16 they do not have the same field width and therefore do not use the same classification field.

16.7.3 Solution to Classifier Mismatch

The way to get the expected behavior is to put the DROP classifier AFTER the other two classifiers when configuring the system. The following shows the corrected configuration.

SHOW CLASSIFIER PORT 4.4

```
--- Classifier Configuration Data -----
```

Port	Rank	Name	Field Match(es)	Action(s)
4.4	51	ipfilt1	IPSOURCE=1.0.0.0/8	FORWARD COUNT
	55	ipmark	IPSOURCE=1.0.0.0/8	FORWARD SETIPTOS=5
	69	ipdeadend	IPSOURCE=ANY	DROP COUNT

The following shows two sets of classifiers that give the expected results

SHOW CLASSIFIER PORT=4.4

```
--- Classifier Configuration Data -----
```

Port	Rank	Name	Field Match(es)	Action(s)
4.4	51	ipfilt1	IPSOURCE=1.0.0.0/8	FORWARD COUNT
	55	ipmark1	IPSOURCE=1.0.0.0/8	FORWARD SETIPTOS=5
	59	ipfilt2	IPSOURCE=2.0.0.0/16	FORWARD COUNT

60	ipmark2	IPSOURCE=2.0.0.0/16	FORWARD SETIPTOS=9
69	ipdeadend	IPSOURCE=ANY	DROP COUNT

16.7.4 QoS Classifier Capacity for FE10/FX10 Cards

With the, FE10/FX10, cards, there are certain limits to the number of classifiers that can be provisioned. These are based on the card's functionality as well as when other features (IGMP and HVLAN) are used.

The FE10/FX10 card uses **masks**; a mask is used to select which parts of a packet to examine for classification. Each GE port can be provisioned with up to 16 masks. On the FE10/FX10 card, 16 masks are shared over multiple ports (one mask for ports 0-5 and another mask for port 6-9). Each classifier with a unique mask (unique fields and/or number of bits within a field) will use up a mask of the set of 16 available.

Note: When the user has HVLANS on the port, then classifiers that match fields beyond the tag(s) will use double the classifier resources.

16.7.5 Adding Classifiers to Service Module Ports/Interfaces

When adding classifiers to multiple ports, certain rules must be followed to ensure the precedence order reflects what the user intended. For example, inputting the following commands:

```
ADD CLASSIFIER=ipfilt1 PORT=4.0-4.1 PRECEDENCE=51
ADD CLASSIFIER=ipfilt2 PORT=4.0-4.1 PRECEDENCE=52
ADD CLASSIFIER=ipfilt3 PORT=4.0-4.1 PRECEDENCE=53
ADD CLASSIFIER=IpDeadEnd PORT=4.0-4.1 PRECEDENCE=69
```

may or may not result in the order of precedence (ipfilt1, ipfilt2, ipfilt3, IpDeadEnd), being configured.

To ensure the precedence order is configured as entered, one of these methods may be used:

1. Disable all ports that will be affected by the Classifier/Precedence. After inputting the commands to configure the ports, enable them.
2. Input the commands, but then reboot the affected card.
3. Configure the classifiers in reverse Precedence order. In the example above, the user would first enter the IpDeadEnd classifier and continue until ipfilt1. (If the commands are entered as in the example rebooting the card will still ensure the precedence order is followed.)

Note: For users, method 3 (configuring the classifiers in reverse precedence order) is preferable the first time the classifiers are configured. Later, if the user wants to modify the classifier list, the user can make the changes and then reboot the card.

16.7.6 Classifiers and Feature Interaction for the fMAP Product

The following rules apply to classifiers and other fMAP features:

- HVLAN - If a classifier for a VLAN ID is set for an upstream port, the match will apply to the outer tag only.
- VLAN translation - If VLAN translation is being performed on a port, the translation is performed first, and then a classifier will match on the translated VLAN ID.
- The fMAP cannot remark the TOS on double tagged (HVLAN) packets. On the subscriber side, a request to remark an IP field (e.g. IPTOS or IPDSCP) is rejected if TAGALL is ON. However, on the network side, because it can be a member of an HVLAN and a “normal” VLAN, the fMAP does not reject this kind of request. Any packet that comes in with only one tag will receive the remarking, but packets that come in with two tags will not.
- IGMP Snooping - If IGMP Snooping is enabled, there are limits to the number of classifiers that can be configured. If a classifier cannot be entered because it exceeds the system’s resources, an alarm is raised.

16.8 LAG

16.8.1 Overview

In the fMAP, there are up to three GE ports that can be used in a LAG configuration. [Figure 16-2](#) shows how the fMAP can be configured. Note the following from the figure:

- All ports in the LAG have same VLAN configuration (tagged and untagged VLANs)
- If one link goes down, traffic switches quickly to other link(s), if bandwidth is available.
- All ports have the same STP state
- There can be up to two LAGs, one per GE3 card. Do not create a LAG that spans across the GE3 cards.
- If one link goes down, all traffic switches to other link in same LAG group

Note: While it is possible to have a LAG group span across the GE3 cards, flooding may occur and therefore this is not recommended. Therefore, since members of the LAG cannot be on different GE3 cards, there can be only one or two LAG groups on the fMAP, with each group having two or three ports on the same card. This configuration gives port level rather than card-level protection.




FIGURE 16-2 Upstream Protection Using LAG in the fMAP

16.8.2 LAG and UFO Incompatibility

UFO mode is set on a VLAN basis. On the fMAP, however, a VLAN can have **only one** associated upstream port. As a result, the two features cannot work together; the VLANs associated with a LAG must not be in UFO mode.

16.8.3 LAG and 9100 GE2RJ Interfaces

As mentioned in 13.3.2, all interfaces being placed into a LAG group must have:

- AUTONEGOTIATE=ON (to allow for negotiation of settings with the far end)
- DUPLEX=FULL
- SPEED=same across all ports

Once a link is placed in a LAG, these attributes cannot be changed. (The attributes can be changed once the interface is removed from the LAG.)

For the GE1, GE3, and GE8 interfaces, the default values set the parameters to the values listed above (with SPEED always set at 1000K), and so associating these interfaces to a LAG does not involve having to change the default settings.

However, for the GE2RJ interface, which supports variable speeds, using the default settings can lead to a loss of LAG functionality. Therefore, when configuring the GE2RJ interface for LAG, the user must **manually** set the interface values as listed above; the speed can be any allowed value as long as it is the same across the interfaces in the LAG.

Moreover, if users have LAGs using the GE2RJ and wish to upgrade the system to 8.0, they should ensure the following before performing the upgrade:

- AUTONEGOTIATE is set to ON
- DUPLEX is set to FULL
- SPEED is set to a value that will be the same for the GE2RJ interfaces in this LAG.

Once the upgrade is performed, the GE2RJ interface will have the proper values.

Following are some sample scenarios in which the values have not been manually set to the above values.

16.8.3.1 Attempt to add port to LAG with auto settings (Interface DN-DN-Offline)

The user attempts to add a port to a LAG with auto settings. Note that Auto Negotiation cannot be “Off” if any of speed, duplex, or flow control are set to “On”. Also note that the interface state is DN-DN-Offline. This is required for the interface configuration commands to follow.

```
officer SEC>> show int 4.1
--- GE Interfaces ---
Interface..... 4.1
Type..... GE
State..... DN-DN-Offline
Description..... <none>

Provisioning
Provisioning Profile..... AutoProv
Direction..... Network
Auto Negotiation..... On
Speed..... Auto
```

```
Duplex..... Auto
Flow Control..... Auto
Remote Monitoring..... Off
```

Actual

```
Direction..... Network
Physical Address..... 00:0C:25:03:91:64
```

VLAN Information

```
Acceptable Frame Types..... All
Ingress Filtering..... On
TPID..... 0x8100
TAGALL..... Off
Untagged VLAN..... 1
Tagged VLAN(s)..... 10
```

Show which LAGs are configured:

```
officer SEC>> show lag
```

```
--- LAG Info Data ---
LAG Name          Provisioned Ports      Mode   Select      Admin Index
                  5.0-5.3              on     macsrc &   1      0
                  4.0                  on     macdest
                  macsrc &
                  macdest

LAG Name          Enabled Ports        Speed  Oper State  Oper  Interface
                  5.0-5.3             4 Gpbs Up         N/A   LAG:0
LAG Name          Enabled Ports        Speed  Oper State  Oper  Interface
                  4.0                 1 Gpbs Up         N/A   LAG:1
```

```
Info (010017): Operation Successful
```

Now attempt to add int 4.1 to the LAG 'toF27too'. First it fails for "speed" setting:

```
officer SEC>> add lag toF27too int 4.1
Error (040342): Interface 4.1 speed cannot be set to 'AUTO' for LAG
```

Change speed to match LAG

```
officer SEC>> set int 4.1 GE speed 1000
Info (020186): Successfully modified interface(s) 4.1
```

Try again:

```
officer SEC>> add lag toF27too int 4.1

Error (040313): Interface 4.1 duplex mode must be full for LAG
```

Now set duplex to match LAG:

```
officer SEC>> set int 4.1 GE duplex full
Info (020186): Successfully modified interface(s) 4.1
```

Look at the interface to see that settings are what we want:

```
officer SEC>> show int 4.1
```

```
--- GE Interfaces ---
```

```
Interface..... 4.1
Type..... GE
State..... DN-DN-Offline
Description..... <none>

Provisioning
  Provisioning Profile..... AutoProv (*)
  Direction..... Network
  Auto Negotiation..... On
  Speed..... 1 Gbps
  Duplex..... Full
  Flow Control..... Auto
  Remote Monitoring..... Off

Actual
  Direction..... Network
  Physical Address..... 00:0C:25:03:91:64

VLAN Information
  Acceptable Frame Types..... All
  Ingress Filtering..... On
  TPID..... 0x8100
  TAGALL..... Off
  Untagged VLAN..... 1
  Tagged VLAN(s)..... 10
```

Try to add the interface to the LAG again, and show the results

```
officer SEC>> add lag toF27too int 4.1
Info (010017): Operation Successful
officer SEC>> show lag
--- LAG Info Data -----
```

LAG Name	Provisioned Ports	Mode	Select Criteria	Admin Key	Index
toF27	5.0-5.3	on	macsrc & macdest	1	0
toF27too	4.0-4.1	on	macsrc & macdest	2	1

```

--- LAG Info Data -----
```

LAG Name	Enabled Ports	Speed	Oper State	Oper Key	Interface ID
toF27	5.0-5.3	4 Gbps	Up	N/A	LAG:0
toF27too	4.0	1 Gbps	Up	N/A	LAG:1

```
Info (010017): Operation Successful
```

The interface was added successfully, since it is in the Provisioned Ports list. This means that all semantic checks passed.

But the interface is still not carrying traffic in the LAG, because it is disabled. (The interface had to be disabled in order to configure the speed and duplex settings earlier.) The user needs to enable the interface.

```
officer SEC>> enable int 4.1
Info (039512): Operation Successful (GE2RJ Slot 4 Port 1)
```

Now check LAG again:

```
officer SEC>> show lag
--- LAG Info Data -----
LAG Name          Provisioned Ports    Mode   Select      Admin Index
                  Criteria           Key
-----
toF27             5.0-5.3            on     macsrc &   1    0
                  macdest
toF27too         4.0-4.1            on     macsrc &   2    1
                  macdest

LAG Name          Enabled Ports       Speed  Oper State  Oper  Interface
                  Key               ID
-----
toF27             5.0-5.3            4 Gpbs Up         N/A  LAG:0
toF27too         4.0-4.1            2 Gpbs Up         N/A  LAG:1
Info (010017): Operation Successful
```

The interface is now carrying traffic in the LAG.

16.8.3.2 Attempt to change an interface setting, when an interface is already a member of a LAG

Note that to change any of the interface settings, the interface must first be disabled. This will remove the interface from the LAG “Enabled Ports” list. (The interface will still however be provisioned against the LAG.)

```
officer SEC>> disable int 4.1
Service may be affected, are you sure (Y/N)? y
Info (039512): Operation Successful (GE2RJ Slot 4 Port 1)
```

Look at interface parameters and note that state is DN-DN-Offline:

```
officer SEC>> show int 4.1
--- GE Interfaces ---

Interface..... 4.1
Type..... GE
State..... DN-DN-Offline
Description..... <none>

Provisioning
  Provisioning Profile..... AutoProv (*)
  Direction..... Network
  Auto Negotiation..... On
  Speed..... 1 Gbps
  Duplex..... Full
```

```
Flow Control..... Off
Remote Monitoring..... Off
```

Actual

```
Direction..... Network
Physical Address..... 00:0C:25:03:91:64
```

VLAN Information

```
Acceptable Frame Types..... All
Ingress Filtering..... On
TPID..... 0x8100
TAGALL..... Off
Untagged VLAN..... 1
Tagged VLAN(s)..... 10
```

Look at LAG and note that interface 4.1 is no longer in the Enabled Ports List

```
officer SEC>> show lag
```

```
--- LAG Info Data -----
```

LAG Name	Provisioned Ports	Mode	Select Criteria	Admin Key	Index
toF27	5.0-5.3	on	macsrc & macdest	1	0
toF27too	4.0-4.1	on	macsrc & macdest	2	1

LAG Name	Enabled Ports	Speed	Oper State	Oper Key	Interface ID
toF27	5.0-5.3	4 Gpbs	Up	N/A	LAG:0
toF27too	4.0	1 Gpbs	Up	N/A	LAG:1

```
Info (010017): Operation Successful
```

Now attempt to change int 4.1 speed and duplex settings:

```
officer SEC>> set int 4.1 GE speed 100
Error (035515): Interface(s) 4.0 not modified because they are members of a LAG.
officer SEC>> set int 4.1 GE duplex half
Error (035515): Interface(s) 4.1 not modified because they are members of a LAG.
```

Because the interface is still a provisioned member of LAG ‘toF27too’, semantic checks will prevent the user from changing any interface settings.

16.8.4 Changes to Feature Support from Release 6.0

In release 6.0, voice, data, and video (multicast traffic) services are all supported with an FE10 port, FX10 port, or an FE2 port as an upstream interface that will support data and multicast services. An FE/FX card can have a mixture of network and customer interfaces.

When a system has FE/FX network interfaces, it may also have enabled GE interfaces.

16.8.4.1 Feature Support

The FE/FX network interfaces **can** support:

- Inband Management - no restrictions
- STP/RTSP - no restrictions
- Traffic Management, including Egress rate limiting - no restrictions
- IGMP Snooping - an FE/FX port that is configured as a network interface will be snooped, as opposed to a GE network interface that is filtered. The maximum number of multicast groups supported by a system with an FE/FX port that is configured as a network interface is limited to 255 minus the number of reserved multicast groups.
- UFO VLAN - an FE/FX can be statically designated as the upstream interface for a UFO VLAN. An FE/FX interface can be configured as a “ring” and STP/RSTP can be used (but not EPSR) to determine the upstream interface dynamically.

Note: Is there a restriction on mixing upstream and downstream interfaces on a FE/FX10 card when using UFO vlans?

- DHCP Relay - an FE/FX interface can be the interface to the DHCP Server; provisioning rules for DHCP will need to be changed.

All other protocols are still **not** supported in release 5.0, including:

- EPSR - an FE/FX interface cannot be added to an EPSR domain as primary or secondary interface via CLI or NMS, but a VLAN associated with an FE interface CAN be added to an EPSR domain
- LAG - an FE/FX port cannot be added to a LAG via the CLI or NMS
- HVLAN - an FE/FX that is configured as a network interface cannot be an HVLAN (cannot set to double tag)

16.8.4.2 Slot Restrictions

None. The FE10 can support network interfaces in **any** service module slot in the fMAP 9000.

16.8.5 Feature Change Details for 6.0

Note: Many of the changes to features apply to all fMAP products and were explained in [14.2](#). These are summarized briefly here and then attributes for the FE/FX upstream interfaces are highlighted.

16.8.5.1 IGMP Snooping

In release 6.0 to support an FE/FX port as a multicast router port, IGMP Snooping needs to know whether an interface's DIRECTION is for a network or customer interface; this is determined by data filling the DIRECTION attribute for an interface as either NETWORK or CUSTOMER.

To control the IGMP snooping of the FE/FX as an upstream port, the snooping mode can be set as follows:

- `MCPASSTHROUGH` - IGMP Snooping will filter IGMP packets and will flood all multicast traffic that is received from the multicast router to the Network Interfaces that are set at `INTERNAL` and are members of the VLAN
- `INTERNAL` (snooping) - The interface will reconfigure the hardware to limit the forwarding of multicast packet only to the ports that have expressed interest in the multicast group.
- `EXTERNAL` (snooping) - Behind the interface (towards the customer) is a device that has a snooping function; The upstream device will only snoop for the first IGMP report and last IGMP leave message, when it knows there is no more interest for a specific multicast stream on the downstream device.

16.8.5.2 Profiles

The FE10 and FEPORT autoprov profiles are not changed by this feature. The `DIRECTION` field is not present in the FEPORT profile; it can only set on a per port basis. Setting profile attributes to a FE port (using the `SET PORT PROFILE` command) does not change the `DIRECTION` attribute of the port.

16.8.6 Software Upgrade

When the fMAP device is upgraded to 6.0, the `DIRECTION` attribute default is set to `CUSTOMER` for FE10 ports.

16.8.7 Traffic Provisioning

The FE uplink port connected to the TN-1000 should have its egress port rate limiter set to 44Mbps, but the bucket size should be the lowest possible which is 4KB for the FE10.

16.8.8 Maintenance/Alarms

16.8.8.1 Card Fault Conditions

Alarms for FE10 cards normally have `MINOR` severity when all ports have `DIRECTION` set to `CUSTOMER`. However, when an FE10 card has at least one port with `DIRECTION` set to `NETWORK`, the card alarms have `MAJOR` severity, similar to alarm severities for other upstream supporting card types such as GE ports.

16.8.8.2 Port Fault Conditions

Alarms for FE ports normally have `INFO` severity when the `DIRECTION` is set to `CUSTOMER`. However, when `DIRECTION` is set to `NETWORK`, the alarms have `MAJOR` severity, similar to alarm severities for other upstream port types such as GE ports

16.8.8.3 Port Outage Threshold

Once an FE port is designated as a network link, a fault on the link will cause a critical port outage threshold alarm if there are no other network links available in the system. This behavior is the same as for GE network links.

16.8.8.4 LEDs

There is no change to the handling of card and link LEDs. Refer to the fMAP Log Manual for details.

16.8.8.5 Logs

A generic log is produced when an FE port direction attribute is changed. Following is an example.

```
PORT008 2004-11-09 11:13:00 3258 INFO
Location: Slot: 3 Port: 1
Description: Provisioning applied to the port database

CLI002 2004-11-09 11:13:00 3257 INFO
User: "officer" on system console entered CLI command:
SET PORT=3.1 FE DIRECTION=net
```

16.8.9 Example Configuration

[Figure 16-3](#) shows a more detailed example and includes sample values for the cards, ports, interfaces, and VLANs. This example configuration is especially useful in understanding many of the features available in 6.0. [Table 16-6](#) shows the key values used in engineering the configuration for the video VLAN.

Note the following key concepts:

- One RG (RG_1) supports IGMP Snooping; the other (RG_2) does not.
- There are three types of interfaces, GE, FE/FX, and ADSL.

TABLE 16-6 Key Parameters / Values in [Figure 16-3](#) for VLAN 80 (V_Video)

Interface (device - link) ^a	DIRECTION	IGMP Snooping Type	FORWARDING
A-1	NETWORK	MCPASSTHROUGH	PRIMARYUPSTREAM
B-2	NETWORK	MCPASSTHROUGH	SECONDARYUPSTREAM
A-3	NETWORK	MCPASSTHROUGH	STP
B-3	NETWORK	MCPASSTHROUGH	STP
A-4	NETWORK	MCPASSTHROUGH	STP
B-4	NETWORK	MCPASSTHROUGH	STP
B-5	NETWORK	MCPASSTHROUGH	STP
C-5	NETWORK	MCPASSTHROUGH	STP

TABLE 16-6 Key Parameters / Values in Figure 16-3 for VLAN 80 (V_Video) (Continued)

Interface (device - link) ^a	DIRECTION	IGMP Snooping Type	FORWARDING
C-6	NETWORK	EXTERNAL - Device (9000) “behind” interface does support IGMP Snooping)	STP
D-6	NETWORK	MCPASSTHROUGH	STP
D-7	NETWORK	MCPASSTHROUGH	STP
E-7	NETWORK	MCPASSTHROUGH - Doesn't care, since not going to another device	STP
E-8	CUSTOMER	INTERNAL - Device (RG) “behind” interface does not support IGMP Snooping	DOWNSTREAM
B-9	CUSTOMER	EXTERNAL - Device (RG) “behind” interface does support IGMP Snoop- ing)	DOWNSTREAM

a. VLAN on every device has a FORWARDINGMODE of UPSTREAMONLY (UFO mode).

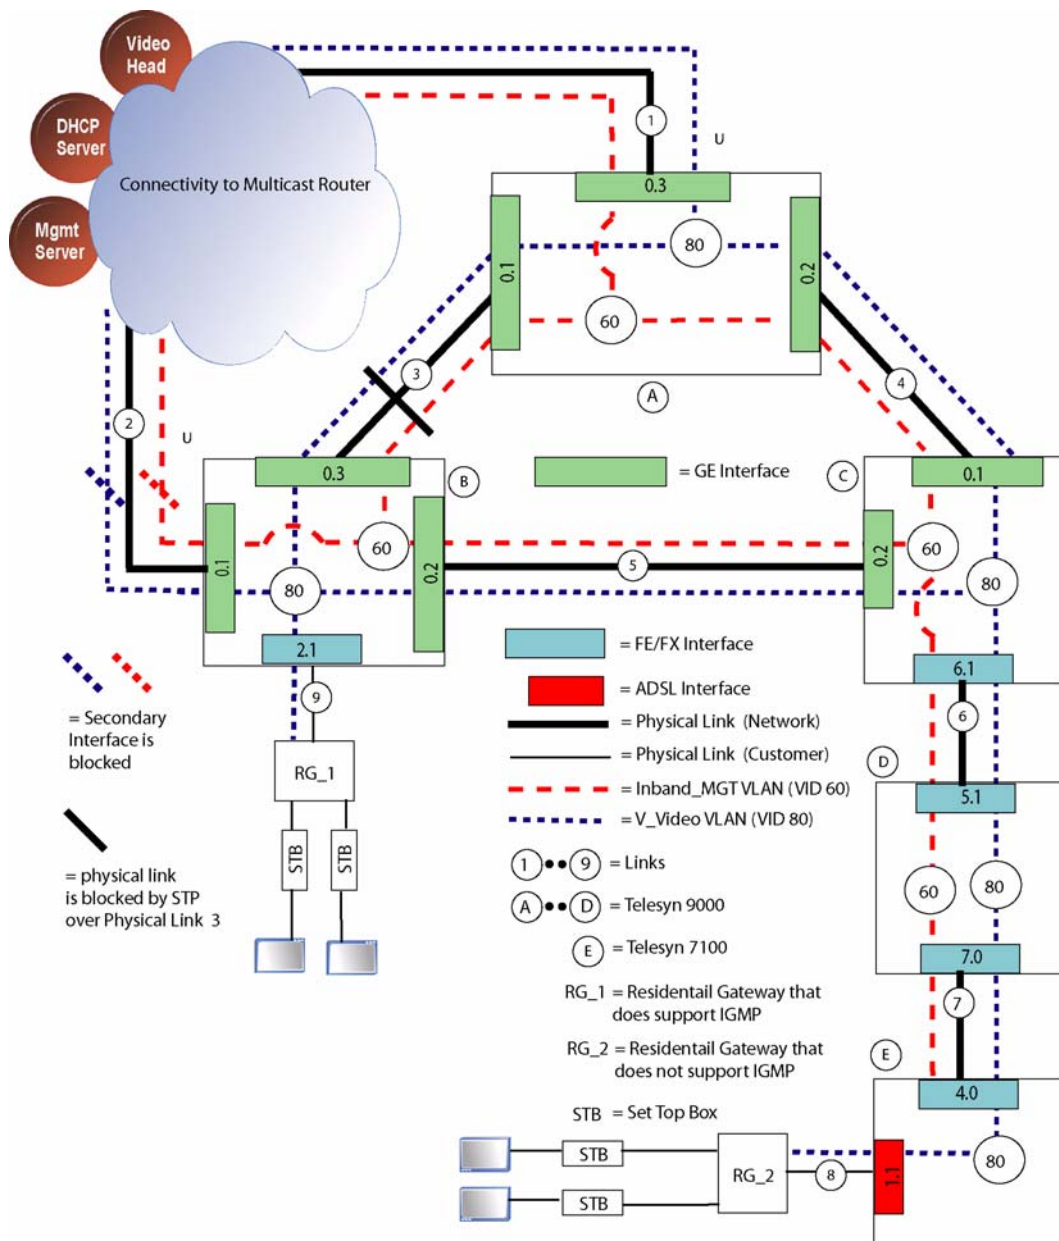


FIGURE 16-3 FE/FX Uplink Configuration Using STP - Example

The inband mgmt VLAN is non-UFO [STANDARD], and has a DIRECTION of NETWORK. The other attributes are not applicable; the management application uses learning to find and query a device.

16.8.10 Convert Subtending Ring in D43 from GE3 to GE8

Note: Before performing this procedure, ensure you understand the VLAN configuration for the subtending ring, since you must configure the GE8 ports to have the same VLAN configuration.

TABLE 16-7 Convert D43 Subtending Ring to GE8

Step	State or Action	Details
1.	Configuration Notes	<p>Before performing this procedure, ensure you understand the EPSR/VLAN configuration for the subtending ring, since you must configure the GE8 ports to have the same EPSR/VLAN configuration. This ensures that service will be disrupted on the ring only as long as it takes to swap the physical links.</p> <p>The GE8 must be inserted in a Resource Module (RM) slot. (These are slots 6/7 in 9700 and slots 6/8 in 9400). If there is already an SM card in one or both slots, the card(s) will have to be moved.</p>
2.	Insert the GE8 card into the Resource Module Slot 7.	For this example, there is only one GE8 card being used for both interfaces (7.0 and 7.1).
3.	Copy the GE8 load from the network server onto CFC Flash.	Do not restart the GE8 card yet.
4.	<p>Move the VLAN configuration from the GE3 ports to the GE8 ports</p> <p>Move 10.1 -> 7.0</p> <p>Move 11.0 -> 7.1</p>	Refer to the before and after figures. This ensures that when the subtending ring is associated with the GE8ports, there is no disruption of service
5.	Set the GE8 card to the PREF-LOAD that you copied over in Step 3.	
6.	Restart the GE8 card.	The card will come up with the correct load and VLAN configuration, , but will not be configured for EPSR yet.
7.	Disable the primary interface for the subtending ring. (10.1)	Since the GE3 subtending ring is still working, data traffic is switched to the secondary port (11.0)
8.	Disable the EPSR domain	Although the EPS domain is now disabled, there are no loops since you disabled the primary interface in Step 7.
9.	Delete the Primary Interface from the domain.	

TABLE 16-7 Convert D43 Subtending Ring to GE8

Step	State or Action	Details
10.	Add the association between the EPSR domain primary interface and the GE8 interface.	In the example, 7.1 is added as the Primary.
11.	(Re)enable the EPSR domain	The logical EPSR configuration is now complete. All that remains is to swap the physical links so that traffic will flow over the GE8 Primary Interface.
12.	Swap the physical link (GE3 with GE8)	This is for the Primary Interface only.
13.	Repeat Steps through 7. through 12. for the secondary interface.	

16.8.11 Upgrade the D42 System from the CFC24 to CFC56

TABLE 16-8 Simplex Software Upgrade - Next System (D42)

Step	State or Action	Details
1.	Pre-Upgrade Notes	With the upgrade of the D43 system , the media card now has the correct CFC56 and FE10 loads for 8.0. The media card can now be used during the D42 upgrade.
2.	Back up the current database: BACKUP DATABASE	For network reliability purposes, backup the existing configuration database to the external network server using the BACKUP DATABASE command.
3.	Prepare the config file to be used for reconfiguring the system after the CFC upgrade. Run backup config.	BACKUP CONFIG FILE=<file name>.txt
4.	Upload the config file to a network server.	Using the PUT FILE command, send the config file to the local network server.
5.	Save a copy of the config file just uploaded for downgrade purposes, if required.	Save a copy of the file.
6.	Edit the config file so that it can be used to configure a CFC56 system	Make the following changes: <ol style="list-style-type: none"> 1. Change all references for CFC24 to CFC56 2. Set the PREFLOAD for the CFC56 8.0 software load. 3. Set the PREFLOAD for the SM card (in this case the FE10) to the correct 8.0 load for the card. 4. Save the file

TABLE 16-8 Simplex Software Upgrade - Next System (D42)

Step	State or Action	Details
7.	Insert the media card into the CFC24	The media card already has the correct CFC56 and FE10 load.
8.	ftp the edited config file onto the CFC Flash for the CFC24.	This is done so that it may copied onto the media card.
9.	ftp the edited config file from the CFC Flash to the media card.	This will be used when the media card is inserted into the CFC56.
10.	Copy the FE10 load from the media card to the CFC Flash	
11.	Copy the FE10 load from the CFC Flash to the FE10 card Flash., but not reboot the card.	
12.	Remove the media card from the CFC 24 card and insert it into the CFC56 card that will going into the D42 system	The CFC56 card should be in an anti-static environment.
13.	In preparation for replacing the hardware, power down the system	Turn both A and B power circuit breakers OFF .
14.	Replace the CFC24. with the CFC56.	Using anti-static procedures: <ol style="list-style-type: none"> 1. Remove the CFC cables for the CFC24. 2. Remove the CFC24 and replace with the CFC56. 3. Replace the CFC cables. 4. Remove/replace any SM cards that will be part of the upgrade.
15.	Power up the system	Turn both A and B power circuit breakers ON . <i>Note: When the system comes up, you now have a CFC56 with no loads, but all relevant loads on the media card. Also, the system will have no IP address, and so you must use the console connection to communicate with the system</i>
16.	List the files that are on the media card.	
17.	Copy the 8.0 CFC56 load and config file from the media card to the CFC Flash	
18.	Set the preferred load on the CFC56	Set the new CM load as the PREF load for the CFC56 (TN-407-A). SET CARD=ACTCFC PREFLOAD=<latest CM load>

TABLE 16-8 Simplex Software Upgrade - Next System (D42)

Step	State or Action	Details
19.	Run restore config to purge the database and reboot the system.	RESTORE CONFIG FILE=<file name>.txt OUT=<file name>.log Note: The system will reboot after the RESTORE command completes execution.
20.	After the restart and the system recovers, the restoration of the config file will be in progress. To check the progress of the restore use the SHOW CONFIG STATUS command.	SHOW CONFIG STATUS <i>Note: This step may take a long time to complete depending on the number of commands included in the config file.</i> Monitor the progress of the restore. Once SHOW CONFIG indicates that the restore is complete, inspect the resulting log file and ensure that all configuration is restored correctly and completely.
21.	Ensure the EPS ring has come back up and that traffic is passing correctly.	SHOW EPS FULL The system should only have been down for as long as it took to replace the CFC24 card, power up the system, copy the files and do a restore config.
22.	Add new SMs if required. At this point, any new SMs can be added.	Using anti-static procedures Add new SMs. Load them with the latest 8.0 release loads. The new load is set as PREFERRED on all cards. SET CARD=<sm slot> PREFLOAD=<8.0 SM load> The user can set the preferred load for each card individually or input a comma-separated slot list of all cards Restart the cards. RESTART CARD=<sm slot-list> The user can restart each card individually or input a comma-separated slot list of all cards.
23.	Back up the new database	Back up the new database: BACKUP DATABASE

16.8.12 Upgrade the D32 System from the CFC24 to CFC56

This procedure includes the same steps as upgrading the D42 system, with the media card being used since it contains the correct CFC56 and FE10 loads.

Note: Since this is the system that connects to the upstream network, all traffic (all subrings) will be lost during the upgrade.

16.8.13 Convert the D32 Uplink to GE8 Interfaces (LAG)

TABLE 16-9 Convert D32 Uplink to GE8 Interfaces with LAG

Step	State or Action	Details
1.	Insert the media card into the CFC56, and ensure the GE8 8.0 load is included on the media card.	The media card makes it unnecessary to retrieve the load from the network server.
2.	Using anti-static procedures, insert the GE8 card into slot 7 (a Resource Module slot)	Since the shelf is powered up, the GE8 card will come up.
3.	Copy the GE8 load from the media card to the CFC Flash	
4.	Set the GE8 load as the PREF-LOAD	
5.	Disable and Enable the GE8 card.	
6.	Create a LAG over the GE8 interfaces that will participate in the LAG	There are only two interfaces on one card that make up this LAG
7.	Add the data VLANs to both GE8 interfaces of the LAG	Also, ensure the IGMP settings for the GE8 interfaces are the same as for the GE3 interfaces.
8.	Move the GE3 uplink (10.2) over to a GE8 uplink (7.0)	Data service will be disrupted only as long as it takes to move the physical link.
9.	Connect the other LAG GE8 port (7.1) with the upstream device	
10.		

17. Routine Administration

17.1 Overview

This section describes some of the routine system administrative procedures performed on the system.

17.2 Database Management

17.2.1 Overview

The system database contains important configuration data. Periodic back ups of the configuration database are important in the event of its corruption or loss. In the event the database is compromised, back ups provide the user with the latest system configuration data. Also, during a software release upgrade, a database back up is performed in case a downgrade is required after the new load files have been committed to. Backing up the database as detailed in the upgrade procedure insures that the system has the most recent data available and will be completely restored.

17.2.2 Database Back Up

During a normal software release upgrade, the database is manually backed up in the event that a downgrade is required after the new load files have been committed to. If a downgrade is required, a database restoration is performed.

Here are the steps involved in performing a database back up:

1. Designate a network server for secure storage of the current configuration database
2. Using the BACKUP DATABASE command, back up the current configuration database to the secure server
3. The user may execute the SHOW TRANSFER command to display the progress of the backup

Following is an example of a database backup:

```
officer SEC> BACKUP DATABASE FILE=F42_DBASE TFTP SERVER=172. 16. 17. 18
Command has been submitted
Info (033758): Database backup submitted with Transfer ID: 0
officer SEC> SHOW TRANSFER ALL
```

ID	CMD	Remote file	Local file	Server	Mode	Status	MB
0	PUT	F42_DBASE		172. 16. 17. 18	TFTP	Progress	0

```
officer SEC> Info (033759): Database backup succeeded
```

17.2.3 Database Purge

The PURGE DATABASE command *ERASES* the current configuration database. This command would be used if for some reason the user wants to reconfigure the system back to factory defaults. What happens, after the command is entered, is that the system reboots and recovers with factory defaults.

Note: TELNET is disabled by default. If the user is connected through TELNET, after the PURGE DATABASE completes, TELNET will be set back to its default setting- disabled. The user will no longer be connected to the system. In order to reconnect, TELNET must be enabled. The user should connect and login in to the system through the CONSOLE port on the CFC card prior to executing the PURGE DATABASE command.

The user can restore the original configuration database using the RESTORE DATABASE command, provided that it had been previously backed up, or by rebuilding it manually using CLI commands.

Example of the PURGE DATABASE command:

```
officer SEC> PURGE DATABASE FORCE
Command has been submitted
PURGE DATABASE - success
```

Note: Use of the PURGE DATABASE command can cause network outages. After the PURGE DATABASE completes, TELNET will be disabled.

17.2.4 Restore Database

The user can replace the current system database with a database that has earlier been backed up to a network server. The user should be aware of backward compatibility criteria for system releases. This operation is performed using the RESTORE DATABASE command.

The normal flow of commands would be:

1. Execute the RESTORE DATABASE command

Note: The system will reboot after the RESTORE DATABASE command is completed.

17.2.5 Database Transaction Failure

If a requested database transaction fails to successfully complete:

- A warning message will be output to the user's CLI session indicating that the command the user entered will be executed, but the transaction will not be recorded in the database.
- A major alarm will be raised.
- A log will be generated. The log message will read "**Database transaction failure**".

In the unlikely event that this situation occurs, the user should capture all management, error, and crash logs, then contact Allied Telesis Technical Support. Note that while this alarm is present any configuration data that is entered by the user may not be stored in the database.

TABLE 17-1 Database Commands Summary

Noun	Verb	Syntax	Description
DATABASE	BACKUP	BACKUP DATABASE FILE={ destinationfile unit:destinationfile serverpath/destinationfile } [{ TFTP SERVER={ ipaddress hostname } ZMODEM FTP SERVER={ ipaddress hostname } USER=userid PASSWORD=password }]	The BACKUP DATABASE command backs up the contents of the system configuration database to a file on an external network server.
TRANSFER	SHOW	SHOW TRANSFER [= { transferid-list ALL }]	Displays the status and progress of a file transfer.

TABLE 17-1 Database Commands Summary (Continued)

Noun	Verb	Syntax	Description
DATABASE	PURGE	PURGE DATABASE [FORCE]	The PURGE DATABASE command purges all contents in the system configuration database and then automatically restarts the control module.
DATABASE	RESTORE	RESTORE DATABASE FILE={ sourcefile unit:sourcefile serverpath/sourcefile } [{ TFTP SERVER={ ipaddress hostname } ZMODEM FTP SERVER={ ipaddress hostname } USER=userid PASSWORD=password } [FORCE]	The RESTORE DATABASE command rewrites the configuration database with contents from a file transferred from an external network server.

17.3 Delete Obsolete Users

For system and network security reasons, obsolete users should be deleted from the system when they are no longer required. Obsolete users can be deleted from the system using the **DELETE USER** command.

17.4 DELETE Obsolete FILES

The system is designed to provide memory for the storage of system files. Obsolete files should be deleted when they are no longer required. The user can display all system files using the **SHOW FILES** and delete obsolete files using the **DELETE FILE** commands.

17.5 Scripting

This product provides the user with the functionality to execute user-defined command (CLI) scripts. The normal flow for the use of scripts includes:

- the user designs and edits a script file using any plain text editor (NOTE: script files cannot be edited from the CLI on the system)
- using the **GET FILE** command the user puts the file onto the control module card into FLASH

- the user can then look at the file and execute it using the SHOW and EXECUTE commands

The following text describes the scripting commands.

The **SHOW SCRIPT** command displays the contents of a Command Line Interface (CLI) script. A script contains CLI command(s) that are executed using the **EXECUTE SCRIPT** command.

The **EXECUTE SCRIPT** command processes all of the commands entered in the specified filename. The script file contains one or more CLI commands. The first line in the file must contain a comment that identifies the file as a script. Other text can also exist on the line, but the word 'script' must appear some place in the line. Comments are identified as a hash(#) as the first character on a line. A CLI command in the script file must occupy a single line. A command cannot span more than one line. If a command requires user interaction like a confirmation, the user response text is included on the line after the command.

Here is a summary of the rules for scripts:

1. The commands in the script file must be syntactically correct.
2. Each command must be on ONE LINE only. In other words, there is no continuation character.
3. The first line of the script file must be a comment line with the word “script” in it. This is used to verify that a file is a valid script file. It is used to prevent the execution of a non-script file (i.e. load file).
4. If a command returns a failure response, the script will continue to process commands following the error. It will not exit due to a parsing error OR command failure.
5. If a command requires a confirmation string, the NEXT LINE must be a ‘Y’ to provide the confirmation response. If something other than a Y or N is provided, the script will quit. Alternatively, the user could disable prompting at the beginning of the script.
6. The user can provide comments and blank lines in script files.
7. The commands used must be within the realm of the user (i.e. Security Officer, Manager, User).

The contents of a script file are played back as written. Any syntax errors in the file are detected as the script is run. If an error is encountered, the device under maintenance is left in an unknown condition.

Following is an example script, between the dashed lines.

```
-----  
# This is an example script file. It must have the word SCRIPT in this line
```

```
#
```

```
# First, list all users already configured in the system
```

```
#
```

```
SHOW USER
```

```
#
```

```
# Now look at the system user information
```

```
#  
SHOW SYSTEM USERCONFIG  
#  
# Now reset all of the users. This command requires a confirmation. Notice the 'Y' on the line following  
#  
RESET USER ALL  
Y  
#  
# Now look at the user information again to verify that the counters were reset  
#  
SHOW USER  
SHOW SYSTEM USERCONFIG  
#  
# This is the end of the script
```

The command flow would be similar to the following:

Create the script using a text editor on the user's workstation. For example, assume a user edits the script above and saves it, naming it ExampleScript.txt.

To load the script file, use the GET FILE command. For example:

```
GET FILE=ExampleScript.txt TFTP SERVER=172.28.11.31
```

To see that the script file exists on the system:

```
SHOW FILES
```

To view the contents of a script file, use the SHOW SCRIPT command:

```
SHOW SCRIPT=ExampleScript.txt
```

The contents of the file will be displayed.

To run the script, use the EXECUTE SCRIPT command

```
EXECUTE SCRIPT=ExampleScript.txt
```

The EXECUTE SCRIPT command will present a confirmation string before executing the script.

When the script is no longer needed, it should be deleted using the DELETE FILE command:

DELETE FILE=ExampleScript.txt

TABLE 17-2 Script Commands

Noun	Verb	Syntax	Description
SCRIPT	SHOW	SHOW SCRIPT=filename	The SHOW SCRIPT command displays the contents of a Command Line Interface (CLI) script.
	EXECUTE	EXECUTE SCRIPT=filename	The EXECUTE SCRIPT command processes all of the commands specified in the specified filename.

18. Alarms / Troubleshooting System

18.1 Overview

This section describes alarms that are part of the system, associated alarm indicators, associated logs and alarm clearing procedures. For any system card, this section assumes that the card is correctly provisioned, is providing service (card Administrative and Operational States are both UP), and a problem has occurred that may affect service.

The system provides certain redundant functionality, for example, redundant network interfaces can be equipped and provisioned. If a failure occurs on one interface, the other interface will continue to provide service through the use of STP, LAG, and EPSR.

The system will automatically recover from severe software faults. If the software can no longer process application data, the system will initiate recovery methods that may involve reloading the system.

The system continuously monitors system cards and their interfaces. When anomalies are discovered, the following indications are generated to alert the user:

- **System logs** - Note that logs can be displayed using the **SHOW LOG** command. System management logs can be output to a server where they can be easily analyzed. Refer to **Configuring Management Interfaces** for more information.
- **Traps** - The system uses SNMP to allow management devices (using an SNMP enabled MIB browser) to query the Managed Information Base (MIB) of a device or a set of devices, and to receive unsolicited messages (traps) from the device for critical events. Refer to section **Configuring Management Interfaces** of this manual for more information.
- **Visual alarms** - The LED(s) of the associated card(s) are illuminated.

Note: This Section provides an overview of the maintenance system for fMAP products. For actual log and alarm output, refer to the Log and Troubleshooting Guide.

18.2 Querying Alarm Status (SHOW ALARMS, CLEAR ALARMS)

18.2.1 Overview

There are three levels of alarm severity on the system, Critical, Major, and Minor. In general, they are described as:

- **Critical** - A critical alarm is used to indicate that a severe, service-affecting condition has occurred and that immediate corrective action is imperative.
- **Major** - A major alarm is used to indicate a serious disruption of service or the malfunctioning or failure of important circuits. These troubles require immediate attention and response by the crafts person to restore or maintain system capability. The urgency is less than critical situations because of lesser immediate or impending effect on service or system performance.
- **Minor** - Minor alarms are used for troubles that do not have serious effect on service to customers or for troubles that do not affect essential system operation.
- **Info** - Represents an informational message. No explicit action is required of the user.

When an anomaly occurs, the system generates management logs. In reality, every time an event occurs on the system, a log is created. Logs are also generated when performance measurement thresholds have been exceeded. For an efficient management configuration, users can configure logs to be filtered, output, and shown on specified devices and formats.

The remainder of this section details the querying of alarms, the CLI commands used to query alarms, and provides an alarm table. The alarm table describes notable alarms, their meaning, and how to clear them.

18.2.2 Displaying alarms

Alarms can be displayed using the SHOW ALARMS command. Depending on the parameters used when the user inputs the SHOW ALARMS command, different information will be provided in the response. The SHOW ALARMS command will be discussed below.

- **SHOW ALARMS ALL** - To display all system alarms, the user inputs this command. The alarm statuses for all system cards will be displayed.
- **SHOW ALARMS CARD** - To display alarms for a specified card
- **SHOW ALARMS PORT** - To display alarms for a specified port
- **SHOW ALARMS SEVERITY** - To display alarms according to their severity

Here are some examples of the use of the SHOW ALARMS command:

```
officer SEC> SHOW ALARMS ALL
```

```
-----  
Shelf                Can Not Read Temperature Major
```

Sensors		
Slot 04, Port 00	No Peer Present	Info
Slot 04, Port 01	No Peer Present	Info
Slot 04, Port 02	No Peer Present	Info
Slot 04, Port 03	No Peer Present	Info
Slot 04, Port 06	No Peer Present	Info
Slot 04, Port 07	No Peer Present	Info
Slot 04, Port 08	No Peer Present	Info
Slot 04, Port 09	No Peer Present	Info
Slot 04, Port 10	No Peer Present	Info
Slot 04, Port 11	No Peer Present	Info
Slot 04, Port 12	No Peer Present	Info
Slot 04, Port 13	No Peer Present	Info
Slot 04, Port 14	No Peer Present	Info
Slot 04, Port 15	No Peer Present	Info
Slot 17, Port 00	No Peer Present	Info
Slot 17, Port 01	No Peer Present	Info
Slot 17, Port 02	No Peer Present	Info
Slot 17, Port 03	No Peer Present	Info
Slot 17, Port 04	No Peer Present	Info
Slot 17, Port 05	No Peer Present	Info
Slot 17, Port 06	No Peer Present	Info
Slot 17, Port 07	No Peer Present	Info
Slot 17, Port 08	No Peer Present	Info
Slot 17, Port 09	No Peer Present	Info
Slot 17, Port 10	No Peer Present	Info
Slot 17, Port 11	No Peer Present	Info
Slot 17, Port 12	No Peer Present	Info
Slot 17, Port 13	No Peer Present	Info
Slot 17, Port 14	No Peer Present	Info
Slot 17, Port 15	No Peer Present	Info
Fan Module	Fan Module Not Present	Major

officer SEC> SHOW ALARMS CARD=2

Slot 02, 04	SB 2.5v ANALOG Power	Critical
	Failed	

officer SEC> SHOW ALARMS SEVERITY=CRITICAL

Shelf	Port Outage Threshold	Critical
Slot 02, 04	SB 2.5v ANALOG Power	Critical
	Failed	

officer SEC> SHOW ALARMS SEVERITY=MAJOR

Shelf	Can Not Read Temperature Sensors	Major
Slot 00, Port 00	Loss Of Signal	Major
Slot 00, Port 00	Loss Of Sync	Major
Slot 00, Port 00	Loss Of Link	Major
Slot 00, Port 01	Loss Of Signal	Major
Slot 00, Port 01	Loss Of Sync	Major
Slot 00, Port 01	Loss Of Link	Major
Slot 00, Port 02	Loss Of Signal	Major
Slot 00, Port 02	Loss Of Sync	Major
Slot 00, Port 02	Loss Of Link	Major
Slot 01, Port 00	Loss Of Sync	Major
Slot 01, Port 00	Loss Of Link	Major
Slot 01, Port 02	Loss Of Signal	Major
Slot 01, Port 02	Loss Of Sync	Major
Slot 01, Port 02	Loss Of Link	Major

officer SEC> SHOW ALARMS SEVERITY=MINOR

Slot 09	Card Not Present	Minor
---------	------------------	-------

officer SEC> SHOW ALARMS SEVERITY=INFO

Slot 04, Port 00	No Peer Present	Info
Slot 04, Port 01	No Peer Present	Info
Slot 04, Port 02	No Peer Present	Info
Slot 04, Port 03	No Peer Present	Info
Slot 04, Port 06	No Peer Present	Info

Slot 04, Port 07	No Peer Present	Info
Slot 04, Port 08	No Peer Present	Info
Slot 04, Port 09	No Peer Present	Info
Slot 04, Port 10	No Peer Present	Info
Slot 04, Port 11	No Peer Present	Info
Slot 04, Port 12	No Peer Present	Info
Slot 04, Port 13	No Peer Present	Info
Slot 04, Port 14	No Peer Present	Info
Slot 04, Port 15	No Peer Present	Info
Slot 17, Port 00	No Peer Present	Info
Slot 17, Port 01	No Peer Present	Info
Slot 17, Port 02	No Peer Present	Info
Slot 17, Port 03	No Peer Present	Info
Slot 17, Port 04	No Peer Present	Info
Slot 17, Port 05	No Peer Present	Info
Slot 17, Port 06	No Peer Present	Info
Slot 17, Port 07	No Peer Present	Info
Slot 17, Port 08	No Peer Present	Info
Slot 17, Port 09	No Peer Present	Info
Slot 17, Port 10	No Peer Present	Info
Slot 17, Port 11	No Peer Present	Info
Slot 17, Port 12	No Peer Present	Info
Slot 17, Port 13	No Peer Present	Info
Slot 17, Port 14	No Peer Present	Info
Slot 17, Port 15	No Peer Present	Info

officer SEC> **SHOW ALARMS PORT=5.0**

Slot 05, Port 00	Loss Of Link	Info
------------------	--------------	------

TABLE 18-1 SHOW ALARMS Command Summary

Noun	Verb	Syntax	Description
ALARMS	SHOW	<pre> SHOW ALARMS [{ ALL CARD={ slot-list ACTCFC INACTCFC ALL } INTERFACE={ type:id- range id-range ifname-list ALL } }] [SEVERITY={ CRITICAL MAJOR MINOR INFO ALL }] [FULL] </pre>	<p>The SHOW ALARMS command displays alarm conditions on system components.</p> <p>The FULL parameter makes SHOW ALARMS show all alarms regardless of whether or not they are masked.</p>
ALARMS CARD MCASTGROUPLIMIT	CLEAR	<pre> CLEAR ALARMS CARD={ slot-list ALL } MCASTGROUPLIMIT </pre>	<p>Clears the Multicast Group Limit alarm.</p>

18.3 Alarm System Features

18.3.1 Overview of Alarm System

To understand the alarm system, the user should first understand the following concepts.

1. Relation of components (hardware and software) in providing service.

For the delivery of services to succeed, a set of components (both hardware and software) must work together, since each provides some attribute of that service. Moreover, these components can be in a hierarchy, with certain components dependent on other components to function correctly. Otherwise the service may not work (or be degraded). In many cases, components have a **parent-child relationship** to each other; a parent must exist and be operational before a child can. (An exception is the MLPPP, which is explained later.) Moreover, a component may have children (and therefore be a parent to these components), and yet also be a child to another component (which is its parent).

Note: This component hierarchy is based on the way components are dependent on each other, and may or may not match how the components are physically or logically connected.

2. Potential customer impact on the loss of service of a component

Since the loss of a component may impact one, several, or even hundreds of customers, the system tries to ensure that alarm indicators for components are appropriate. One key attribute of a component is whether it is **customer-interfacing** or **network-interfacing**, since in most cases a network-interfacing component has a greater impact on loss (or degradation) of service. Also, when a component is higher up in the hierarchy (one that is a parent to other components), the impact of a loss of service will usually be greater.

3. The final status of components the user wishes to have

As explained in 4.2 and 4.3, most components have an **Administrative State**, an **Operational State**, and a **Status**, the result of what the system has done with the component. (This is shown by placing the two states and status together but separated by dashes, such as UP-UP-ONLINE). In most cases, the default for the fMAP products is for the components, when installed, to have an Administrative State of UP, and the system to try to make the Operational State go to UP as well. The component can then provide service and go ONLINE if there are no errors. The alarm system is based on this idea; alarms are produced when the user **intends** to put the customer feature into service (with all related components UP-UP-ONLINE), but one or more components do not function correctly.

Note: In some cases alarms are maintained (but with reduced severity) when the component is disabled; this provides some history of the defect condition.

This third concept works with the previous two in that when a component fails (especially a card), the components further down the hierarchy may no longer be able to provide service, yet they will not produce an alarm since they are dependent on the parent component. These child components will then have an Operational State of DOWN and a status of DEPENDENCY. This is useful in isolating the source of the problem, as explained below.

Refer to [Figure 18-1](#); in this figure, the hierarchy and state of components are shown when the system is in service and there are no errors. This figure will be used to illustrate various error conditions.

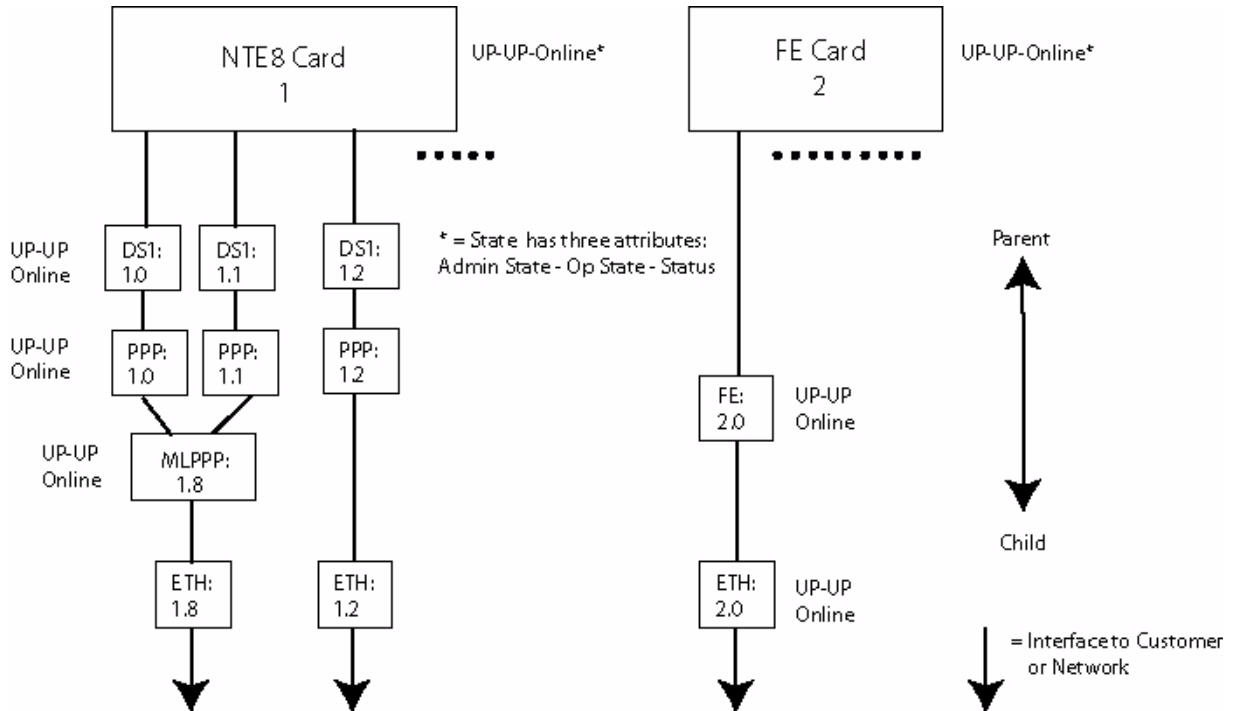


FIGURE 18-1 Component Hierarchy - No Error Conditions

18.3.2 Alarm types - (Interface, Card, System) and the Alarm System

Following the concepts in [18.3.1](#), there are the following types of alarms for the fMAP product:

- Interface** - These are alarms that involve components that face outward from the system. In most cases these are towards the customer, although in some situations these could be towards the network. These alarms usually indicate to the user that service (or a particular service type) cannot be provided to a customer.

When an interface fails in a previously working system, it is usually in one of two states:

- Failed** - The interface cannot provide any service, and the user must assume that all services for that interface are lost. The Op State is DOWN. An alarm is usually produced (the severity depending on whether the component is disabled or enabled), and if the card has Port/Interface LEDs, the Inservice (INS) is OFF while the Error (ERR) is lit.
- Degraded** - The interface can still provide service, but at a degraded level. The OP state is still UP. An alarm may be produced, with severity at a lesser level than an interface in the Failed state. If the card has Port/Interface LEDs, the Inservice (INS) is still ON while the Error (ERR) LED is ON.

Note: When these alarms occur, there is an associated log as well, with a category of PORT/INTF/PSPN. Refer to the fMAP Log Manual, which contains a complete list of all logs, associated alarm messages, component status, LED states, and descriptions.

2. **non-CFC Card** - These are alarms associated with a specific card. Since the card supports multiple interfaces, the loss of a card usually means the potential loss of multiple interfaces. When a card fails in a previously working system, the card is usually in one of two states:
 - Failed - The card cannot provide any service, and the user must assume that any associated interfaces cannot provide service(s) as well. The Op State is DOWN. An alarm is usually produced (the severity depending on whether the alarm is disabled or enabled). For the card's LEDs, the Fault (FLT) LED is ON, and the OK to Pull (OTP) is ON as well, meaning that since service has been lost the card can be pulled without a further degradation of service. Moreover, the CFC severity LED is ON, which is usually MINOR but may be MAJOR, depending on the concepts listed in [18.3.1](#)
 - Degraded - The card can still provide service, but the interfaces it supports may not be able to provide a full level of service. The Op State is still UP. An alarm is usually produced (the severity depending on whether the alarm is disabled or enabled). For the card's LEDs, the Fault (FLT) LED is ON, and the OK to Pull (OTP) is OFF, since removing the card could further degrade service. For the CFC card, its severity LED is ON, which is usually MINOR but may be MAJOR, depending on the concepts listed in [18.3.1](#).
3. **CFC card** - The CFC card must be described separately, since a loss of service on this card can affect the entire system. Moreover, with the fMAP 9700, the CFC is in duplex mode; if the active CFC cannot function, the system will switch activity to the other CFC. With this ability in mind, following is the behavior for the CFC states.
 - Failed, Active - The CFC card cannot provide any service, and so all cards and interfaces are down. The Op State is DOWN. However, since this will trigger a switch of activity, this CFC will become the inactive CFC, and so its state will change to Degraded.

Note: In duplex mode, an active and FAILED CFC will always switch activity to ensure services are not lost. Because of this, there is never an active CFC in a FAILED state for duplex.

- Degraded, Active - The CFC can still provide service, but it is in simplex mode and so may fail if there are further alarms for the CFC. The Op State is still UP. An alarm is produced, and this is usually major or critical, since a further alarm may mean loss of all service. The CFC FLT and severity LEDs are ON, with the severity as MAJOR or CRITICAL that matches the alarm level. Note that the Inservice LED is ON and the OTP LED is OFF, since the card cannot be pulled without the loss of the system.
- Failed, Inactive - The system is in duplex mode, and the inactive CFC is not providing service. The Op State is DOWN. An alarm is produced, and this is usually major or minor, since the loss of the card will not affect service. The CFC FLT and severity LEDs are ON, with the severity as MAJOR or CRITICAL that matches the alarm level. Note that the Inservice LED is OFF and the OTP LED is ON, since the card can be pulled without a loss of service.
- Degraded, Inactive - The system is in duplex mode, and so the CFC is not providing service. The Op State is UP and the OTP LED is OFF, since the card may still be needed in case of a failure of the active CFC.

Note: When these alarms occur, there is an associated log as well, with a category of CARD and CFCP. (CFCP logs are about the state of CFC protection.) Refer to the fMAP Log Manual, which contains a complete list of all logs, associated alarm messages, component status, LED states, and descriptions.

4. **System** - These are errors at the system level, and usually involve the CFC being unable to read/process data correctly. This can be caused by something minor, such as the death of a non-critical task, or critical, such as traffic volume overwhelming the CFC. The states and LEDs are similar to the CFC conditions described above.
5. Also, when the number of Interface alarms has reached a number that affects the overall service level of the fMAP, a system alarm is produced, as explained in [18.3.3](#).

18.3.3 Interface/Port Outage Threshold Feature

When a single customer interface fails, the problem is considered minor for the overall state of the system. However, when a certain number of interfaces have failed, this could indicate a greater problem than just the interface. Therefore, there is a separate alarm that is associated with the number of alarms and a severity associated with the number, as follows. (Note that these ranges can be changed.)

- More than 128=Critical,
- 25 to 128=Major
- Less than 24=Minor

When this occurs, there are usually two areas to investigate:

1. Failed Uplinks
2. Failed Service Modules

Note: When all UPLINK interfaces are out of service, a CRITICAL alarm will be raised regardless of the threshold values.

18.3.4 Configurable Alarm Severity (INTERFACE Alarm)

In release 7.0, the user has the ability to control the severity level of interface alarms. As explained in [18.3.3](#), an interface failure is usually considered a minor alarm, but if a **specific** customer interface is considered a high priority, the user can control the severity to highlight that particular interface. The commands to perform this are:

```
SET
INTERFACE={ type:id-range
            | id-range
            | ifname-list
            | ALL
            }
```

```
ALARM
SEVERITY={ NONE
           | INFO
           | MINOR
           | MAJOR
           | CRITICAL
           }
```

```
[ FORCE ]
```

```
-----
SETDEFAULTS
```

```
INTERFACE={ type:id-range
            | id-range
            | ifname-list
            | ALL
            }
```

```
ALARM
[ SEVERITY ]
```

```
SHOW
```

```
INTERFACE
[ ={ type:id-range
    | id-range
    | ifname-list
    | ALL
    } ]
```

```
ALARM
SEVERITY
[ ={ NONE
    | INFO
    | MINOR
    | MAJOR
    | CRITICAL
    | DEFAULT
    | NONDEFAULT
    | ALL
    } ]
```

When doing `SET INTERFACE ALARM SEVERITY`, the specified severity level is used for any failing alarms on the specified interface(s). It is also used to “cap” severities of any degrading alarms on the interface(s). For exam-

ple, if an interfaces has a failing alarm defined that defaults to MAJOR and a degrading alarm that defaults to MAJOR, set the interface alarm severity to MINOR will cause both the failing and degrading alarms to be raised as MINOR. However, if an interface has a failing alarm of MINOR and a degrading alarm of MINOR as defaults, then setting the interface alarm severity to MAJOR only affects the failing alarm; the degrading alarm would still be raised as MINOR.

Using `SETDEFAULTS INTERFACE ALARM SEVERITY` will set the severities back to the hardcoded defaults.

The user can display the current severity setting for any interface using `SHOW INTERFACE ALARM SEVERITY`.

For example:

```
officer SEC>> show interface=18.0 alarm severity
```

```
--- Interface Alarms Settings ---
```

Interface	Severity
18.0	Info (default)

```
officer SEC>> set interface 18.0 alarm severity minor
officer SEC>> sh interface 18.0 alarm severity
```

```
--- Interface Alarms Settings ---
```

Interface	Severity
18.0	Minor

```
officer SEC>> setdefaults interface 18.0 alarm severity
officer SEC>> show interface 18.0 alarm severity
```

```
--- Interface Alarms Settings ---
```

Interface	Severity
18.0	Info (default)

18.4 Overview of Troubleshooting

18.4.1 Overview

Certain assumptions are made concerning the state of a system card. It is assumed that the card has been provisioned, has been providing service, a problem has occurred, and replacement is required. To replace a card, disable it, replace it, and enable it. The system will reload the software and configuration data that was previously provisioned for the card. No reprovisioning is required.

18.4.2 Card Diagnostics

The system provides the user with manual card diagnostics using the DIAGNOSE command. The DIAGNOSE command allows the user to execute tests on card slots. It runs a series of diagnostic tests on the specified card or list of cards.

Currently, out-of-service diagnostics are the only diagnostics supported. For service module, network module, and the inactive CFC cards, the card must be in the administratively DOWN state for the diagnostics to run (DISABLE CARD). Note that diagnostics are executed each time a service module, network module, or inactive CFC card is enabled.

For the control module, the diagnostics are not run immediately, but are instead scheduled to run during the next restart of the control module. Since diagnostics tend to be destructive to configuration data, they cannot be executed on the control module while the system is in service. Therefore, diagnostics are scheduled and are executed during the next system restart. The restart is requested using the RESTART CARD command. Refer to section **Provisioning Network, Service, and Control Modules** for more information on the ENABLE and DISABLE command.

18.5 DIAGNOSE command

18.5.1 Diagnosing the CFC Control Module

For simplex systems, in order to diagnose the CFC Control Module, the user executes the DIAGNOSE CARD command with the OUTOFSERVICE option. The user then executes a restart of the card. During the enabling sequence for the CFC, diagnostics are executed.

Note: Diagnosing the CFC card should be performed during low traffic periods.

For a duplex configuration, the INACTCFC card can be disabled, diagnosed and enabled. However, the ACTCFC card must be diagnosed in the same manner as a system in the simplex configuration.

Note: The following example is for a CFC6, but it can be applied to any CFC.

Following is an example of the use of the DIAGNOSE command.

```
officer SEC> DISABLE CARD=3 FORCE
```

Info (030079): Operation Successful (ADSL16 Slot 03)

officer SEC> **DIAGNOSE CARD=3 OUTOFSERVICE**

Command has been submitted

Info (030079): Operation Successful (ADSL16 Slot 03)

officer SEC> **SHOW CARD ALL**

```
-----
```

Slot	Prov Card Type	State	Faults
0	not provisioned	-	-
1	not provisioned	-	-
2	not provisioned	-	-
3	ADSL16	DOWN-DOWN-Offline	No Faults
4	ADSL8S	UP-DOWN-NotInstalled	Card Not Present
5	not provisioned	-	-
6	not provisioned	-	-
7	not provisioned	-	-
8	CFC6	INACTIVE UP-UP-Online	No Faults
9	CFC6	INACTIVE UP-UP-Online	No Faults
10	not provisioned	-	-
11	not provisioned	-	-
12	CFC6	ACTIVE UP-UP-Online	No Faults
13	CFC6	ACTIVE UP-UP-Online	No Faults
14	not provisioned	-	-
15	not provisioned	-	-
16	not provisioned	-	-
17	not provisioned	-	-
18	not provisioned	-	-
19	not provisioned	-	-
20	not provisioned	-	-
21	not provisioned	-	-
FAN	FAN8	UP-DOWN-Not Installed	Fan Module Not Present

```
-----
```

officer SEC> **ENABLE CARD=3**

Info (030079): Operation Successful (ADSL16 Slot 03)

officer SEC> **SHOW CARD ALL**

```
-----
```

Slot	Prov Card Type	State	Faults
------	----------------	-------	--------


```

-----
0      not provided      -      -
1      not provided      -      -
2      not provided      -      -
3      ADSL16            UP-UP-Online      No Faults
4      ADSL8S            UP-DOWN-NotInstalled  Card Not Present
5      not provided      -      -
6      not provided      -      -
7      not provided      -      -
8      CFC6              INACTIVE UP-UP-Online  No Faults
9      CFC6              INACTIVE UP-UP-Online  No Faults
10     not provided      -      -
11     not provided      -      -
12     CFC6              ACTIVE  UP-UP-Online  No Faults
13     CFC6              ACTIVE  UP-UP-Online  No Faults
14     not provided      -      -
15     not provided      -      -
16     not provided      -      -
17     not provided      -      -
18     not provided      -      -
19     not provided      -      -
20     not provided      -      -
21     not provided      -      -
FAN    FAN8              UP-DOWN-NotInstalled  Fan Module Not Present
-----

```

officer SEC> **DIAGNOSE CARD=8**

% Invalid or incomplete command

officer SEC> **DIAGNOSE CARD=8 OUTOFSERVICE**

Command has been submitted

Error (030083): Command rejected: Operation requires card to be disabled
(CFC6 Slot 08,09)

officer SEC> **DISABLE CARD=8**

Service may be affected, are you sure (Y/N)? Y

Info (030079): Operation Successful (CFC6 Slot 08,09)

officer SEC> **DIAGNOSE CARD=8 OUTOFSERVICE**

Command has been submitted

Info (030079): Operation Successful (CFC6 Slot 08,09)

officer SEC> **ENABLE CARD=8**

Info (030079): Operation Successful (CFC6 Slot 08,09)

officer SEC> **SHOW CARD ALL**

```
-----
```

Slot	Prov Card Type	State	Faults
0	not provisioned	-	-
1	not provisioned	-	-
2	not provisioned	-	-
3	ADSL16	UP-UP-Online	No Faults
4	ADSL8S	UP-DOWN-NotInstalled	Card Not Present
5	not provisioned	-	-
6	not provisioned	-	-
7	not provisioned	-	-
8	CFC6	INACTIVE UP-DOWN-InTest	No Faults
9	CFC6	INACTIVE UP-DOWN-InTest	No Faults
10	not provisioned	-	-
11	not provisioned	-	-
12	CFC6	ACTIVE UP-UP-Online	No Faults
13	CFC6	ACTIVE UP-UP-Online	No Faults
14	not provisioned	-	-
15	not provisioned	-	-
16	not provisioned	-	-
17	not provisioned	-	-
18	not provisioned	-	-
19	not provisioned	-	-
20	not provisioned	-	-
21	not provisioned	-	-
FAN	FAN8	UP-DOWN-Not Installed	Fan Module Not Present

```
-----
```

officer SEC> **SHOW CARD ALL**

```
-----
```

Slot	Prov Card Type	State	Faults
0	not provisioned	-	-
1	not provisioned	-	-
2	not provisioned	-	-

```
-----
```

3	ADSL16	UP-UP-Onl i ne	No Faul ts
4	ADSL8S	UP-DOWN-NotI nstal l ed	Card Not Present
5	not provi si oned	-	-
6	not provi si oned	-	-
7	not provi si oned	-	-
8	CFC6	I NACTI VE UP-UP-Onl i ne	No Faul ts
9	CFC6	I NACTI VE UP-UP-Onl i ne	No Faul ts
10	not provi si oned	-	-
11	not provi si oned	-	-
12	CFC6	ACTI VE UP-UP-Onl i ne	No Faul ts
13	CFC6	ACTI VE UP-UP-Onl i ne	No Faul ts
14	not provi si oned	-	-
15	not provi si oned	-	-
16	not provi si oned	-	-
17	not provi si oned	-	-
18	not provi si oned	-	-
19	not provi si oned	-	-
20	not provi si oned	-	-
21	not provi si oned	-	-
FAN	FAN8	UP-DOWN-Not I nstal l ed	Fan Modul e Not Present

officer SEC> **ENABLE CARD=8 I NACTCFC**

Info (030013): CFC6 Slot 08,09 is already enabled

18.5.2 Logs

Logs associated with dry contact alarm definitions are:

- **SHLF003:** modification of a dry contact alarm definition
- **SHLF004:** creation of a dry contact alarm definition
- **SHLF005:** destruction of a dry contact alarm definition

Logs associated with dry contact alarm conditions:

- **SHLF008:** dry contact alarm raised
- **SHLF009:** dry contact alarm cleared

18.6 SYSTEM COOLING

The system is engineered to operate in a specified range of environmental temperatures. The **Installation Guide** provides information and details about the temperature range.

The **SHOW SYSTEM COOLING** command provides the user with information regarding the current temperature inside the chassis of the system.

Refer to section 3 for more information on the SHOW SYSTEM command.

TABLE 18-2 SHOW SYSTEM COOLING command

Noun	Verb	Syntax	Description
SYSTEM COOLING	SHOW	SHOW SYSTEM COOLING	The SHOW SYSTEM COOLING command displays various information about shelf temperature and fan conditions.

18.7 Audits

Internal maintenance audits are automatically executed by the system. Periodic audits run on an internally set schedule. System audits search for cards and ports with interrupted maintenance activity or failures and take action accordingly. The action taken depends on whether the audit is performed periodically or after a swap of activity. If a periodic audit runs, action is only taken on the second detection of an interrupted maintenance action or failure.

Following is the schedule for system audits:

1. Periodic.
 - Card presence audit – every 60 minutes.
 - Card or port state audit - every 5 minutes and will take action on the card if its in a stable state for 2 audit intervals or 10 minutes.
 - Card or port defect audit – every 15 minutes.
 - Shelf audit – every 15 minutes.
 - Load file audit - runs every 24 hours.
 - Database audit - runs every 60 minutes.
 - Compares RAM to RAM in both CFCs in Duplex
 - Compares RAM to FLASH in each CFC
2. Post-restart.
 - Immediately runs:
 - Card presence audit.
 - After 30 minutes, runs:

- Load file audit.

TABLE 18-3 Audit Types

Audit Type	Action
Card Presence Audit	Checks and corrects the following situations: <ol style="list-style-type: none"> 1. Card is inserted into the slot, but the software is not aware of its presence 2. Card is not present in the slot, but software incorrectly registers the card as being inserted into the slot.
Shelf Audit	The following attributes are refreshed in software (updated from hardware) <ul style="list-style-type: none"> - Management MAC address - Inband MAC address The following attributes are updated in the hardware (updated from software): <ul style="list-style-type: none"> - Domain name and Host name - Lamp test status - LED settings for cards, system and FAP
Card/port defect audit	Periodically polls the defect status of the hardware and compares it with software status. Any mismatch is handled by adjusting the software value to the hardware value. This audit could result in alarm raising or clearing and accompanying changes to card and port status.
Load file audit	Runs a CRC calculation on system load files (those with <i>.tar</i> extension) in flash and compares it with the CRC value in the file header. If a mismatch is detected, it generates a log and an alarm against the CFC24.

18.8 System Recovery

The user should be aware of the fact that the fMAP system has been designed to automatically recover from certain conditions and situations. For example, if the system is placed under a very heavy traffic load and critical system tasks are jeopardized, it will automatically swap activity in an attempt to recover from the overload and sustain required services. If certain critical system tasks are not allowed to run to completion in their allotted time period, due to the above mentioned network congestion, the system will execute a restart to attempt to recover and re-establish service.

18.9 TRACEROUTE

18.9.1 Overview

The TRACEROUTE command is used to trace all intermediate nodes (or hops) that a packet traverses from the source IP address up to just before it can reach the user specified destination IP address. It displays the DNS resolvable names and IP addresses of these routers' (hops) IP address and the round-trip times to each of these intermediate nodes.

TRACEROUTE can be used as a diagnostic and troubleshooting tool to examine traffic flows or to determine network bottle-necks (places where traffic gets lost or delayed).

18.9.2 Using TRACEROUTE

TABLE 18-4 TRACEROUTE commands

Noun	Verb	Syntax	Description
TRACEROUTE	TRACEROUTE	<pre>TRACEROUTE={ ipAddress hostName } [MINTTL=number] [MAXTTL=number] [TIMEOUT=seconds] [TOS=0..255] [NORESOLVE]</pre>	<p>Displays all intermediate nodes (hops) from source to destination along with round-trip times to each of these nodes. The output also includes the destination.</p> <p>This command expects a valid IP address or a DNS resolvable destination name. All other parameters are independently optional and defaults will be used unless they are explicitly specified by the user.</p> <p>MINTTL - is the minimum Time To Live parameter, MAXTTL - is the maximum Time to Live (default max is 30).</p> <p>TIMEOUT - refers to the maximum allowed time in seconds to receive an ICMP response from a node (3 seconds is the default maximum).</p> <p>TOS - is the type of service field in the IP header. 3 bits are reserved for it. This can be a value from 0..7</p> <p>This command can be executed only by users with a SECURITYOFFICER privilege level or higher.</p> <p>TRACEROUTE can only be executed one user at a time.</p>
TRACEROUTE	STOP	STOP TRACEROUTE	Aborts the current TRACEROUTE command that is running on the device.

18.10 Call Debugging

18.10.1 Address Resolution Protocol (ARP)

Once the IP provisioning for a POTS24 card is completed and the card is ONLINE, the user can display all or individual Address Resolution Protocol (ARP) entries using the **SHOW IP ARP** command. The output displays

three fields per entry, the IP address, MAC address, and the reference count that indicates the number of times that ARP entry was accessed. All or selected entries in the ARP table can be cleared by using the **CLEAR IP ARP** command.

18.10.2 Ping

Once the POTS24 card is in service, the PING command can be used to verify connectivity between the POTS24 card interface and other hosts in the same network. The PING command initiates the process of sending ICMP echo packets from the specified POTS24 card virtual Ethernet interface to the specified host and waits for a response. The user is presented with a summary of the number of packets sent and received along with an indication of the percentage of packets lost.

For example, to initiate ping to IP address 1.1.2.2 from a POTS24 card that has virtual Ethernet interface vlan:123.0 and IP address 1.2.3.4, the user would enter either:

```
PING 1. 1. 2. 2 FROM INTERFACE=VLAN: 123. 0
```

or

```
PING 1. 1. 2. 2 FROM IPADDRESS=1. 2. 3. 4
```

In the event that a user wishes to end a repetitive PING request, the STOP PING command terminates ping operation and presents information regarding the number of packets sent and received.

18.10.3 Unsupported commands

The TRACEROUTE and SHOW IP ROUTE commands are not supported for the POTS24 card.

18.10.4 Call State

When a POTS24 port is in the UP-UP-ONLINE or UP-UP-DEGRADED state, the SHOW PORT command displays the current call processing state of the port. The call processing status is represented by the following components:

- **Hook Status:**
 - **On:** the line is not involved with an active call
 - **Off:** the line is involved with an active call
- **Connection Status**
 - **RxOnly:** receive only from the network to the subscriber
 - **TxOnly:** transmit only from the subscriber to the network
 - **RxTx:** both receive and transmit
- **Active Codec**
 - **None**
 - **G711**

- [G726_32](#)

18.10.5 VOICECALL Trace Logs

Detailed call events pertaining to a port or set of ports can be captured by defining and enabling trace log criteria on the system. The `ADD TRACE VOICECALL` command is used to define trace criteria and has the following parameters:

- **EVENT:** The call event, from among the following:
 - **OPENLOOP:** the subscriber loop has transitioned to open state (on hook)
 - **CLOSELOOP:** the subscriber loop has transitioned to closed state (off hook)
 - **MGCPOFFHOOK:** an off hook signal has been received from the Call Agent
 - **MGCPONHOOK:** an on hook signal has been received from the Call Agent
 - **MODEMDETECT:** modem tone has been detected
 - **ALL:** indicates all of the above call events
- **INTERFACE:** a port or range of ports, or **ALL** for all POTS24 ports on all POTS24 cards

Note that the trace criteria is defined by the explicit combination of event and interface. For example, suppose the user enters the following:

```
ADD TRACE VOICECALL EVENT=OPENLOOP I NTERFACE=0.0
ADD TRACE VOICECALL EVENT=CLOSELOOP I NTERFACE=0.0
```

This creates 2 explicit trace criteria instances. Now suppose the user tries to create a third trace criteria as follows: `TRACE`

```
ADD TRACE VOICECALL EVENT=ALL I NTERFACE=0.0
```

This last command is accepted as a third trace criteria instance, even though it represents a duplicate trace criteria for 2 of the events on the same interface.

The `DELETE TRACE VOICECALL` command is used to remove `TRACE` criteria. If parameters are not specified, **ALL** is assumed. Continuing with the example above, suppose the user enters:

```
DELETE TRACE VOICECALL EVENT=ALL I NTERFACE=0.0
```

This command deletes the trace criteria instance, but the 2 existing trace criteria instances will still exist.

The `SHOW TRACE VOICECALL` command is used to display the status of the trace criteria instances that have been defined. However, this command does not display the actual logs in the trace buffer. Here is an example of the output:

Interface	Event
ALL	MODEMDETECT
7.0	CLOSELOOP
7.1	ALL
7.3	MGCPOFFHOOK
7.3	OPENLOOP
7.4	ALL
7.23	OPENLOOP

Note: The *ENABLE TRACE* command is used to enable logging of defined traces. Once traces are added and enabled, when call events occur that match the criteria, a trace log is created and stored in the trace log buffer. The *SHOW TRACE* command is used to display the logs in the buffer, and the *CLEAR TRACE* command is used to clear the buffer. See the *User Event Logging* section (18.13) for more details.

TABLE 18-5 VOICECALL commands

Noun	Verb	Syntax	Description
VOICECALL	ADD	<pre> ADD TRACE VOICECALL [EVENT={ OPENLOOP CLOSELOOP MGCPPOFFHOOK MGCPONHOOK MODEMDETECT ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	Used to define VOICECALL trace criteria.
VOICECALL	DELETE	<pre> DELETE TRACE VOICECALL [EVENT={ OPENLOOP CLOSELOOP MGCPPOFFHOOK MGCPONHOOK MODEMDETECT ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	Used to remove TRACE criteria.

TABLE 18-5 VOICECALL commands (Continued)

Noun	Verb	Syntax	Description
VOICECALL	SHOW	<pre> SHOW TRACE VOICECALL [EVENT={ OPENLOOP CLOSELOOP MGCPOFFHOOK MGCPONHOOK MODEMDETECT ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	Used to display the logs in the buffer.
TRACE	ENABLE	<pre> ENABLE TRACE [OUTPUT={ CLI } [FORMAT={ FULL SUMMARY }]] </pre>	Enables logging of defined traces.
TRACE	SHOW	<pre> SHOW TRACE [{ STATUS BUFFER [DATE=[op] yyyy-mm-dd [-yyyy-mm-dd]] [FORMAT={ FULL SUMMARY }] [REVERSE] [SEQUENCE=seqnum [-seqnum]] [TAIL [=count]] [TIME=[op] hh:mm:ss [-hh:mm:ss]] }] </pre>	Displays the logs in the buffer.
TRACE	CLEAR	<pre> CLEAR TRACE [FORCE] </pre>	Used to clear the trace buffer.

18.11 IGMP Trace

18.11.1 Overview

The fMAP provides the user with the ability to display and view IGMP parameters. From the CLI, the user can perform the following functions in real time:

- Display reports
- Display queries
- Display events (timers, router ports, cleanup, etc.)

The user can also configure an output filter on the displays for:

- IP address (unicast client, or MCAST address)
- MAC address (unicast client, or MCAST)
- Display known MCAST routers

See section, **Configuring Management Interfaces** for information on configuring output filters.

18.11.2 Command Examples

Examples of the commands follow:

```
SHOW IGMP Snooping
```

Display IGMP information on the system.

```
SHOW IGMP Snooping Card 4.4 Full
```

Display IGMP information for a specific card.

```
SHOW IGMP Snooping Counter Card 4
```

Display IGMP counters for a specific card.

```
SHOW IGMP Snooping Counter Interface=4.4
```

Display IGMP counters for a specific interface.

```
SHOW IGMP Snooping Counter MessageResponse
```

Displays the IGMP snooping 'message response' counters and statistics. This data can help in determining channel change response time.

```
SHOW IGMP Snooping Counter Standard
```

Displays the IGMP snooping standard counters and statistics.

```
SHOW IGMP Snooping Interface=4.4 Full
```

Display complete IGMP information for a specific interface.

```
SHOW IGMP SNOOPING MCASTGROUPS FULL
```

Displays the current multicast group(s) that are subscribed on the system.

Examples and explanations of the counters follow. The first example displays IGMP counters for a specified card.

```
officer SEC> SHOW IGMP COUNTER CARD 5
--- IGMP Card-Level Message Counters -----
Card   Message Type   Good Count   Error Count
-----
5      Report         230          0
      Leave         110          15
      General Query 0            0
      OSPF          0            0
      DVMRP         0            0
      PIMV1         0            0
      PIMV2         0            0
-----
Info (010017): Operation Successful
```

This example displays IGMP counters for a specified interface.

```
officer SEC> SHOW IGMP COUNTER INTERFACE 5.4
--- IGMP Interface-Level Message Counters -----
Interface Message Type   Good Count   Error Count
-----
5.4       Report         0            0
      Leave         0            0
      General Query 0            0
      OSPF          0            0
      DVMRP         0            0
      PIMV1         0            0
      PIMV2         0            0
-----
Info (010017): Operation Successful
```

An example of displaying IGMP MESSAGERESPONSE follows. Here is an explanation of the fields:

- Counter - Eight counters (buckets) are supported.
- Response Range - Each counter tracks an individual response range from 0-4+ seconds.
- Last Updated - Timestamp indicating the last time the count was incremented.
- Message Count - The number of messages that have experienced <Response Range> delay across IGMP message queue.

```
officer SEC> SHOW IGMP COUNTER MESSAGERESPONSE
--- IGMP System-Level Message Response Counters -----
Counter Response Range   Last Updated   Message Count
-----
0      0-249 (msec)      2004-04-06 06:08:56   38412
1      250-499 (msec)    2004-04-05 19:25:03    38
2      500-749 (msec)    2004-04-05 19:24:59    0
3      750-1000 (msec)   2004-04-05 19:24:59    0
4      1-2 (sec)         2004-04-05 19:24:59    0
5      2-3 (sec)         2004-04-05 19:24:59    0
6      3-4 (sec)         2004-04-05 19:24:59    0
7      4+ (sec)          2004-04-05 19:24:59    0
-----
Info (010017): Operation Successful
```

An example of displaying IGMP COUNTER STANDARD follows. Here is an explanation of the fields:

- Message Type - Seven message types currently processed by IGMP.
- Good Count - Number of <Message Type> messages that were processed to completion.

- Error Count - Number of <Message Type> messages that were not processed to completion.

```

officer SEC> SHOW IGMP COUNTER STANDARD
--- IGMP System-Level Message Counters ---
Message Type      Good Count      Error Count
-----
Report            230              0
Leave              110              15
General Query     6                0
OSPF              1                0
DVMRP             0                0
PIMV1             0                0
PIMV2            0                0
Info (010017): Operation Successful

```

Information can be displayed about IGMP joins, leaves, and events, etc. using the SHOW IGMP SNOOPING TRACE command.

Note: *The ENABLE TRACE command is used to enable logging of defined traces. Once traces are added and enabled, when call events occur that match the criteria, a trace log is created and stored in the trace log buffer. The SHOW TRACE command is used to display the logs in the buffer, and the CLEAR TRACE command is used to clear the buffer. See the User Event Logging section (18.13) for more details.*

TABLE 18-6 IGMP TRACE commands

Noun	Verb	Syntax	Description
TRACE IGMPS- NOOPING	ADD	<pre> ADD TRACE IGMPSNOOPING MESSAGESTYPE={ REPORTV1 REPORTV2 LEAVE GENERALQUERY LASTMEMBERQUERY ALL } [INTERFACE={ type:id-range id-range ifname-list ALL }] [MACADDRESS={ macaddress ALL }] [GROUPADDRESS={ ipaddress ALL }] </pre>	Adds an IGMPSNOOPING event trace.

TABLE 18-6 IGMP TRACE commands (Continued)

Noun	Verb	Syntax	Description
TRACE IGMPS- NOOPING	DELETE	<pre>DELETE TRACE IGMPSNOOPING [MESSAGEATYPE={ REPORTV1 REPORTV2 LEAVE GENERALQUERY LASTMEMBERQUERY ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] [MACADDRESS={ macaddress ALL }] [GROUPADDRESS={ ipaddress ALL }]</pre>	Deletes an IGMPSNOOPING event trace.
TRACE IGMPS- NOOPING	SHOW	<pre>SHOW TRACE IGMPSNOOPING [MESSAGEATYPE={ REPORTV1 REPORTV2 LEAVE GENERALQUERY LASTMEMBERQUERY ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] [MACADDRESS={ macaddress ALL }] [GROUPADDRESS={ ipaddress ALL }]</pre>	The SHOW IGMPSNOOPING TRACE command is used to display joins, leaves, queries, and reports for the IGMP snooping feature.

18.12 EPSR Trace

18.12.1 Overview

The fMAP provides the user with the ability to display, in real time, the following EPSR events.

- Hello packet generation
- Hello packet reception
- Failover Time Expiry
- Failover Time Reset
- LINK_UP/LINK_DOWN message generation
- LINK_UP/LINK_DOWN message reception
- RING_DOWN Flush generation
- RING_DOWN Flush reception and corresponding flush action
- RING_UP Flush generation
- RING_UP Flush reception and corresponding flush action
- Secondary port unblocking
- Port into pre forwarding state

See section, **Configuring Management Interfaces** for information on configuring output filters.

18.12.2 Command Examples

Examples of the commands follow:

Note: *The **ENABLE TRACE** command is used to enable logging of defined traces. Once traces are added and enabled, when call events occur that match the criteria, a trace log is created and stored in the trace log buffer. The **SHOW TRACE** command is used to display the logs in the buffer, and the **CLEAR TRACE** command is used to clear the buffer. See the User Event Logging section ([18.13](#)) for more details.*

TABLE 18-7 IGMP TRACE commands

Noun	Verb	Syntax	Description
TRACE EPSR	ADD	<pre> ADD TRACE EPSR [={ epsrdomain-list ALL }] MESSAGETYPE={ HEALTH RINGUPFLUSH RINGDOWNFLUSH LINKDOWN ALL } [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	Adds an EPSR event trace.
TRACE EPSR	DELETE	<pre> DELETE TRACE EPSR [={ epsrdomain-list ALL }] [MESSAGETYPE={ HEALTH RINGUPFLUSH RINGDOWNFLUSH LINKDOWN ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	Deletes an EPSR event trace.

TABLE 18-7 IGMP TRACE commands (Continued)

Noun	Verb	Syntax	Description
TRACE EPSR	SHOW	<pre> SHOW TRACE EPSR [={ epsrdomain-list ALL }] [MESSAGEATYPE={ HEALTH RINGUPFLUSH RINGDOWNFLUSH LINKDOWN ALL }] [INTERFACE={ type:id-range id-range ifname-list ALL }] </pre>	The SHOW EPSR TRACE command is used to display EPSR events.

18.13 User Event Logging

18.13.1 Overview

Event Logs are logs that are enabled and disabled by the user in order to trace a series of events. The user enables Event Logging for a specific application and is able to control what triggering events will generate logs. This can be useful for user level debugging. The user also has the ability to control whether or not Event Logs are sent to the CLI or just to the log buffer.

18.13.2 Usage Notes

Event Log settings do not persist over restarts.

18.13.3 Overview of Commands

Event Logs are stored in a buffer. All application logs share a common buffer. The size of this buffer can be set by the user.

The order of commands to set up Event Logging is:

1. (Optional) User sets the log buffer size to store the logs. There is only one buffer for all logs. This buffer always exists, but can be resized by the user. Resizing the buffer clears it of all current logs.
2. User sets up any application specific settings that control what triggering events will generate logs.

3. User enables application event logging. Event Logs now go into the buffer and to the CLI (optional).

Each application has unique arguments for controlling which of its logs go into the buffer. For example, call processing (CALLP) might have an argument TRIGGER for actions such as ONHOOK and OFFHOOK, as well as which slot and ports to trace events. For instance, the user could log only ONHOOK events on slot 12 port 5.

The user has the ability to suspend and resume logging (using DISABLE LOG and ENABLE LOG commands) without clearing the buffer. The buffer is only cleared with the CLEAR LOG command or when the Event Log buffer is resized using the SET LOG BUFFERSIZE command).

Note: *The user should take into consideration that event logging or tracing could affect system performance during heavy network traffic periods.*

18.13.4 Event Logging commands

18.13.4.1 Changing buffer size

```
SET TRACE [ BUFFERSIZE=events [ FORCE ] ]
```

This command sets general options for all logs of a given type. The only available option now is the Log Event buffersize. Resizing the buffer purges all current logs and creates a new buffer for storing logs. This will also generate an Event Log so that the user can see that the logging buffer was cleared at this point.

18.13.4.2 Enabling Event Logging

```
ENABLE TRACE [ OUTPUT={ CLI } [ FORMAT={ FULL | SUMMARY } ] ]
```

This allows Event Logs to be stored (in Event Log buffer) and to be displayed to the CLI (optional). This will also generate an Event Log so that the user can see that logging was enabled at this point.

This command can be called more than once. For example, the user could input ENABLE TRACE to turn on tracing, then decide that the CLI needs to be enabled as well and input ENABLE TRACE OUTPUT CLI.

18.13.4.3 Disabling Event Logging

```
DISABLE TRACE
```

This stops Event Logs for this application from being stored or displayed. Existing Event Logs in the buffer are NOT cleared out, but new logs will not be added. This will also generate an Event Log so that the user can see that logging was disabled at this point.

18.13.4.4 Clearing Event Logs

```
CLEAR TRACE [ FORCE ]
```

This clears the logs from the buffer. If logging is enabled new logs can still enter the buffer. A log will be generated saying that the logs were cleared.

18.13.4.5 Resetting Event Logging

```
SETDEFAULTS TRACE [ FORCE ]
```

This resets all the Event Logging settings to their default state. Logging is disabled and the buffer is reset to its original size (and cleared). This will also generate an Event Log so that the user can see that logging was reset at this point.

18.13.4.6 Viewing Event Logs

```
SHOW TRACE [ { STATUS | BUFFER [ DATE=[ op ] yyyy-mm-dd [ -yyyy-mm-dd ] ] [ FORMAT={ FULL  
| SUMMARY } ] [ REVERSE ] [ SEQUENCE=seqnum [ -seqnum ] ] [ TAIL [ =count ] ] [ TIME=[ op ]  
hh:mm:ss [ -hh:mm:ss ] ] } ]
```

This shows the Event Logging status (STATUS) or the selected logs in the buffer (BUFFER). Filters can be used to determine what type of logs to display.

The FORMAT option controls how much data is displayed to the CLI. Unlike MGMT logs, which use MSGONLY to show the whole multiple line message and SUMMARY to show just the header, Event Logs will use SUMMARY to show a reduced info header and only the first line of application data.

FULL log:

```
CALLP Slot:12 2004-03-11 16:48:12.480 0400
```

```
OFFHOOK detected
```

```
Location: Port: 3
```

SUMMARY log:

```
CALLP 16:48:12.480 OFFHOOK detected
```

```
Displaying the buffer:
```

```
offl cer SEC> SHOW TRACE BUFFER
```

```
CALLP Slot:12 2004-03-11 16:49:33.500 0401
```

```
ONHOOK detected
```

```
Location: Port: 3
```

```
CALLP Slot:12 2004-03-11 16:48:12.480 0400
```

```
OFFHOOK detected
```

```
Location: Port: 3
```

```
CALLP Slot:12 2004-03-11 16:45:00.280 0399
```

```
ONHOOK detected
```

```
Location: Port: 3
```

CALLP Slot:12 2004-03-11 16:44:12.000 0388

ONHOOK detected

Location: Port: 4

CALLP 2004-03-11 16:43:33.900 0387

ONHOOK detected

Location: Slot: 12 Port: 5

18.13.4.7 Offloading Event Logs

```
PUT LOG FILE={ destinationfile | unit:destinationfile | serverpath/destinationfile } [ { TFTP SERVER={ ipaddress | hostname } | ZMODEM | FTP SERVER={ ipaddress | hostname } USER=userid PASSWORD=password } ] [ TYPE={ MGMT | ERROR | TRACE | CRASH } ] [ CARD={ ACTCFC | INACTCFC } ]
```

Offloads the Event logs to a remote file using one of the supported transfer methods.

A sample session follows:

```
officer SEC>> SET TRACE
This will clear Trace logs. Continue (Y/N)? Y
Info (010017): Operation Successful
officer SEC>> ADD TRACE VOICECALL EVENT ALL
Info (010017): Operation Successful
officer SEC>> ENABLE TRACE OUTPUT=CLI
Info (010017): Operation Successful
officer SEC>> SHOW TRACE STATUS
--- Trace Settings -----
```

Trace System is ENABLED

CLI Streaming is ENABLED

CALLP Call Traces for POTS Interfaces
Interface: ALL Event: ALL

```
-----
Info (010017): Operation Successful
```

The telephone connected to a POTS24 card, port 19.0 went OFF, then ON HOOK several times. The resulting traces are displayed here:

```
officer SEC>> CALLP 2004-07-23 08:13:41.000 slot(19) 1
Off-hook detected on port 0

CALLP 2004-07-23 08:13:45.000 slot(19) 2
On-hook detected on port 0

CALLP 2004-07-23 08:13:47.000 slot(19) 3
Off-hook detected on port 0

CALLP 2004-07-23 08:13:50.000 slot(19) 4
On-hook detected on port 0

CALLP 2004-07-23 08:13:58.000 slot(19) 5
Off-hook detected on port 0

CALLP 2004-07-23 08:14:17.000 slot(19) 6
On-hook detected on port 0
```

```
CALLP 2004-07-23 08:14:19.000 slot(19) 7
Off-hook detected on port 0
```

```
CALLP 2004-07-23 08:14:21.000 slot(19) 8
On-hook detected on port 0
```

```
officer SEC>> DISABLE TRACE
Info (010017): Operation Successful
```

Now, a trace was set on interface 19.0 of the POTS24 card.

```
officer SEC>> ADD TRACE VOICECALL INTERFACE=19.0
```

```
Info (010017): Operation Successful
```

```
officer SEC>> SHOW TRACE VOICECALL
```

Interface	Event
19.0	ALL

```
officer SEC>> ENABLE TRACE OUTPUT=CLI
```

```
Info (010017): Operation Successful
```

```
officer SEC>> SHOW TRACE STATUS
```

```
--- Trace Settings ---
```

```
Trace System is ENABLED
```

```
CLI Streaming is ENABLED
```

```
CALLP Call Traces for POTS Interfaces
```

```
Interface: 19.0 Event: ALL
```

```
Info (010017): Operation Successful
```

The telephone connected to POTS24 card, interface 19.0 went OFF, then ON HOOK several times. The resulting traces are displayed here:

```
officer SEC>> CALLP 2004-07-23 08:19:33.000 slot(19) 18
Off-hook detected on port 0
```

```
CALLP 2004-07-23 08:19:35.000 slot(19) 19
On-hook detected on port 0
```

```
CALLP 2004-07-23 08:19:36.000 slot(19) 20
Off-hook detected on port 0
```

```
CALLP 2004-07-23 08:19:38.000 slot(19) 21
On-hook detected on port 0
```

```
officer SEC>> SHOW TRACE BUFFER
```

```
CALLP 2004-07-23 08:19:38.000 slot(19) 21
On-hook detected on port 0
```

```
CALLP 2004-07-23 08:19:36.000 slot(19) 20
Off-hook detected on port 0
```

```
CALLP 2004-07-23 08:19:35.000 slot(19) 19
On-hook detected on port 0
```

```
CALLP 2004-07-23 08:19:33.000 slot(19) 18
Off-hook detected on port 0
```

```
TRACE_SYS 2004-07-23 08:17:58.867 slot(13) 17
```

```
Trace Buffer resized (new size 200 logs)
```

```
Info (010017): Operation Successful
```

TABLE 18-8 Event Logging (TRACE) commands

Noun	Verb	Syntax	Description
TRACE	SET	SET TRACE [BUFFERSIZE=events [FORCE]]	This command sets general options for all logs of a given type. The only available option now is the Log Event buffer-size. Resizing the buffer purges all current logs and creates a new buffer for storing logs. This will also generate an Event Log so that the user can see that the logging buffer was cleared at this point. The maximum size of the buffer is 10,000 logs. The default is 200 logs.
TRACE	ENABLE	ENABLE TRACE [OUTPUT={ CLI } [FORMAT={ FULL SUMMARY }]]	This allows Event Logs to be stored (in Event Log buffer) and to be displayed to the CLI (optional). This will also generate an Event Log so that the user can see that logging was enabled at this point.
TRACE	DISABLE	DISABLE TRACE	This stops Event Logs for this application from being stored or displayed. Existing Event Logs in the buffer are NOT cleared out, but new logs will not be added. This will also generate an Event Log so that the user can see that logging was disabled at this point.
TRACE	CLEAR	CLEAR TRACE [FORCE]	This clears ALL logs from the buffer. If logging is enabled new logs can still enter the buffer. A log will be generated saying that the logs were cleared.

TABLE 18-8 Event Logging (TRACE) commands (Continued)

Noun	Verb	Syntax	Description
TRACE	SETDEFAULTS	SETDEFAULTS TRACE [FORCE]	This resets all the Event Logging settings to their default state. Logging is disabled and the buffer is reset to its original size (and cleared). This will also generate an Event Log so that the user can see that logging was reset at this point.
TRACE	SHOW	SHOW TRACE [{ STATUS BUFFER [DATE=[op] yyyy-mm-dd [-yyyy-mm-dd]] [FORMAT={ FULL SUMMARY }] [REVERSE] [SEQUENCE=seqnum [-seqnum]] [TAIL [=count]] [TIME=[op] hh:mm:ss [-hh:mm:ss] }]	This shows the Event Logging status (STATUS) or the selected logs in the buffer (BUFFER). Filters can be used to determine what type of logs to display. The FORMAT option controls how much data is displayed to the CLI. Unlike MGMT logs (which use MSGONLY to show the whole multiple line message and SUMMARY to show just the header) Event Logs will use SUMMARY to show a reduced info header and only the first line of application data.

18.14 CES Troubleshooting

Note: Many of the concepts covered here are the same for the NTE8. However, the information will be complete in both subsections since users may have only one of the configurations.

18.14.1 Overview

The fMAP provides the user with certain tools that can be used to troubleshoot the CES. These tools are discussed here.

18.14.2 Card and fMAP Chassis Support

The CES8 service module is equipped with PULL, FAULT, and INSRV LEDs. They behave consistently with existing service modules.

18.14.3 Concepts and Terms

- Loopback - A configuration of the framer hardware where the data sent in one direction is looped back in the direction from where it was received. This can be done at the path level or the line level. Loopback is typically used to diagnose or troubleshoot problems with an interface.
- AIS - Alarm Indication Signal. The DS1 AIS is an unframed all ones signal.
- RAI - Remote Alarm Indication. When a DS1 sink detects Loss of Signal (LOS) or it receives a signal for which framing cannot be found (e.g. AIS), it forces a zero into the second bit of each channel of the 24 channel structure. RAI is a framed signal. In release 5.0, RAI is passed through to the distant end.

18.14.4 Loopback

There are 3 types of loopback, Line Loopback, Payload Loopback, and Inward Loopback

Note: Payload loopback is not supported in release 7.0.

- Line Loopback - The signal received from the DS1/E1 port is returned back out the DS1/E1 port, consisting of the full 1.544 Mbps/2.048 Mbps signal with:
 - Bit sequence integrity maintained
 - No change in framing
 - No removal of bipolar violations
- Payload Loopback - The signal from the DS1/E1 port is returned back out the DS1/E1 port, consisting only of the DS0 payload. Not supported in release 7.0.
- Inward Loopback - The data that was destined for a particular DS1/E1 port is looped back toward the network. (Also known as CSU loopback)




FIGURE 18-2 Loopbacks

In release 7.0, loopbacks may be initiated only by Management interfaces (CLI/SNMP). Loopbacks set using the Management interfaces are persistent. The persistence is maintained by a loopback configuration attribute, which is independent of the administrative state of the DS1/E1 port. The DS1/E1 port is considered operationally DOWN when a loopback is configured, because no “thru service” can be provided. Therefore, when a loopback is configured on the interface, an administratively UP interface would be “UP-DOWN-Loopback”, while an administratively DOWN interface would be “DOWN-DOWN-Loopback”.

Defect monitoring is disabled on the interface while a loopback is present.

When a connected interface is placed in the line loopback or payload loopback state, the interworking function will send AIS out the packet side.

When a connected interface is placed in the internal loopback, the system will send AIS out the DS1/E1 port side.

An interface in loopback mode may be used as a timing reference. The timing signal is passed through even when the loopback is in place.

In release 7.0, packet-level loopback on the pseudo-span side is not supported. The internal loopback on a connected DS1/E1 port effectively loops the pseudo-span data back up the pseudo-span.

In release 7.0, signal injection within the system to be looped back is not supported (e.g. by external equipment).

18.14.5 Port State Management

When a physical port is in the DISABLE state (i.e. administratively down), the signal will be turned off on the physical interface. The port LEDs reflect the current state of the interface as follows:

TABLE 18-9 Port LED Meanings

Green PORT LED (INS)	Orange PORT LED (ERR)	Meaning
OFF	ON	Port provisioned in service, but is in a LOS failed state
ON	Blink @ .5Hz	Port provisioned in service, but degraded, (e.g. CV threshold crossed)
ON	OFF	Port provisioned in service, in sync, no errors (normal quiescent state)
Blink @ .5Hz	Blink @ .5Hz	Port is in diagnostic state, Loopback is set
OFF	OFF	Port is provisioned but out of service administratively

18.14.6 Fault Management

Fault management is supported on the DS1/E1 port up to the Line level. Generally, line level faults (such as coding violations) are caught at the DS1/E1 port interface, but path level faults (such as Loss of Frame) are passed through the Pseudo-span to the emulated circuit.

Due to the number of subscribers that could be affected by a DS1/E1 port fault, failure of a DS1/E1 port will be a **Major** alarm (unlike other SM ports, which normally have minor or info severity). The DS1/E1 port still counts as 1 port against the port outage threshold when the CES8 card is in a failed condition (similar to ports on other cards).

The following alarms can be raised on the CES8 card:

- Card not present and all the common alarms currently supported. However, note that the CES8 card alarm will itself be **Major**.
- Configuration Failure

The following alarms can be raised on the DS1/E1 port entity:

- Configuration failure (Severity=Major, failed)
- LOS (Severity=Major, failed)
- Degraded timing reference (when the interface is timed off the internal oscillator and there is a *better* timing reference provisioned) (Severity=Minor, degraded)

The following “run-time status attributes” about the DS1/E1 port entity. These are qualifiers concerning the current conditions on the interface that do not affect the administrative or operational state of the interface. Changes in these run-time attributes do not result in any alarm indications, traps, management logs, etc. However, the values can be displayed at the CLI (or via the NMS) to support remote fault isolation:

- AIS received

The following alarms can be raised on the pseudo-span entity:

- Configuration failure (Severity=Major, failed)
- Loss of Packet Stream (LOPS) (Severity=Major, failed)
- Remote LOPS (Severity=Major, degraded)
- Degraded timing reference (when the interface is timed off the internal oscillator and there is a “better” timing reference provisioned) (Severity=Minor, degraded)
- Failed timing reference (when the interface is timed off the internal oscillator and the internal oscillator has failed) (Severity=Major, failed)

The following “run-time status attributes” about the Pseudo-span entity:

- Remote Carrier Fault
- Local Carrier Fault

18.14.6.1 Faults from the DS1/E1 port

Refer to [Figure 18-3](#). When LOS is detected on the DS1/E1 port, the local interworking function (IWF) will immediately set Remote Carrier Fault for packets going out the pseudo-span where that DS1/E1 port is connected. Packet transmission continues at the regular period, because the clocking is not interrupted, although it may have switched to the internal oscillator. An AIS will be indicated in the outgoing packet. An LOS defect sustained throughout a soak period will result in an LOS failure and an alarm generated on the port.




FIGURE 18-3 Data Flow during LOS on DS1/E1 Port

Referring to [Figure 18-4](#), when the system detects AIS on the DS1/E1 port, the local interworking function does nothing to the packets going out the pseudo-span where that DS1/E1 port is connected. Thus, AIS is sent through the emulated circuit to the far end. If the user queries the state of the DS1/E1 port, for example at the Network Management System (NMS), it will show a run-time attribute indicating that the system is receiving AIS. There is no indication on the pseudo-span. Note that the system will not return RAI back if it receive AIS. This is a path-level signal which is the responsibility of the remote DS1 sink.

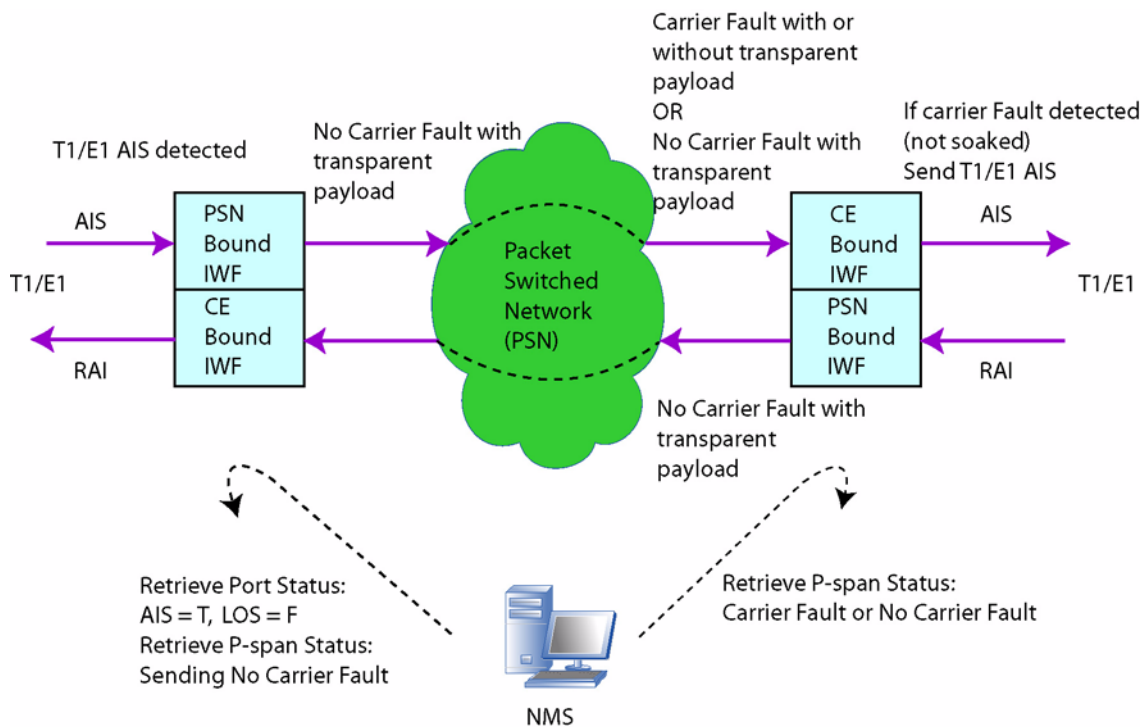


FIGURE 18-4 Data Flow During AIS on DS1/E1 Port

18.14.6.2 Transient Defects from the Pseudo-span Side

Referring to [Figure 18-4](#), when the system receives packets from the pseudo-span (because the remote end has LOS on their DS1/E1 port), the interworking function will output the “all ones” data pattern rather than any payload in the packet.

When the system misses a packet (e.g., the packet was lost or delayed past its useful life), the system interworking function will output the “all ones” data pattern and the appropriate counter is incremented.

18.14.6.3 Failures from the Pseudo-span Side

Referring to [Figure 18-5](#), the only failure that can occur on the pseudo-span itself is the LOPS failure. The LOPS fault, as defined by the SAToP standard, is not supported in release 5.0. LOPS faults are detected by counting the number of packets “lost” on each pseudo-span. A packet is considered “lost” if it was not there when it was time to play it out. This may be because the packet was lost in transit, arrived too late, or was corrupted along the way, etc. If the count of packets “lost” on a particular pseudo-span in a second is > 10 packets, then a LOPS defect is declared on that pseudo-span. At that time, the R bit is set in packets that are sent on that pseudo-span (i.e., in the reverse direction). If the LOPS defect is sustained for 2.5 seconds, we declare a LOPS failure and post an alarm on the pseudo-span entity. If the system detects a second with 0 packets “lost”, the

LOPS defect is cleared. At that time, the R bit is cleared in the reverse direction. If the absence of LOPS defect is sustained for 10 seconds, the LOPS failure is cleared.

Posting and clearing the LOPS failure results in Management Log entries and traps to the NMS. The user can also display the alarms (or absence of alarm) using the SHOW ALARM and SHOW PSPAN commands.

If the local interworking function detects a LOPS (Loss Of Packet Stream) defect in the downstream direction (i.e. loss of packets received from the pseudo-span), the local interworking function will:

- Output the “all ones” data pattern to the DS1/E1 port in place of all lost packets.
- Set the “R” bit in the packets sent upstream (i.e., into the pseudo-span).

Note that this is analogous to the RDI indication on a DS1.

When the system receives packets from the pseudo-span that have the “R” bit set (e.g. because the remote end has LOPS on their end of the pseudo-span), the system soaks the indication for 2.5 seconds. If the system is still receiving packets with the “R” bit set after the 2.5 seconds, it will post a “Remote LOPS” condition (a “degraded” alarm, not a “failed” alarm) on the pseudo-span. After the system stops receiving packets with the “R” bit set for 10 seconds, it clears the “Remote LOPS” condition.

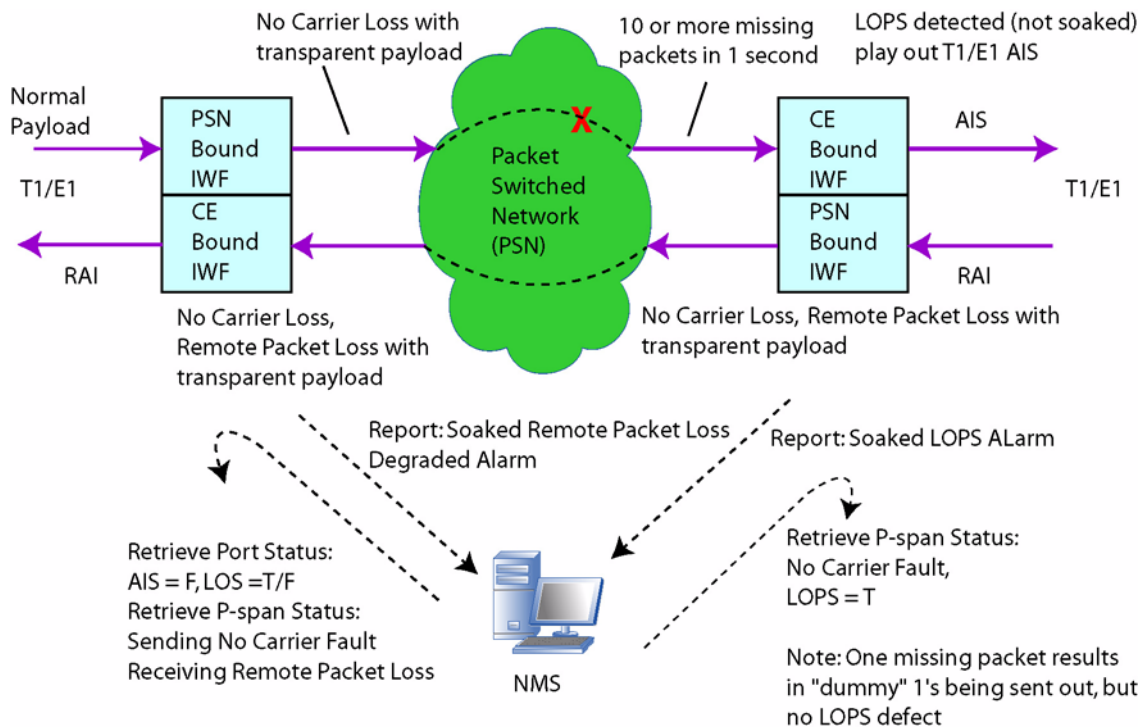


FIGURE 18-5 Data Flow during a Uni-Directional packet loss condition

18.14.7 CSE Troubleshooting Examples

This subsection identifies possible trouble scenarios that might occur in the field and explains how to use specific tools to isolate and resolve the situation.

18.14.7.1 Mis-configuration of the Pseudo-span

1. **Incorrect Destination UDP Port** - At both the local and remote ends, the user will see the counter for packets received for an invalid pseudo-span incrementing (because the system validates on both destination and source UDP address) and the target pseudo-span will be in the LOPS state.
2. **Incorrect Destination IP address** - At the local end, the user will see the counter for packets received for an invalid pseudo-span incrementing (because the system validates on both destination and source IP address) and the target pseudo-span will be in the LOPS state. At the remote end, the user will see the received packets count on the pseudo-span not incrementing, and it will be in the LOPS state. On the CES8 card who's IP address the user entered, they should see the counter for packets received for an invalid pseudo-span incrementing (because the system validates on both destination and source IP address).

18.14.8 CES/CES8 IP Debugging

Existing CLI commands, such as TRACEROUTE, PING, ARP table management, etc., can be used to perform IP debugging on the CES8.

18.15 EPON2 Troubleshooting

The EPON2 configuration is integrated into the overall maintenance system, as described in the beginning of this section. The general rules for the components are as follows:

- Alarms on the EPON2 card - The EPON2 card alarm is Major due to the number of subscribers that are potentially impacted.
- EPON interface - The alarms will have MAJOR severity due to the number of potential subscribers on the physical port. When the interface is disabled, the alarms will be cleared, as is the case with other interfaces in the system.
- ONU alarms - These have a MINOR severity due to the fact that it is just one subscriber. When the interface is disabled, the alarms will be cleared, as is the case with other interfaces in the system.
- ONU Interface:
 - Local faults - These are reported by the OLT about the ONU.
 - Remote faults - reported by the ONU about the ONU.

Refer to the Log Manual for a description of the log and alarm messages that are specific for the EPON2 and ONU.

18.16 Technical Support Scripts

18.16.1 Overview for 7.0

When there is a field problem that needs a detailed analysis, there is in release 7.0 a mechanism that allows the customer to run scripts; these scripts produce a screen output that is captured by the customer and sent to Technical Support. Technical Support and Software Design can then analyze the output to further diagnose the problem.

These scripts have the following attributes for release 7.0:

- The files are encrypted.
- They are in the /loads directory, being placed there using the GET FILE command. (In getting and using this feature, the customer uses the set of commands for retrieving, renaming, deleting, and placing files.)
- They are not stored in Compact Flash.
- They are run using the command `SHOW TECHSUPPORT FILE=<script filename> [FORCE]`.
- Output from the script execution is displayed directly to the user's session. The user must capture the output and place it where Allied Telesis Technical Support can retrieve the file.
- Only one script may be executed at one time.

When the 'SHOW TECHSUPPORT' is invoked and the FORCE option is not supplied, the user is presented with the following prompt:

```
This command may affect customer service. Are you sure you want to proceed (Y/N)?
```

The customer will work with Allied Telesis Technical Support to use this feature.

19. Specifications / Standards

19.1 Overview

[Table 19-1](#) lists the RFCs supported by the fMAP products and the specific objects that are used for this release.

Note: References to RFCs adhere to the following copyright convention.

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

19.2 MIB Objects Supported

TABLE 19-1 MIB Objects Supported for this release

RFC	Object	Compliance (Complete, or if not complete, what is supported)
RFC1213	snmp	Complete
	system	Complete
	ip	ipForwarding
	ipAddrTable	ipAdEntAddr ipAdEntNetMask
RFC1493 (Bridge MIB)	dot1dBase	Complete
	dot1dTpPortTable	dot1dBasePort dot1dBasePortIfIndex dot1dBasePortCircuit
	dot1dStp	Complete
	dot1dStpPortTable	Complete
	dot1dTp	Complete
RFC1450 (SNMPv2 MIB)	coldStart Trap	Complete
	authenticationFailure Trap	Complete
RFC1514 (Host-Resources MIB)	hrSystem	hrSystemUptime hrSystemDate hrSystemNumUsers

TABLE 19-1 MIB Objects Supported for this release (Continued)

RFC	Object	Compliance (Complete, or if not complete, what is supported)
RFC1573 (IF MIB)	interfaces	Complete
	ifTable	ifIndex ifDescr ifType ifPhysAddress ifAdminStatus ifOperStatus ifLastChange ifInOctets ifInUcastPkts ifInDiscards ifInErrors ifInUnknownProtocol ifOutOctets ifOutUcastPkts ifOutDiscards ifOutErrors
	linkDown Trap	Complete
	linkUp Trap	Complete
RFC2037 (Entity MIB)	entPhysicalTable	Complete
	entLogicalTable	Complete
	entLPMappingTable	Complete
	entPhysicalContainsTable	Complete
	entLastChangeTime	Complete
	entConfigChange Trap	Complete
RFC2515 (ATM MIB)	atmInterfaceTCTable	atmInterfaceTCAlarmState
	atmVclTable	atmVclVpi atmVclVci atmVclAdminStatus atmVclOperStatus atmVccAalType atmVccAal5EncapsType

MIB Objects Supported

TABLE 19-1 MIB Objects Supported for this release (Continued)

RFC	Object	Compliance (Complete, or if not complete, what is supported)
RFC2662 (ADSL MIB)	adslLineTable	Complete
	adslAtucPhysTable	adslAtucInvSerialNumber adslAtucCurrSnrMgn adslAtucCurrAtn adslAtucCurrStatus adslAtucCurrAttainableRate
	adslAtucPerfDataTable	Complete
	adslAturPerfDataTable	Complete
	adslAtucIntervalTable	Complete
	adslAturIntervalTable	Complete
RFC2662 (ADSL MIB) - Cont.	adslLineConfProfileTable	adslLineConfProfileName adslAtucConfRateMode adslAtucConfTargetSnrMgn adslAtucChanConfFastMinTxRate adslAtucChanConfInterleaveM- inTxRate adslAtucChanConfFastMaxTxRate adslAtucChanConfInterleaveMax- TxRate adslAtucChanConfMaxInterleaveDelay adslAturConfRateMode
	adslLineAlarmConfProfileTable	Complete
	adslAtucPerfLofsThreshTrap	Complete
	adslAtucPerfLossThreshTrap	Complete
	adslAtucPerfLprsThreshTrap	Complete
	adslAtucPerfESsThreshTrap	Complete
	adslAtucPerfLolsThreshTrap	Complete
	adslAtucInitFailureTrap	Complete
	adslAturPerfLofsThreshTrap	Complete
	adslAturPerfLossThreshTrap	Complete
	adslAturPerfLprsThreshTrap	Complete
	adslAturPerfESsThreshTrap	Complete

TABLE 19-1 MIB Objects Supported for this release (Continued)

RFC	Object	Compliance (Complete, or if not complete, what is supported)
RFC2674 (VLAN MIB)	dot1qBase	dot1qVlanVersionNumber dot1qMaxVlanId dot1qMaxSupportedVlans dot1qNumVlans dot1qGvrpStatus
	dot1qVlan	dot1qNextFreeLocalVlanIndex dot1qVLANNumDeletes
	dot1qVlanCurrentTable	Complete
	dot1qVlanStaticTable	Complete
	dot1qPortVlanTable	dot1qPvid dot1qPortAcceptableFrameTypes dot1qPortIngresFiltering dot1qPortGvrpStatus
RFC2819 (RMON MIB)	etherStatsTable	Complete
	historyControlTable	Complete
	etherHistoryTable	Complete
	alarmTable	Complete
	eventTable	Complete
	risingAlarm Trap	Complete
	fallingAlarm Trap	Complete
RFC3273 (HC RMON MIB)	etherStatsHighCapacityTable	Complete
RFC3440 (ADSL-LINE EXT MIB)	adslAtucPerfDataExtTable	Complete
	adslAturPerfDataExtTable	Complete
	adslConfProfileExtTable	Complete
	adslAlarmConfProfileExtTable	Complete

*

TABLE 19-2 fMAP Tables

Table	Object	Notes
snmpCacheTable	snmpUseCache snmpCacheAging snmpCacheSize	
ATNLagMibTable	lagID lagName lagAdminState lagOperState lagDataRate lagPortList	
AtnMgcpStatsEntry (indexed table)	mgcpSentMessages	Unsigned32, Read Only
	mgcpRcvdMessages	Unsigned32, Read Only
	mgcpLostMessages	Unsigned32, Read Only
	mgcpCmdsRetransmitted	Unsigned32, Read Only
	mgcpRcvdBadVersionMessages	Unsigned32, Read Only
	mgcpUnrecognizedMessages	Unsigned32, Read Only
AtnRtpStatsTable (indexed by AtnRtpStatsEntry)	rtpLocalIpAddr	IpAddress
	rtpLocalPort	INTEGER (0..65535)
	rtpRemoteIpAddr	IpAddress
	rtpRemotePort	INTEGER (0..65535)
	rtpSentPacketCount	Unsigned32, Read Only
	rtpSentOctetCount	Unsigned32, Read Only
	rtpRecvPacketCount	Unsigned32, Read Only
	rtpRecvOctetCount	Unsigned32, Read Only
	rtpExtHighSeqNumRecv	Unsigned32, Read Only
	rtpFractionLost	Unsigned32, Read Only
	rtpPacketLostCount	Unsigned32, Read Only
	rtpLatePacketCount	Unsigned32, Read Only
	rtpInterarrivalJitter	Unsigned32, Read Only
rtpSentTimestamp	Unsigned32, Read Only	

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnPortFaultTable (indexed table)	atnPortLossofLinkCounter	Read Only
	atnPortLossofSignalCounter	Read Only
	atnPortLossofFrameCounter	Read Only Does not apply to GE1
	atnPortPeerNotPresentCounter	Read Only Does not apply to GE1
	atnPortFaultResetFlag	Read Write Does not apply to GE1 When set to T, all counters go to 0 and Flag set back to F
atnAdslCountersEntry (indexed table) for QOS	adslEgressQueue0DropPkts	Unsigned32, Read Only
	adslEgressQueue0SentPkts	Unsigned32, Read Only
	adslEgressQueue1DropPkts	Unsigned32, Read Only
	adslEgressQueue1SentPkts	Unsigned32, Read Only
	adslEgressQueue2DropPkts	Unsigned32, Read Only
	adslEgressQueue2SentPkts	Unsigned32, Read Only
	adslEgressQueue3DropPkts	Unsigned32, Read Only
	adslEgressQueue3SentPkts	Unsigned32, Read Only
adslQosResetFlag	AtnBoolean, Read Write (When set to T, all counters go to 0 and Flag set back to F)	

MIB Objects Supported

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnCXECounterObjs (scaler table)	atnInQOverFlow	Read Only
	atnALFIFOOverFlow	Read Only
	atnBufMemOverFlow	Read Only
	atnALFIFOOverFlow	Read Only
	atnBufMemOverFlow	Read Only
	atnRliDiscard	Read Only
	atnRleDiscard	Read Only
	atnMCBCLimitDiscard	Read Only
	atnTTLScoping	Read Only
	atnWFHBD	Read Only
	atnMACError	Read Only
	atnCXECountersObjsResetFlag	Read Write (When set to T, all counters go to 0 and Flag set back to F)
atnCardTable (indexed by AtnCardEntry)	atnCardAdminState atnCardOperState atnCardAvailabilityStatus atnCardProceduralStatus atnCardStandbyStatus atnCardType atnCardPorts atnCardPrefLoad atnCardTempLoad atnCardFlashLoad atnCardRunningLoad atnCardSpecial atnCardPort atnCardRowStatus	
atnShdslCardTable (indexed by atnShdslCardEntry)	atnShdslCardAnnex atnShdslCardWettingCurrent atnShdslCardWireMode	
atnCesCardTable (indexed by AtnCesCardEntry)	atnCesCardPortType atnCesCardTimingReference atnCesCardTimingReferenceInterface	

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnCfcCardTable (indexed by atnCfcCardEntry)	atnCfcCardAltLoad	
atnGePortTable (indexed by AtnGePortEntry)	atnGePortAdminState atnGePortOperState atnGePortAvailStatus atnGePortAutoNegotiate atnGePortFlowControl atnGePortDescription atnGePortDirection	

MIB Objects Supported

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnAdslPortTable (indexed by atnAdslPortEntry)	atnAdslPortAdminState atnAdslPortOperState atnAdslPortAvailStatus atnAdslPortMode atnAdslPortLineType atnAdslPortInterleaveDelay atnAdslPortEC atnAdslPortMaxUpstreamRate atnAdslPortMinUpstreamRate atnAdslPortMaxDownstreamRate atnAdslPortMinDownstreamRate atnAdslPortTargetSNRMargin atnAdslPortLineQualityMonitor atnAdslPortVpi atnAdslPortVci atnAdslPortDescription atnAdslPortConnectionState atnAdslPortActualLineStandard atnAdslPortActualLineType atnAdslPortTrellisCoding atnAdslPortAtucConnectRate atnAdslPortAtucMaxAttainableRate atnAdslPortAtucSNR atnAdslPortAtucAttenuation atnAdslPortAtucOutputPower atnAdslPortAturConnectRate atnAdslPortAturMaxAttainableRate atnAdslPortAturSNR atnAdslPortAturAttenuation atnAdslPortAturOutputPower atnAdslActualSoftwareAnnex atnAdslPortDirection	

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnFePortTable (indexed by atnFePortEntry)	atnFePortAdminState atnFePortOperState atnFePortAvailStatus atnFePortSpeed atnFePortAutoNegotiate atnFePortDuplex atnFePortFlowControl atnFePortDescription atnFePortActualSpeed atnFePortActualDuplex atnFePortActualFlowControl atnFePortDirection	
atnPotsPortTable (indexed by atnPotsPortEntry)	atnPotsPortAdminState atnPotsPortOperState atnPotsPortAvailStatus atnPotsPortCapability atnPotsPortMinPacketization atnPotsPortMaxPacketization atnPotsPortBufferDelay atnPotsPortBufferMode atnPotsPortEchoCancellation atnPotsPortVoiceActivityDetection atnPotsPortComfortNoiseGeneration atnPotsPortPacketLossConcealment atnPotsPortTransmitPowerGain atnPotsPortReceivePowerGain atnPotsPortLineSupervision atnPotsPortBufferLength atnPotsPortHookStatus atnPotsPortConnectionStatus atnPotsPortActiveCodec atnPotsPortDescription atnPotsPortDirection	

MIB Objects Supported

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnShdslPortTable (indexed by atnShdslPortEntry)	atnShdslPortAdminState atnShdslPortOperState atnShdslPortAvailStatus atnShdslPortMaxConnectRate atnShdslPortMinConnectRate atnShdslPortTargetSNRMargin atnShdslPortLineQualityMonitor atnShdslPortVpi atnShdslPortVci atnShdslPortDescription atnShdslPortConnectionState atnShdslPortActualConnectRate atnShdslPortStucSNR atnShdslPortStucAttenuation atnShdslPortStucOutputPower atnShdslPortSturSNR atnShdslPortSturAttenuation atnShdslPortSturOutputPower atnShdslPortReceiverGain atnShdslPortPSDMask atnShdslPortTipRingConfiguration atnShdslPortWirepairSwapped atnShdslPortDirection	
atnCesPortTable (indexed by atnCesPortEntry)	atnCesPortAdminState atnCesPortOperState atnCesPortAvailStatus atnCesPortLineBuildOut atnCesPortLineEncoding atnCesPortLoopbackState atnCesPortLineType atnCesPortStuffByte atnCesPortTimingReference atnCesPortDescription atnCesPortDirection	

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnFxPortTable (indexed by atnFxPortEntry)	atnFxPortAdminState atnFxPortOperState atnFxPortAvailStatus atnFxPortSpeed atnFxPortAutoNegotiate atnFxPortDuplex atnFxPortFlowControl atnFxPortDescription atnFxPortActualSpeed atnFxPortActualDuplex atnFxPortActualFlowControl atnFxPortDirection	
atnCardAlarmTable (indexed by AtnCardAlarmEntry)	atnCardDefect atnCardAlarmRaisedTime atnCardAlarmMasked atnCardAlarmSeverity	
atnPortAlarmTable (indexed by AtnPortAlarmEntry)	atnPortDefect atnPortAlarmRaisedTime atnPortAlarmMasked atnPortAlarmSeverity	

MIB Objects Supported

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnPSpanPerfDataTable (indexed by AtnPSpanPerfDataEntry)	(Day Counts)	
	atnPSpanPerfES	
	atnPSpanPerfLOPSS	
	atnPSpanPerfLatePackets	
	atnPSpanPerfEarlyPackets	
	atnPSpanPerfLostPackets	
	atnPSpanPerfPacketsReceived	
	atnPSpanPerfBytesReceived	
	atnPSpanPerfPacketsSent	
	atnPSpanPerfBytesSent	
	atnPSpanPerfInterarrivalJitterMin	
	atnPSpanPerfInterarrivalJitterMax	
	atnPSpanPerfInterarrivalJitterAvg	
	(15 min Counts)	
	atnPSpanPerfCurr15MinES	
	atnPSpanPerfCurr15MinLOPSS	
	atnPSpanPerfCurr15MinLatePackets	
atnPSpanPerfCurr15MinEarlyPackets		
atnPSpanPerfCurr15MinLostPackets		
atnPSpanPerfCurr15MinInterarrivalJitterMin		
atnPSpanPerfCurr15MinInterarrivalJitterMax		
atnPSpanPerfCurr15MinInterarrivalJitterAvg		
atnPSpanPerfValidIntervals	Integer	
atnPSpanPerfInvalidIntervals	Integer	
atnPSpanThresholdTable (indexed by AtnPSpanThresholdEntry)	atnPSpanCurr15MinThreshES	Integer
	atnPSpanCurr15MinThreshLOPSS	Integer
	atnPSpanCurr15MinThreshLatePackets	Integer
	atnPSpanCurr15MinThreshEarlyPackets	Integer
	atnPSpanCurr15MinThreshLostPackets	Integer

TABLE 19-2 fMAP Tables (Continued)

Table	Object	Notes
atnPSpanIntervalTable (indexed by AtnPSpanIntervalEntry)	atnPSpanIntervalNumber atnPSpanIntervalES atnPSpanIntervalLOPSS atnPSpanIntervalLatePackets atnPSpanIntervalEarlyPackets atnPSpanIntervalLostPackets atnPSpanIntervalInterarrivalJitterMin atnPSpanIntervalInterarrivalJitterMax atnPSpanIntervalInterarrivalJitterAvg	
atnDsITrapCurrentThresholdTable (indexed by AtnDsITrapCurrentThresholdEntry)	atnDsITrapCurrentThresholdLineIndex atnDsITrapCurrentThresholdESs atnDsITrapCurrentThresholdSESs atnDsITrapCurrentThresholdSEFSs atnDsITrapCurrentThresholdUASs atnDsITrapCurrentThresholdCSSs atnDsITrapCurrentThresholdPCVs atnDsITrapCurrentThresholdLESs atnDsITrapCurrentThresholdLCVs atnDsITrapCurrentThresholdLOSSs	
atnDsITrapFarEndCurrentThresholdTable (indexed by AtnDsITrapFarEndCurrentThresholdEntry)	atnDsITrapFarEndCurrentThresholdLineIndex atnDsITrapFarEndCurrentThresholdESs atnDsITrapFarEndCurrentThresholdSESs atnDsITrapFarEndCurrentThresholdSEFSs atnDsITrapFarEndCurrentThresholdUASs atnDsITrapFarEndCurrentThresholdCSSs atnDsITrapFarEndCurrentThresholdPCVs atnDsITrapFarEndCurrentThresholdLESs atnDsITrapFarEndCurrentThresholdLCVs	
atnDsILOSSCurrentCountTable (indexed by AtnDsILOSSCurrentCountEntry)	atnDsICurrentLOSs	
atnDsILOSSDayCountTable (indexed by AtnDsILOSSDayCountEntry)	atnDsIDayLOSs	

19.3 Physical Standards

The following table lists the standards supported by the fMAP devices, and includes supporting documentation.

19.4 Protocol / Software Standards

TABLE 19-3

Protocol / Standard	Reference
ARP	RFCs 826, 395
DHCP	RFC 1700
IEEE 802.2	
IEEE 802.3	
IGMP	