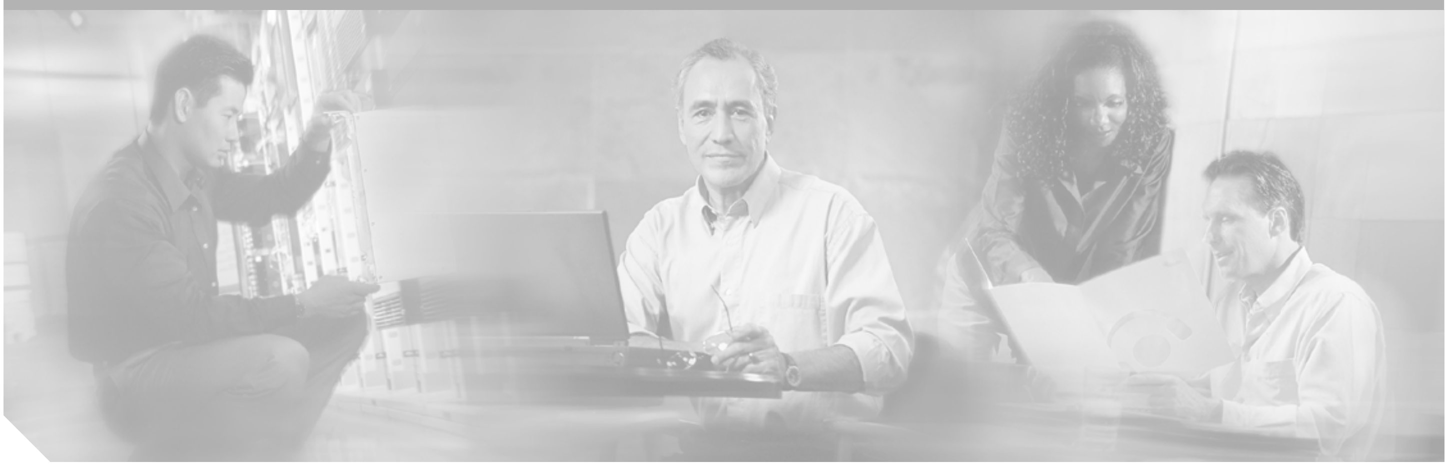

Cisco Gigabit-Ethernet Optimized VoD Solution Design and Implementation Guide, Release 1.1

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number:
Text Part Number: OL-4921-01


CONTENTS

Preface

This preface explains the objectives, intended audience, and organization of the *Cisco Gigabit-Ethernet Optimized VoD Solution Design and Implementation Guide*. The section also defines the conventions used to convey instructions and information, available related documentation, and the process for obtaining Cisco documentation and technical assistance.

This preface presents the following major topics:

Document Version and Solution Release

This is the first version of this document, which covers Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution.

Document History

Document Version	Date	Notes
1	10/30/2003	This document was first released. Release 1.0 documentation was released on 07/30/2003. The current document (Release 1.1), though substantially similar to the previous version, discusses support for nCube video servers and the Cisco uMG9820 QAM Gateway.

Document Objectives and Scope

This guide describes the architecture, the components, and the processes necessary for the design and implementation of the Cisco Gigabit-Ethernet Optimized VoD Solution.

This guide presents the fundamental design and configuration information that is required to establish the various services provided by the Cisco Gigabit-Ethernet Optimized VoD Solution. MSO (multiple system operator) and service provider networks may have additional requirements that are beyond the scope of this document.

Note This document is primarily for Cisco products. To establish and maintain the third-party products and applications that may be a part of the Cisco Gigabit-Ethernet Optimized VoD Solution, refer to the documentation provided by the vendors of those products.

Audience

The target audience for this document is assumed to have basic knowledge of and experience with the installation and acceptance of the products covered by this solution. See Chapter 1, "Solution Overview."

In addition, it is assumed that the user understands the procedures required to upgrade and troubleshoot optical transport systems and Ethernet switches, with emphasis on Cisco Catalyst 4500 series switches).

Note This document addresses Cisco components only. It does not discuss how to implement third-party optical components, VoD servers, or QAM devices, or how to enable service between QAM devices and HFC distribution.

Document Organization

The major sections of this document are as follows:

Section	Title	Major Topics
Chapter 1	Solution Overview	Introduces applications, example scenarios, and components.
Chapter 2	Designing the Solution	Provides detailed requirements of various scenarios.
Chapter 3	Implementing and Configuring the Solution	Describes the configuration and implementation of the solution and provides example implementations.
Chapter 4	Providing Redundancy and Reliability	Describes failure scenarios and their remedies.
Chapter 5	Monitoring and Troubleshooting	Provides remedies to possible types of failure and loss of quality of service.

Chapter 6	Deploying the Cisco Gigabit-Ethernet Optimized VoD Solution in Fiber Ring Topologies	Compares ring vs. Gigabit Ethernet topologies, and describes how to convert existing fiber ring topologies to the hub-and-spoke architecture required by the solution.
Appendix A	Sample Switch Configurations for Symmetric 10-GE Topology	Provides headend and Dhub switch configurations for single EtherChannel, multiple EtherChannels, and subtended control path.
Appendix B	No Switch in Dhub: Sample Configurations for Cisco Catalyst 4500 Series Switches	Provides headend and Dhub switch configurations for single VLAN, split VLAN, and multiple VLANs.
Appendix C	Switch in Dhub: Sample Redundancy Configurations for Cisco Catalyst 4500 Series Switches	Provides active and standby headend switch configurations, as well as Dhub configurations, for HSRP without UDLR, HSRP with UDLR, HSRP with interface tracking, and Layer 3 1+1 protection.

Related Documentation

Solution Documentation

This document, and *Release Notes for Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1*, are available under Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1, at the following URL:

http://www.cisco.com/univercd/cc/td/doc/solution/vodsols/geopt1_1/index.htm

Switch Documentation

Documentation resources for the Cisco Catalyst 4500 series switches are available at the following URLs:

http://www.cisco.com/en/US/products/hw/switches/ps4324/index.html

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/index.htm

For the Cisco IOS software used in this release, refer to the following:

- Switch Documentation, Cisco IOS Software Release 12.1(19)EW

  http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_19/index.htm

- Cisco IOS Software Release 12.1(19)EW for Sup II-Plus, III, IV

  http://www.cisco.com/en/US/products/hw/modules/ps2797/prod_bulletin09186a008019741d.html

- Release Notes for the Catalyst 4500 Series Switch, Cisco IOS Release 12.1(19)EW

  www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/relnotes/ol_2170.htm

Optical Component Documentation

Cisco ONS 15216 DWDM Filters

- *Cisco ONS 15216 FlexLayer User Guide*
  www.cisco.com/univercd/cc/td/doc/product/ong/15216/flxlyr10/index.htm

Cisco DWDM GBICs

- *Cisco DWDM Gigabit Interface Converter Installation Guide*
  www.cisco.com/univercd/cc/td/doc/product/gbic_sfp/gbic_doc/78_15574.htm
- *Cisco Dense Wavelength Division Multiplexing GBICs Compatibility Matrix*
  www.cisco.com/univercd/cc/td/doc/product/gbic_sfp/gbic_doc/ol_4604.htm

QAM Gateway Documentation

- Cisco uMG9820 QAM Gateway
  http://www.cisco.com/univercd/cc/td/doc/product/cable/vod/umg9820/index.htm

Note Other references are provided as appropriate throughout this document.

Document Conventions

Command descriptions use the following conventions:

boldface font	Commands and keywords are in **boldface**.
italic font	Arguments for which you supply values are in *italics*.
[ ]	Elements in square brackets are optional.
{ x | y | z }	Alternate keywords are grouped in braces and separated by vertical bars.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Screen examples use the following conventions:

`screen` font	Terminal sessions and information the system displays are in `screen` font.
`boldface screen` font	Information you must enter is in **`boldface screen`** font.[1]
`italic screen` font	Arguments for which you supply values are in *`italic screen`* font.

→	This pointer highlights an important line of text in an example.
^	The symbol ^ represents the key labeled Control. For example, the key combination ^D in a screen display means hold down the Control key while you press the D key.
< >	Nonprinting characters, such as passwords, are in angle brackets in contexts where italic font is not available.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

1. As this document makes use of annotated configurations, the rigorous use of boldface type to indicate what the user must enter is relaxed.

Notes use the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the publication.

Timesavers use the following conventions:

Timesaver This symbol means *the described action saves time*. You can save time by performing the action described in the paragraph.

Cautions use the following conventions:

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Tips use the following conventions:

Tip This symbol means the following information *will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information, similar to a Timesaver.

Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

http://www.cisco.com/univercd/home/home.htm

You can access the Cisco website at this URL:

http://www.cisco.com

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpck/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:

  http://www.cisco.com/en/US/partner/ordering/index.shtml

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit e-mail comments about technical documentation to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, the Cisco Technical Assistance Center (TAC) provides 24-hour-a-day, award-winning technical support services, online and over the phone. Cisco.com features the Cisco TAC website as an online starting point for technical assistance. If you do not hold a valid Cisco service contract, please contact your reseller.

Cisco TAC Website

The Cisco TAC website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The Cisco TAC website is available 24 hours a day, 365 days a year. The Cisco TAC website is located at this URL:

http://www.cisco.com/tac

Accessing all the tools on the Cisco TAC website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a login ID or password, register at this URL:

http://tools.cisco.com/RPF/register/register.do

Opening a TAC Case

Using the online TAC Case Open Tool is the fastest way to open P3 and P4 cases. (P3 and P4 cases are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Case Open Tool automatically recommends resources for an immediate solution. If your issue is not resolved using the recommended resources, your case will be assigned to a Cisco TAC engineer. The online TAC Case Open Tool is located at this URL:

http://www.cisco.com/tac/caseopen

For P1 or P2 cases (P1 and P2 cases are those in which your production network is down or severely degraded) or if you do not have Internet access, contact Cisco TAC by telephone. Cisco TAC engineers are assigned immediately to P1 and P2 cases to help keep your business operations running smoothly.

To open a case by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)
EMEA: +32 2 704 55 55
USA: 1 800 553-2447

For a complete listing of Cisco TAC contacts, go to this URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

TAC Case Priority Definitions

To ensure that all cases are reported in a standard format, Cisco has established case priority definitions.

Priority 1 (P1)—Your network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Priority 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Priority 3 (P3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Priority 4 (P4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Go to this URL to visit the company store:

  http://www.cisco.com/go/marketplace/

- The Cisco *Product Catalog* describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the Cisco Product Catalog at this URL:

  http://cisco.com/univercd/cc/td/doc/pcat/

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

  http://www.ciscopress.com

- *Packet* magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access Packet magazine at this URL:

  http://www.cisco.com/packet

- *iQ Magazine* is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:

  http://www.cisco.com/go/iqmagazine

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

  http://www.cisco.com/ipj

- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:

  http://www.cisco.com/en/US/learning/index.html

Solution Overview

This chapter presents the following major topics:

Solution Description

The Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1 enables cable operators and MSOs (multiple system operators) to offer Video on Demand (VoD) services to consumer customers over their existing hybrid fiber coax (HFC), with existing next-generation digital set-top boxes. The solution leverages a Gigabit Ethernet (GE) transport network between the headend, where the video servers reside, to the distribution hub (Dhub), where the HFC network terminates. Cisco switches are used in the headend and, optionally, in the Dhub. This solution uses Cisco Catalyst 4500 series switches, and its high-level architecture is depicted in Figure 1-1 on page 1-3.

MPEG-2 video is (1) unicast from the video server to the switch; (2) transported from the switch through Gigabit Ethernet interface converters (GBICs; not shown) over an optical Ethernet network between Cisco ONS 15216 Metro FlexLayer DWDM 100-GHz terminal filters at each end; (3) delivered to GE quadrature amplitude modulation (QAM) devices (often referred to simply as QAMs); (4) and then modulated onto the HFC plant for viewing by the authorized user, through a set-top box (STB). See also Architecture and Scope, page 1-2. This solution uses both third-party QAM devices and the Cisco uMG9820 QAM Gateway.

Note For documentation on the Cisco uMG9820 QAM gateway, see Cisco uMG9820 QAM Gateway at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/cable/vod/umg9820/index.htm

The intelligent switched infrastructure provided in this solution reduces the overall complexity and operational costs of the network, while permitting scalability for future growth.

Key Applications

Some key applications for this solution include the following:

- Movies on Demand (pay-per-view model)

  In this application, both new and older releases are offered by the MSO for on-demand viewing.

- Subscription Video on Demand

  This application provides on-demand viewing of content from premium cable stations. Available content is managed by the premium network, not the MSO. The MSOs store the content locally on their VoD servers, but the network chooses what content is stored there, and how often it is refreshed. The content consists of movies, as well as original programming from the network.

- Video on Demand—Specific Content

  This application uses content specifically developed for Video on Demand. This content is typically targeted at specific demographics, such as enthusiasts or consumers with a specific area of interest.

- Network Personal Video Recorder (Time-Shifted TV)

  For this application, traditional linear network programming is put into the on-demand application. Consumers can view programming when they want to view it. This application can be the highest-volume form of VoD.

Architecture and Scope

The Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1, consists of Cisco transport and edge components as identified in Figure 1-1.

Note References in this document to third-party components such as VoD servers, QAM devices, HFC plant, or set-top boxes are solely for the purpose of illustrating an overall VoD architecture. They are not intended to imply support of those components by Cisco. Although some third-party components have been tested for interoperability, those components will be supported by their respective vendors.

Throughout this document, “QAM device” refers to both third-party devices and the Cisco uMG9820 QAM Gateway.

For a listing of third-party components and vendors referenced in this solution, see the following:

- Video Servers, page 1-8
- Gigabit Ethernet QAM Devices, page 1-8

Figure 1-1 High-Level Architecture of the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1

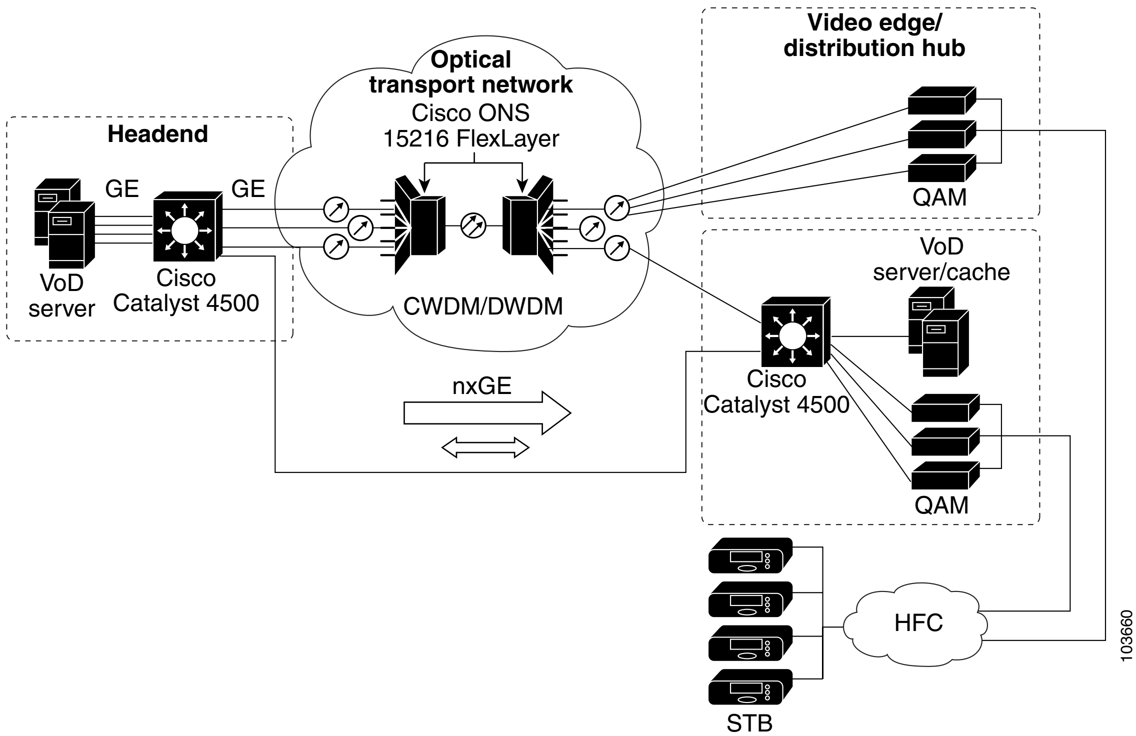

Note the following:

- The above illustrates GE feeds to Cisco uMG9820 QAM gateways, which are capable of using the full capacity of the feed. Where third-party devices cannot support that bandwidth, use daisy chaining where possible to distribute the bandwidth to multiple QAM gateways. See QAM Device Connectivity, page 2-8.
- Depending on the design of the network, the feeds to the Dhubs can use either 10 GE or DWDM optical networking. DWDM is used where distances are greater than those supported by 10 GE modules.

Basic Solution Scenarios

All Ethernet topologies in this release of the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1 include an Ethernet switch/router in the video headend. There are two main Ethernet topology choices for video transport. The topology choices depend on whether or not the Dhub has Ethernet switching capability.

The following sections introduce two solution scenarios, and provide a high-level comparison between the two:

- Switch in Dhub
- No Switch in Dhub
- Scenario Feature Comparison

Switch in Dhub

This basic scenario uses a switch at the distribution hub (Dhub), with either a dedicated GE bidirectional wavelength for provisioning and control information, or an arbitrary IP return path that already exists. Figure 1-2 and Figure 1-3 show optical and IP topologies, respectively, of the switch-in-Dhub scenario with a dedicated GE return path. The optical topologies are variations of the physical connections between the switches. The configurations on the switches do not change.

Figure 1-2 Switch in Dhub: Optical Topology—Dedicated GE Return Path

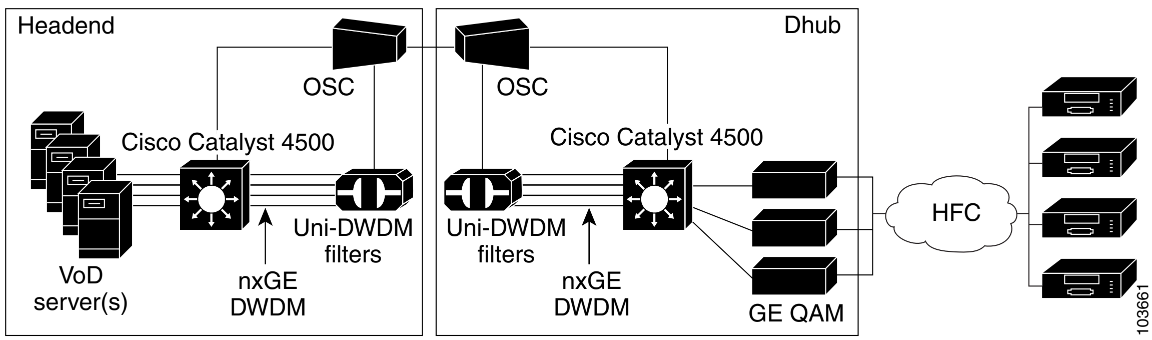

Figure 1-3 Switch in Dhub: IP Topology—Dedicated GE return path

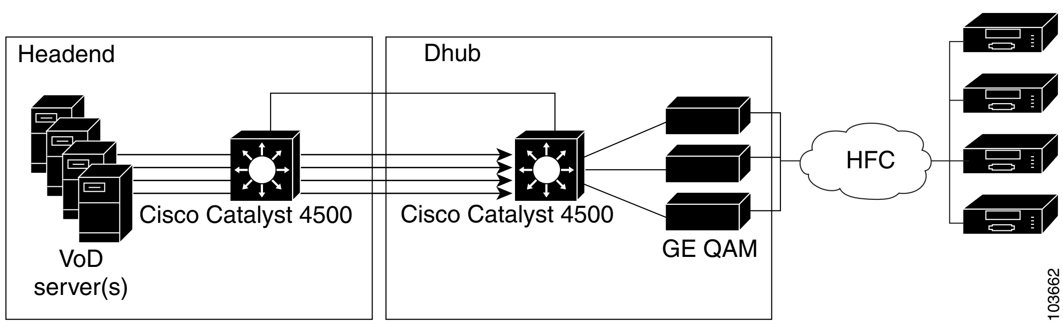

No Switch in Dhub

The other basic scenario does not include a switch at the Dhub. In this case, the GE QAM device has some basic Ethernet switching capability to aggregate responses from multiple, daisy-chained GE QAM devices over a single return path. In the no-switch-in-Dhub scenario, the GE QAM devices can be connected directly to the headend switch, either individually or in groups of three. Since each of the Harmonic NSG QAM devices used in this release of the solution handles about 300 Mbps worth of video data, two NSG QAM devices can be daisy-chained off another NSG QAM device that is connected directly to the headend switch. The master QAM device will forward any traffic not addressed to itself to other members of the daisy-chain. Figure 1-4 and Figure 1-5 show optical and Ethernet topologies, respectively, of the no-switch-in-Dhub scenario.

For detailed configuration information, see Chapter 3, "Implementing and Configuring the Solution.".

Figure 1-4 Optical Topology

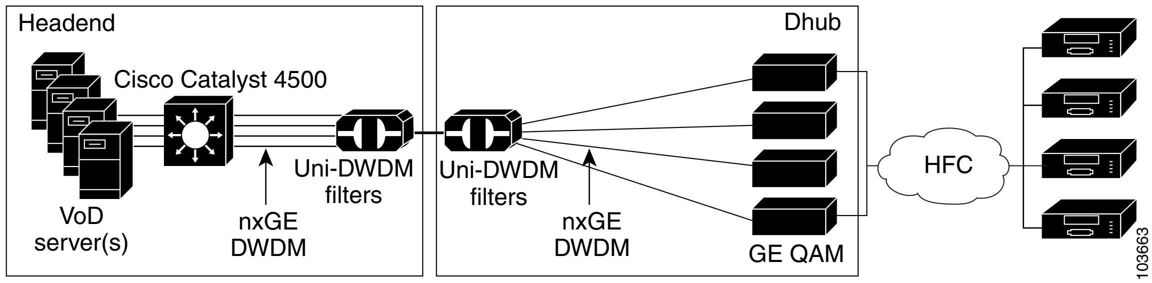

Figure 1-5 Ethernet Topology

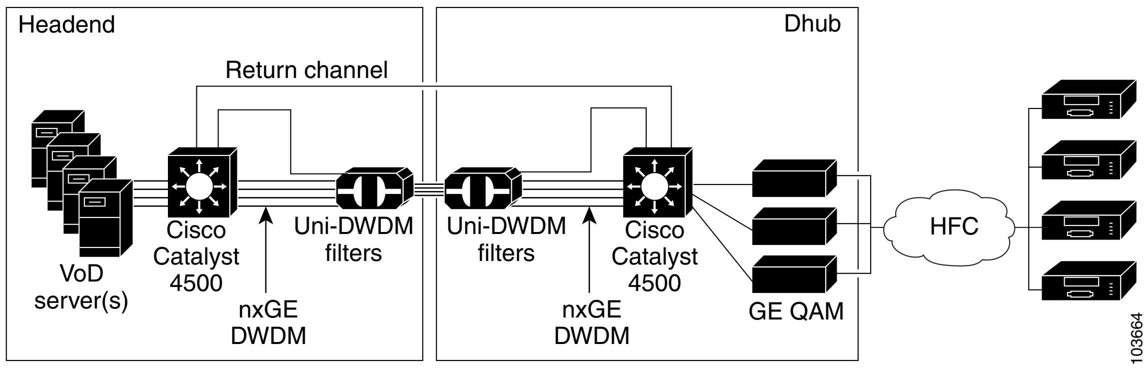

In customer deployments, all solution components are located in either a video headend site or a Dhub site. The basic solution topology is an Ethernet hub-and-spoke topology between the headend site and multiple Dhub sites. The Ethernet hub-and-spoke topology can be built in either physical hub-and-spoke or physical- ring fiber environments. When the solution is deployed in networks that use physical ring topologies, the physical ring networks must be converted to an Ethernet hub-and-spoke network at the optical layer.

Note For instructions for converting a ring network to a hub-and-spoke network, using physical dark-fiber media and *x*WDM methods, see Chapter 6, "Deploying the Cisco Gigabit-Ethernet Optimized VoD Solution in Fiber Ring Topologies."

Scenario Feature Comparison

Table 1-1 provides a high-level comparison of the features of the two scenarios.

Table 1-1 High-Level Comparison of Scenario Features

Feature	Switch in Dhub	No Switch in Dhub	Comments
Optical redundancy	yes	yes	See Chapter 4, "Providing Redundancy and Reliability."
Dhub return channel	yes	no	See Optical Topology, page 2-9.

Table 1-1 High-Level Comparison of Scenario Features (continued)

Feature	Switch in Dhub	No Switch in Dhub	Comments
Automatic population of MAC bridging table	yes	no	See Manual Population of ARP and Bridging Tables on the Dhub Switch, page 2-14.
Automatic population of ARP table	yes	no	
Dynamic IP routing	yes	no	See Ethernet Topology, page 2-1.
Dynamic load balancing	yes	no	See Load-Balancing Strategies, page 2-4.

Solution Components

This solution uses the following basic components:

- Optical Transport Network
- Ethernet Switches
- Video Servers
- Gigabit Ethernet QAM Devices
- Element Management (provisioning and fault management through the Cisco IOS CLI)

Optical Transport Network

The optical portion of the solution consists primarily of a unidirectional optical network to support the video streams. This consists of dense wavelength-division multiplexing (DWDM) gigabit interface converters (GBICs) in the headend switch connecting to a multiplexing unidirectional filter. On the Dhub side is a demultiplexing unidirectional filter that breaks out the individual wavelengths (content streams) and hands them either to a Cisco Ethernet switch, or directly to the GE QAM devices. If individual streams are handed off to a switch, this hand-off would occur to a receive-only GBIC. (For more information, see DWDM GBICs in the Headend Switch, page 1-6.)

The Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1 uses 32 wavelengths of the C-Band's 100-GHz ITU DWDM channels. There are eight bands of four wavelengths, allowing a bandgap of a single 100-GHz channel to remain unused between each four-wavelength group. Thus, to cover all 32 wavelengths, 32 individual DWDM GBICs, each representing an individual color of the total transmitted spectrum, are used. (For more information, see Filters, page 1-7.)

Note Cisco 15216 Optical Supervisory Channel (OSC) modules are used in some topologies to multiplex a bidirectional CWDM GE channel onto the same fiber as the DWDM link.

DWDM GBICs in the Headend Switch

Many MSOs have a single unidirectional optical link between the headend and the Dhub. DWDM (as apposed to CWDM) is typically chosen as the optical multiplexing technology, because the distance between the headend and Dhub often exceeds the CWDM range.

Cisco GBICs are hot-swappable devices that plug into a GE port, linking the port with the network. The solution uses DWDM optical components that include DWDM lasers on switches used for transmission. The DWDM GBICs used to terminate the unidirectional links in the Dhub switch have receive-only capability.

GE Bidirectional Links

CWDM or DWDM can potentially be used to handle the return path if done over a dedicated GE wavelength. In either scenario, a segmented bidirectional link can either be parallel to the unidirectional video links, or can use any bidirectional network. In fact, where other networks participate in locations of the solution components, as in a DOCSIS network, the bidirectional link does not have to be parallel.

This link can be used for the return path for provisioning information from switches and other devices connected to the unidirectional downstream link. It is assumed that all downstream wavelengths share the same fibers whether they are CWDM or DWDM. Upstream wavelengths use a second physical layer.

Filters

Cisco offers 16 two-channel filters, and 4 eight-channel filters. Any combination of these filters can be used to support multiple GE channels (or “colors” in an optical context). The filter in the headend is configured as a DWDM multiplexer, and the filter in the Dhub is configured as a DWDM demultiplexer. The use of unidirectional filters saves the cost of return path optics in both locations, as well as of a dedicated laser for each channel deployed in the Dhub.

All filters have an incoming or outgoing fiber port (for pass-through). There is also a monitor, or tap, port, which can be used for optical-performance monitoring of all wavelengths on the fiber when it is attached to an external optical monitoring system. Finally, there are two or eight individual-wavelength ports that filter specific wavelengths in or out.

When the need for more than a gigabit of traffic is apparent, or when the need for additional DWDM wavelengths is envisioned for multiple forward channels or one or more return channels from the Dhub to the headend, additional unidirectional filters can be used to provide asymmetric DWDM connectivity between the headend and the Dhub. This scenario also uses a set of bidirectional DWDM GBICs in the headend and Dhub switches. The DWDM output of the GBIC in the headend switch is driven as one wavelength in the filter complex that is used between the headend and Dhub sites. The DWDM output of the GBIC in the Dhub switch is also driven as one wavelength in a unidirectional filter complex that is used between Dhub and headend sites. The two unidirectional filters allow asymmetric bandwidth to be provisioned between the sites.

Ethernet Switches

The switches that have been tested and are supported in this release of the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1 are members of the Cisco Catalyst 4500 series. The Catalyst 4507R is used for the headend switch to gain maximum density and redundancy capabilities. Any of the other members of the series, such as the Cisco Catalyst 4503, can be used as the switch in the Dhub, where redundancy is not of great concern.

To ensure that the GE ports can be driven close to 100% link utilization for video, the Ethernet switches in the solution must be capable of operating their Gigabit Ethernet interfaces in a unidirectional, nonblocking manner, in both transmit and receive directions. For example, the switch must behave normally (not set off alarms, for instance) when only receiving Ethernet frames on a given interface, or only transmitting frames on a given interface. This functionality can be limited to nonblocking GE ports only on the Cisco 4500 series switch used in the solution.

Note For switch-related documentation, see Cisco 4000 Family Switches at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/index.htm

Video Servers

For this solution, VoD servers will be supplied by third-party VoD server vendors. Table 1-2 lists the capabilities of the VoD servers used in Release 1.0 of the solution.

Table 1-2 Capabilities of VoD Servers Used in the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1, Release 1.0

VoD Server Vendor/ Product	Daisy Chain?	Optical Failover?	Nonstandard 802.x?	Bidirectional GE?	GE Failover?	ARP Sender?	ARP Receiver?	Management Port
Concurrent/ MediaHawk www.ccur.com	No	No	No	No	No	Yes	Yes	Yes
SeaChange/ ITV Media Cluster www.schange.com	No	No	No	No	No	No	No	Yes
nCube/ n4x On-Demand Server www.ncube.com	No[1]	No[2]	No	Yes[3]	No[4]	Yes	Yes	Yes

1. Servers are interconnected by means of an internal scalable network.
2. No failover at optical layer, but failover is supported at a higher layer.
3. Optimized for unidirectional video delivery.
4. No failover at the GE physical link layer, but failover is supported at a higher layer.

Note While third-party servers will not be supported as part of this solution, configurations relevant to the transport network are documented in this guide.

Gigabit Ethernet QAM Devices

These devices take MPEG streams arriving over GE transport in UDP packets and use QAM modulation and RF upconversion to introduce the video onto RF channels on the HFC plant. Video edge equipment for this release of the solution will be supplied by third-party QAM vendors. Table 1-3 lists the capabilities of the QAM devices used in Release 1.1 of the solution. (Versions of the Harmonic NSG earlier than version 2.1 do not support GE redundancy, and provide unreliable ARP responses.)

Table 1-3 Capabilities of QAM Devices Used in the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.1

QAM Device Vendor/ Product	Daisy Chain?	Optical Failover?	Nonstandard 802.x?	Bidirectional GE?	ARP Sender?	ARP Receiver?	Management Port
Harmonic/ NSG www.harmonicinc.com	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cisco/ Cisco uMG9820 QAM Gateway www.cisco.com/univercd/cc/td/doc/product/cable/vod/umg9820/index.htm	No	Yes	Yes	Yes	No	Yes	Yes

Note While third-party QAM devices will not be supported as part of this solution, configurations relevant to the transport network are documented in this guide.

Element Management

Provisioning and fault management of Cisco Catalyst 4500 series switches for the solution are performed through the command line interface (CLI) of the Cisco IOS. Fault management is not provided for the passive optical components of the solution, including multiplexers, demultiplexers, splitters, and the OSC (optical supervisory channel).

Note The OSC needs a bidirectional network between two components, such as a switch and a filter. (See CWDM Multiplexing Using the Cisco 15216 OSC-1510, page 2-10.) Although the OSC can offer network management features, performance management is not included in this release.

CHAPTER 2

Designing the Solution

In customer deployments, all Cisco Gigabit-Ethernet Optimized VoD Solution components will be located in either a video headend site or a distribution hub (Dhub) site. The basic topology is an Ethernet hub-and-spoke topology between the headend site and multiple Dhub sites. The Ethernet hub and spoke topology can be built in either physical hub-and-spoke or physical fiber-ring environments. When the solution is deployed in networks that use physical ring topologies, the ring networks must be converted to an Ethernet hub-and-spoke network at the optical layer. (Different optical layer topologies have no effect on either the operation or performance of Gigabit Ethernet.)

Note For instructions for converting a ring network to a hub-and-spoke network, using physical dark-fiber media and *x*WDM methods, see Chapter 6, "Deploying the Cisco Gigabit-Ethernet Optimized VoD Solution in Fiber Ring Topologies."

All Ethernet topologies in the initial releases (Release 1.0 and Release 1.1) of the solution include an Ethernet switch in the video headend. There are two main Ethernet topology choices for video transport. The topology choices depend on whether or not there is an Ethernet switching capability in the Dhub.

This chapter presents the following major topics:

- Switch in the Dhub, page 2-1
- No Switch in the Dhub, page 2-11

Switch in the Dhub

This section presents the following topics:

- Ethernet Topology
- IP Topology
- Optical Topology

Ethernet Topology

Figure 2-1 on page 2-2 illustrates the Ethernet topology used in the Cisco Gigabit-Ethernet Optimized VoD Solution for Ethernet switching in the Dhub. This section presents the following topics:

- Designing for Growth
- UDLR

- QoS
- OSPF
- Subtended Dhubs
- QAM Device Connectivity

Figure 2-1 Switch in Dhub: Ethernet Topology—Dedicated GE Bidirectional Wavelength

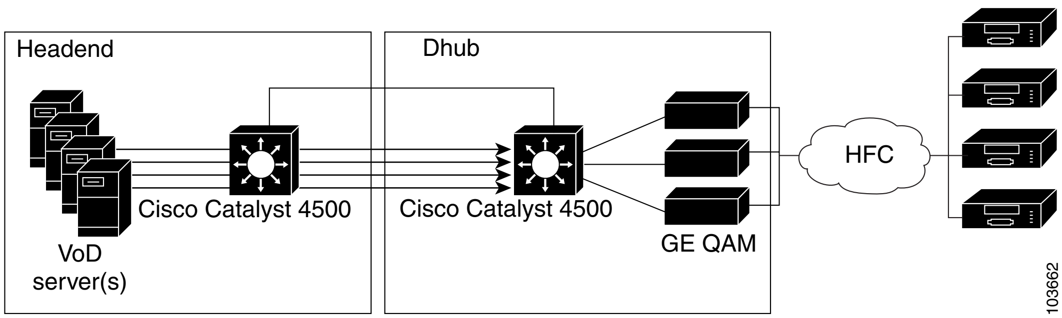

This topology allows the switching path to be broken into three sets of component links:

- The links between the VoD servers and the headend switch
- The links between the headend switch and the Dhub switch
- The links between the Dhub switch and the QAM devices

Since multiple equal-cost links connect the headend and Dhub switches, load balancing will be applied. Both Layer 3 IP and EtherChannel load balancing are used in this solution. Depending on the number of ports used to connect the headend switch to the Dhub switch, one or both of these load-balancing techniques are used. The use of load balancing simplifies provisioning, because the links between the headend and Dhub switches appear as one large pipe to both routing protocols, as well as to the IP forwarding logic. (See Enabling Load Balancing, page 3-4.)

This topology also allows for a control-plane back channel between the Dhub and headend switches. The video transport interfaces between the headend and Dhub switches are unidirectional, helping to reduce cost. Because of this, a separate return path link must be used.

To allow IP routing and other protocols to work transparently, it is important to bind the return path link There are two ways to do this:

- The first method, referred to here as an *asymmetric EtherChannel*, includes one bidirectional link in the EtherChannel group between the headend and Dhub switches. This method essentially creates an asymmetric bidirectional EtherChannel interface.
- The second method, Cisco *UniDirectional Link Routing* (UDLR), binds a unidirectional interface to a GRE (generic routing encapsulation) tunnel that is used as the return path. The two interfaces are bound below the IP layer to create one logical bidirectional interface at the IP layer.

In Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution, either an asymmetric EtherChannel or UDLR are used for the return-channel connectivity between the headend and Dhub switches. Which method to use will depend on the specifics of the connectivity between the switches. An asymmetric EtherChannel will be used when the headend and Dhub switches are directly connected with a single EtherChannel group. In more-complex connectivity scenarios, UDLR must be used in place of the asymmetric EtherChannel. These scenarios are described below.

IP Topology

Figure 2-2 on page 2-3 illustrates the IP topology of the Cisco Gigabit-Ethernet Optimized VoD Solution for the switch-in-Dhub scenario.

Figure 2-2 Switch in Dhub: IP Topology—Dedicated GE Return Path

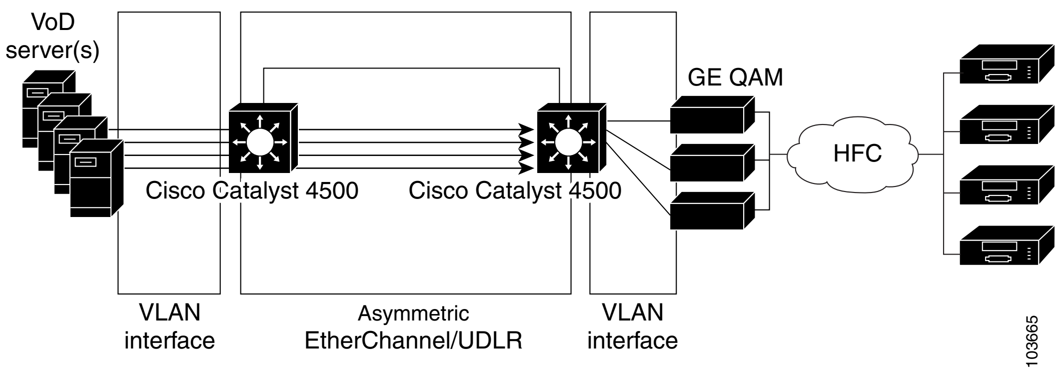

There are three logical IP segments:

- The links between the VoD servers and the headend switch. These links are terminated through a MAC layer bridge group into a single IP interface on the headend switch.
- The links between the Dhub switch and the QAMs. These links are terminated through a MAC layer bridge group into a single IP interface on the Dhub switch.
- One or more EtherChannel groups or physical Layer 3 interfaces between the headend and Dhub switches. How this IP segment is broken into IP interfaces depends on the specifics of the connectivity between the headend and Dhub switches.

  If the headend switch is connected to the Dhub switch with a single EtherChannel group, an asymmetric EtherChannel will be used. As a result, there will be a single IP interface on the headend and Dhub switches for this EtherChannel group.

The use of an asymmetric EtherChannel means that flows from the headend to the Dhub will be distributed across all of the links in the EtherChannel group. Traffic from the headend to the Dhub will consist of video flows, but may also include nonvideo data such as network management traffic and IP routing updates.

For configuration details, see Establishing Interfaces on the Headend Switch, page 3-5, and Establishing Interfaces on the Dhub Switch, page 3-10.

Designing for Growth

The following sections provide recommendations for network designers to accommodate future growth in capacity between the headend and the Dhub:

- Load-Balancing Strategies
- Strategies to Facilitate Expansion

Load-Balancing Strategies

This section introduces issues and strategies related to EtherChannel and IP load balancing. It is valid to configure EtherChannel, IP load balancing, or a combination of EtherChannel and IP load balancing across a set of ports from a headend to a Dhub switch. In theory, it should be possible to use any of the above three load-balancing techniques, along with any combination of EtherChannel and IP load balancing, for any number of ports between those switches. However, limitations in the load-balancing architecture and its implementation on the Cisco Catalyst 4500 series switches force restrictions on how load balancing can be used for this solution. The restrictions—and strategies to deal with them—are presented below.

The Cisco Catalyst 4500 series switches implement a stateless, equal-cost load-balancing algorithm in hardware for both EtherChannel and IP load balancing. Because the hardware supports only equal-cost load balancing, all EtherChannel groups should have the same number of ports assigned to them when EtherChannel and IP load balancing are used in combination. This restricts valid configurations of EtherChannel in combination with IP load balancing to cases where the number of links in each EtherChannel group, multiplied by the number of EtherChannel interfaces per Dhub, equals the number of ports between the headend and Dhub. Consider also the following:

- The EtherChannel load-balancing algorithm on the Cisco Catalyst 4500 series switches is nonoptimal when the number of ports in an EtherChannel group does not divide evenly into 8.
- Interactions between EtherChannel and IP load balancing cause nonoptimal load-balancing behavior for specific combinations of EtherChannel and IP load-balancing group size.

Table 2-1 lists the combinations of EtherChannel and IP load balancing that have been shown, through both simulation and system testing, to exhibit optimal load-balancing behavior for Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution.

Caution To ensure optimal load-balancing behavior, it is strongly recommended that EtherChannel and IP load balancing be configured within the limits of Table 2-1.

Notice that there is more than one potential load-balancing configuration for some values of headend-to-Dhub ports. For example, load balancing for 4 ports between a headend and Dhub switch can be configured by using either 1 EtherChannel group of 4 ports or 4 separate physical Layer 3 interfaces. The best configuration to use in these cases will depend on a number of factors, including IP addressing scheme, ease of configuration, and flexibility in adding ports as capacity demands increase.

Table 2-1 Approved EtherChannel Assignment and IP Load-Balancing Combinations

<table>
<tr><td>Headend ports per Dhub</td><td>1</td><td colspan="2">2</td><td>3</td><td colspan="2">4</td><td>5</td><td>6</td><td>7</td><td colspan="3">8</td><td>12</td><td colspan="2">16</td></tr>
<tr><td>Ports per EtherChannel</td><td></td><td>2</td><td></td><td></td><td>4</td><td></td><td></td><td></td><td></td><td>8</td><td>4</td><td></td><td>4</td><td>8</td><td>4</td></tr>
<tr><td>EtherChannel interfaces per Dhub</td><td></td><td>1</td><td></td><td></td><td>1</td><td></td><td></td><td></td><td></td><td>1</td><td>2</td><td></td><td>3</td><td>2</td><td>4</td></tr>
<tr><td>Physical Layer 3 interfaces per Dhub</td><td>1</td><td></td><td>2</td><td>3</td><td></td><td>4</td><td>5</td><td>6</td><td>7</td><td></td><td></td><td>8</td><td></td><td></td><td></td></tr>
</table>

Note Table 2-1 documents only load-balancing configurations of up to 16 ports between a headend and Dhub switch. This is because the Cisco Catalyst 4500 series switches do not support more than 16 GE ports running at line rate. Also, note that table does not show any load-balancing configurations for 9, 10, 11, 13, 14, or 15 ports. The EtherChannel and IP load-balancing restrictions described above result in nonoptimal load-balancing behavior in these cases. See also Understanding the Load-Balancing Behavior of Asymmetric EtherChannels, page 5-12.

To make EtherChannel load balancing more deterministic, only the destination IP port number is used as the input to the load-balancing function. The Harmonic NSG QAM device (along with other GE QAM devices) uses a static mapping to derive the QAM channel and MPEG program number from the destination port number of an incoming MPEG stream. Consequently, a fully loaded QAM will have a known set of destination IP port values for all the MPEG streams that it services.

Strategies to Facilitate Expansion

This section describes strategies that can be used to facilitate the growth in traffic to a Dhub by increasing the number of ports between a headend and Dhub switch. These strategies use the information in Table 2-1 to define different strategies for increasing capacity, depending on the requirements of the network. Other expansion strategies may be derived from the information in Table 2-1 to match best the requirements of network designs that do not match the objectives of the strategies listed in this section.

In some network designs, it is important to simplify provisioning by avoiding the renumbering of IP interfaces as capacity requirements increase. In these cases, it is best to use a combination of EtherChannel and IP load balancing to increase port capacity. This results in a strategy where port capacity is increased in units of EtherChannel groups. With the information in Table 2-1, this strategy can be implemented by using a combination of 2- and 4-port EtherChannel group sizes. Table 2-2 on page 2-5 illustrates the subset of the combinations of Table 2-1 that could be used to implement this strategy. While this strategy simplifies provisioning by not requiring the renumbering of IP interfaces as capacity increases, it does not allow as much flexibility in port number increments as is shown in Table 2-1.

Table 2-2 Load-Balancing Strategies to Avoid Renumbering IP Interfaces

Headend ports per Dhub	1	2	4	8	12	16
Ports per EtherChannel		2	4	4	4	4
EtherChannel interfaces per Dhub		1	1	2	3	4
Physical Layer 3 interfaces per Dhub	1					

If the renumbering of IP interfaces is not an issue for a particular network design, an alternative strategy could be used in which (1) IP load balancing alone is used for headend-to-Dhub port combinations of 1 to 8, and (2) a combination of IP and EtherChannel load balancing is used to support more than 8 ports. Table 2-2 illustrates the subset of the combinations of Table 2-1 on page 2-4 that could be used to implement this strategy. The strategy illustrated in Table 2-3 provides the full flexibility of Table 2-1 on page 2-4, while minimizing the renumbering of IP subnets as capacity requirements increase.

Table 2-3 Modularized EtherChannel Assignment and IP Load-Balancing Combinations

Headend ports per Dhub	1		3	4	5	6	7	8	12	16
Ports per EtherChannel									4	4
EtherChannel interfaces per Dhub									3	4
Physical Layer 3 interfaces per Dhub	1	2	3	4	5	6	7	8		

UDLR

When multiple EtherChannel groups are configured between the headend and Dhub switches, it is inefficient to require an Ethernet return channel for each interface. Consequently, Cisco UniDirectional Link Routing (UDLR) is used on all of the unidirectional IP interfaces configured between the two switches. (This is discussed further in Subtended Dhubs, page 2-7.) The configuration of the UDLR interface creates a second IP interface on the Dhub switch.

When UDLR is configured, the GRE tunnel for the UDLR interface will end up using the bidirectional Ethernet port that was configured to be part of the asymmetric EtherChannel. OSPF cost metrics are configured on the Dhub switch to direct upstream IP packets through the bidirectional IP interface, rather than through the UDLR back-channel interface. This is needed because the UDLR back channel is process-switched, and will consequently have very low throughput.

Note When UDLR is used in this scenario, there will be very few protocols that will send packets over the UDLR return channel. The only protocols that will end up using the return channel in this scenario will be ARP, and potentially the hello protocol in OSPF.

QoS

To enable video streams to be serviced ahead of nonvideo traffic, QoS must be enabled on the headend switch. To achieve this, a **policy-map** statement is mapped to an access control list (ACL) on the headend switch interface that is connected to the VoD servers. These statements will mark video traffic with "*DSCP EF*" and nonvideo traffic with "*DSCP 0*." If a management port is connected to the headend switch, all packets arriving on this port will be marked with "DSCP 0." Egress QoS will be enabled on the Ethernet ports of asymmetric EtherChannel groups connected to Dhubs.

Packets marked with "DSCP EF" will be serviced by a priority queue. Packets marked with any other DSCP value will be serviced by a weighted queue configured for 80% of the physical link bandwidth. This configuration allows the video load to be distributed evenly cross all members of an EtherChannel group, while ensuring that video is always transmitted in a timely manner.

Note High video-link utilization may adversely affect convergence times when IP routing protocols are used.

For configuration details, see Establishing Quality of Service (QoS), page 3-2.

OSPF

OSPF is used as the Internet routing protocol in Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution. OSPF is enabled on all IP interfaces in the headend and Dhub switches. OSPF populates the routing table on the headend switch with routes to the QAM devices, and also enables equal-cost load balancing when multiple IP interfaces are configured between the headend and Dhub switches.

Subtended Dhubs

Some topologies for the Cisco Gigabit-Ethernet Optimized VoD Solution will include subtended Dhubs. A subtended Dhub is a Dhub whose bidirectional link is connected to another Dhub, and whose unidirectional link(s) are connected to the headend switch. Refer to Figure 2-3 on page 2-8, which illustrates the architecture for subtended Dhubs and return channels, with support for bidirectional and unidirectional interfaces. When subtended Dhubs are used, the upstream Dhubs will have more than one bidirectional link connected to them. On the upstream Dhubs, each of the bidirectional links will appear as a separate IP interface.

Since the bidirectional interface to the subtended Dhub is not directly connected to the headend switch, but the unidirectional interface(s) are, an asymmetric EtherChannel cannot be used. Instead, UDLR is used for the return channel for all IP interfaces between the Dhub and the headend switch, as discussed in UDLR, page 2-6. When UDLR is configured, the GRE tunnel for the UDLR interface will use the least-cost IP path back to the headend switch. This will turn out to be the combination of the bidirectional (asymmetric) Ethernet link (EtherChannel) between (1) the subtended Dhub and upstream Dhub, and (2) the Ethernet port that was configured to be part of the asymmetric interface between the upstream Dhub and headend switch. As noted earlier, the configuration of the UDLR interface ends up creating a second IP interface on the subtended Dhub switch. Figure 2-3 on page 2-8 illustrates this second path from the headend to subtended Dhub B.

Caution This second path should never be used for VoD delivery, because the bidirectional link between Dhub A and Dhub B will typically not have as much bandwidth as the directly connected unidirectional link between the headend and Dhub B.

Normally, the second path between the headend and subtended Dhub B will not be selected by routing protocols, because it will have a higher cost than the directly connected link between the switches. However, if the directly connected link fails, the second path will be chosen for VoD traffic. The use of the second path by VoD streams would likely cause congestion and result in degraded video quality for VoD traffic going from the headend to the subtended Dhub. In addition, the use of the second path would cause VoD streams from the headend to both Dhubs A and B to be routed through the EtherChannel group from the headend to Dhub A. Because of this, a link failure of the directly connected link from the headend to Dhub B would result in degraded video quality for customers of both Dhub B and Dhub A.

Figure 2-3 Subtended Dhubs and Return Channel Architecture

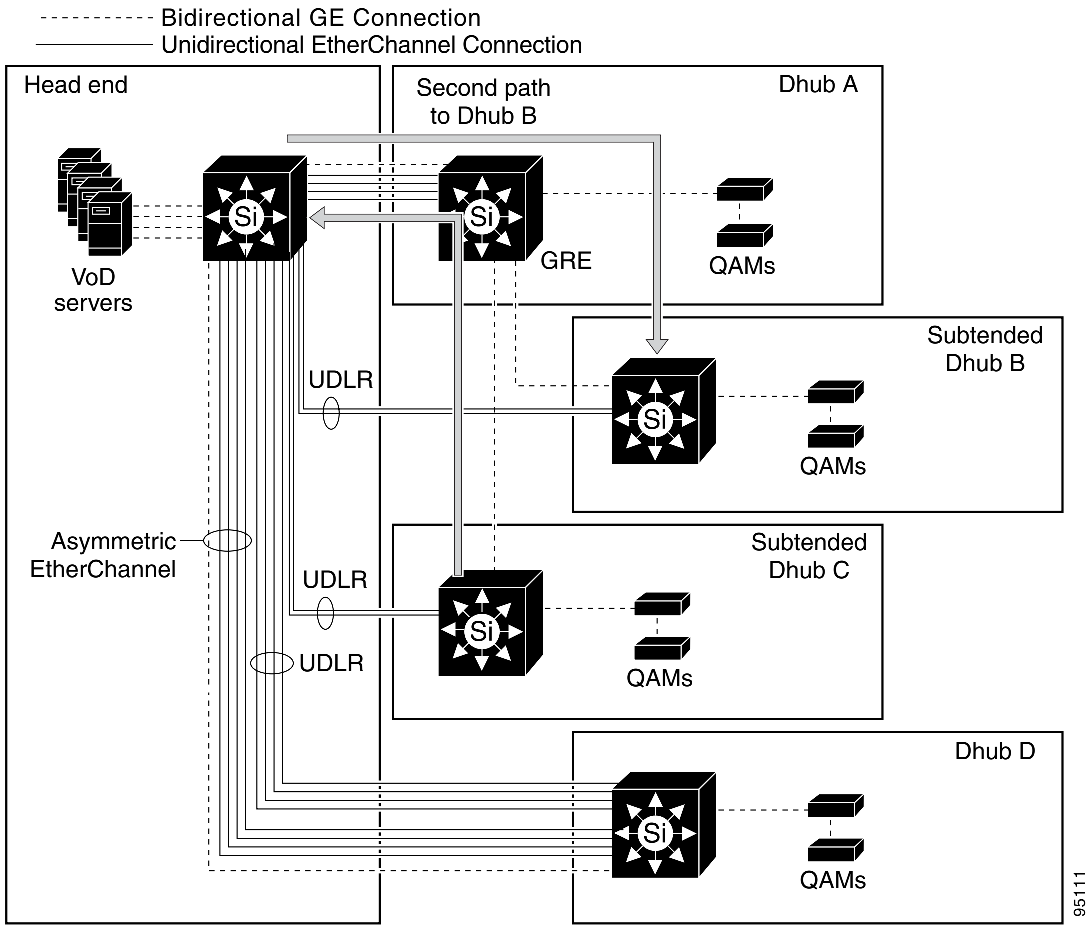

To prevent this scenario from occurring, the IP routing configuration must be modified to prevent the second path from the headend to Dhub B from being used by the headend switch. (Though the second path is still advertised and is entered into the switch's database, it is blocked from the routing table.) This can be accomplished by configuring OSPF distribution lists on the headend switch. A distribution list filters OSPF routes received on a particular interface to include or exclude routes to specific subnets. Since the headend switch is supposed to use only the directly connected EtherChannel to send packets to the QAMs in a particular Dhub, a "permit" access control list can be configured on each IP interface in the headend switch. The ACL for each interface will contain one permit entry for the QAM subnet on the directly connected Dhub switch. In this way, the headend switch will learn only a single route to the QAMs in each Dhub. That route will use the directly connected link to the Dhub.

Note For example configurations, see Establishing Subtended Dhubs, page 3-14.

QAM Device Connectivity

Some third-party QAM devices can be daisy chained. These devices support basic packet switching to allow IP packets not destined for one QAM device to be passed to the other QAM devices in the chain. Even though the Dhub switch can provide a separate fiber to each QAM device, daisy chaining has the advantage of reducing the number of GE ports needed in the Dhub. However, the risk of failure is increased, because a failure on the link between the switch and the first QAM device causes all devices in the chain to fail. In most situations, more than a single GE worth of traffic will be required between the QAM device complex and the Dhub switch. In these cases, each "master" GE QAM device should have its own connection to the Dhub switch.

When QAM dvices that are not capable of daisy chaining are used in this solution, each QAM device must be connected to the Dhub switch through a separate bidirectional GE port. The capabilities of the QAM devices used in the Cisco Gigabit-Ethernet Optimized VoD Solution are shown in Table 1-3 on page 1-9.

Note For example configurations, see Connecting to QAM Devices, page 3-13.

Optical Topology

Many multiple system operators (MSOs) have a single unidirectional optical link between the headend and a Dhub. DWDM (as opposed to CWDM) is typically chosen as the optical multiplexing technology, because the distance between the headend and Dhub often exceeds the CWDM range.

The Cisco Gigabit-Ethernet Optimized VoD Solution uses DWDM optical components that are cost-reduced to include DWDM lasers on nodes that are used for transmission. For example, the DWDM GBICs used to terminate the unidirectional links in the Dhub switch have receive-only capability. This saves the cost of the DWDM laser on each receive port. The Cisco 15216 FlexLayer DWDM filters have also been cost-reduced, so that the filter in the headend is capable only of DWDM multiplexing, and the filter in the Dhub is capable only of DWDM demultiplexing. This eliminates the need for passive optics in both components and for a DWDM laser for the filter in the Dhub—with attendant cost savings.

The back channel between the headend and Dhub switches requires a second fiber between the headend and the Dhub. This second fiber may be part of an existing SONET infrastructure that is being redeployed for GE transport, or it may be part of a side-by-side data network already deployed by the MSO. In many cases, however, a second fiber will not already be in use and will have to be allocated for use in the video network.

There are several options for implementing the GE back channel:

- DWDM over Dark Fiber
- DWDM Multiplexing Using a Second Set of Uni-DWDM Filters
- CWDM Multiplexing Using the Cisco 15216 OSC-1510

The best choice for a particular design depends on the existing infrastructure and future applications that may be required.

DWDM over Dark Fiber

DWDM over dark fiber will typically be used when a return fiber between the headend and Dhub sites is not already in use. DWDM over dark fiber uses a set of bidirectional DWDM GBICs in the headend and Dhub switches. The DWDM output of the GBIC in the headend switch is driven as one wavelength in the Uni-DWDM filter complex between the headend and Dhub sites. The DWDM output of the GBIC in the Dhub switch directly drives a return fiber connected to the receive side of bidirectional GBIC in the headend switch. This option does not need additional optical components to provide the back-channel functionality. Figure 2-4 illustrates both the Ethernet and optical layers used for the DWDM-over-dark-fiber GE return path option for the switch-in-Dhub scenario.

Figure 2-4 Switch in Dhub: Optical Topology—DWDM over Dark Fiber GE Return Path

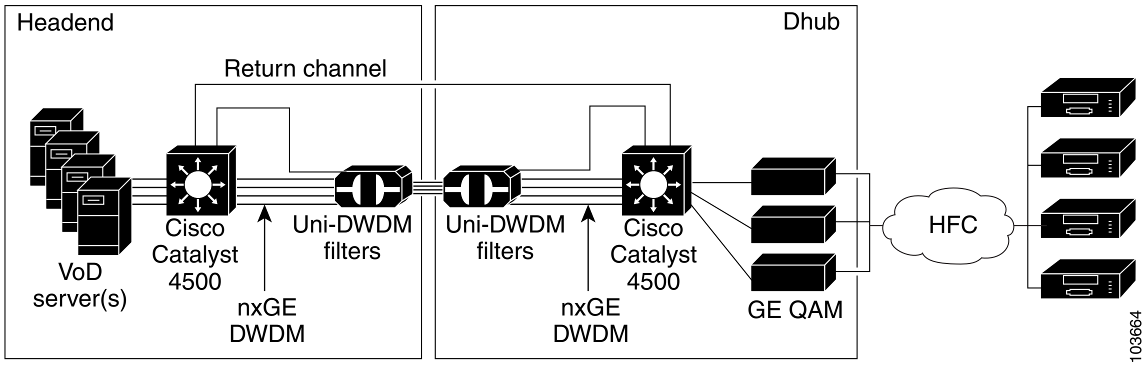

DWDM Multiplexing Using a Second Set of Uni-DWDM Filters

When DWDM fiber is already available, or when the need for multiple DWDM wavelengths is envisioned for the return channel from Dhub to headend, a second unidirectional uni-DWDM filter can be used to provide asymmetric DWDM connectivity between headend and Dhub. This scenario also uses a pair of bidirectional DWDM GBICs in the headend and Dhub switches. The DWDM output of the GBIC in the headend switch is driven as one wavelength in the filter complex between headend and Dhub sites. The DWDM output of the GBIC in the Dhub switch is also driven as one wavelength in a unidirectional filter complex between Dhub and headend sites. The two unidirectional filters allow asymmetric bandwidth to be provisioned between headend and Dhub sites. Figure 2-5 on page 2-10 illustrates the asymmetric DWDM return path option for the switch-in-Dhub scenario.

Figure 2-5 Switch in Dhub: Optical Topology—DWDM with Return DWDM MUX

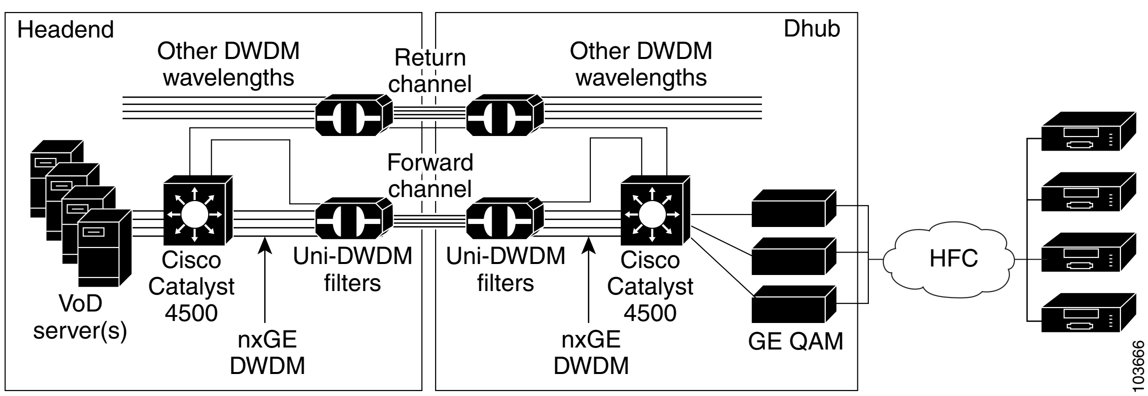

CWDM Multiplexing Using the Cisco 15216 OSC-1510

When the fiber distance between the headend and Dhub sites allows it, the Cisco 15216 OSC-1510 can be used to multiplex a bidirectional CWDM link into the downstream DWDM wavelengths. In this scenario, standard bidirectional CWDM GBICs can be used in the headend and Dhub switches. The downstream CWDM channel is multiplexed/demultiplexed onto the DWDM fiber after and before, respectively, the uni-DWDM filters. The upstream CWDM channel is driven on a return fiber dedicated for this purpose. The use of the OSC (optical supervisory channel) and CWDM filters can reduce the cost of implementing the optical return path over the DWDM scenarios described above. However, because of the distance limitations in multiplexing CWDM into the optical path, this approach is recommended only for fiber topologies where long-distance connections are not needed. Figure 2-6 illustrates the CWDM return path option for the switch-in-Dhub scenario.

Note For further information, see Implementing the Cisco ONS 15216 FlexLayer, page 3-18.

Figure 2-6 Switch in Dhub: Optical Topology—CWDM Return Path

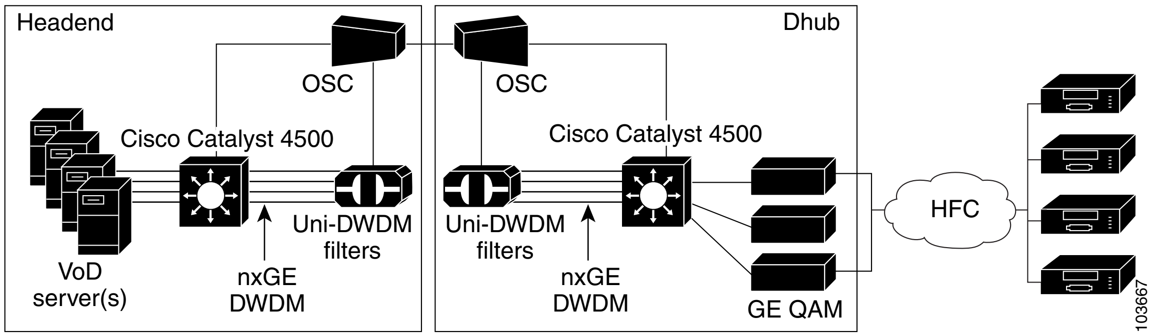

No Switch in the Dhub

This section presents the following major topics:

- Ethernet Topology
- Optical Topology

Ethernet Topology

Figure 2-7 on page 2-11 illustrates the Ethernet topology when the Cisco Gigabit-Ethernet Optimized VoD Solution is deployed with no switching capability in the Dhub. This topology is extremely simple, because there are essentially only two segments in the Ethernet switching path.

Figure 2-7 No Switch in Dhub: Ethernet Topology

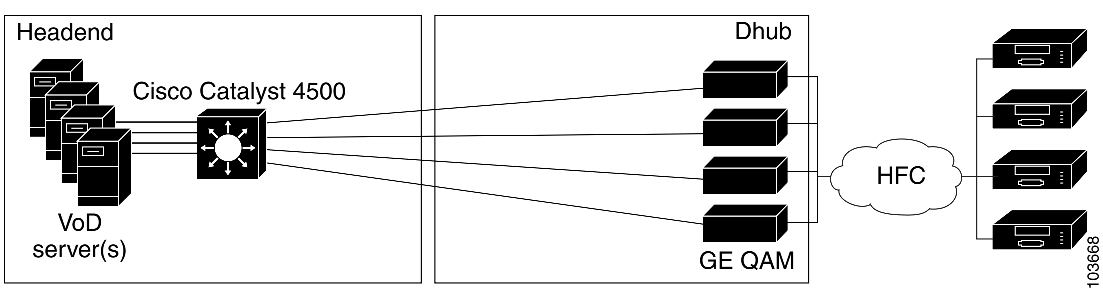

Note This topology does not support a backchannel mechanism in the first phase of the solution. This is because Cisco UDLR (UniDirectional Link Routing) is a proprietary protocol that requires a Cisco switch at each end. Because there is only one Cisco switch in this topology, UDLR cannot be supported. A future release will include support for an RFC 3077 session as a back-channel mechanism between the QAM device and the headend switch.

Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution support two different types of Layer 2 and Layer 3 switching configurations for this topology:

- Layer 2 Switching
- IP Switching Between the Headend and All Dhubs
- IP Switching Between the Headend and Each Dhub Link

These are described in the following sections.

Layer 2 Switching

In Layer 2 switching, all of the VoD servers and QAM devices will appear in the same IP subnet. The Cisco Catalyst 4500 switch must be configured to perform MAC layer bridging between all of its connected interfaces. Since no back channel will be available between the QAM devices and the switch, the MAC layer bridge table entries for the interfaces between the headend and Dhub must be populated manually. In addition, because there is no back channel between the headend and the Dhub, ARP requests from the VoD server to the QAM devices will not work. This means that the ARP table in the VoD server must also be populated manually. Figure 2-8 on page 2-12 illustrates the Ethernet topology for Layer 2 switching for the no-switch-in-Dhub scenario.

Figure 2-8 No Switch in Dhub: Ethernet Topology for Layer 2 Switching

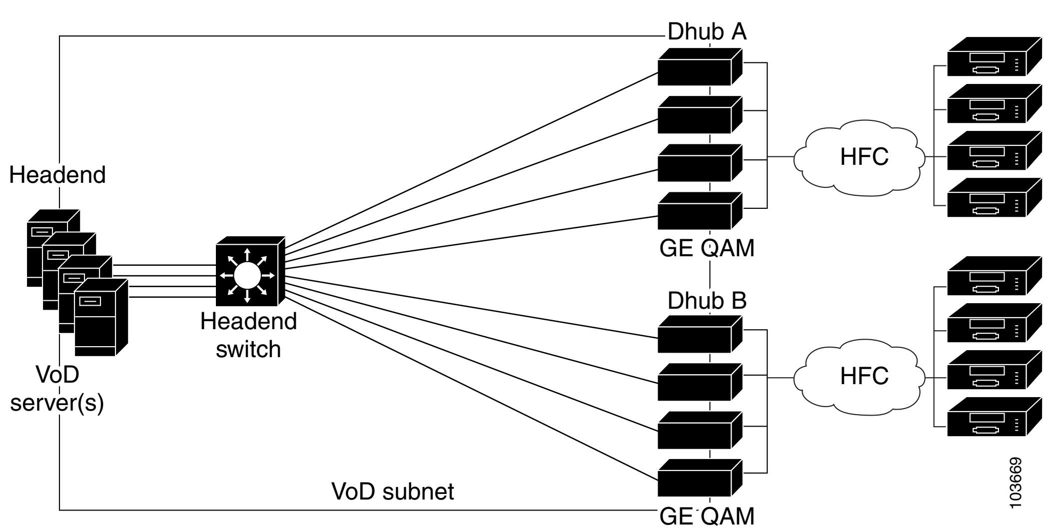

IP Switching Between the Headend and All Dhubs

When IP switching is used between the headend and all the Dhubs, two interfaces must be configured on the headend switch:

- An interface consisting of all of the links between the VoD servers and the headend switch

  These links will all be terminated through a MAC layer bridge group into a single IP interface on the headend switch.

- An interface consisting of all of the links between the headend switch and the QAM devices in all Dhubs

These links will also be terminated into a single IP interface on the headend switch. Figure 2-9 illustrates the topology for switching between the headend and all Dhubs for the no-switch-in-Dhub scenario.

Figure 2-9 No Switch in Dhub: Topology for IP Switching Between the Headend and All Dhubs

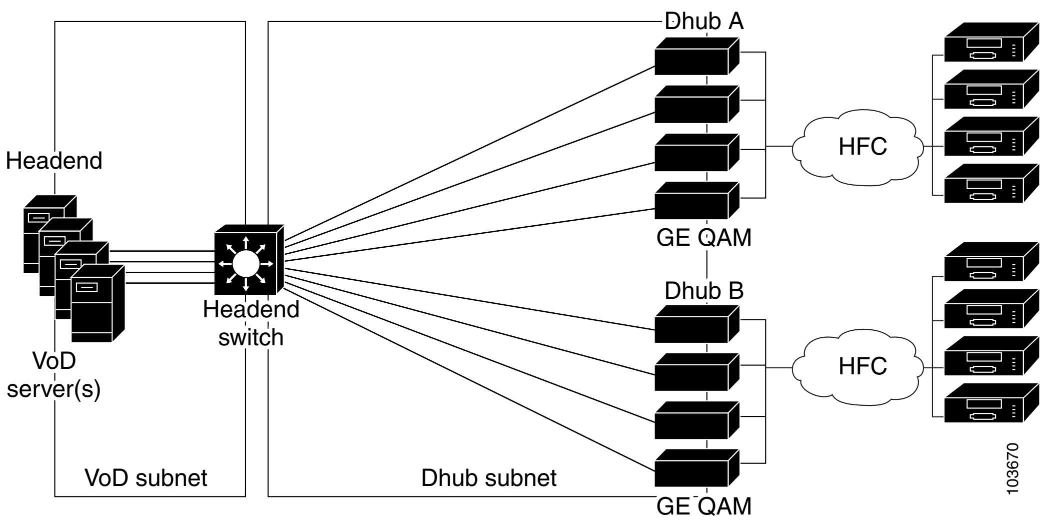

While this also requires manual configuration, the configuration will be reduced, because IP discovery protocols can be used on the bidirectional link between the VoD servers and the headend switch. For example, ARP is used to populate the ARP table on VoD servers that support ARP. The limitations of VoD servers in regards to ARP functionality are shown in Table 1-2 on page 1-8.

For VoD servers that do not support an ARP sender (such as the SeaChange VoD server), a subset of HSRP (Hot Standby Routing Protocol) functionality is configured on the Ethernet ports connected to the VoD servers. In addition to implementing a failover protocol, HSRP supports the configuration of a virtual MAC address that can be shared between the active and standby HSRP nodes. For VoD servers that do not support ARP, HSRP is configured with only one interface in the HSRP group. As a result, HSRP failover will not be enabled. Instead, an HSRP virtual MAC address will be configured to replace the native MAC address of the VLAN interface. The HSRP virtual MAC address allows ARP table configuration on the VoD server to be made independently of the headend switch to which it is connected.

Because the Dhubs are directly connected to the headend switch, the IP address and subnet mask assigned to the Dhub VLAN interface on the headend switch will populate the IP switching table with forwarding information needed to forward packets to the Dhubs.

IP Switching Between the Headend and Each Dhub Link

When IP switching is used between the headend and each Dhub, there will be one interface configured on the headend switch for all of the ports connected to VoD servers and one interface configured for each Dhub that the headend switch is connected to. This configuration trades off increased complexity in the configuration of static routes on the headend switch for decreased complexity in configuration of the bridging table as compared to the scenario described in IP Switching Between the Headend and All Dhubs, page 2-12. Figure 2-10 illustrates the topology for IP switching between the headend and each Dhub link for the no-switch-in-Dhub scenario.

Figure 2-10 No Switch in Dhub: Topology for IP Switching Between the Headend and Each Dhub Link

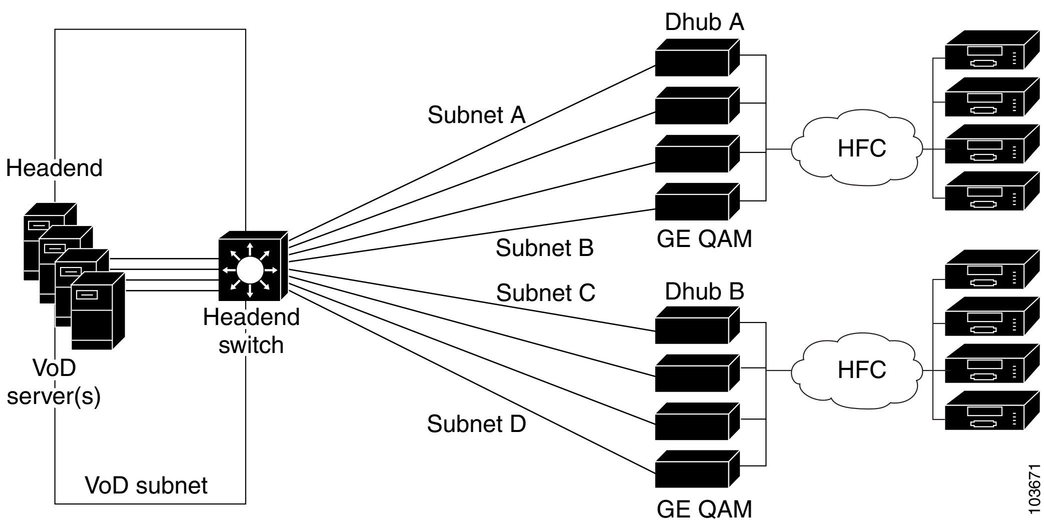

This configuration also provides a straightforward migration path to the first implementation phase of RFC 3077. RFC 3077 describes a protocol that emulates a logical bidirectional link between two nodes that are directly connected with a unidirectional link and indirectly connected with a separate bidirectional IP path. While RFC 3077 can be used to emulate a bidirectional multipoint network (a multipoint MAC layer bridged network in 802.3 environments), it is simpler to emulate a bidirectional point-to-point link using RFC 3077. Because of this, the first implementation of RFC 3077 in the Cisco IOS will emulate a bidirectional point-to-point link whose endpoints terminate as IP layer interfaces. Since the configuration described in this section breaks the IP topology into one subnet per link connected to Dhubs, the initial implementation of RFC 3077 can be used to provide a back-channel capability for each of the links connected to Dhubs. The RFC 3077 back channel can be used to carry ARP requests between the headend switch and the QAM devices to dynamically populate the ARP table on the headend switch and on the QAM devices.

In this configuration, each link from the headend switch to each Dhub resides in a separate subnet. As a result, there will be one physical interface configured on the headend switch for each link connected to a Dhub. In cases where there are multiple links to a Dhub, the outbound interfaces may be included in a single VLAN to use network addresses more efficiently. Otherwise, each interface will require a unique subnet.

Since the QAM devices are directly connected to the headend switch, the IP address and subnet mask configured for each VLAN interface will populate the IP switching table with the forwarding information needed to forward packets to the Dhubs connected on that link.

Manual Population of ARP and Bridging Tables on the Dhub Switch

Normally, the switch discovers the IP and MAC addresses of the destination interface (for the QAM devices, the GE port for Layer 2 switching). The switch then populates (1) its bridging table with pairs consisting of a MAC address and the associated physical port, and (2) its ARP table with pairs consisting of a MAC address and an IP address.

To accommodate QAM devices that do not send ARP replies, you will need to configure the ARP tables manually on the Dhub switch. If multiple QAM devices are connected to the Dhub switch in a single VLAN and they do not send ARP replies, you should manually populate the MAC address table to avoid possible port flooding on the VLAN. If the QAM devices are connected through Layer 3 ports, or if each QAM device is in its own VLAN, then the MAC address table does not have to be populated manually.

Note Versions of the Harmonic NSG earlier than V2.1 provide unreliable ARP responses. For example configurations, see Manually Configuring the ARP and MAC Address Tables on the Dhub Switch (Optional), page 3-13. See also Understanding ARP and MAC Address Table Timeouts, page 5-12.

Optical Topology

Figure 2-11 illustrates both the Ethernet and optical layers used for the no-switch-in-Dhub option in Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution.

Note For a discussion of optical components, see Optical Topology, page 2-9.

Figure 2-11 No Switch in Dhub: Optical Topology

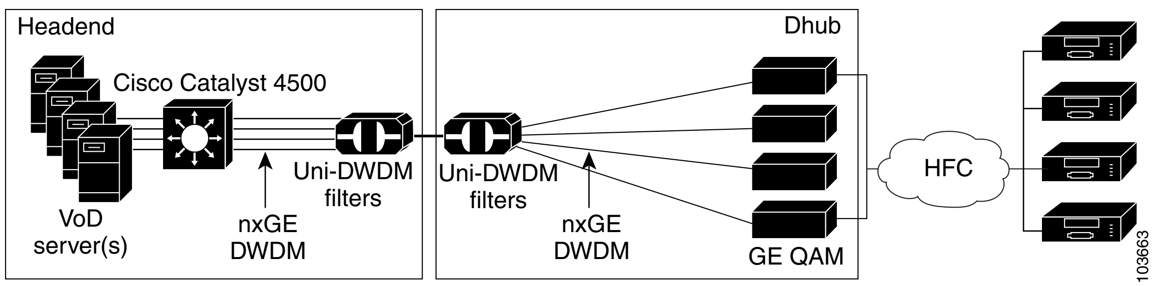

CHAPTER 3

Implementing and Configuring the Solution

Overview

This chapter addresses the two fundamental Ethernet/IP architectures:

- Switch in the Dhub, page 3-1
- No Switch in the Dhub, page 3-15

In addition, resources for implementing the optical architecture are provided below:

- Implementing Optics, page 3-17
- Implementing the Cisco uMG9820 QAM Gateway, page 3-18

Switch in the Dhub

This section addresses the following topics:

- Configuring the Headend: Cisco Catalyst Switches (Basic Configuration)
- Configuring the Dhub: Cisco Catalyst Switches
- Establishing Multiple EtherChannels
- Establishing Subtended Dhubs

For background, see Switch in the Dhub, page 2-1. Figure 3-1 illustrates an example network architecture for the switch-in-Dhub scenario. Extrapolate to add additional servers and switches.

Figure 3-1 Example Network Architecture: Switch in Dhub

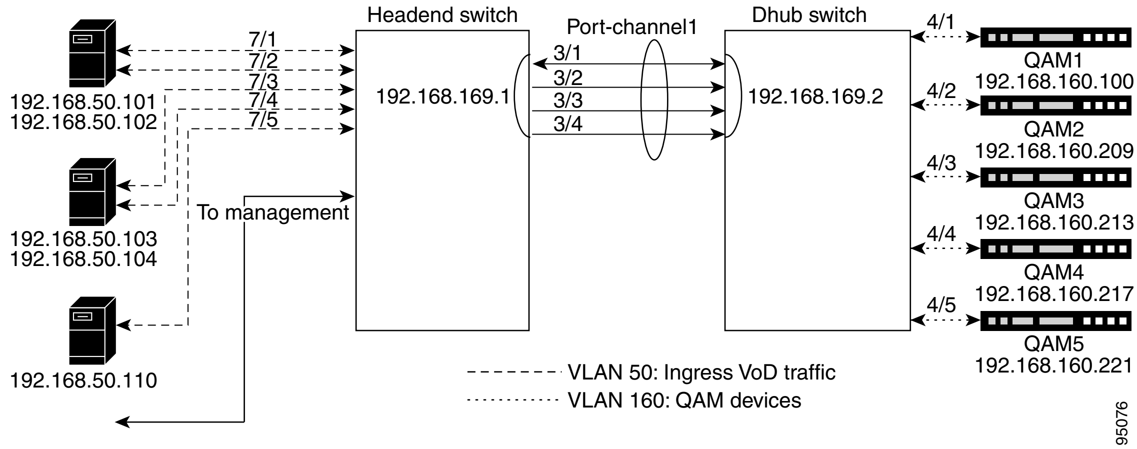

Configuring the Headend: Cisco Catalyst Switches (Basic Configuration)

This section illustrates the nondefault provisioning that is required to establish payload channels between a Cisco Catalyst switch and the VoD servers and QAMs where a switch is used in the Dhub. Switch defaults are not discussed here, but are available in Appendix A, "Switch in Dhub: Sample Configurations for Cisco Catalyst 4500 Series Switches."

Note For switch-related documentation, see Cisco 4000 Family Switches at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/index.htm

The following topics are illustrated:

- Establishing Quality of Service (QoS)
- Enabling OSPF
- Enabling Load Balancing
- Establishing Interfaces on the Headend Switch
- Using and Monitoring HSRP
- Example Configuration on a Headend Switch

Establishing Quality of Service (QoS)

For background, see QoS, page 2-6. Cisco recommends that you tag all video traffic so that it is queued properly and efficiently. The following procedure illustrates the key interrelated components of establishing QoS for video.

Step 1 Establish QoS. This is a global command.

```
qos
```

Step 2 Establish a video class. This requires an access list.

```
class-map match-all class_video
  match access-group name acl_video
```

The extended access-list name **acl_video** is arbitrary. The option **match-all** ensures that all statements in the list are matched.

Step 3 Establish an extended access list.

```
ip access-list extended acl_video
 remark Identify video traffic (UDP ports 257-33023).
 permit udp any any range 257 33023
```

Above port 256, ranges will vary, depending upon the installation.

Note The Harmonic NSG QAM device uses the 2-byte UDP port number to specify the output QAM and program number for an incoming video stream. One of the leftmost eight bits is turned on to specify one of the eight output QAMs. The least-significant bit (LSB) for QAM 1, and the most-significant bit (MSB) is for QAM 8. The rightmost eight bits specify the output program number (range is 1–255). Thus, the access list will use the valid NSG UDP port range to identify video streams (range is 257–33023).

Step 4 Establish a policy map. This exploits the Type of Service (TOS) bits in the video IP packet. You will need to establish Differentiated Services Code Point (DSCP) values.

```
policy-map setDSCP <---This value is arbitrary, and is applied to service-policy on a VLAN
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
```

The parameter setDSCP is arbitrary, and provides an association set by the **service-policy** command, in the VLAN(s) for video traffic. The value 46 turns on the EF (Expedited Forwarding) flag.

The following helps explain the DSCP values and Tx-Queue (transmit queue) mappings, and illustrates how **policy-map** changes the ToS bits for all incoming packets that are in **class_video**. The DSCP bits are the first 6 bits, and the last two (the ECT and CE bits) are ignored.

The incoming video packets from the VoD servers have a DSCP of 7, mapped from hex to binary as follows:

```
!
! Type of service = 1C
!    0 0 0 .  . . . .     = routine
!    . . . 1  . . . .     = low delay
!    . . . .  1 . . .     = high throughput
!    . . . .  . 1 . .     = high reliability
!    . . . .  . . 0 .     = ECT bit - transport protocol will ignore the CE bit
!    . . . .  . . . 0     = CE bit - no congestion
!
```

The following policy-map changes the DSCP to 46.

```
!
! Type of service = B8
!    1 0 1 .  . . . .     = Critic/ECP
!    . . . 1  . . . .     = low delay
!    . . . .  1 . . .     = high throughput
!    . . . .  . 0 . .     = normal reliability
!    . . . .  . . 0 .     = ECT bit - transport protocol will ignore the CE bit
!    . . . .  . . . 0     = CE bit - no congestion
!
```

Here is the default DSCP mapping. The **policy-map** changes the video packets from the DSCP value of 7, which maps to Tx-Queue 1, to a DSCP value of 46, which maps to Tx-Queue 3. Use the command **show qos maps** to view the mapping table and confirm this:

```
! Headend#show qos maps
! DSCP-TxQueue Mapping Table (dscp = d1d2)
! d1 : d2  0  1  2  3  4  5  6  7  8  9
! ------------------------------------
!  0 :    01 01 01 01 01 01 01 01 01 01
!  1 :    01 01 01 01 01 01 02 02 02 02
!  2 :    02 02 02 02 02 02 02 02 02 02
!  3 :    02 02 03 03 03 03 03 03 03 03
!  4 :    03 03 03 03 03 03 03 03 04 04
!  5 :    04 04 04 04 04 04 04 04 04 04
!  6 :    04 04 04 04
```

Enabling OSPF

Enable OSPF, as in the following example.

```
router ospf 100
 log-adjacency-changes
 passive-interface Vlan50
 network 1.1.1.0 0.0.0.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
 distribute-list 1 in Port-channel1
 distribute-list 2 in Port-channel2
 distribute-list 3 in Port-channel3
 distribute-list 4 in Port-channel4
 distribute-list 4 in Port-channel5
```

OSPF will advertise all the headend switch interfaces to the DHubs. This includes the GRE tunnel endpoints (loopbacks) and the VoD server VLAN. The **distribute-list** commands are used to block alternative routes to DHub QAM VLANs from being added to the routing table. In other words, the only path to VLAN160 on DHub A should be through Port-channel1.

Area 0, the default, represents a backbone.

Enabling Load Balancing

Both Layer 3 and EtherChannel load balancing should be used on the headend switch. Default hashing algorithms XOR the source and destination IP addresses. In the VoD application, this produces a small set to be distributed across multiple equal-cost paths. To increase this set and achieve a better distribution, Cisco recommends that you include the Layer 4 destination ports in the hashing algorithm.

Layer 3 Load Balancing

For Layer 3 (CEF IP) load balancing, the Layer 4 ports are included with the following command:

```
ip cef load-sharing algorithm include-ports destination
```

EtherChannel Load Balancing

For EtherChannel load balancing, the Layer 4 ports are used with the following command:

```
port-channel load-balance dst-port
```

Note See Load-Balancing Strategies, page 2-4.

Establishing Interfaces on the Headend Switch

Traffic is segregated naturally through the assignment of interfaces. Corresponding interfaces must also be configured on any Dhubs in the switch-in-Dhub scenario. Interfaces are assigned to VLANs on the switch.

The following interface-related topics are illustrated, in the general order in which they would be applied:

- Establishing a VLAN and Addresses for VoD Server Traffic to the Switch
- Establishing the GE Interfaces to the VoD Servers
- Establishing GE Interfaces from the Headend Switch to the Dhub Switch
- Establishing an Asymmetric EtherChannel to the Dhub Switch
- Establishing Tunnels to a Dhub Switch (Optional)
- Establishing Interfaces for Management (Optional)

Establishing a VLAN and Addresses for VoD Server Traffic to the Switch

The following illustrates the establishment of a VLAN and range of IP addresses for video traffic to the switch. (See Figure 3-1 on page 3-2.) Addressing will vary for multiple servers in disparate networks.

Step 1 First assign a VLAN, then assign an address range to the VLAN.

```
interface Vlan50
 description VoD Servers
 ip address 192.168.50.2 255.255.255.0
 no ip redirects
 no ip unreachables
```

Note The last two commands above prevent the switch from returning ICMP notification messages to servers if a destination IP address is unreachable or has an erroneous destination.

Step 2 Establish a service policy for the VLAN. The following applies values set by the parameter policy-map for TOS. See Establishing Quality of Service (QoS), page 3-2.

```
 service-policy input setDSCP <---Applies values set by policy-map
```

Step 3 Establish virtual addresses.

```
 standby 50 ip 192.168.50.1 <---See Note below
```

Note The line above establishes a virtual IP and MAC address for the VoD VLAN. By convention, the virtual IP address is the first host address in the subnet; the physical interfaces or SVIs (switch VLAN interfaces) are subsequent host addresses.

Be sure to include a group number on this and all standby commands. Otherwise, the group number defaults to 0, resulting in improper operation.

Establishing the GE Interfaces to the VoD Servers

Refer to Figure 3-1 on page 3-2. The following establishes GE interfaces for the video traffic to the VoD servers, interfaces GigabitEthernet7/1, GigabitEthernet7/5, GigabitEthernet7/9, GigabitEthernet7/13, and GigabitEthernet7/17, assigning them to VLAN 50.

Note Slot 7 holds a 24-port linecard that supports 10/100/1000 Mbps over copper. While the optical ports on the switch support only 1000 Mbps, the copper ports have to be set explicitly to this rate.

Caution When using VoD servers that have copper GE interfaces, be sure to space the VoD server connections appropriately. On a 24-port 10/100/1000 Mbps linecard, use one VoD server connection for every four ports on the linecard. On a 14-port 1000 Mbps linecard, use one server connection for every four copper GE ports, and one for each optical GE port.

```
interface GigabitEthernet7/1
 description VODserver1 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 duplex full
 speed 1000 <---See Note above
!
interface GigabitEthernet7/5
 description VODserver1 GigE1
 switchport access vlan 50
 load-interval 30
 speed 1000
!
interface GigabitEthernet7/9
 description VODserver2 GigE0
 switchport access vlan 50
 load-interval 30
 speed 1000
!
interface GigabitEthernet7/13
 description VODserver2 GigE1
 switchport access vlan 50
 load-interval 30
 speed 1000
!
interface GigabitEthernet7/17
 description VODserver3 GigE0
 switchport access vlan 50
 load-interval 30
 speed 1000
```

Establishing GE Interfaces from the Headend Switch to the Dhub Switch

Refer to Figure 3-1 on page 3-2. The following illustrates the establishment of interfaces GigabitEthernet3/1, GigabitEthernet3/2, GigabitEthernet3/3, and GigabitEthernet3/4. GigabitEthernet3/1 is bidirectional, with all others being unidirectional. Corresponding channels will need to be established on the Dhub switch. Once you establish the interfaces, assign them to an EtherChannel as described in Establishing an Asymmetric EtherChannel to the Dhub Switch, page 3-8.

Caution When adding send-only links to EtherChannels that are already carrying video traffic, be sure to configure the receive-side switch and all physical connections before configuring the send-side switch. See Preventing Lost Traffic: Adding Send-Only GE Links to Operational EtherChannels, page 5-12.

```
interface GigabitEthernet3/1
 description Bidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 tx-queue 3 <---See first Note below
   priority high
 channel-group 1 mode on <---See second Note below
!
```

Note There are 64 DSCP (Differentiated Services Code Point) values, which are mapped to four queues on the switch. This ensures that traffic marked EF (video), which is mapped to queue 3, has a higher priority than traffic mapped to the other three queues.

Note The command **channel-group** binds the interface to Port-channel1.

```
interface GigabitEthernet3/2
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only <---See Note below
 channel-group 1 mode on
!
```

Note This forces the GE port up with transmit-only capabilities. This relieves the CPU and improves throughput.

```
interface GigabitEthernet3/3
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/4
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
```

For the corresponding configuration on the Dhub switch, see Establishing GE Interfaces from the Dhub Switch to the Headend Switch, page 3-11.

Establishing an Asymmetric EtherChannel to the Dhub Switch

Once the GE interfaces are configured, they can be assigned to an EtherChannel. The switch-in-Dhub scenario requires the establishment of an asymmetric EtherChannel between the switches, as illustrated below. Refer to Figure 3-1 on page 3-2.

```
interface Port-channel1
 description asymmetric EtherChannel to DHub_A
 ip address 192.168.169.1 255.255.255.252
 load-interval 30
```

The **load-interval** is the measurement interval for interface statistics. (The default **load-interval** is 300.)

For the corresponding configuration on the Dhub switch, see Establishing an Asymmetric EtherChannel from the Dhub Switch to the Headend Switch, page 3-11.

Establishing Tunnels to a Dhub Switch (Optional)

As noted in Ethernet Topology, page 2-1, with this UDLR application tunnels provide the mechanism for creating a return path for a unidirectional interface or EtherChannel. Create loopback interfaces on the headend and Dhub switches to serve as endpoints for the UDLR tunnel. Loopback interfaces (rather than physical interfaces) are used as endpoints of the tunnel, because loopback interfaces never go down, and each tunnel requires it own unique set of endpoints. In this example, the return path for a unidirectional EtherChannel (Port-channel2) will be configured, beginning at the headend switch. For the corresponding configuration on the Dhub switch, see Establishing Tunnels to a Headend Switch (Optional), page 3-12.

Tunnels are used in the following scenarios:

- Establishing Multiple EtherChannels, page 3-14
- Establishing Subtended Dhubs, page 3-14

To establish a tunnel from the headend switch to a Dhub switch:

Step 1 Create a loopback interface for each tunnel. (For numbering recommendations, see Tip below.)

```
interface Loopback2
 description Endpoint of Tunnel2
 ip address 1.1.1.5 255.255.255.255
```

You must create a corresponding destination address to communicate to the same loopback on the Dhub switch, as shown in Step 2.

Tip Loopback1 is not used as Port-channel1 is asymmetrical and does not need a tunnel. For convenience in provisioning and following configurations, it is useful to match the numbers for loopback, tunnel, and Port-channel.

Step 2 Create the tunnel interface:

```
interface Tunnel2
 description Return Path for Unidirectional EtherChannel to DHub_B
 no ip address
 tunnel source 1.1.1.5
 tunnel destination 2.2.2.6
 tunnel udlr receive-only Port-channel2
```

You also must create Tunnel2 on the Dhub switch. See Establishing Tunnels to a Headend Switch (Optional), page 3-12.

Note UDLR tunnels are configured as **receive-only** on the headend switch, and as **send-only** on the Dhub switch.

Step 3 Repeat Step 1 andStep 2 with unique tunnel IDs and IP addresses for each tunnel configured between the headend and the Dhubs. There can be multiple GE interfaces or EtherChannels between the headend and a single Dhub.

Step 4 Configure an OSPF network statement on the headend switch.

```
router ospf 100 network 1.1.1.0 0.0.0.255 area 0
```

Note Because loopback interfaces are configured on both the headend and Dhub switches for use as tunnel endpoints, the Dhub switches need to know how to route to the tunnel endpoints on the headend switch. However, the headend switch does not need to know how to route to the endpoints of the Dhub switches, bacause packets travel only from the Dhub to the headend. You must therefore include the tunnel endpoints in a network statement under the OSPF process.

Establishing Interfaces for Management (Optional)

If your installation supports a management application, simply assign Layer 3 addresses to dedicated northbound and southbound management channels.

Using and Monitoring HSRP

If redundant switches are used, it is important to monitor the status of the failover protocol. You should also use this command to determine the virtual MAC address that should be configured on VoD servers that do not support ARP.

Note Redundancy is discussed in detail in Chapter 4, "Providing Redundancy and Reliability."

The following example illustrates the result of a **show standby** command:

```
Headend#show standby
Vlan50 - Group 50
Local state is Active, priority 100

  Hellotime 3 sec, holdtime 10 sec

  Next hello sent in 0.346
Virtual IP address is 192.168.50.1 configured
Active router is local
Standby router is unknown
Virtual mac address is 0000.0c07.ac32
2 state changes, last state change 3d22h
IP redundancy name is "hsrp-Vl50-50" (default)
```

Example Configuration on a Headend Switch

For a complete example configuration of a headend switch connected to a single Dhub switch over a single EtherChannel, see Single EtherChannel: Headend Switch to Dhub Switch, page A-1.

Configuring the Dhub: Cisco Catalyst Switches

In the switch-in-Dhub architecture, the configuration on the Dhub switch is considerably simpler. Simply maintain the correspondences and VLANs that were established at the headend. This section addresses the following topics:

- Establishing QoS
- Enabling OSPF
- Enabling Load Balancing
- Establishing Interfaces on the Dhub Switch
- Manually Configuring the ARP and MAC Address Tables on the Dhub Switch (Optional)

Establishing QoS

No configuration is required on Dhub switches to perform QoS. For issues that pertain to the headend switch, see QoS, page 2-6.

Enabling OSPF

Cisco recommends that you enable OSPF routing on the switch, to advertise to the headend switch the route to the QAM VLAN, as well as to receive routes from the headend switch. Here we set the PID to 100, the same as in Enabling OSPF, page 3-4.

```
router ospf 100
 log-adjacency-changes
 passive-interface Vlan160 <---See Note below
 network 192.168.0.0 0.0.255.255 area 0
```

Note The QAM devices do not need to participate in OSPF messaging.

Enabling Load Balancing

No configuration is required on the Dhub switch to perform load balancing. For issues that pertain the to the headend switch, see Load-Balancing Strategies, page 2-4.

Establishing Interfaces on the Dhub Switch

The interfaces established on the Dhub switch must naturally correspond to those on the headend switch. Below we illustrate the following:

- Establishing a VLAN and Addresses for Video Traffic to the QAM Devices
- Establishing GE Interfaces from the Dhub Switch to the Headend Switch
- Establishing an Asymmetric EtherChannel from the Dhub Switch to the Headend Switch

Establishing a VLAN and Addresses for Video Traffic to the QAM Devices

The following illustrates the establishment of a VLAN and range of IP addresses for video traffic to the QAM devices. Refer to Figure 3-1 on page 3-2.

```
interface Vlan160
 description QAM Devices
 ip address 192.168.160.1 255.255.255.0
 no ip redirects
 no ip unreachables
```

Note The last two commands above prevent the switch from returning ICMP notification messages to servers if a destination IP address is unreachable or has an erroneous destination.

Establishing GE Interfaces from the Dhub Switch to the Headend Switch

Below we establish Gigabit Ethernet interfaces GigabitEthernet3/1, GigabitEthernet3/2, GigabitEthernet3/3, and GigabitEthernet3/4. Interface GigabitEthernet3/1 is bidirectional, with all others being unidirectional. (Refer to Figure 3-1 on page 3-2.) This corresponds to Establishing GE Interfaces from the Headend Switch to the Dhub Switch, page 3-6. Once you establish the interfaces, assign them to an EtherChannel as described in Establishing an Asymmetric EtherChannel from the Dhub Switch to the Headend Switch, page 3-11.

```
interface GigabitEthernet3/1
 description Bidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/4
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
```

Establishing an Asymmetric EtherChannel from the Dhub Switch to the Headend Switch

Once GE interfaces have been configured, they can be aggregated in an EtherChannel. The switch interface used is the same as that on the headend switch. (Refer to Figure 3-1 on page 3-2.) This corresponds to Establishing an Asymmetric EtherChannel to the Dhub Switch, page 3-8.

Note To accommodate additional traffic, see Load-Balancing Strategies, page 2-4. See also Establishing Multiple EtherChannels, page 3-14.

See Establishing an Asymmetric EtherChannel to the Dhub Switch, page 3-8.

```
interface Port-channel1
 description asymmetric EtherChannel from Headend
 ip address 192.168.169.2 255.255.255.252
 load-interval 30
```

The **load-interval** is the measurement interval for interface statistics. The default **load-interval** is 300.

Establishing Tunnels to a Headend Switch (Optional)

See Establishing Tunnels to a Dhub Switch (Optional), page 3-8. This configuration on the Dhub switch must correspond to that established on the headend switch.

To establish a tunnel from the Dhub switch to the headend switch:

Step 1 Create a loopback interface for each tunnel, corresponding to what is established on the headend switch..

```
interface Loopback2
 description Endpoint of Tunnel2
 ip address 2.2.2.6 255.255.255.255
```

A corresponding destination address, as illustrated below, will need to be created for communicating to the same loopback on the headend switch.

Step 2 Create the tunnel interface:

```
interface Tunnel2
 description Return Path for Unidirectional EtherChannel to Dhub_B
 no ip address
 tunnel source 2.2.2.6
 tunnel destination 1.1.1.5
 tunnel udlr send-only Port-channel2
 tunnel udlr address-resolution
```

This corresponds to Tunnel2 on the headend switch.

Note UDLR tunnels are configured as **send-only** on the Dhub switch, and as **receive-only** on the headend switch. The command **tunnel udlr address-resolution** enables Dhub_B to send ARP responses for Port-channel2 through Tunnel2 to the headend switch.

Step 3 Repeat Step 1 and Step 2 with unique tunnel IDs and IP addresses for each tunnel configured between the Dhub and the headend. There can be multiple GE interfaces or EtherChannels between the headend and a single Dhub.

Manually Configuring the ARP and MAC Address Tables on the Dhub Switch (Optional)

Where the tables must be populated manually, the following lines illustrate the correspondence that will have to be entered, through a text editor, into the configuration file. (See Manual Population of ARP and Bridging Tables on the Dhub Switch, page 2-14.) The addresses of the QAM devices being supported are shown in Figure 3-1 on page 3-2.

ARP Example

The following lines will find their MAC address correspondence in MAC Address Table Example, below:

```
arp 192.168.160.100 0020.a300.92aa ARPA
arp 192.168.160.209 000b.4627.0001 ARPA
arp 192.168.160.213 000b.4627.0002 ARPA
arp 192.168.160.217 000b.4627.0003 ARPA
arp 192.168.160.221 000b.4627.0004 ARPA
```

MAC Address Table Example

The following lines correspond to the ARP lines in ARP Example, above:

```
mac-address-table static 0020.a300.92aa vlan 160 interface GigabitEthernet4/1
mac-address-table static 000b.4627.0001 vlan 160 interface GigabitEthernet4/2
mac-address-table static 000b.4627.0002 vlan 160 interface GigabitEthernet4/3
mac-address-table static 000b.4627.0003 vlan 160 interface GigabitEthernet4/4
mac-address-table static 000b.4627.0004 vlan 160 interface GigabitEthernet4/5
```

Connecting to QAM Devices

The following example illustrates the assigning of interfaces in the Dhub switch to the QAM devices, in a single VLAN.

```
interface GigabitEthernet4/1
 description QAM1
 switchport access vlan 160
 switchport mode access
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet4/2
 description QAM2
 switchport access vlan 160
 switchport mode access
 load-interval 30
!
interface GigabitEthernet4/3
 description QAM3
 switchport access vlan 160
 switchport mode access
 load-interval 30
```

Example Configuration on Dhub Switch

For a complete example configuration of a Dhub switch connected to a single headend switch over a single EtherChannel, see Single EtherChannel: Dhub Switch to Headend Switch, page A-5.

Establishing Multiple EtherChannels

When more than one EtherChannel interface is used between a headend and a Dhub switch, one EtherChannel will need to be asymmetric, to support bidirectional control traffic. The return path for the second EtherChannel uses UDLR over the bidirectional path on the asymmetric EtherChannel, and is established as discussed in Establishing Tunnels to a Dhub Switch (Optional), page 3-8. See Figure 3-2 on page 3-14. See also Load-Balancing Strategies, page 2-4.

Figure 3-2 Switch in Dhub: Multiple EtherChannels

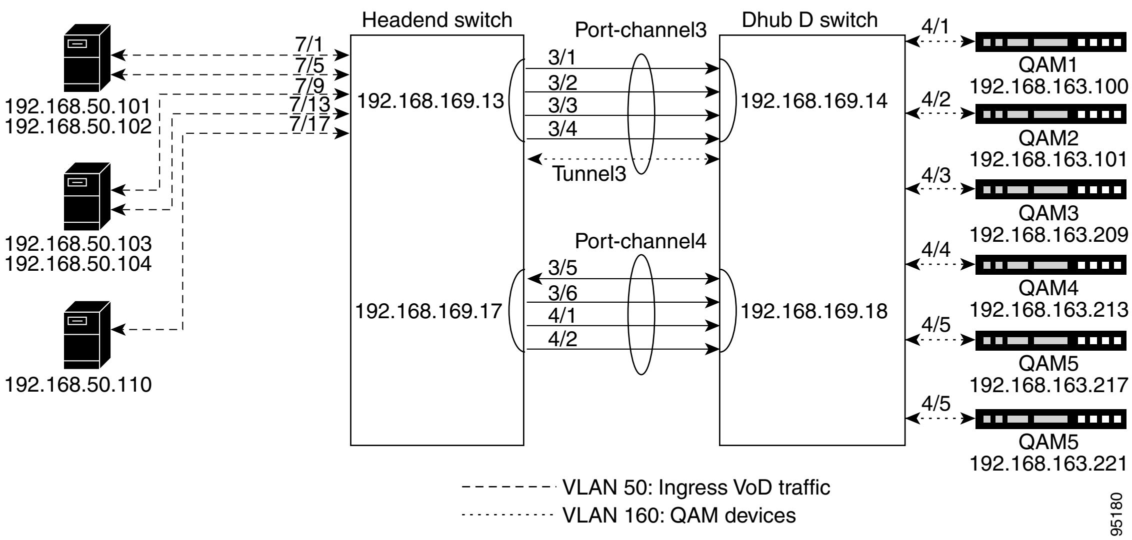

Configuration is as in Configuring the Headend: Cisco Catalyst Switches (Basic Configuration), page 3-2, and Configuring the Dhub: Cisco Catalyst Switches, page 3-10, but with the following differences:

- A tunnel is configured on the headend switch.
  - The tunnel source is configured as a loopback interface on the headend switch.
  - The tunnel destination is configured with the same loopback interface on the Dhub switch.
  - The tunnel is associated with the receive port of an EtherChannel.
- A tunnel is configured on the Dhub switch.
  - The tunnel source is configured as a loopback interface on the Dhub switch.
  - The tunnel destination is configured with the same loopback interface on the Dhub switch.
  - The tunnel is associated with the send port of the same EtherChannel selected above.

For example configurations of the headend and Dhub switch, see the following:

- Multiple EtherChannels: Headend Switch to Dhub Switch, page A-8
- Multiple EtherChannels: Dhub Switch to Headend Switch, page A-12

Establishing Subtended Dhubs

With subtended Dhubs (see Subtended Dhubs, page 2-7), a unidirectional EtherChannel from one Dhub switch is returned through another Dhub switch. See Figure 3-3.

Figure 3-3 Switch in Dhub: Subtended Control Path

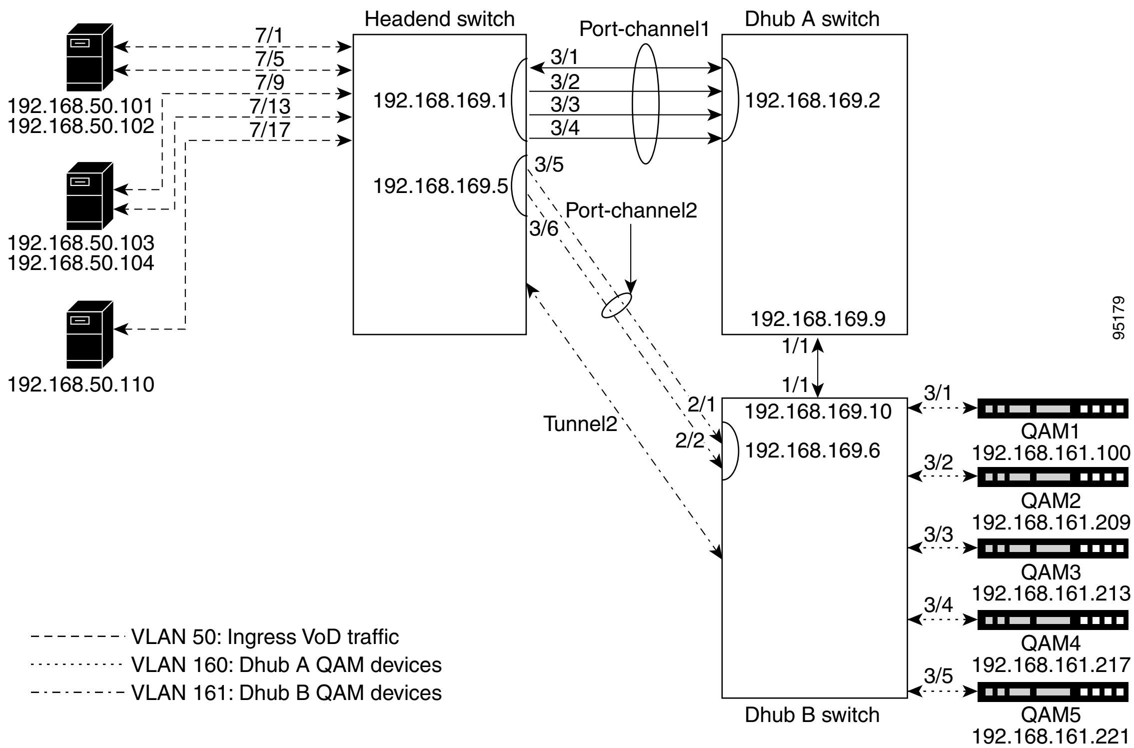

Configuration is as in Configuring the Headend: Cisco Catalyst Switches (Basic Configuration), page 3-2, and Configuring the Dhub: Cisco Catalyst Switches, page 3-10, but with the following differences:

- A tunnel is configured on the headend switch.
  - The source is configured as a loopback interface on the headend switch.
  - The destination is configured with the same loopback interface on the Dhub switch.
  - The tunnel is associated with the receive port of an EtherChannel.
- A tunnel is configured on the Dhub switch.
  - The source is configured as a loopback interface on the Dhub switch.
  - The destination is configured with the same loopback interface on the Dhub switch.
  - The tunnel is associated with the send port of the same EtherChannel selected above.
  - The tunnel is configured to send ARP replies.

For example configurations of the headend and Dhub switch, see the following:

- Subtended Control Path: Headend Switch to Two Dhubs, page A-15
- Subtended Control Path: Dhub A to Headend Switch, page A-19
- Subtended Control Path: Dhub B to Headend Switch, page A-22

No Switch in the Dhub

For background, see No Switch in the Dhub, page 2-11. This section addresses the following topics:

- Establishing a Single VLAN
- Establishing a Split VLAN
- Establishing Multiple VLANs

For sample configurations that illustrate the above, see Appendix B, "No Switch in Dhub: Sample Configurations for Cisco Catalyst 4500 Series Switches."

Note For all of the following VLAN examples, ARP and bridging tables are populated manually, as discussed in Manually Configuring the ARP and MAC Address Tables on the Dhub Switch (Optional), page 3-13.

Establishing a Single VLAN

Figure 3-4 illustrates the implementation of a single VLAN, the simplest case. Here all ports are members of the same VLAN. For a sample switch configuration, see Single VLAN, page B-1.

Figure 3-4 No Switch in Dhub: Single VLAN

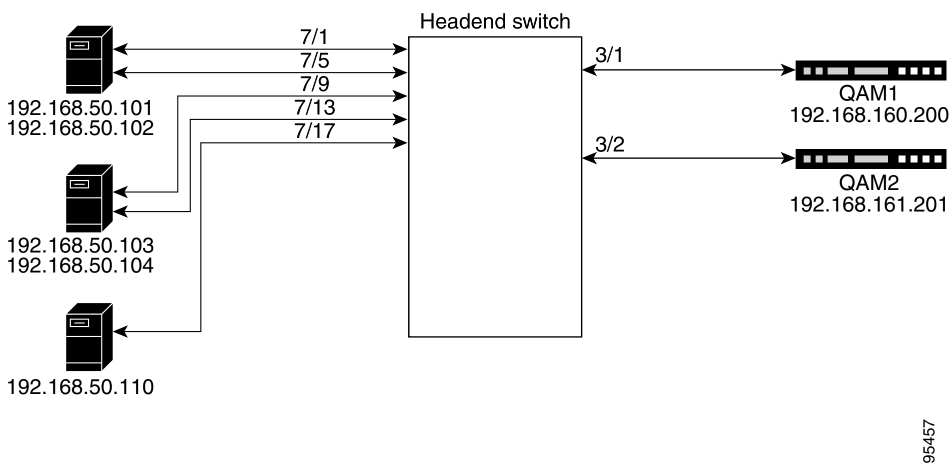

Establishing a Split VLAN

Figure 3-5 illustrates the implementation of a split VLAN, where one VLAN is reserved for ingress VoD traffic (source ports) and another for the QAM devices (destination ports). For a sample switch configuration, see Split VLAN, page B-4.

Figure 3-5 No Switch in Dhub: Split VLAN

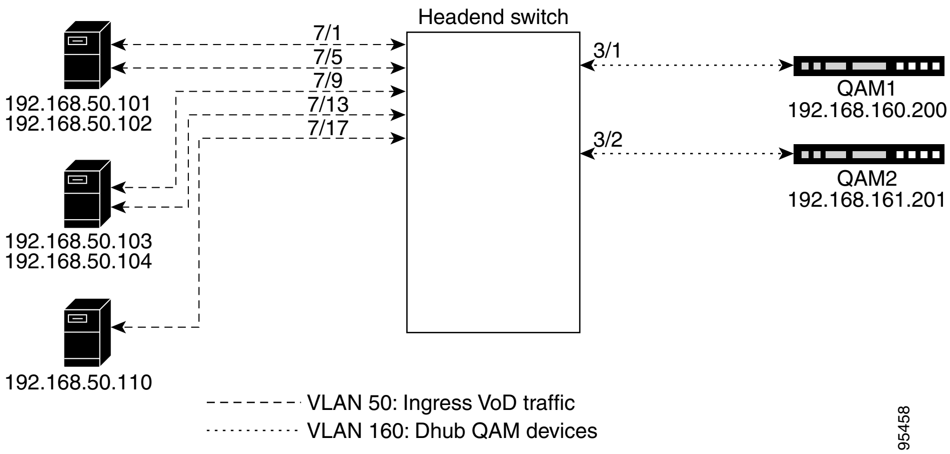

Establishing Multiple VLANs

Figure 3-6 illustrates the implementation of a VLAN for ingress VoD traffic as well as multiple VLANs to the QAM devices. Here, source ports are members of one VLAN, and destination ports are grouped (usually by region) into discrete VLANs. For a sample switch configuration, see Multiple VLANs, page B-7.

Figure 3-6 No Switch in Dhub: Multiple VLANs

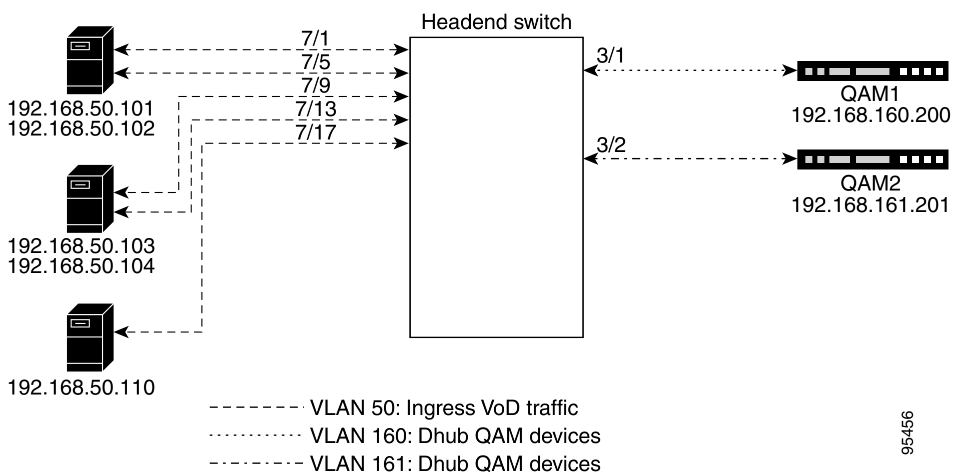

Implementing Optics

The following discussions present a variety of options for implementing the various optics and supervisory channels for the Cisco Gigabit-Ethernet Optimized VoD Solution:

- Implementing the Cisco ONS 15216 FlexLayer
- Implementing the Cisco 15216 OSC-1510

Implementing the Cisco ONS 15216 FlexLayer

The Cisco Gigabit-Ethernet Optimized VoD Solution uses the Cisco ONS 15216 FlexLayer solution to provide modular support for a variety of optical functions. A single chassis accommodates multiplex/demultiplex filters, combiner or splitter assemblies, and optical attenuators, providing for easy and cost-effective expansion.

Note For more about features and the various modules, refer to "Data Sheet: Cisco ONS 15216 Metropolitan Dense Wavelength Division Multiplexing 100-GHz FlexLayer Filter Solution," at the following URL:

http://www.cisco.com/warp/public/cc/pd/olpl/metro/15200/prodlit/flexp_ds.htm

To install and use the Cisco ONS 15216 FlexLayer and its various components, refer to *Cisco ONS 15216 FlexLayer User Guide, Release 1.0*, at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/ong/15216/flxlyr10/

Implementing the Cisco 15216 OSC-1510

The Cisco ONS 15216 OSC-1510 can be used in the Cisco Cisco Gigabit-Ethernet Optimized VoD Solution to provide optical supervisory channel (OSC) communication to a site without the need for a OADM (optical add/drop multiplexer), EDFA (erbium-doped fiber amplifier), or multiplexing/demultiplexing at that site. This passive single-channel 100-GHz device allows you to add or drop a protected OSC wavelength in each direction at any point of a DWDM link. The dropped OSC channel is then sent to the receive GBIC port on the switch.

Note For a description of the Cisco ONS 15216 OSC-1510 and installation instructions, refer to *Cisco ONS 15216 OSC-1510 User Guide* at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/ong/15216/osc.htm

Although that document refers to the Cisco Catalyst 2950 switch, the Cisco ONS 15216 OSC-1510 is compatible with the Cisco Catalyst 4500 series switches.

Implementing the Cisco uMG9820 QAM Gateway

For information on preparing, installing, starting, and configuring the Cisco uMG9820 QAM Gateway, refer to the *Cisco uMG9820 QAM Gateway Installation and Configuration Guide*, at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/cable/vod/umg9820/index.htm

Providing Redundancy and Reliability

The architecture of the network determines the approaches used to ensure service. This chapter addresses remedies for the following topologies:

- Switch in Dhub, page 4-2
- No Switch in Dhub, page 4-12

Fundamental Failure Scenarios

There are five fundamental failure scenarios:

1. A DWDM fiber cut
2. A failure of a passive optical element
3. A failure of a headend switch linecard
4. A failure of a Dhub switch linecard
5. A failure of a headend switch

These scenarios are referred to by number in the following discussion. The redundancy options discussed in this chapter were selected for their ability to deal with any one of these failure modes. Multiple simultaneous failures are not considered. For a discussion of particular link failures and remedies, see Link Failure Modes, page 4-3.

Note Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution provide different redundancy options for the switch-in-Dhub and no-switch-in-Dhub topologies. This is because some of the redundancy options depend on the ability for Layer 3 routing protocols to detect a failure and route around it.

For an illustration of link failure modes and their responses, see Switch in Dhub: Ethernet Troubleshooting Examples, page 5-8.

Note All of the Ethernet topologies tested for Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution are hub and spoke. For a discussion of ring topologies, as well as converting from rings to hub and spoke, see Chapter 6, "Deploying the Cisco Gigabit-Ethernet Optimized VoD Solution in Fiber Ring Topologies."

Switch in Dhub

There are two redundancy options for the switch-in-Dhub topology:

- Headend Switch Protection
- Layer 3 1+1 Protection

These are discussed in detail below.

Headend Switch Protection

Headend switch protection is the most redundant—but most costly— failover mechanism for protection in this solution. This protection mechanism relies on two fibers between headend and Dhub sites, two headend switches, and two sets of optical multiplexers/OSCs. Depending on the capabilities of the VoD server, the VoD streams from the VoD server will either be(1) sent from one of two GE ports on the VoD server, or (2) replicated at the optical layer by means of an optical splitter. On VoD servers that are capable of GE failover, two GE ports will be used on the VoD server, with video being transmitted on one of the two ports. Figure 4-1 illustrates the topology for headend switch protection for VoD servers that support GE failover. Table 1-2 on page 1-8 shows which VoD servers are capable of GE port failover.

Note Headend switch protection applies to all five of the failure modes listed in Fundamental Failure Scenarios, page 4-1.

Figure 4-1 Headend Switch Protection: GE Failover on the VoD Server

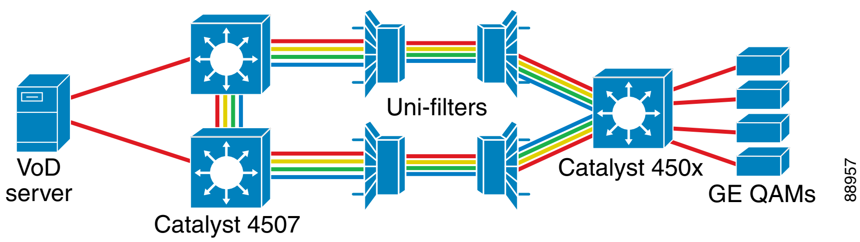

On VoD servers that are not capable of GE failover, an optical splitter can be used to split the optical signal between the VoD server and headend switches. Figure 4-2 illustrates the topology for headend switch protection with optical splitting.

Figure 4-2 Headend Switch Protection: Optical Splitting

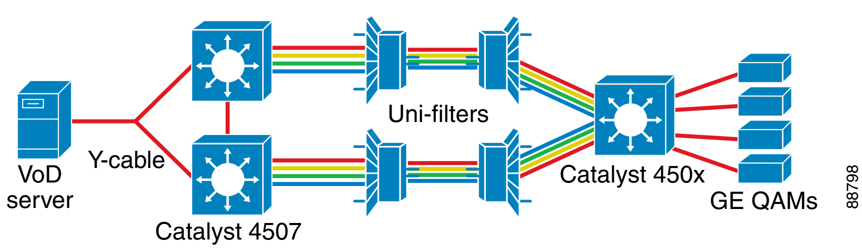

In this failover scenario, the two headend switches are connected to each other through an EtherChannel group that has the same number of optical links as the headend-to-Dhub connections. The same combination of EtherChannel groups and IP interfaces is configured on both the links between the headend switches and the headend-to-Dhub connections. On VoD servers that support failover, the GE ports facing the VoD servers on each of the two headend switches are configured as separate routable interfaces, and an IP routing protocol (OSPF in Release 1.0 and Release 1.1 of the Cisco Gigabit-Ethernet Optimized VoD Solution) is enabled on these interfaces. The VoD servers must be capable of supporting multiple IP subnets; the servers must also support OSPF on the linecards, as OSPF must also be enabled on the VoD server line cards. (See OSPF, page 2-7.)

When either a GE link is cut or a headend switch fails, OSPF routes traffic from one set of GE links to another. The OSPF cost metric must be configured to be higher for one VoD server's GE link than for the other. The lower-cost GE link will be the primary GE link, and the higher-cost link will be the standby link.

On VoD servers that do not support GE failover, the two headend switches must run HSRP between the GE interfaces connected to the VoD server. (See Using and Monitoring HSRP, page 3-9.) One switch will be the active HSRP node, and the other will be the standby node. The standby HSRP node will have its GE port disabled, so no video will be passed through it.

OSPF must be enabled on each EtherChannel interface on the headend switches, as well as on the Dhub switch. The headend switches and the Dhub switches for all Dhubs connected to the headend must share the same OSPF process ID (PID). On VoD servers that support GE failover, the GE interfaces on the VoD server must also share the same OSPF PID.

Link Failure Modes

The following sections show how the system works under the five scenarios listed in Fundamental Failure Scenarios, page 4-1. Those modes can be expanded to take into account two basic types of link failure:

- Failures that affect all links between a headend and a Dhub switch
- Failures that affect a subset of links between those switches, as shown in Table 4-1.

Table 4-1 Detailed Summary of Link Failure Modes

Failure Category	Link Failure Mode	Notes
Affects all links between headend and Dhub switch	**1.** DWDM fiber cut	Loss of the bidirectional link brings down the entire EtherChannel.
	2. Failure of some passive optical network elements	
	3. Failure of some Dhub switch linecards	
	4. Failure of some headend switch linecards	See Link Failure Mode 4, page 4-4.
Affects a subset of links between headend and Dhub switches	**5.** Failed GBIC or GE fiber cut	See Link Failure Modes 5, 6, and 7, page 4-4.
	6. Failure of some passive optical network elements	
	7. Failure of some Dhub switch linecards	
Affects some links between headend and Dhub switches	**8.** Failure of some headend switch linecards	See Link Failure Mode 8, page 4-4.
Failure of headend switch	**9.** Failure of headend switch	See Link Failure Mode 9, page 4-4.

Details are addressed in the following sections.

Link Failure Mode 4

If the all of the ports connecting the headend and Dhub switches fail, the headend switch immediately learns that the ports are down. OSPF reroutes traffic through the alternate headend switch to the Dhub switch.

Link Failure Modes 5, 6, and 7

The behavior of the network in these three scenarios depends on whether the GE ports that are affected (1) are part of an EtherChannel group, or (2) are physical Layer 3 interfaces that are part of an IP load-balancing group.

In case (1), the streams carried on those ports will be lost. EtherChannel will not redistribute the load if a unidirectional link is cut, because the port-layer keep-alive protocol is disabled for unidirectional links. Because of this, the headend EtherChannel logic will not recognize the loss of a member EtherChannel link, and will therefore not redistribute the EtherChannel load across the remaining member links. When new video sessions are initiated, sessions whose packets are hashed to the failed port(s) will be dropped at the headend switch. There is no automated failover for this type of scenario.

In case (2), the Dhub switch will learn immediately that the interfaces are down. However, UDLR will not relay this information to the feed side of the interface. (See UDLR, page 2-6.) Since OSPF is a unidirectional routing protocol, OSPF link-state updates will not cause the headend switch to remove the failed interfaces from the switch's RIB (routing information base). In OSPF, the mechanism used to inform the headend switch that its unidirectional feed interface is down is the hello protocol. In normal OSPF operation, the hello protocol is not used to detect link failure. To get the OSPF hello protocol to detect a link failure in the headend switch in a timely manner, OSPF fast hello must be enabled in Cisco IOS. OSPF fast hellos cause the feed interface to send a hello packet periodically to the Dhub switch, which then returns the packet to the feed interface on any directly connected return link. When the fiber is cut, OSPF hellos are not received by the Dhub switch, and so are not be returned to the headend switch.

The failure of the headend switch to receive hello packets causes the feed interface to go down on the that switch. Consequently, IP load balancing will redistribute the load across the remaining interfaces in the IP load-balancing group. For reasons of economy, the traffic load on an interface is often at or near 100%, and a link failure with this level of oversubscription can result in a loss of video on all interfaces. To prevent this type of failure from affecting video, the capacity of the IP load-balancing group can be overprovisioned by using *N* additional GE links, to provide for *N* port failures.

Link Failure Mode 8

If a subset of the ports connecting the two headend switches fail at the headend switch, either the EtherChannel logic or OSPF detects the failure immediately—depending on whether the GE ports that are affected are part of an EtherChannel group or are physical Layer 3 interfaces that are part of an IP load-balancing group. In either case, the EtherChannel or IP load-balancing logic will redistribute the load across the remaining links in the load-balancing group. To prevent this type of failure from affecting video, the capacity of the EtherChannel or IP load-balancing group can be over provisioned by using *N* additional GE links to take into account *N* port failures.

Link Failure Mode 9

In topologies where the VoD servers do not support GE failover, HSRP will detect the failure by means of the HSRP keep alive protocol. (See HSRP Redundancy Scenarios, page 4-5.) When HSRP detects the failure on the standby switch, it switches the GE interface(s) on the standby switch from standby to active. IP routing on this headend switch should have determined that the lowest-cost path to the QAM devices is through a directly connected switch interface. As a result, video streams will be routed to the Dhub switch through this interface.

In topologies where VoD servers support GE failover, a headend switch failure is detected as a set of interface failures by OSPF on the standby headend switch and the VoD server. OSPF reroutes video streams from the primary headend switch to the standby headend switch.

HSRP Redundancy Scenarios

This section presents a detailed illustration of three switch-in-Dhub HSRP redundancy scenarios tested as part of the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.0 and Realease 1.1.

- Switch in Dhub: HSRP without UDLR
- Switch in Dhub: HSRP with UDLR
- Switch in Dhub: HSRP with Interface Tracking

Note For sample configurations for the above, see Appendix C, "Switch in Dhub: Sample Redundancy Configurations for Cisco Catalyst 4500 Series Switches."

Switch in Dhub: HSRP without UDLR

Figure 4-3 illustrates a scenario where HSRP is used without UDLR.

Figure 4-3 Switch in Dhub: HSRP without UDLR

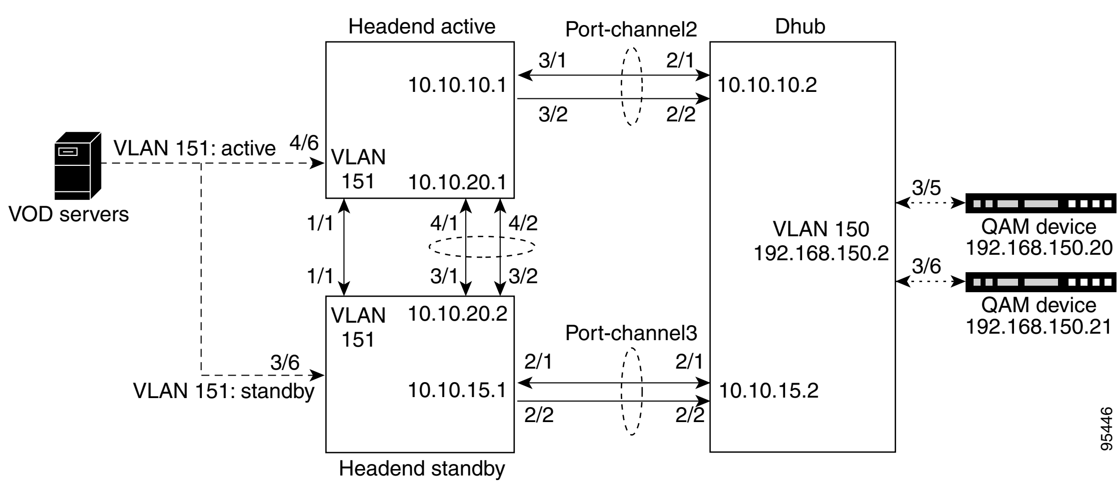

Scenario Characteristics

Note the following key characteristics of this scenario:

- An asymmetric EtherChannel (Port-channel) is required between the primary (active) headend switch and the Dhub, as well as between the standby headend switch and the Dhub.
  - There is a bidirectional EtherChannel between both the active and standby headend switches.

- A return fiber is required for both headend switches.

- OSPF is used to provide link redundancy.
- VoD server input is to a receive-only interface on both headend switches. (A 2:1 optical splitter can be used to provide identical feeds to both switches.)
- A separate VLAN is used (151 in this example), on which HSRP is active. Configuration excerpts are show below for both headend switches. (The virtual **standby ip** address is shared.)

```
Primary Headend Switch: Configuration Excerpt

ip address 192.168.151.100 255.255.255.0
 no ip redirects
 service-policy input setDSCP
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 100
 standby preempt

Standby Headend Switch: Configuration Excerpt

interface Vlan151
 ip address 192.168.151.101 255.255.255.0
 no ip redirects
 logging event link-status
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 10
```

Note For more information, refer to Using HSRP for Fault-Tolerant IP Routing at the following URL:

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs009.htm

- An additional link (GE 1/1 on both switches) is required between the primary and standby headend switches on VLAN 151, so they can pass ICMP standby messages. An access control list (ACL) is used to ensure that video traffic is not passed on this link, allowing only ICMP messages to go out. The following is an ACL example:

```
ip access-list extended acl_standby_link
 remark Stop video traffic on standby_link (g1/1)
 permit icmp any any
```

Failure Recovery Time

Table 4-2 on page 4-7 lists a variety of failures and time to recover from them for this scenario.

Note See also Switch in Dhub: Ethernet Troubleshooting Examples, page 5-8.

Table 4-2 Switch in Dhub: Recovery Times for HSRP without UDLR Scenario

Failed Components	Action	Recovery Time, sec	Comments
DWDM fiber	Pulled DWDM fiber between headend and Dhub without interface tracking activated. (See Switch in Dhub: HSRP with Interface Tracking, page 4-8.)	5–7	This is the time for OSPF to switch to backup path.
Dhub switch linecard	Pulled linecard serving Port-channel1 to Dhub.		
Headend switch linecard			
GBIC or GE fiber on bidirectional link	Pulled GBIC or GE fiber (same effect).	5–7	
GBIC or GE fiber on unidirectional link		No recovery.	
Headend switch	Turned power off.	2	Preempting the headend switch resulted in a 7-sec outage.

Configuration Examples

For sample configurations for this scenario, see HSRP without UDLR, page C-1.

Switch in Dhub: HSRP with UDLR

Figure 4-4 illustrates where HSRP is used with UDLR (unidirectional link routing).

Note For a background and configuration examples, refer to Configuring Unidirectional Link Routing at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcpt3/1cfudlr.htm

For UDLR command syntax, refer to Unidirectional Link Routing Commands at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122sup/122csum/csum1/122csip3/1sfudlr.htm

Figure 4-4 Switch in Dhub: HSRP with UDLR

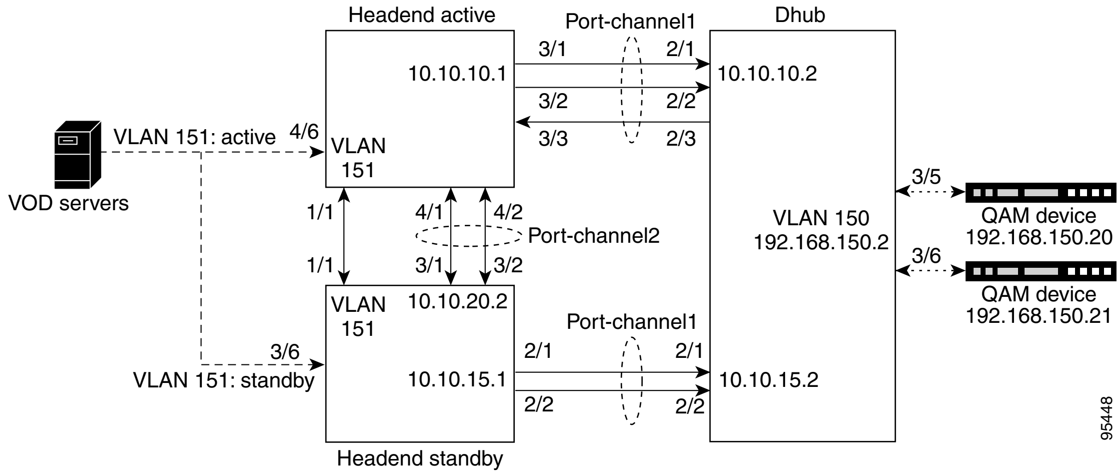

Scenario Characteristics

Note the following key characteristics of this scenario:

- This configuration requires one return fiber.
  - UDLR provides a return path for the transmit-only EtherChannel (Port-channel1 between the standby headend switch and the Dhub switch).
  - OSPF is used for link redundancy.
- Port-channel1 between the primary (active) headend switch and the Dhub switch consists of three physical ports:
  - Two are transmit-only.
  - One is receive-only (from the point of view of the headend switch).
  - A loss of the transmit fiber will not take down the EtherChannel to the Dhub (Port-channel1 from the standby headend switch).
- If the headend receive fiber fails, routing goes down for both active and standby headend switches.

Failure Recovery Time

Time to recover from a link outage is 5–7 seconds.

Configuration Examples

For sample configurations for this scenario, see HSRP with UDLR, page C-9.

Switch in Dhub: HSRP with Interface Tracking

Figure 4-5 illustrates HSRP between the headend and Dhub switches, with the HSRP interface tracking feature applied.

Figure 4-5 Switch in Dhub: HSRP with Interface Tracking

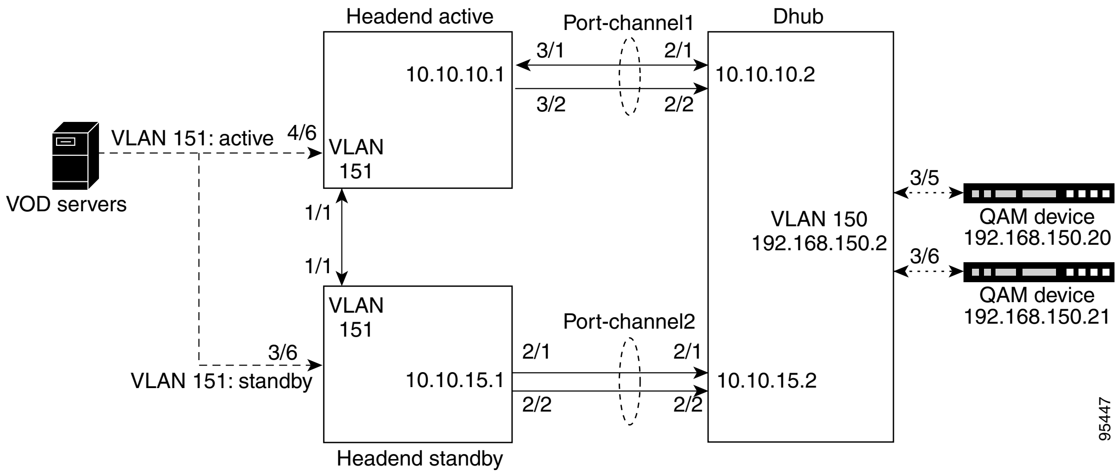

Scenario Characteristics

Note the following key characteristics of this scenario:

- The HSRP interface tracking feature detects a link failure between the headend switch and the Dhub switch. This reduces recovery time in the event of a link outage.
- Interface tracking is applied to interface GE 3/1, the bidirectional link of Port-channel1.
- OSPF is still used to learn routes to the Dhub switch. However, it is not used for redundancy.
- There is no EtherChannel between the primary (active) and standby headend switches.
- An ACL on the primary switch prevents video traffic from going down GE 1/1 to the standby switch. The same ACL is also configured on the standby switch.

Failure Recovery Time

If Port-channel1 fails, recovery time is 1–2 seconds.

Configuration Examples

For sample configurations for this scenario, see HSRP with Interface Tracking, page C-17.

Layer 3 1+1 Protection

Layer 3 1+1 protection covers four of the five failure scenarios, but is less costly than headend switch protection. Failure scenarios 1 through 4 are covered, but the failure of a headend switch itself is not covered.

This protection mechanism is identical to headend switch protection, except that it has a single headend switch. Because of this, the redundant interfaces connecting the headend and Dhub switches are both connected to the same headend switch.

Figure 4-6 Switch in Dhub: Layer 3 1+1 Protection

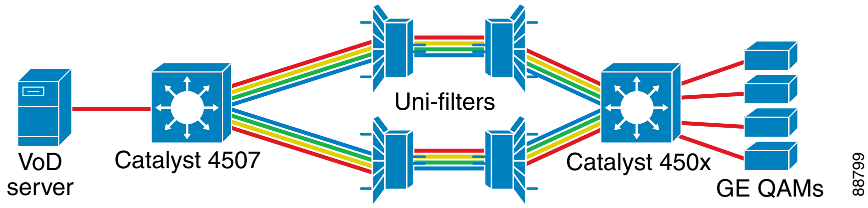

The configuration of the headend and Dhub switches, and the network operation for each scenario, are identical to that in Headend Switch Protection, page 4-2, except that HSRP redundancy is not enabled (or needed) on the single headend switch.

Table 4-3 summarizes the protections available, with failover times, for switch-in-Dhub failure scenarios.

Table 4-3 Switch in Dhub: Summary of Protections

	Headend Switch Protection	**Layer 3 1+1 Protection**
Failure Scenarios covered	Fiber cut	Fiber cut
	Failure of a passive optical element	Failure of a passive optical element
	Failure of a headend switch linecard	Failure of a headend switch linecard
	Failure of a Dhub switch linecard	Failure of a Dhub switch linecard
	Failure of a headend switch	
FailoverTime Range	1 msec – 1 sec	1 msec – 1 sec

Layer 3 1+1 Protection Scheme

This section illustrates an example Layer 3 1+1 protection scheme for use in switch-in-Dhub scenarios. Figure 4-7 illustrates the optical topology for the specific case of two Cisco Catalyst 4507 switches (running Supervisor Engine IV) with WS-4306 interface cards and transmit and receive GBICS. Figure 4-8 illustrates the IP topology for the above.

Figure 4-7 Switch in Dhub: Layer 3 1+1 Protection—Optical Topology

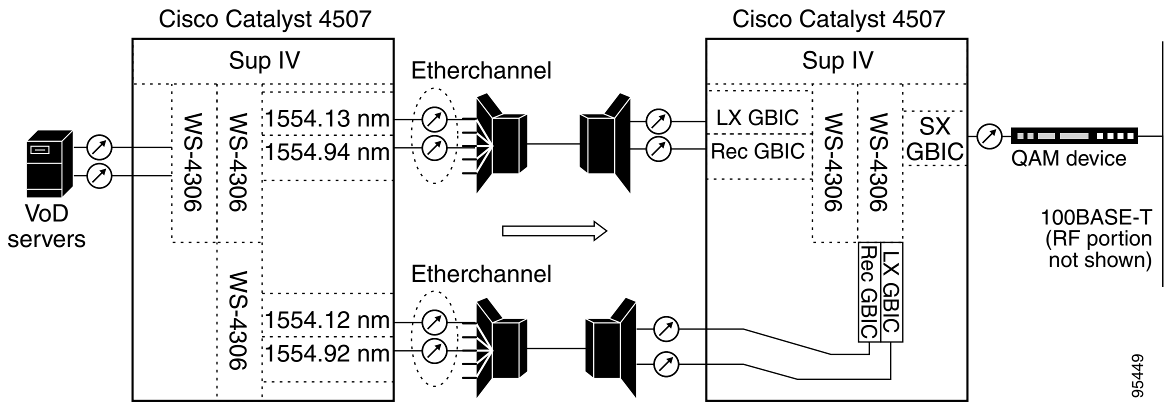

Figure 4-8 Switch in Dhub: Layer 1 1+1 Protection—IP Topology

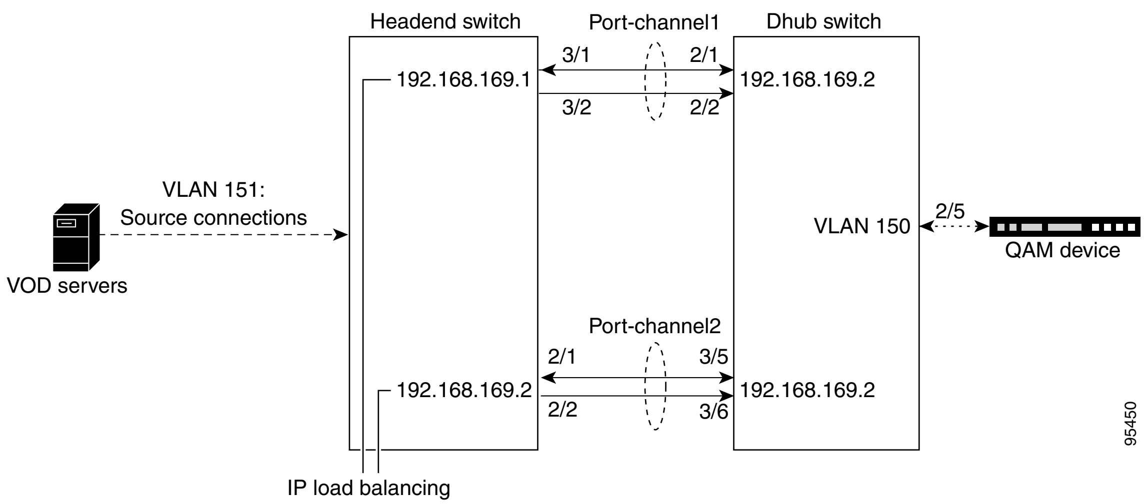

Scenario Characteristics

Note the following key characteristics of this scenario:

- OSPF provides redundancy between Port-channel1 and Port-channel2.
- IP load balancing is used on the above EtherChannels.

Failure Recovery Time

In the event of a link failure, recovery time is 5–7 seconds.

Configuration Examples

For sample configurations for this scenario, see Layer 3 1+1 Protection, page C-22.

No Switch in Dhub

There are two redundancy options for the no-switch-in-Dhub topology:

- Y-cable Optical Protection
- Fiber Splitter Redundancy

These are discussed in detail below.

Note Recovery time for redundant link failures is less than 1 second.

Y-cable Optical Protection

Y-cable optical protection uses an optical splitter in the Dhub to send the same optical signal through a redundant DWDM network to two GE interfaces at the QAM module. To support this method of protection, the QAM must support optical failover between the two interfaces. (See Gigabit Ethernet QAM Devices, page 1-8.)

Because this failover method duplicates all of the components of the DWDM network, it protects against any failure in the network. The failover time depends on how quickly the QAMs can detect LOS (loss of signal) on the primary GE port and switch to the backup GE port as a result. The detection of LOS on a standard GE port is typically on the order of 5 msec.

Figure 4-9 No Switch in Dhub: Y-cable Optical Protection

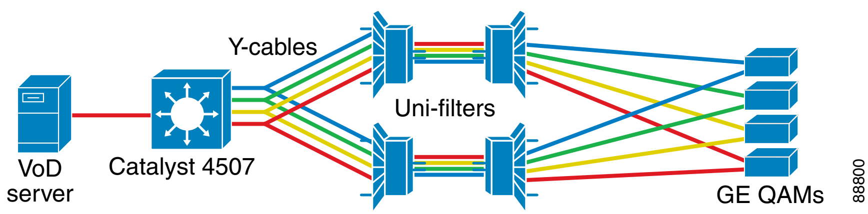

Fiber Splitter Redundancy

Fiber splitter redundancy works identically to Y-cable redundancy, except that it splits the optical signal after the DWDM multiplexer at the headend. Since the signal is split after the DWDM multiplexer, only one DWDM multiplexer is needed at the headend. This redundancy scheme trades off the ability to deal with failures in optical components at the headend for the savings of not having to duplicate optical components in the headend.

This method relies on the same optical failover mechanism in the QAMs that Y-cable redundancy uses. As a result, the recovery times will be the same for these two protection schemes.

Table 4-4 on page 4-13 summarizes the protections available, with failover times, for no-switch-in-Dhub failure scenarios.

Table 4-4 No Switch in Dhub: Summary of Protections

	Y-cable Optical Protection	**Fiber Splitter Redundancy**
Failure Scenarios Covered	Fiber cut	Fiber cut
	Failure of a passive optical element	
Failover Time Range	< 1 sec	< 1 sec

CHAPTER 5

Monitoring and Troubleshooting

This chapter provides an introduction to monitoring and troubleshooting the Cisco Ethernet switches in the Cisco Gigabit-Ethernet Optimized VoD Solution, Release 1.0 and Release 1.1, and presents the following major topics:

- Using Basic CLI Commands to Troubleshoot the Switch, page 5-1
- Switch in Dhub: Ethernet Troubleshooting Examples, page 5-8
- More Troubleshooting Tips, page 5-11

Using Basic CLI Commands to Troubleshoot the Switch

The following commands are addressed:

- show interface (and variations: gigabitethernet, vlan, port-channel, and tunnel)
- logging event link-status

show interface

The command **show interface**, without a specific interface as an option, retrieves information from every active interface. Below are the counters that are relevant to troubleshooting:

- *input errors*—This is a count of any errors that occurred while the switch is trying to receive packets from the referenced port.

  The counter includes both **CRC** (cyclic redundancy check) and **frame errors**, but does not include ignored packets. CRC errors occur when the received packets fail the CRC. Frame errors occur when the receiving frame is not complete. The **ignored** counter counts the number of frames dropped on input because of resource exhaustion in the switch fabric. **Overruns** occur when inter frame gaps (IFGs) are so short that a new Ethernet frame arrives before the previous frame has been completely stored in shared memory.

- *output errors*—This is a count of any errors that occurred while the switch is trying to transmit packets from the referenced port.

  Collisions shows the number of times a collision occurred while the switch is trying to transmit a packet from the referenced port. This counter should be 0 for a port operating in full-duplex mode. The **interface resets** counter counts the number of times the port reset itself, generally the result of link up or link down transitions. **Underruns** occur when packets are not retrieved quickly enough from shared memory to be transmitted.

- *babbles* and *late collision*—A **babble** is an error caused by the transmission of frames in excess of 1518 bytes in size.

  A **late collision** is a collision that occurs outside of the collision window, which is typically caused by a duplex mismatch or a wire length that exceeds the distance limitations (100 meters for 10/100BASE-T ports). The **deferred** counter tabulates the number of times the port had to wait to transmit as a result of traffic on the wire.

- *lost carrier* and *no carrier*—The carrier is an electrical signal that Ethernet devices use to detect whether the wire is currently being used by another transmitting station.

  The **lost carrier** counter increases each time a carrier sense loss occurs. This happens when the hardware is transmitting a frame onto the wire and does not see its own carrier wave on the Ethernet. The absence of the carrier signal increments the **no carrier** counter.

Sample Output for show interfaces on a Headend Switch

The following basic CLI **show interface** commands can provide a great deal of information of assistance in troubleshooting Cisco Catalyst 4500 series switches. The examples below show nominal values for **show interfaces** on a headend switch for a single GE port, a VLAN, and an EtherChannel, respectively:

- show interfaces gigabitethernet
- show interfaces vlan
- show interfaces port-channel

Note The VLAN, VLAN50, is in Port-channel2, which contains two GE interfaces—3/5 and 3/6 in this example.

show interfaces gigabitethernet

```
Headend#show interfaces gigabitethernet 3/5
GigabitEthernet3/5 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet Port, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: Unidirectional Link of EtherChannel 2
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 41/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is force-up, media type is SX
  output flow-control is off, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 162032000 bits/sec, 14826 packets/sec
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts (0 multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     835379461 packets output, 1137771157239 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
```

```
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out

Headend#show interfaces gigabitethernet 3/6
GigabitEthernet3/6 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet Port, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: Unidirectional Link of EtherChannel 2
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 40/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is force-up, media type is SX
  output flow-control is off, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 158527000 bits/sec, 14498 packets/sec
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts (0 multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     835489311 packets output, 1137915654913 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

show interfaces vlan

```
Headend#show interfaces vlan50
Vlan50 is up, line protocol is up
  Hardware is Ethernet SVI, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: VoD Servers
  Internet address is 192.168.50.2/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 155/255
  Encapsulation ARPA, loopback not set
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:06:04, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 610762000 bits/sec, 56805 packets/sec
  30 second output rate 0 bits/sec, 0 packets/sec
  L3 in Switched: ucast: 3340981987 pkt, 4490274590976 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 4012 pkt, 192576 bytes - mcast: 0 pkt, 0 bytes
     3340981999 packets input, 4490274591912 bytes, 0 no buffer
     Received 12 broadcasts (0 IP multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     24215 packets output, 1163224 bytes, 0 underruns
     0 output errors, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

show interfaces port-channel

```
Headend#show interfaces port-channel 2
Port-channel2 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: Unidirectional EtherChannel to DHub_B
  Internet address is 192.168.169.5/30
  MTU 1500 bytes, BW 2000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 40/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Unknown duplex, Unknown Speed, media type is unknown media type
  output flow-control is unsupported, input flow-control is unsupported
  Members in this channel: Gi3/5 Gi3/6
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 314028000 bits/sec, 28505 packets/sec
  L3 in Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 1668713589 pkt, 2242751063616 bytes - mcast: 0 pkt, 0 bytes
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     1668774184 packets output, 2272833746238 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

show interfaces tunnel

This command is illustrated for both headend and Dhub switches in Troubleshooting Tunnels, page 5-4.

logging event link-status

This command is useful in providing the up/down status of links, sending messages to the console when the status of a link changes. This command can be used in conjunction with management applications such as Cisco Info Center (CIC) (which can pick up console messages and perform a notification process), as there is less latency with this command than there is with MIBs and traps. However, take into account that **logging event link-status** should not be used unless it is necessary, as logging can burden the CPU, especially if traps are also being used.

Troubleshooting Tunnels

This section covers both working and nonworking unidirectional EtherChannels:

- Working Unidirectional EtherChannel Between Headend and DHub
- Nonworking Unidirectional EtherChannel Between Headend and DHub

Note Lines of interest are highlighted in **bold**.

Working Unidirectional EtherChannel Between Headend and DHub

For reference, see Figure 3-3 on page 3-15. The packets on Tunnel2 flow from DHub_B to Headend. Use the command **show interfaces tunnel2** on both switches to verify that packets are sent by the DHub_B switch through Tunnel2 and received by the Headend switch. Outputs from the following commands are illustrated:

- show interfaces tunnel
- debug tunnel
- show ip ospf neighbor
- show arp

show interfaces tunnel

At the Headend

```
Headend#show interfaces Tunnel2
Tunnel2 is up, line protocol is down
  Hardware is Tunnel
  Description: Return Path for Unidirectional EtherChannel to DHub_B
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 1.1.1.5, destination 2.2.2.6
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
  Checksumming of packets disabled, fast tunneling enabled
  Tunnel is receive-only UDLR tunnel for send-only interface Port-channel2
  Last input 00:00:01, output never, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     2 packets input, 212 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

At the Dhub

```
DHub_B#show interfaces Tunnel2
Tunnel2 is up, line protocol is up
  Hardware is Tunnel
  Description: Return Path for Unidirectional EtherChannel from Headend
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 2.2.2.6, destination 1.1.1.5
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
```

```
  Checksumming of packets disabled, fast tunneling enabled
  Tunnel is send-only UDLR tunnel for receive-only interface Port-channel2
  Last input never, output 00:00:00, output hang never
  Last clearing of "show interface" counters 00:00:17
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     2 packets output, 184 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

debug tunnel

Use the command **debug tunnel** on both switches to see packets travel through the tunnel.

At the Headend

```
Headend#debug tunnel
Tunnel Interface debugging is on
Headend#
3d18h: Tunnel2: GRE/IP to classify 2.2.2.6->1.1.1.5 (len=92 ttl=253 tos=0x0) ok,
retval=0x0
3d18h: Tunnel2: GRE/IP to decaps 2.2.2.6->1.1.1.5 (len=92 ttl=252)
```

At the Dhub

```
DHub_B#debug tunnel
Tunnel Interface debugging is on
DHub_B#
2w5d: Tunnel2: GRE/IP encapsulated 2.2.2.6->1.1.1.5 (linktype=7, len=92)
```

show ip ospf neighbor

If the unidirectional EtherChannel and tunnel are working properly, DHub_B will be recognized as Headend's OSPF neighbor.

At the Headend

```
Headend#show ip ospf neighbor Port-channel2

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.6           1   FULL/DR         00:00:30    192.168.169.6   Port-channel2
```

show arp

There will also be an ARP entry for DHub_B's side of Port-channel2 in Headend's ARP table.

At the Headend

```
Headend#show arp | include Port-channel2
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  192.168.169.5           -   000c.850e.80bf  ARPA   Port-channel2
Internet  192.168.169.6          13   0009.e89a.2dff  ARPA   Port-channel2
```

Nonworking Unidirectional EtherChannel Between Headend and DHub

For reference, see Figure 3-3 on page 3-15. When Tunnel2 is down because of a physical connection problem, the line protocol on the DHub_B side will be "down" and the number of output drops should increment. Outputs from the following commands are illustrated:

- show interfaces tunnel
- debug tunnel
- show ip ospf
- show arp

show interfaces tunnel

At the Dhub

```
DHub_B#show interfaces Tunnel2
Tunnel2 is up, line protocol is down
  Hardware is Tunnel
  Description: Return Path for Unidirectional EtherChannel from Headend
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 2.2.2.6, destination 1.1.1.5
  Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
  Checksumming of packets disabled, fast tunneling enabled
  Tunnel is send-only UDLR tunnel for receive-only interface Port-channel2
  Last input never, output 00:01:12, output hang never
  Last clearing of "show interface" counters 00:14:49
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 30
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     135 packets output, 13584 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

debug tunnel

If you use **debug tunnel** again, you will see that no packets are being sent through the tunnel.

At the Headend

```
Headend#debug tunnel
Tunnel Interface debugging is on
Headend#
```

Note There is no output for **debug tunnel**.

At the Dhub

```
DHub_B#debug tunnel
Tunnel Interface debugging is on
DHub_B#
```

Note There is no output for **debug tunnel**.

show ip ospf

If you look at Headend's OSPF neighbors, you will no longer see DHub_B listed as a neighbor.

At the Headend

```
Headend#show ip ospf neighbor Port-channel2
```

Note DHub_B is no longer listed.

show arp

The ARP entry for DHub_B's side of Port-channel2 might still be in Headend's ARP table, but it will be aging. If you execute a **clear arp**, the ARP entry will be removed from the ARP table and will not be added until the tunnel is reestablished.

At the Headend

```
Headend#show arp | include Port-channel2
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  192.168.169.5           -   000c.850e.80bf  ARPA   Port-channel2
Internet  192.168.169.6          21   0009.e89a.2dff  ARPA   Port-channel2

Headend#clear arp

Headend#show arp | include Port-channel2
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  192.168.169.5           -   000c.850e.80bf  ARPA   Port-channel2
```

Troubleshooting Cisco Catalyst 4500 Series Switches

For additional tips on troubleshooting Cisco Catalyst 4500 series switches, see "Troubleshooting Tips" at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/trbl_ja.htm

Switch in Dhub: Ethernet Troubleshooting Examples

The following example illustrates a configuration in which links are broken between a headend and a Dhub switch connected by two asymmetric EtherChannels. This scenario only supports 2 Gbps of video traffic, as Port-channel2 is used for link redundancy.

Figure 5-1 illustrates the basic topology. The response of the switches to the various failure modes provides useful information about anticipated behavior, as well as clues to troubleshooting should such failures occur.

Note For a discussion of link failure modes, see Chapter 4, "Providing Redundancy and Reliability."

Figure 5-1 Basic Ethernet Troubleshooting Topology

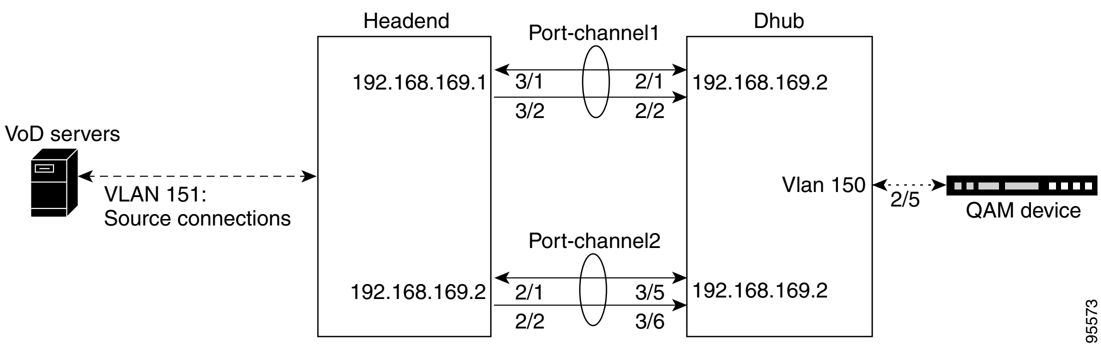

Preconditions

- There are two EtherChannels: Port-channel1 (P1) and Port-channel2 (P2).
- OSPF is enabled, with equal-cost OSPF routes for P1 and P2.
- Load balancing is applied to the destination ports for both EtherChannels with CEF (Cisco Express Forwarding). This causes video flows to be distributed evenly across all four GE ports.

IP Routes

Route tables for the headend and Dhub switches are presented below.

Headend Route Table

```
C    192.168.151.0/24 is directly connected, Vlan151
O    192.168.150.0/24 [110/2] via 192.168.170.2, 00:00:32, Port-channel2
                      [110/2] via 192.168.169.2, 00:00:32, Port-channel1
     192.168.170.0/30 is subnetted, 1 subnets
C       192.168.170.0 is directly connected, Port-channel2
     192.168.169.0/30 is subnetted, 1 subnets
C       192.168.169.0 is directly connected, Port-channel1
```

Dhub RouteTable

```
O    192.168.151.0/24 [110/2] via 192.168.170.1, 00:01:44, Port-channel2
                      [110/2] via 192.168.169.1, 00:01:44, Port-channel1
C    192.168.150.0/24 is directly connected, Vlan150
     192.168.170.0/30 is subnetted, 1 subnets
C       192.168.170.0 is directly connected, Port-channel2
     192.168.169.0/30 is subnetted, 1 subnets
C       192.168.169.0 is directly connected, Port-channel1
```

Actions and Responses

1. Pull the fiber on the send-only GE of P1 (GE 2/2 on the Dhub and GE 3/1 on the headend).
   - Sessions on that port are lost and never recover. EtherChannel load balancing does not rebalance on a loss of a send-only link.
   - GE 2/2 shows down/down on the Dhub:

```
2d21h: %EC-5-UNBUNDLE: Interface GigabitEthernet2/2 left the port-channel
Port-channel1
```

   - P1 on the Dhub now only shows GE 2/1 as a member.
   - No event is shown on the headend. The disabled send-only link is still reported as up/up.
   - No alarms are noted on the QAM device, although the missing sessions are removed from the output extraction page.
2. Replace the fiber on the send-only GE of P1.
   - Sessions recover within 2 seconds.
3. Pull the fiber on the bidirectional GE of P1 (GE 2/1 on the Dhub and GE 3/1 on the headend).
   - Video is momentarily rebalanced to the unidirectional link (evidenced by “blinking”), but after a few seconds the entire EtherChannel goes down as OSPF rerouting takes over.
   - No sessions are lost.
   - On the Dhub, GE 2/1 is down/down; P1 is up/up and shows GE 2/2 as member.
   - On the headend, GE 3/1 down/down; P1 is up/up and shows G3/2 as member
   - The Dhub log reads as follows:

```
2d21h: %EC-5-UNBUNDLE: Interface GigabitEthernet2/1 left the port-channel
Port-channel1
2d21h: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet2/1, changed
state to down
2d21h: %LINK-3-UPDOWN: Interface GigabitEthernet2/1, changed state to down
```

   - The headend log reads as follows:

```
Jun 13 05:19:52.174: %EC-5-UNBUNDLE: Interface GigabitEthernet3/1 left the
port-channel Port-channel1
*Jun 13 05:19:53.174: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet3/1, changed state to down
*Jun 13 05:19:54.174: %LINK-3-UPDOWN: Interface GigabitEthernet3/1, changed state
to down
*Jun 13 05:19:54.302: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.170.2 on
Port-channel1 from FULL to DOWN, Neighbor Down: Dead timer d
*Jun
```

   - P1 is no longer in the **sh ip route** for 192.168.150.0.
   - No alarms are on the QAM device.
   - All traffic now flows through P2.

Note A **sh int g2/2** on the Dhub shows a large output rate for this receive-only link. This is erroneous, as there is no output, as shown below:

```
30 second output rate 4084883554000 bits/sec, 1651326426 packets/sec
```

4. Pull the P1 DWDM fiber.
   - Recovery time is 7 seconds
   - P1 is down/down in the Dhub, but is up/up in the headend.
   - The IP route now has P2 as the only link for 192.168.150.0.
   - The Dhub log reads as follows:

```
00:09:45: %EC-5-UNBUNDLE: Interface GigabitEthernet2/1 left the port-channel
Port-channel1
00:09:45: %EC-5-UNBUNDLE: Interface GigabitEthernet2/2 left the port-channel
Port-channel1
00:09:46: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet2/1,
changed state to down
00:09:46: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel1, changed
state to down
00:09:47: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.170.1 on Port-channel1 from
FULL to DOWN, Neighbor Down: Dead timer expired
00:09:47: %LINK-3-UPDOWN: Interface GigabitEthernet2/1, changed state to down
00:09:47: %LINK-3-UPDOWN: Interface Port-channel1, changed state to down
```

   - The headend log reads as follows:

```
*Jun 13 05:51:29.331: %EC-5-UNBUNDLE: Interface GigabitEthernet3/1 left the
port-channel Port-channel1
*Jun 13 05:51:30.335: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet3/1, changed state to down
*Jun 13 05:51:31.331: %LINK-3-UPDOWN: Interface GigabitEthernet3/1, changed state
to down
*Jun 13 05:51:31.955: %OSPF-5-ADJCHG: Process 100, Nbr 192.168.170.2 on
Port-channel1 from FULL to DOWN, Neighbor Down: Dead timer d
```

5. Pull the P2 DWDM fiber.
   - The results are similar to those for pulling the P1 DWDM fiber.
6. Remove the receive-only GBIC or pull the associated fiber (same effect).
   - All video streams on the failed port are lost.
   - There is no recovery from streams lost on the failed port.
   - The Dhub log reads as follows:

```
2w1d: %EC-5-UNBUNDLE: Interface GigabitEthernet2/2 left the port-channel
Port-channel1
```

More Troubleshooting Tips

This following duscussions can be of assistance in troubleshooting:

- Example Link Utilization and Related Traffic Statistics
- Understanding the Load-Balancing Behavior of Asymmetric EtherChannels
- Preventing Lost Traffic: Adding Send-Only GE Links to Operational EtherChannels
- Understanding ARP and MAC Address Table Timeouts
- Understanding Buffers on Cisco Catalyst 4500 Series Switches

Example Link Utilization and Related Traffic Statistics

As a benchmark for troubleshooting links that support video streams, it can be useful to see some nominal traffic statistics for various degrees of link utilization. As the number of video streams increases, so does link utilization—and also latency. Table 5-1 lists example link utilization and related traffic statistics for a Concurrent VoD server. Numbers for other servers may vary.

Table 5-1 Example Link Utilization and Related Traffic Statistics for a Concurrent VoD Server

Link Utilization, percent	Number of Video Streams	Arrival Time, msec			Latency, microsec		Packets	
		Minimum	Average	Maximum	Minimum	Maximum	Out of Order	Dropped
88	240	0.84	2.82	4.41	21.8	968.5	0	0
90	245	0.84	2.82	4.41	21.8	1369.7	0	0
92	250	0.84	2.82	4.41	21.8	1068.0	0	0
93	255	0.84	2.82	4.41	—	—	—	—
95	260	—	—	—	—	—	—	—

Understanding the Load-Balancing Behavior of Asymmetric EtherChannels

The load balancing behavior of an EtherChannel with both bidirectional and send-only links is as follows:

- If a bidirectional link is lost, the flows on that link are rebalanced to the other links. If this link is the only bidirectional link in the EtherChannel, then it is possible that OSPF will lose this route and force the traffic to a new route.
- If a send-only link is lost, the flows on that link are not rebalanced to the other links.

Preventing Lost Traffic: Adding Send-Only GE Links to Operational EtherChannels

When adding send-only links to EtherChannels that are already carrying video traffic, be sure to configure the receive-side (Dhub) switch and all physical connections before the configuring the send-side (headend) switch. If the send-side switch is configured first, then the send-only link is forced up (even though the receive-side switch is not ready to receive), and traffic will be load-balanced onto that link. This causes that traffic to be lost.

Understanding ARP and MAC Address Table Timeouts

Because the video data flow in VoD is unidirectional, the MAC address table for each GE port connected to an edge QAM device eventually times out, causing data bound for that device to be flooded to all ports. To prevent this, set the ARP timeout for that port (or VLAN) to a value less than the value for mac-address-table aging-time. Because the default value of **aging-time** is 300 (seconds), the ARP timeout can be set to 150, as follows: **arp timeout 150**.

Understanding Buffers on Cisco Catalyst 4500 Series Switches

The shared-memory switching buffer for Cisco Catalyst 4500 series switches running Supervisor Engines 2, 3, and 4 is 16 Mbytes. Unlike with many Cisco platforms, there is no command to change or tune the buffers on Cisco Catalyst 4500 series switches. The various transmit-queue buffer sizes are shown in Table 5-2.

Table 5-2 Transmit-Queue Buffer Sizes for Cisco Catalyst 4500 Series Switches

Paquets per Queue	Packets per Queue on a Subport	Packets on a Nonblocking GE Port	
		Per Queue	Per Port
240	960	1920	7680

Deploying the Cisco Gigabit-Ethernet Optimized VoD Solution in Fiber Ring Topologies

This chapter describes how the hub-and-spoke architecture of the Cisco Gigabit-Ethernet Optimized VoD Solution fits into networks where existing fiber has already been deployed in ring topologies. Fiber rings are very common in MSOs (multiple system operators) that have historically used SONET for video transport and now want to upgrade their transport infrastructure using the same fiber rings.

A hub-and-spoke model based on Gigabit Ethernet (GE) cannot be used in a fiber-ring topology unless the fiber ring is converted into a logical hub-and-spoke topology at the optical layer. Because of the additional optical complexity of using GE, it may seem simpler and more cost-effective to deploy a ring-based packet switching technology such as DPT (Dynamic Packet Transport) directly on the fiber ring.

The information presented in this chapter addresses only the use of physical dark-fiber media and *x*WDM methods.

Note For more information about DPT, refer to Dynamic Packet Transport at the following URL:

http://www.cisco.com/warp/public/cc/techno/wnty/dpty/index.shtml

This chapter presents the following major topics:

- Ring-Based Transport vs. Gigabit Ethernet, page 6-2

  Discusses the benefits and drawbacks of ring-based packet-switching technologies such as DPT compared to GE in VoD and video broadcast networks.

- Converting Fiber Rings to Hub-and-Spoke Gigabit Ethernet, page 6-6

  Describes the basic principles behind converting a fiber ring into a logical hub-and-spoke topology at the optical layer.

Note The information presented in this chapter is outside the scope of solution testing for Release 1.0 and Release 1.1 of the solution.

Ring-Based Transport vs. Gigabit Ethernet

It is straightforward to deploy ring-based packet-transport technologies such as DPT in dual fiber-ring topologies, because the DPT interfaces can be used to drive dark fiber directly. As noted previously, with GE the fiber ring must be converted into a logical hub-and-spoke topology at the optical layer. This is an advantage for ring-based packet-transport technologies over GE, because less passive and active optical equipment needs to be deployed.

In the event of a fiber cut, the DPT network layer also has the ability to reroute traffic on the ring in less than 50 msec. Since the Cisco Gigabit-Ethernet Optimized VoD Solution uses Gigabit EtherChannel, the network-layer convergence times in the event of a fiber cut are similar to those of DPT. An EtherChannel interface is capable of detecting a fiber cut and re-load-balancing the traffic on the remaining members of an EtherChannel group in less than 10 msec. Consequently, both DPT and GE have similar network-layer convergence times in the event of a fiber cut.

Switching Efficiency

Ring-based packet transport technologies also be compared to Gigabit Ethernet by comparing the relative switching efficiency of each technology using different traffic patterns. The switching efficiency for a particular traffic pattern in a transport network can be expressed in terms of the ratio of the traffic entering or exiting the network to the total switching capacity of the network. To determine the relative cost of a particular transport technology with a given traffic pattern (assuming the cost to switch a unit of bandwidth in a particular switching technology is linear), the switching efficiency can be multiplied by the cost to switch a unit of bandwidth in a given transport technology. The analysis used in this section assumes that the cost to switch a unit of bandwidth in ring-based packet transport technologies is the same as that for Gigabit Ethernet. Given this assumption, the relative switching cost can be determined by using the relative switching efficiency of each technology used for video distribution.

Ring-Based Packet-Transport Technologies

Ring-based packet-transport technologies are most bandwidth-efficient when traffic distribution around the ring is even. Traffic distribution around a ring will be most even when the distribution pattern of traffic entering and leaving the ring at each ring node is random. When traffic is distributed evenly around a ring, each switching node around the ring (the DPT interfaces) will be switching the same amount of traffic on the ring, off the ring, and around the ring. Since the switching capacity of each node in ring-based packet-transport technologies such as DPT must be the same, a random traffic distribution pattern is the best in terms of switching efficiency. The switching efficiency of a ring network that uses a random traffic-distribution pattern is approximately 1.

The most inefficient traffic-distribution pattern for ring-based packet-transport technologies occurs when all traffic originates or terminates at one node on the ring. Unfortunately, the traffic in typical video distribution systems originates from a single network node—the headend. Consequently, the traffic pattern associated with video distribution matches this model.

With a video-distribution traffic pattern, the headend node will be switching all traffic onto the ring. This traffic will then be switched off the ring by one of the other nodes (that is, a Dhub). Because the switching capacity of each node must be the same for packet transport technologies such as DPT, this traffic-distribution pattern uses switch resources inefficiently. If the number of nodes is expressed by the variable N, the switching efficiency E of a ring network for video distribution can be expressed (in a common spreadsheet or C-based routine) by Eq. (1):

Eq. (1) $E = (N - 1) / (\text{Roundup}((N - 1) / 2, 1) * N)$

where Roundup() is Roundup(number, num_digits). Taking the value of the first parameter to the value of the second parameter (the value to round up to), this rounds a number upwards (away from zero).

Eq. (1) shows that as the number of nodes in a ring network used for video distribution goes up, the switching efficiency of the network goes down. For video distribution, a ring network with three nodes on the ring is 66% efficient, while a ring network with five nodes on the ring is only 40% efficient. Note that Eq. (1) takes into account the fact that ring networks such as DPT networks are built to transport bidirectional traffic, and so have equal switching capacity in both directions around the ring.

Figure 6-1 illustrates how traffic flows, with the resulting switching efficiency, when a DPT ring is used for video distribution. Here the ring must be provisioned to carry two units of bandwidth in each direction, or four units of switching capacity per node. This is because the headend switch will switch two units of bandwidth in both directions. Because there are five nodes in the ring, the total switching capacity of the ring is 4*5 or 20 units of bandwidth. The amount of traffic that enters or leaves the network can be determined by whether traffic is sourced or dropped at each node on the ring. This amounts to eight units of bandwidth in our example. Therefore, the switching efficiency *E* of a five-node network used for video distribution is 8/20 = 40%.

Figure 6-1 Switching Efficiency for a Five-Node Video Distribution Ring

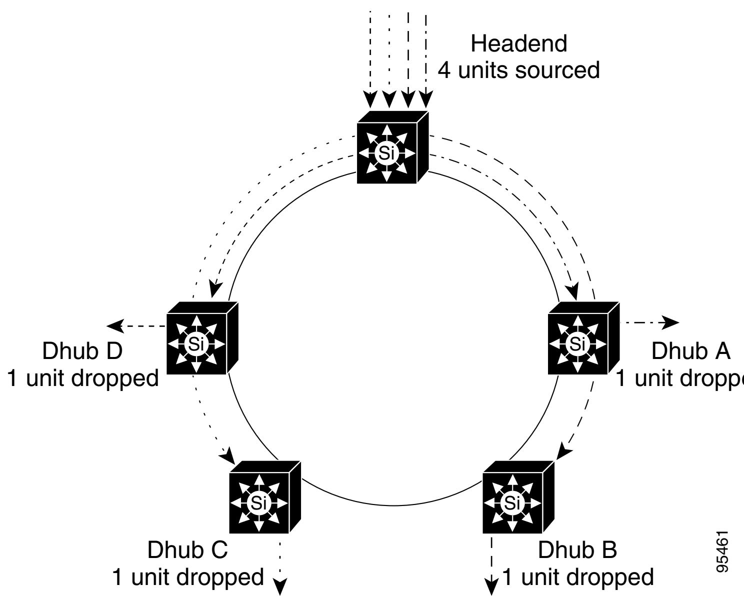

Note that the switching efficiency of a ring-based packet-transport network used for video distribution may change, depending on the distribution of traffic from the headend to the Dhubs. When the traffic from the headend is evenly split in both directions, the ring must be provisioned for the worst-case amount of traffic originating from the headend. A totally uneven traffic distribution can result in all traffic from the headend flowing in one direction only. Eq. (1) made the assumption that the amount of traffic sent from the headend to each Dhub is the same. Removing this assumption, Eq. (2) shows the worst-case switching efficiency *E* for *N* nodes in a ring-based packet-network for video:

Eq. (2) $E = 1 / (N - 1)$

Gigabit Ethernet

When Gigabit Ethernet is deployed in a hub-and-spoke topology, it is most efficient when all traffic originates or terminates at the hub source. This matches the traffic pattern for video distribution, and is the reason why hub-and-spoke GE is well-suited for video distribution—although the switching efficiency for the network will still be less than 1.

As shown in Eq. (3), switching efficiency E in a single-tier hub-and-spoke GE network can be determined by taking the ratio of the bandwidth sourced at the headend, B_H, to the sum of the bandwidths B_{Dn} that are provisioned between the headend and each Dhub:

Eq. (3) $$E = B_H / (B_{D1} + B_{D2} + B_{D3} + \ldots B_{Dn})$$

The bandwidth provisioned at the headend and on each Dhub link will be determined by the MSO, and is typically based on a probability model that takes into account the following:

- The number of customers in the Dhub who subscribe to a particular service such as VoD
- A mathematical model that takes into account customer tuning patterns.

Both Zipf and binomial distributions have been used to model the tuning behavior of customers. A general law of probability states that as the size of a sample population increases, the ratio of peak usage to average usage within that population decreases. This is illustrated below.

Figure 6-2 shows a binomial distribution with an average probability of 10% and a population size of 50. Here the ratio of the maximum (12.5) to the average (5) for a population of 50 subscribers is approximately 2.5.

Figure 6-2 Binomial Probability Distribution for a Population Sample of 50

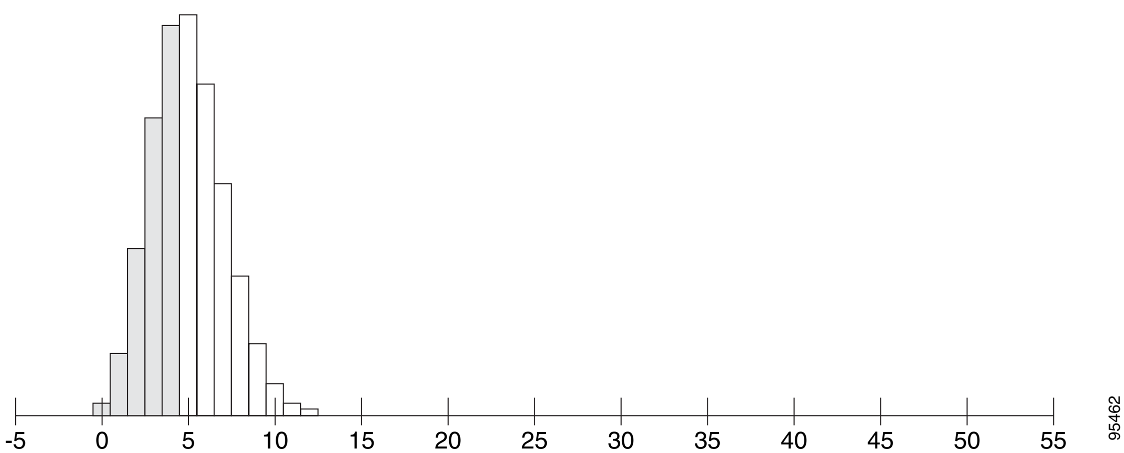

Figure 6-3 on page 6-5 shows that the ratio of the maximum (25) to the average (15) for a population of 150 subscribers is approximately 1.66.

Figure 6-3 Binomial Probability Distribution for a Population Sample of 150

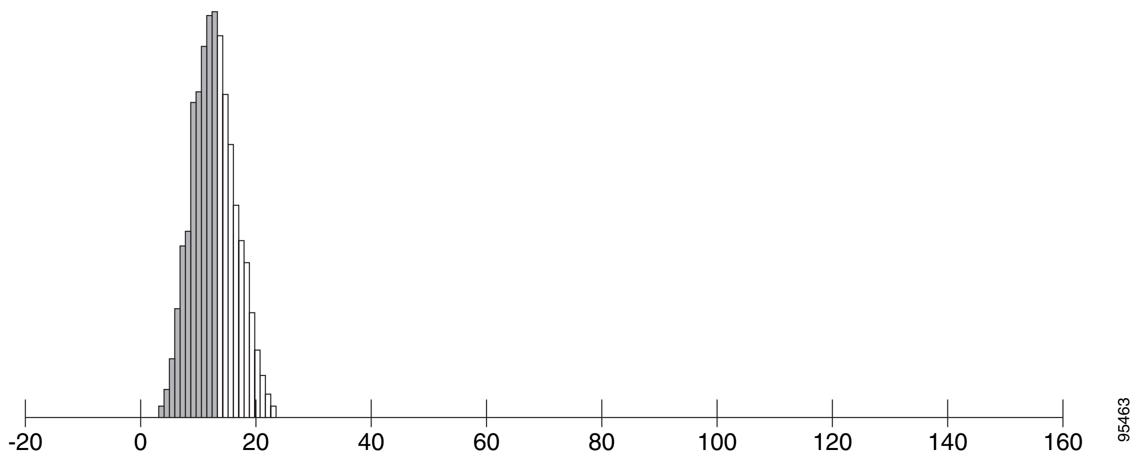

When MSOs provision VoD bandwidth in a headend, they apply this probability analysis to a subscriber population that is the sum of all subscribers in all Dhubs connected to the headend. To provision the bandwidth for each Dhub link, they will use the probability analysis described above on the subscriber population of that Dhub.

If analysis results of the above examples were used for a headend serving three Dhubs of 50 subscribers each, then the bandwidth that would be provisioned on each of the Dhub links from Figure 6-2 on page 6-4 would be 12.5 units, while the amount of bandwidth that would be provisioned in the headend from Figure 6-3 would be 25 units. Using Eq. (3), the switching efficiency of this network would be 25 / (12.5 + 12.5 + 12.5) = 66%. Figure 6-4 illustrates the traffic flows and resulting provisioning numbers that result from this analysis.

Figure 6-4 Hub-and-Spoke Efficiency Model

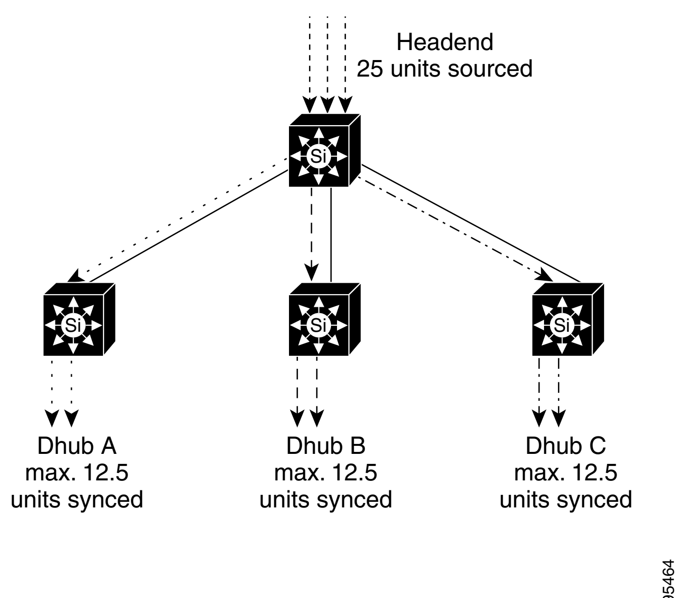

Note that as the population of subscribers using a hub-and-spoke model increases, the bandwidth efficiency of the network will increase. This is because of the nature of binomial distributions. As the sample size of a binomial distribution increases, the ratio of average to peak with that distribution decreases.

Converting Fiber Rings to Hub-and-Spoke Gigabit Ethernet

In order to run hub-and-spoke Gigabit Ethernet in an environment where existing fiber is already deployed in rings, the fiber ring must be converted to a logical hub-and-spoke network by means of optical splitters and combiners and DWDM filters. This section provides an example of how to provide optical hub-and-spoke connectivity over a physical fiber laid out in a ring.

Figure 6-5 shows an example Cisco Gigabit-Ethernet Optimized VoD Solution (Release 1.0 and Release 1.1) GE topology that can be built over a fiber ring. Figure 6-6 on page 6-7 shows an example mapping of the headend and Dhub switches shown in Figure 6-5 to nodes on a physical fiber ring. Fiber rings typically consist of two or more fibers. This section shows how to convert one fiber into the downstream hub-and-spoke optical path.

Figure 6-5 Cisco Gigabit-Ethernet Optimized VoD Solution (Release 1.0 and Release 1.1) Gigabit Ethernet Topology

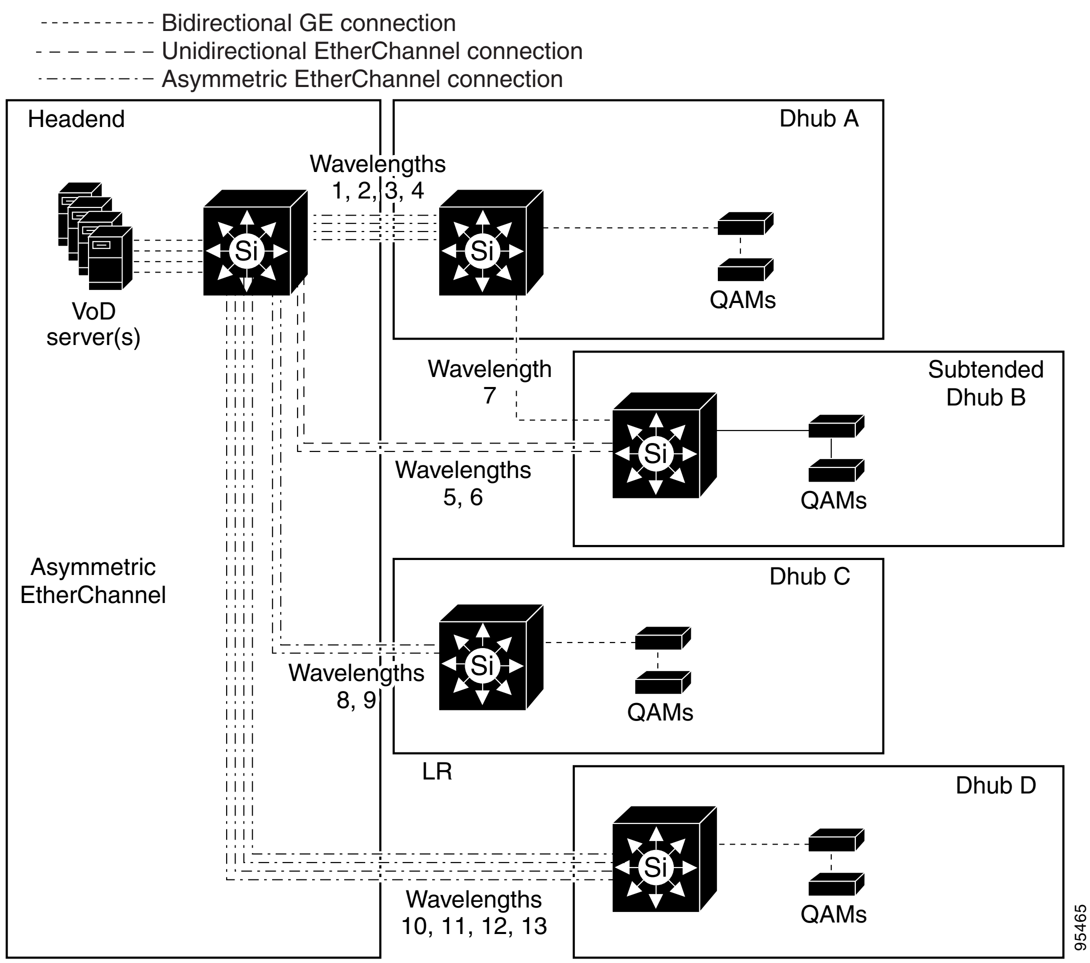

Figure 6-6 Cisco Gigabit-Ethernet Optimized VoD Solution 1.0 Fiber Topology

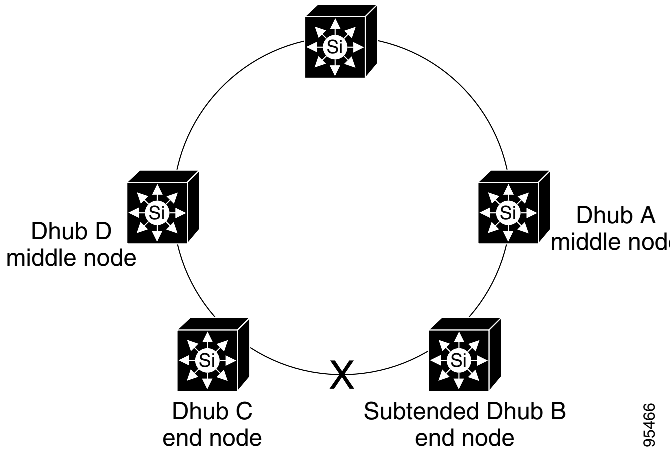

To convert the fiber topology shown above to a logical hub-and-spoke topology, the ring must first be broken at a point that is topologically farthest from the headend switch: that location is the fiber that connects Dhub B to Dhub C. Once this fiber is broken, the nodes on the ring can be labeled in relation to their position on the ring. Nodes that have a fiber going in and coming out are labeled middle nodes, while the two nodes that are adjacent to the broken fiber are labeled end nodes. The optical configuration for a middle node is slightly different from that for an end node, because a middle node splits the optical signal in order to pass it to downstream nodes.

Once the nodes are labeled, an optical wavelength plan must be drawn based on the Ethernet topology. Figure 6-5 on page 6-6 labels each group of GE links with a set of DWDM wavelengths that will be used to carry the optical signals from the DWDM GBICs in the headend to the Dhubs. With the wavelength plan in place, the optical designs for each node can be created.

Note The wavelength allocations shown in Figure 6-5 on page 6-6 do not take into account the fact that wavelengths can be reused when they are run on separate physical fibers. This was done to simplify the example. An optical deployment that is short on DWDM wavelengths can reuse wavelengths whenever possible.

Figure 6-7 on page 6-8 shows how the wavelength plan from Figure 6-5 on page 6-6 maps to the fiber topology.

Figure 6-7 Converting a Fiber Ring to Hub-and-Spoke Topology

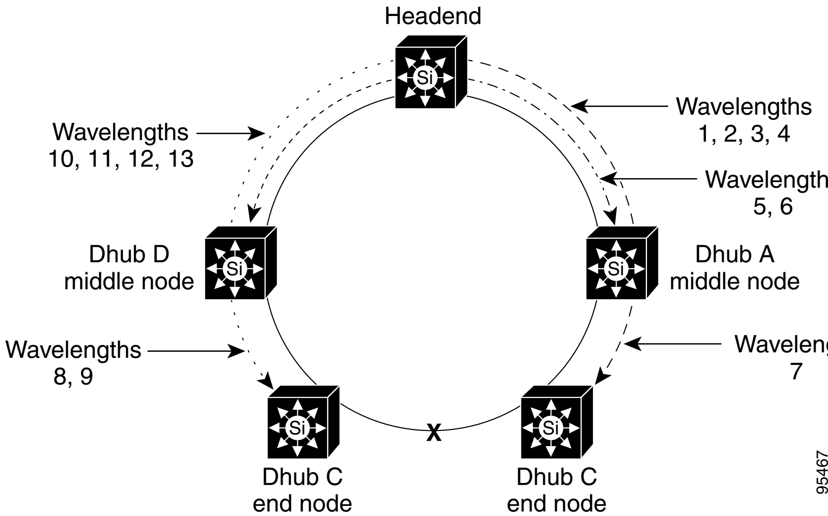

Figure 6-8 shows the optical design for Dhub A. Because Dhub A is a middle node, an optical splitter must be used to split the optical signal into one fiber for the Dhub switch and another fiber that feeds downstream nodes. From Figure 6-5 on page 6-6, the switch in Dhub A terminates four GE links, which are carried in DWDM wavelengths 1, 2, 3, and 4. These wavelengths are filtered out of the fiber from the splitter by means of two 2-port filters. The switch in Dhub A has an outgoing GE link that feeds subtended Dhub B over wavelength 7. That link is fed to the outgoing fiber by means of an optical combiner.

Figure 6-8 Middle Node Details: Dhub A

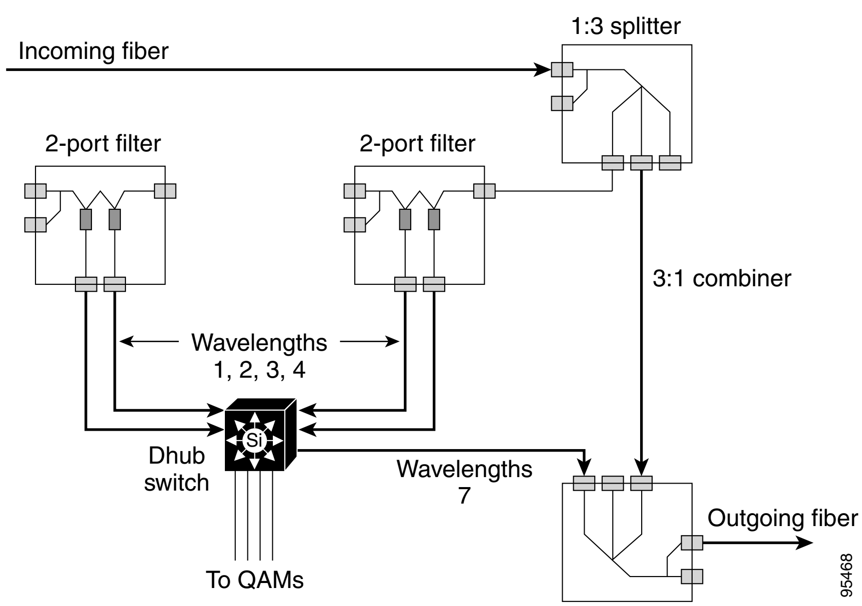

Figure 6-9 shows the optical design for subtended Dhub B. Since Dhub B is an end node, no optical splitters are needed. All that is needed are two 2-port filters that extract wavelengths 5, 6, and 7 from the incoming fiber. Wavelength 7 is the GE link from Dhub A, while wavelengths 5 and 6 are the unidirectional links from the headend.

Figure 6-9 End Node Details: Subtended Dhub B

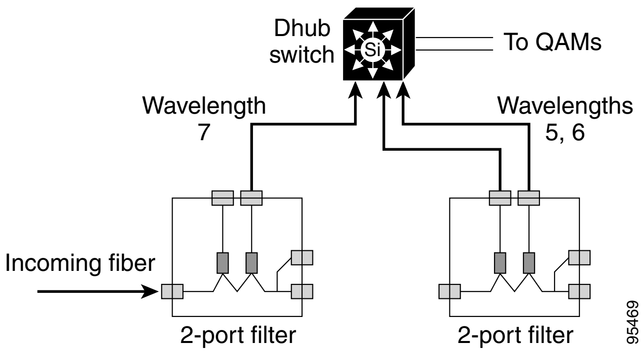

Figure 6-10 shows the optical design for Dhub D. Dhub D is similar to Dhub A, except that it does not feed any subtended Dhubs. Because of this, the optical design for Dhub D is identical to that for Dhub A—with the exception that the combiner from Dhub A is not needed.

Figure 6-10 Middle Node Details: Dhub D

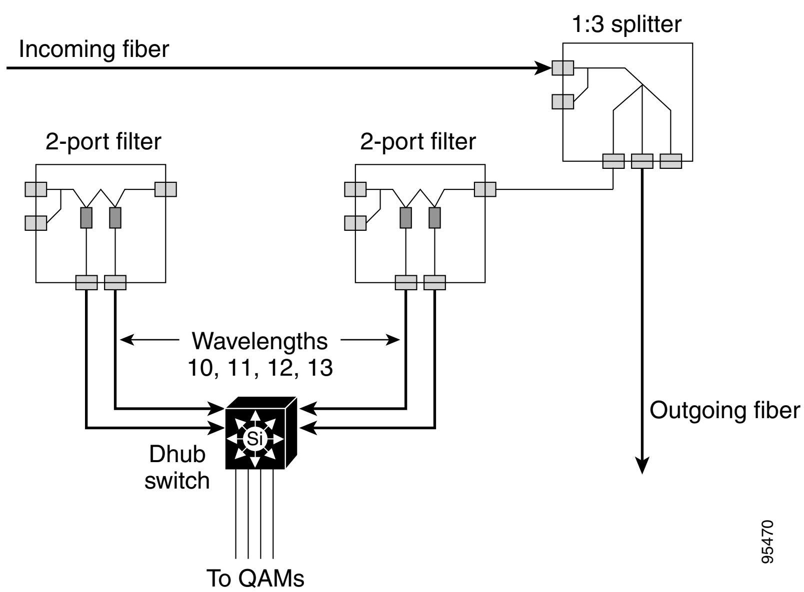

Figure 6-11 shows the optical design for Dhub C. Because Dhub C is an end node, it uses a single 2-port filter to extract the wavelengths that terminate on the Dhub switch.

Figure 6-11 End Node Details: Dhub C

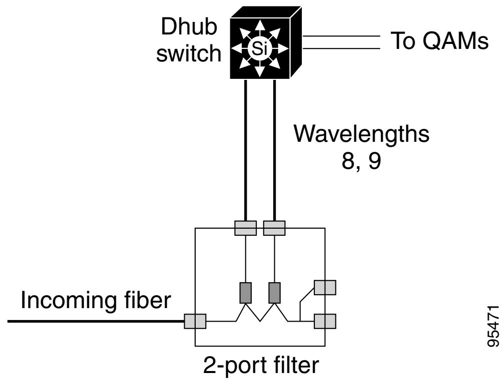

APPENDIX A

Switch in Dhub: Sample Configurations for Cisco Catalyst 4500 Series Switches

This appendix illustrates the following switch-in-Dhub configuration examples:

Single EtherChannel

Single EtherChannel: Headend Switch to Dhub Switch

This switch-in-Dhub example illustrates a configuration on a Cisco Catalyst 4500 series switch in the headend connected to a Dhub switch through a single EtherChannel (Port-channel). Refer to Figure 3-1 on page 3-2. For the corresponding configuration, see Single EtherChannel: Dhub Switch to Headend Switch, page A-5.

```
Current configuration : 8157 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
```

```
hostname Headend
!
!
qos
ip subnet-zero
no ip domain-lookup
!
!
class-map match-all class_video
  match access-group name acl_video
!
!
policy-map setDSCP
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
!
spanning-tree extend system-id
port-channel load-balance dst-port
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Port-channel1
 description Asymmetric EtherChannel to DHub_A
 ip address 192.168.169.1 255.255.255.252
 load-interval 30
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Bidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
```

```
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/4
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
!
interface GigabitEthernet4/2
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
<---output omitted--->
!
interface GigabitEthernet6/1
!
interface GigabitEthernet6/2
!
interface GigabitEthernet6/3
!
interface GigabitEthernet6/4
!
interface GigabitEthernet6/5
!
interface GigabitEthernet6/6
!
interface GigabitEthernet7/1
 description VODserver1 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 no cdp enable
!
interface GigabitEthernet7/2
!
interface GigabitEthernet7/3
!
interface GigabitEthernet7/4
!
interface GigabitEthernet7/5
 description VODserver1 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
```

```
interface GigabitEthernet7/6
!
interface GigabitEthernet7/7
!
interface GigabitEthernet7/8
!
interface GigabitEthernet7/9
 description VODserver2 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/10
!
interface GigabitEthernet7/11
!
interface GigabitEthernet7/12
!
interface GigabitEthernet7/13
 description VODserver2 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/14
!
interface GigabitEthernet7/15
!
interface GigabitEthernet7/16
!
interface GigabitEthernet7/17
 description VODserver3
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/18
!
<---output omitted--->
!
interface GigabitEthernet7/22
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan50
 description VoD Servers
 ip address 192.168.50.2 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
 service-policy input setDSCP
 standby 50 ip 192.168.50.1
!
```

```
router ospf 100
 log-adjacency-changes
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
ip access-list extended acl_video
 remark Identify video traffic (UDP ports 257-33023).
 permit udp any any range 257 33032
!
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 login
!
end
```

Single EtherChannel: Dhub Switch to Headend Switch

This switch-in-Dhub example illustrates a configuration on a Cisco Catalyst 4500 series switch in the Dhub connected to a headend switch through a single EtherChannel (Port-channel). Refer to Figure 3-1 on page 3-2. For the corresponding configuration, see Single EtherChannel: Headend Switch to Dhub Switch, page A-1.

```
Current configuration : 5752 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Dhub_A
!
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Port-channel1
 description Asymmetric EtherChannel from Headend
 ip address 192.168.169.2 255.255.255.252
 load-interval 30
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
```

```
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Bidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/4
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
 description QAM1
 switchport access vlan 160
 switchport mode access
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet4/2
 description QAM2
 switchport access vlan 160
 switchport mode access
 load-interval 30
!
interface GigabitEthernet4/3
 description QAM3
 switchport access vlan 160
 switchport mode access
 load-interval 30
!
interface GigabitEthernet4/4
 description QAM4
 switchport access vlan 160
 switchport mode access
 load-interval 30
```

```
!
interface GigabitEthernet4/5
 description QAM5
 switchport access vlan 160
 switchport mode access
 load-interval 30
!
interface GigabitEthernet4/6
!
interface GigabitEthernet5/1
!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
interface GigabitEthernet5/5
!
interface GigabitEthernet5/6
!
interface GigabitEthernet6/1
!
interface GigabitEthernet6/2
!
interface GigabitEthernet6/3
!
interface GigabitEthernet6/4
!
interface GigabitEthernet6/5
!
interface GigabitEthernet6/6
!
interface GigabitEthernet7/1
!
interface GigabitEthernet7/2
!
<---output omitted--->
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan160
 description QAM Devices
 ip address 192.168.160.1 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan160
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
arp 192.168.160.100 0020.a300.92aa ARPA
!
```

```
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 login
!
mac-address-table static 0020.a300.92aa vlan 160 interface GigabitEthernet4/1
!
end
```

Multiple EtherChannels

Multiple EtherChannels: Headend Switch to Dhub Switch

This switch-in-Dhub example illustrates a configuration on a Cisco Catalyst 4500 series switch in the headend connected to a Dhub switch through two EtherChannels (Port-channels). Refer to Figure 3-2 on page 3-14. For the corresponding configuration, see Multiple EtherChannels: Dhub Switch to Headend Switch, page A-12.

```
Current configuration : 8157 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Headend
!
qos
ip subnet-zero
ip cef load-sharing algorithm include-port destination
no ip domain-lookup
!
!
class-map match-all class_video
  match access-group name acl_video
!
!
policy-map setDSCP
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
!
spanning-tree extend system-id
port-channel load-balance dst-port
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Loopback3
 description Endpoint of Tunnel3
```

```
 ip address 1.1.1.13 255.255.255.255
!
interface Port-channel3
 description Unidirectional EtherChannel to DHub_D
 ip address 192.168.169.13 255.255.255.252
 load-interval 30
!
interface Port-channel4
 description Asymmetric EtherChannel to DHub_D
 ip address 192.168.169.17 255.255.255.252
 load-interval 30
!
interface Tunnel3
 description Return Path for Unidirectional EtherChannel 3
 no ip address
 tunnel source 1.1.1.13
 tunnel destination 2.2.2.14
 tunnel udlr receive-only Port-channel3
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Unidirectional Link of EtherChannel 3
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 3 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of EtherChannel 3
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 3 mode on
!
interface GigabitEthernet3/3
 description Unidirectional Link of EtherChannel 3
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 3 mode on
!
interface GigabitEthernet3/4
 description Unidirectional Link of EtherChannel 3
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
```

```
 unidirectional send-only
 channel-group 3 mode on
!
interface GigabitEthernet3/5
 description Bidirectional Link of EtherChannel 4
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 channel-group 4 mode on
!
interface GigabitEthernet3/6
 description Unidirectional Link of EtherChannel 4
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 4 mode on
!
interface GigabitEthernet4/1
 description Unidirectional Link of EtherChannel 4
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 4 mode on
!
interface GigabitEthernet4/2
 description Unidirectional Link of EtherChannel 4
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 4 mode on
!
<---output omitted--->
!
interface GigabitEthernet7/1
 description VODserver1 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 no cdp enable
!
interface GigabitEthernet7/2
!
interface GigabitEthernet7/3
!
interface GigabitEthernet7/4
!
interface GigabitEthernet7/5
 description VODserver1 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
```

```
!
interface GigabitEthernet7/6
!
interface GigabitEthernet7/7
!
interface GigabitEthernet7/8
!
interface GigabitEthernet7/9
 description VODserver2 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/10
!
interface GigabitEthernet7/11
!
interface GigabitEthernet7/12
!
interface GigabitEthernet7/13
 description VODserver2 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/14
!
interface GigabitEthernet7/15
!
interface GigabitEthernet7/16
!
interface GigabitEthernet7/17
 description VODserver3 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/18
!
<---output omitted--->
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan50
 description VoD Servers
 ip address 192.168.50.2 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
 service-policy input setDSCP
 standby 50 ip 192.168.50.1
!
router ospf 100
```

```
 log-adjacency-changes
 passive-interface Vlan50
 network 1.1.1.0 0.0.0.255 area 0
 network 192.168.0.0 0.0.0.255 area 0
!
ip classless
no ip http server
!
!
ip access-list extended acl_video
 remark Identify video traffic (UDP ports 257-33023).
 permit udp any any range 257 33032
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 password cisco123
 login
!
end
```

Multiple EtherChannels: Dhub Switch to Headend Switch

This switch-in-Dhub example illustrates a configuration on a Cisco Catalyst 4500 series switch in the Dhub connected to a headend switch through two EtherChannels (Port-channels). Refer to Figure 3-2 on page 3-14; this configuration is on Dhub D. For the corresponding configuration, see Multiple EtherChannels: Headend Switch to Dhub Switch, page A-8.

```
Current configuration : 6812 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname DHub_D
!
no logging console
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Loopback3
 description Endpoint of Tunnel3
 ip address 2.2.2.14 255.255.255.255
!
interface Port-channel3
 description Unidirectional EtherChannel from Headend
```

```
 ip address 192.168.169.14 255.255.255.252
 ip ospf cost 666
 load-interval 30
!
interface Port-channel4
 description Asymmetric EtherChannel from Headend
 ip address 192.168.169.18 255.255.255.252
 load-interval 30
!
interface Tunnel3
 description Return Path for Unidirectional EtherChannel 3
 no ip address
 tunnel source 2.2.2.14
 tunnel destination 1.1.1.3
 tunnel udlr send-only Port-channel3
 tunnel udlr address-resolution
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Unidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 3 mode on
!
interface GigabitEthernet3/2
 description Unidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 3 mode on
!
interface GigabitEthernet3/3
 description Unidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 3 mode on
!
interface GigabitEthernet3/4
 description Unidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 3 mode on
!
interface GigabitEthernet3/5
 description Bidirectional Link of Uni+Bidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 channel-group 4 mode on
!
```

```
interface GigabitEthernet3/6
 description Unidirectional Link of Uni+Bidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 4 mode on
!
interface GigabitEthernet4/1
 description Unidirectional Link of Uni+Bidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 4 mode on
!
interface GigabitEthernet4/2
 description Unidirectional Link of Uni+Bidirectional EtherChannel from Headend
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 4 mode on
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
interface GigabitEthernet4/6
!
interface GigabitEthernet5/1
 description QAM1
 switchport access vlan 163
 switchport mode access
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet5/2
 description QAM2
 switchport access vlan 163
 switchport mode access
 load-interval 30
 speed nonegotiate
!
<---output omitted--->
!
interface GigabitEthernet7/1
!
interface GigabitEthernet7/2
!
interface GigabitEthernet7/3
!
interface GigabitEthernet7/4
!
interface GigabitEthernet7/5
!
interface GigabitEthernet7/6
!
interface GigabitEthernet7/7
!
interface GigabitEthernet7/8
!
```

```
interface GigabitEthernet7/9
!
interface GigabitEthernet7/10
!
interface GigabitEthernet7/11
!
interface GigabitEthernet7/12
!
interface GigabitEthernet7/13
!
interface GigabitEthernet7/14
!
interface Vlan1
 no ip address
!
interface Vlan163
 description QAM VLAN
 ip address 192.168.163.1 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan163
 network 192.168.0.0 0.0.0.255 area 0
!
ip classless
no ip http server
!
arp 192.168.163.100 0020.a300.92aa ARPA
!
!
line con 0
 exec-timeout 360 0
 stopbits 1
line vty 0 4
 login
!
mac-address-table static 0020.a300.92aa vlan 163 interface GigabitEthernet5/1
!
end
```

Subtended Control Path

Subtended Control Path: Headend Switch to Two Dhubs

This switch-in-Dhub example illustrates a configuration on a Cisco Catalyst 4500 series switch in the headend connected to two Dhub switches through two EtherChannels (Port-channels) and a tunnel. Refer to Figure 3-3 on page 3-15. For the corresponding configurations, see Subtended Control Path: Dhub A to Headend Switch, page A-19 and Subtended Control Path: Dhub B to Headend Switch, page A-22.

```
Current configuration : 7852 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
```

```
no service password-encryption
service compress-config
!
hostname Headend
!
enable password cisco123
!
qos
ip subnet-zero
no ip domain-lookup
!
!
class-map match-all class_video
  match access-group name acl_video
!
!
policy-map setDSCP
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
!
spanning-tree extend system-id
port-channel load-balance dst-port
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Loopback2
 description Endpoint of Tunnel2
 ip address 1.1.1.5 255.255.255.255
!
interface Port-channel1
 description Asymmetric EtherChannel to DHub_A
 ip address 192.168.169.1 255.255.255.252
 load-interval 30
!
interface Port-channel2
 description Unidirectional EtherChannel to DHub_B
 ip address 192.168.169.5 255.255.255.252
 load-interval 30
 no keepalive
!
interface Tunnel2
 no ip address
 tunnel source 1.1.1.5
 tunnel destination 2.2.2.6
 tunnel udlr receive-only Port-channel2
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Bidirectional Link of EtherChannel 1
 no switchport
 no ip address
```

```
 load-interval 30
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of EtherChannel 1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
 description Unidirectional Link of EtherChannel 1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/4
 description Unidirectional Link of EtherChannel 1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/5
 description Unidirectional Link of EtherChannel 2
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 2 mode on
!
interface GigabitEthernet3/6
 description Unidirectional Link of EtherChannel 2
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 2 mode on
!
interface GigabitEthernet4/1
!
<---output omitted--->
!
interface GigabitEthernet6/4
!
interface GigabitEthernet6/5
!
interface GigabitEthernet6/6
```

```
!
interface GigabitEthernet7/1
 description VODserver1 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 no cdp enable
!
interface GigabitEthernet7/2
!
interface GigabitEthernet7/3
!
interface GigabitEthernet7/4
!
interface GigabitEthernet7/5
 description VODserver1 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/6
!
interface GigabitEthernet7/7
!
interface GigabitEthernet7/8
!
interface GigabitEthernet7/9
 description VODserver22 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/10
!
interface GigabitEthernet7/11
!
interface GigabitEthernet7/12
!
interface GigabitEthernet7/13
 description VODserver2 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/14
!
interface GigabitEthernet7/15
!
interface GigabitEthernet7/16
!
interface GigabitEthernet7/17
 description VODserver3
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
!
<---output omitted--->
```

```
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan50
 description VoD Servers
 ip address 192.168.50.2 255.255.255.0
 no ip redirects
 no ip unreachables
 service-policy input setDSCP
 standby 50 ip 192.168.50.1
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan50
 network 1.1.1.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
 distribute-list 1 in Port-channel1
 distribute-list 2 in Port-channel2
!
ip classless
no ip http server
!
ip access-list extended acl_video
 remark Identify video traffic (UDP ports 257-33023).
 permit udp any any range 257 33032
!
access-list 1 remark Only Allow Incoming VLAN160 OSPF Advertisements
access-list 1 permit 192.168.160.0 0.0.0.255
access-list 2 Only Allow Incoming VLAN161 OSPF Advertisements
access-list 2 permit 192.168.161.0 0.0.0.255
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 password cisco123
 login
!
end
```

Subtended Control Path: Dhub A to Headend Switch

This switch-in-Dhub example illustrates a configuration on a Cisco Catalyst 4500 series switch in the Dhub connected to a headend switch through a single EtherChannel (Port-channel). Refer to Figure 3-3 on page 3-15. Dhub A also has a Layer 3 connection to Dhub B. For the corresponding configuration, see Subtended Control Path: Headend Switch to Two Dhubs, page A-15, and Subtended Control Path: Dhub B to Headend Switch, page A-22.

```
Current configuration : 6002 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
```

```
service compress-config
!
hostname Dhub_A
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Port-channel1
 description Asymmetric EtherChannel from Headend
 ip address 192.168.169.2 255.255.255.252
 load-interval 30
!
interface GigabitEthernet1/1
 description Bidirectional Link for Tunnel
 no switchport
 ip address 192.168.169.9 255.255.255.252
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Bidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/4
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/5
```

```
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
 description QAM1
 switchport access vlan 160
 switchport mode access
 load-interval 30
 speed nonegotiate
!
<---output omitted--->
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
interface GigabitEthernet5/5
!
interface GigabitEthernet5/6
!
interface GigabitEthernet6/1
!
interface GigabitEthernet6/2
!
<---output omitted--->
!
interface GigabitEthernet7/1
!
interface GigabitEthernet7/2
!
<---output omitted--->
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan160
 description QAM Devices
 ip address 192.168.160.1 255.255.255.0
 no ip redirects
 no ip unreachables
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan160
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
arp 192.168.160.100 0020.a300.92aa ARPA
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 password cisco123
 login
!
```

```
mac-address-table static 0020.a300.92aa vlan 160 interface GigabitEthernet4/1
end
```

Subtended Control Path: Dhub B to Headend Switch

This switch-in-Dhub example illustrates a configuration on a Cisco Catalyst 4500 series switch in the Dhub connected to a headend switch through a single EtherChannel (Port-channel) and a tunnel. Refer to Figure 3-3 on page 3-15. Dhub B also has a Layer 3 connection to Dhub A. For the corresponding configurations, see Subtended Control Path: Headend Switch to Two Dhubs, page A-15, and Subtended Control Path: Dhub A to Headend Switch, page A-19.

```
Current configuration : 3281 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Dhub_B
!
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
process-max-time 20
!
interface Loopback2
 description Endpoint of Tunnel2
 ip address 2.2.2.6 255.255.255.255

!
interface Port-channel2
 description Unidirectional EtherChannel from Headend
 ip address 192.168.169.6 255.255.255.252
 ip ospf cost 666
 load-interval 30
 no keepalive
!
interface Tunnel2
 no ip address
 tunnel source 2.2.2.6
 tunnel destination 1.1.1.5
 tunnel udlr send-only Port-channel2
 tunnel udlr address-resolution
!
interface GigabitEthernet1/1
 description Bidirectional Link for Tunnel
 no switchport
 ip address 192.168.169.10 255.255.255.252
 load-interval 30
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
```

```
 load-interval 30
 unidirectional receive-only
 channel-group 2 mode on
!
interface GigabitEthernet2/2
 description Unidirectional Link of EtherChannel
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 2 mode on
!
interface GigabitEthernet2/3
!
interface GigabitEthernet2/4
!
interface GigabitEthernet2/5
!
interface GigabitEthernet2/6
!
interface GigabitEthernet3/1
 description QAM1
 switchport access vlan 161
 switchport mode access
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet3/2
 description QAM2
 switchport access vlan 161
 switchport mode access
 load-interval 30
!
interface GigabitEthernet3/3
 description QAM3
 switchport access vlan 161
 switchport mode access
 load-interval 30
!
interface GigabitEthernet3/4
 description QAM4
 switchport access vlan 161
 switchport mode access
 load-interval 30
!
interface GigabitEthernet3/5
 description QAM5
 switchport access vlan 161
 switchport mode access
 load-interval 30
!
interface GigabitEthernet3/6
!
interface Vlan1
 no ip address
!
interface Vlan161
 description QAM Devices
 ip address 192.168.161.1 255.255.255.0
 no ip redirects
 no ip unreachables
!
router ospf 100
 log-adjacency-changes
```

```
 passive-interface Vlan161
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
arp 192.168.161.100 0020.a300.92aa ARPA
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 login
!
mac-address-table static 0020.a300.92aa vlan 161 interface GigabitEthernet3/1
!
end
```

APPENDIX B

No Switch in Dhub: Sample Configurations for Cisco Catalyst 4500 Series Switches

This appendix presents example configurations, using VLANs, for the following redundancy scenarios for no-switch-in-Dhub scenarios.

The following configurations are presented:

Single VLAN

For the network diagram that represents this configuration, see Establishing a Single VLAN, page 3-16.

```
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Headend
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
```

```
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description QAM1
 switchport access vlan 50
 load-interval 30
 speed nonegotiate
 unidirectional send-only
!
interface GigabitEthernet3/2
 description QAM2
 switchport access vlan 50
 load-interval 30
 speed nonegotiate
 unidirectional send-only
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
!
interface GigabitEthernet4/2
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
interface GigabitEthernet4/6
!
interface GigabitEthernet5/1
!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
interface GigabitEthernet5/5
!
interface GigabitEthernet5/6
!
interface GigabitEthernet6/1
!
interface GigabitEthernet6/2
!
interface GigabitEthernet6/3
!
interface GigabitEthernet6/4
!
interface GigabitEthernet6/5
!
interface GigabitEthernet6/6
!
interface GigabitEthernet7/1
 description VODserver1 GigE0
 switchport access vlan 50
```

```
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/5
 description VODserver1 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/9
 description VODserver2 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/13
 description VODserver2 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/17
 description VODserver3
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/6
!
interface GigabitEthernet7/7
!
<---output omitted--->
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan50
 description VoD Servers
 ip address 192.168.50.2 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
 standby 50 ip 192.168.50.1
!
ip classless
no ip http server
!
!
!
line con 0
```

```
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 password cisco123
 login
!
mac-address-table static 0020.a300.92aa vlan 50 interface GigabitEthernet3/1
mac-address-table static 0090.f000.39dd vlan 50 interface GigabitEthernet3/2
!
end
```

Split VLAN

For the network diagram that represents this configuration, see Establishing a Split VLAN, page 3-16.

```
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Headend
!
ip subnet-zero
no ip domain-lookup
!
!
spanning-tree extend system-id
no spanning-tree vlan 50
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description QAM1
 switchport access vlan 160
 switchport mode access
 load-interval 30
 speed nonegotiate
 unidirectional send-only
!
interface GigabitEthernet3/2
 description QAM2
 switchport access vlan 160
 switchport mode access
 load-interval 30
 speed nonegotiate
 unidirectional send-only
```

```
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
!
interface GigabitEthernet4/2
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
interface GigabitEthernet4/6
!
interface GigabitEthernet5/1
!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
interface GigabitEthernet5/5
!
interface GigabitEthernet5/6
!
interface GigabitEthernet6/1
!
interface GigabitEthernet6/2
!
interface GigabitEthernet6/3
!
interface GigabitEthernet6/4
!
interface GigabitEthernet6/5
!
interface GigabitEthernet6/6
!
interface GigabitEthernet7/1
 description VODserver1 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/2
!
interface GigabitEthernet7/3
!
interface GigabitEthernet7/4
!
interface GigabitEthernet7/5
 description VOSserver1 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
```

```
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/6
!
interface GigabitEthernet7/7
!
interface GigabitEthernet7/8
!
interface GigabitEthernet7/9
 description VODserver2 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/10
!
interface GigabitEthernet7/11
!
interface GigabitEthernet7/12
!
interface GigabitEthernet7/13
 description VODserver2 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/14
!
interface GigabitEthernet7/15
!
interface GigabitEthernet7/16
!

interface GigabitEthernet7/17
 description VODserver3 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
! !
interface GigabitEthernet7/18
!
interface GigabitEthernet7/19
!
interface GigabitEthernet7/20
!
interface GigabitEthernet7/21
!
interface GigabitEthernet7/22
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan50
```

```
 description VoD Servers
 ip address 192.168.50.2 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
 standby 50 ip 192.168.50.1
!
interface Vlan160
 description QAM Devices
 ip address 192.168.160.1 255.255.255.0
 no ip redirects
 no ip unreachables
!
ip classless
no ip http server
!
!
arp 192.168.160.200 0020.a300.92aa ARPA
arp 192.168.160.201 0090.f000.39dd ARPA
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 password cisco123
 login
!
mac-address-table static 0020.a300.92aa vlan 160 interface GigabitEthernet3/1
mac-address-table static 0090.f000.39dd vlan 160 interface GigabitEthernet3/2
!
end
```

Multiple VLANs

For the network diagram that represents this configuration, see Establishing Multiple VLANs, page 3-17.

```
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Headend
!
ip subnet-zero
no ip domain-lookup
!
!
spanning-tree extend system-id
no spanning-tree vlan 50
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
```

```
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description QAM1
 switchport access vlan 160
 switchport mode access
 load-interval 30
 speed nonegotiate
 unidirectional send-only
!
interface GigabitEthernet3/2
 description QAM2
 switchport access vlan 161
 switchport mode access
 load-interval 30
 speed nonegotiate
 unidirectional send-only
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
!
interface GigabitEthernet4/2
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
interface GigabitEthernet4/6
!
interface GigabitEthernet5/1
!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
interface GigabitEthernet5/5
!
interface GigabitEthernet5/6
!
interface GigabitEthernet6/1
!
interface GigabitEthernet6/2
!
interface GigabitEthernet6/3
!
interface GigabitEthernet6/4
!
```

```
interface GigabitEthernet6/5
!
interface GigabitEthernet6/6
!
interface GigabitEthernet7/1
 description VODserver1 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/2
!
interface GigabitEthernet7/3
!
interface GigabitEthernet7/4
!
interface GigabitEthernet7/5
 description VODserver1 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/6
!
interface GigabitEthernet7/7
!
interface GigabitEthernet7/8
!
interface GigabitEthernet7/9
 description VODserver2 GigE0
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/10
!
interface GigabitEthernet7/11
!
interface GigabitEthernet7/12
!
interface GigabitEthernet7/13
 description VODserver2 GigE1
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/14
!
interface GigabitEthernet7/15
!
interface GigabitEthernet7/16
!
interface GigabitEthernet7/17
 description VODserver3 GigE0
 switchport access vlan 50
 switchport mode access
```

```
 load-interval 30
 speed 1000
 no cdp enable
!
interface GigabitEthernet7/18
!
interface GigabitEthernet7/19
!
interface GigabitEthernet7/20
!
interface GigabitEthernet7/21
!
interface GigabitEthernet7/22
!
interface GigabitEthernet7/23
!
interface GigabitEthernet7/24
!
interface Vlan1
 no ip address
!
interface Vlan50
 description VoD Servers
 ip address 192.168.50.2 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
 standby 50 ip 192.168.50.1
!
interface Vlan160
 description QAM Devices
 ip address 192.168.160.1 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
!
interface Vlan161
 description 2nd QAM VLAN
 ip address 192.168.161.1 255.255.255.0
 no ip redirects
 no ip unreachables
 load-interval 30
!
ip classless
no ip http server
!
!
arp 192.168.161.201 0090.f000.39dd ARPA
arp 192.168.160.200 0020.a300.92aa ARPA
!
!
line con 0
 exec-timeout 0 0
 stopbits 1
line vty 0 4
 password cisco123
 login
!
mac-address-table static 0020.a300.92aa vlan 160 interface GigabitEthernet3/1
mac-address-table static 0090.f000.39dd vlan 161 interface GigabitEthernet3/2
!
end
```

APPENDIX C

Switch in Dhub: Sample Redundancy Configurations for Cisco Catalyst 4500 Series Switches

This appendix presents sample redundancy configurations for the following switch-in-Dhub scenarios:

HSRP without UDLR	• Active Headend: HSRP without UDLR, page C-2 • Standby Headend: HSRP without UDLR, page C-4 • Dhub: HSRP without UDLR, page C-7
HSRP with UDLR	• Active Headend: HSRP with UDLR, page C-9 • Standby Headend: HSRP with UDLR, page C-12 • Dhub: HSRP with UDLR, page C-14
HSRP with Interface Tracking	• Active Headend: HSRP with Interface Tracking, page C-17 • Standby Headend: HSRP with Interface Tracking, page C-19 • Dhub: HSRP with Interface Tracking, page C-22
Layer 3 1+1 Protection	• Headend: Layer 3 Protection, page C-22 • Dhub: Layer 3 Protection, page C-25

Note See Using and Monitoring HSRP, page 3-9, and HSRP Redundancy Scenarios, page 4-5.

HSRP without UDLR

The following configurations are related in this series of examples:

- Active Headend: HSRP without UDLR
- Standby Headend: HSRP without UDLR
- Dhub: HSRP without UDLR

For topology and discussion, see Switch in Dhub: HSRP without UDLR, page 4-5.

Active Headend: HSRP without UDLR

```
Current configuration : 3484 bytes
!
version 12.1
service nagle
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname HeadEnd
!
!
clock timezone pst -8
qos
ip subnet-zero
ip cef load-sharing algorithm original
no ip domain-lookup
!
ip multicast-routing
!
class-map match-all class_video
  match access-group name acl_video
!
!
policy-map setDSCP
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
!
!
spanning-tree extend system-id
no spanning-tree vlan 151
port-channel load-balance dst-ip
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Port-channel1
 description Primary Asymmetric EtherChannel to DHub_Z
 ip address 10.10.10.1 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
 no keepalive
!
interface Port-channel2
 description Backup Link to Standby Switch
 ip address 10.10.20.1 255.255.255.252
 ip ospf cost 200
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface GigabitEthernet1/1
 description Trunk to Headend for VLAN151
```

```
 switchport access vlan 151
 speed nonegotiate
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Bidirectional Link of P1
 no switchport
 no ip address
 logging event link-status
 load-interval 30
 speed nonegotiate
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 no keepalive
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
 description Bidirectional Link to P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 tx-queue 3
   priority high
 channel-group 2 mode on
!
interface GigabitEthernet4/2
 description Bidirectional Link to P2
 no switchport
 no ip address
 speed nonegotiate
 tx-queue 3
   priority high
 channel-group 2 mode on
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
interface GigabitEthernet4/6
 description VOD server input
```

```
 switchport access vlan 151
 switchport mode access
 load-interval 30
 speed nonegotiate
 unidirectional receive-only
!
interface Vlan1
 no ip address
!
interface Vlan151
 description VOD Server Inputs
 ip address 192.168.151.100 255.255.255.0
 no ip redirects
 service-policy input setDSCP
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 100
 standby preempt
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan151
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
ip access-list extended acl_video
 remark Identify video traffic (UDP ports 257-33023 for QAM)
 permit udp any any range 257 33032
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
!
monitor session 1 source interface Gi2/1
monitor session 1 destination interface Gi3/5
end
```

Standby Headend: HSRP without UDLR

```
Current configuration : 3198 bytes
!
version 12.1
service nagle
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname HE_Standby
!
logging buffered 65536 debugging
!
clock timezone pst -8
```

```
ip subnet-zero
no ip domain-lookup
!
ip multicast-routing
!
spanning-tree extend system-id
no spanning-tree vlan 151
!
process-max-time 20
!
interface Port-channel1
 description Asymmetric EC to DHub_Y
 ip address 10.10.15.1 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Port-channel2
 description Backup Link to Primary
 ip address 10.10.20.2 255.255.255.252
 ip ospf cost 200
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Tunnel0
 no ip address
!
interface GigabitEthernet1/1
 description Link for VLAN 151 standby messages
 switchport access vlan 151
 ip access-group acl_standby_link out
 speed nonegotiate
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
 description Bidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet2/2
 description Bidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet2/3
 load-interval 30
!
interface GigabitEthernet2/4
 load-interval 30
 speed nonegotiate
!
```

```
interface GigabitEthernet2/5
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet2/6
 description Standby interface for VOD Servers
 switchport access vlan 151
 load-interval 30
 speed nonegotiate
 unidirectional receive-only
!
interface GigabitEthernet3/1
 description Bidirectional Link of P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 channel-group 2 mode on
!
interface GigabitEthernet3/2
 description Bidirectional Link of P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 channel-group 2 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
 description Bidirectional Link of P2
 no switchport
 no ip address
!
interface GigabitEthernet3/6
 description Bidirectional link of p2
 no switchport
 no ip address
!
interface Vlan1
 no ip address
!
interface Vlan150
 no ip address
!
interface Vlan151
 ip address 192.168.151.101 255.255.255.0
 no ip redirects
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 10
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan151
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
```

```
ip access-list extended acl_standby_link
 remark Stop video traffic on standby_link (g1/1)
 permit icmp any any
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
!
monitor session 1 source interface Gi2/5
end
```

Dhub: HSRP without UDLR

```
Current configuration : 2482 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Dhub_Y
!
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
process-max-time 20
!
interface Port-channel1
 description Asymmetric EC from Primary
 ip address 10.10.10.2 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Port-channel2
 description Asymmetric EC from Secondary
 ip address 10.10.15.2 255.255.255.252
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
 description Bidirectional Link for P1
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 tx-queue 3
```

```
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet2/2
 description Unidirectional link for P1
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 tx-queue 3
   priority high
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet2/3
!
interface GigabitEthernet2/4
!
interface GigabitEthernet2/5
!
interface GigabitEthernet2/6
!
interface GigabitEthernet3/1
 description Bidirectional link for P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 channel-group 2 mode on
!
interface GigabitEthernet3/2
 description Unidirectional link for P2
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 2 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
 description Output to Marquis
 switchport access vlan 150
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet3/6
 description QAM1
 switchport access vlan 150
 load-interval 30
 speed nonegotiate
!
interface Vlan1
 no ip address
!
interface Vlan150
 description VLAN for Edge QAM devices
 ip address 192.168.150.2 255.255.255.0
 arp timeout 150
!
router ospf 100
 log-adjacency-changes
```

```
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
end
```

HSRP with UDLR

The following configurations are related in this series of examples:

- Active Headend: HSRP with UDLR
- Standby Headend: HSRP with UDLR
- Dhub: HSRP with UDLR

For topology and discussion, see Switch in Dhub: HSRP with UDLR, page 4-7.

Active Headend: HSRP with UDLR

```
Current configuration : 3716 bytes
!
version 12.1
service nagle
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname HeadEnd
!
!
clock timezone pst -8
qos
ip subnet-zero
ip cef load-sharing algorithm original
no ip domain-lookup
!
ip multicast-routing
!
class-map match-all class_video
  match access-group name acl_video
!
!
policy-map setDSCP
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
!
```

```
!
spanning-tree extend system-id
no spanning-tree vlan 151
port-channel load-balance dst-ip
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Port-channel1
 description Primary Asymmetric EtherChannel to DHub_Z
 ip address 10.10.10.1 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
 no keepalive
!
interface Port-channel2
 description Backup Link to Standby Switch
 ip address 10.10.20.1 255.255.255.252
 ip ospf cost 200
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface GigabitEthernet1/1
 description Trunk to Headend for VLAN151
 switchport access vlan 151
 speed nonegotiate
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Send Only Link of P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 no keepalive
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
 description Rec Only link of P1
 no switchport
 no ip address
 load-interval 30
```

```
 tx-queue 3
   priority high
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
 description Bidirectional Link to P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 tx-queue 3
   priority high
 channel-group 2 mode on
!
interface GigabitEthernet4/2
 description Bidirectional Link to P2
 no switchport
 no ip address
 speed nonegotiate
 tx-queue 3
   priority high
 channel-group 2 mode on
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
!
interface GigabitEthernet4/6
 description VOD server input
 switchport access vlan 151
 switchport mode access
 load-interval 30
 speed nonegotiate
 unidirectional receive-only
!
interface Vlan1
 no ip address
!
interface Vlan151
 description VOD Server Inputs
 ip address 192.168.151.100 255.255.255.0
 no ip redirects
 service-policy input setDSCP
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 100
 standby preempt
 standby track GigabitEthernet3/3 100
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan151
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
```

```
!
ip classless
no ip http server
!
!
ip access-list extended acl_video
 remark Identify video traffic (UDP ports 257-33023 for QAM)
 permit udp any any range 257 33032
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
!
monitor session 1 source interface Gi2/1
monitor session 1 destination interface Gi3/5
end
```

Standby Headend: HSRP with UDLR

```
Current configuration : 3321 bytes
!
version 12.1
service nagle
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname HE_Standby
!
!
clock timezone pst -8
ip subnet-zero
no ip domain-lookup
!
ip multicast-routing
!
spanning-tree extend system-id
no spanning-tree vlan 151
!
process-max-time 20
!
interface Port-channel1
 description Asymmetric EC to DHub_Y
 ip address 10.10.15.1 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Port-channel2
 description Backup Link to Primary
 ip address 10.10.20.2 255.255.255.252
 ip ospf cost 200
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
```

```
!
interface Tunnel0
 no ip address
 tunnel source 10.10.20.2
 tunnel destination 10.10.10.2
 tunnel udlr receive-only Port-channel1
!
interface GigabitEthernet1/1
 description Link for VLAN 151 standby messages
 switchport access vlan 151
 ip access-group acl_standby_link out
 speed nonegotiate
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
 description Bidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet2/2
 description Bidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet2/3
 load-interval 30
!
interface GigabitEthernet2/4
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet2/5
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet2/6
 description Standby interface for VOD Servers
 switchport access vlan 151
 load-interval 30
 speed nonegotiate
 unidirectional receive-only
!
interface GigabitEthernet3/1
 description Bidirectional Link of P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 channel-group 2 mode on
!
interface GigabitEthernet3/2
 description Bidirectional Link of P2
```

```
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 channel-group 2 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
 description Bidirectional Link of P2
 no switchport
 no ip address
!
interface GigabitEthernet3/6
 description Bidirectional link of p2
 no switchport
 no ip address
!
interface Vlan1
 no ip address
!
interface Vlan150
 no ip address
!
interface Vlan151
 ip address 192.168.151.101 255.255.255.0
 no ip redirects
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 10
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan151
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
ip access-list extended acl_standby_link
 remark Stop video traffic on standby_link (g1/1)
 permit icmp any any
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
!
monitor session 1 source interface Gi2/5
end
```

Dhub: HSRP with UDLR

```
Current configuration : 2836 bytes
!
```

```
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname Dhub_Y
!
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
process-max-time 20
!
interface Port-channel1
 description Uni+Bidirectional EC from Primary
 ip address 10.10.10.2 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Port-channel2
 description Asymmetric EC from Secondary
 ip address 10.10.15.2 255.255.255.252
 ip ospf cost 300
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Tunnel0
 no ip address
 tunnel source 10.10.10.2
 tunnel destination 10.10.20.2
 tunnel udlr send-only Port-channel2
 tunnel udlr address-resolution
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
 description Bidirectional Link for P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet2/2
 description Unidirectional link for P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional receive-only
 channel-group 1 mode on
!
```

```
interface GigabitEthernet2/3
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet2/4
!
interface GigabitEthernet2/5
!
interface GigabitEthernet2/6
!
interface GigabitEthernet3/1
 description Bidirectional link for P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 unidirectional receive-only
 channel-group 2 mode on
!
interface GigabitEthernet3/2
 description Unidirectional link for P2
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 2 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
 description Output to Marquis
 switchport access vlan 150
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet3/6
 description QAM1
 switchport access vlan 150
 load-interval 30
 speed nonegotiate
!
interface Vlan1
 no ip address
!
interface Vlan150
 description VLAN for Edge QAM devices
 ip address 192.168.150.2 255.255.255.0
 arp timeout 150
!
router ospf 100
 log-adjacency-changes
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
```

```
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
end
```

HSRP with Interface Tracking

The following configurations are related in this series of examples:

- Active Headend: HSRP with Interface Tracking
- Standby Headend: HSRP with Interface Tracking
- Dhub: HSRP with Interface Tracking

For topology and discussion, see Switch in Dhub: HSRP with Interface Tracking, page 4-8.

Active Headend: HSRP with Interface Tracking

```
Current configuration : 3899 bytes
!
version 12.1
service nagle
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
service compress-config
!
hostname HeadEnd
!
!
clock timezone pst -8
qos
ip subnet-zero
ip cef load-sharing algorithm original
no ip domain-lookup
!
ip multicast-routing
!
class-map match-all class_video
  match access-group name acl_video
!
!
policy-map setDSCP
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
!
!
spanning-tree extend system-id
no spanning-tree vlan 151
port-channel load-balance dst-ip
!
redundancy
```

```
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Port-channel1
 description Primary Asymmetric EtherChannel to DHub_Z
 ip address 10.10.10.1 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
 no keepalive
!
interface Port-channel2
 description Backup Link to Standby Switch
 ip address 10.10.20.1 255.255.255.252
 ip ospf cost 200
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30

!
interface GigabitEthernet1/1
 description Trunk to Headend for VLAN151
 switchport access vlan 151
 ip access-group acl_standby_link out
 speed nonegotiate
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Send Only Link of P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 no keepalive
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
!
```

```
interface GigabitEthernet4/2
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
interface GigabitEthernet4/6
 description VOD server input
 switchport access vlan 151
 switchport mode access
 load-interval 30
 speed nonegotiate
 unidirectional receive-only
!
interface Vlan1
 no ip address
!
interface Vlan151
 description VOD Server Inputs
 ip address 192.168.151.100 255.255.255.0
 no ip redirects
 service-policy input setDSCP
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 100
 standby preempt
 standby track GigabitEthernet3/1 100
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan151
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
ip access-list extended acl_standby_link
 remark Stop video traffic on standby_link (g1/1)
 permit icmp any any
ip access-list extended acl_video
 remark Identify video traffic (UDP ports 257-33023 for QAM)
 permit udp any any range 257 33032
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
!
monitor session 1 source interface Gi2/1
monitor session 1 destination interface Gi3/5
end
```

Standby Headend: HSRP with Interface Tracking

```
Current configuration : 3180 bytes
```

```
!
version 12.1
service nagle
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname HE_Standby
!
!
clock timezone pst -8
ip subnet-zero
no ip domain-lookup
!
ip multicast-routing
!
spanning-tree extend system-id
no spanning-tree vlan 151
!
process-max-time 20
!
interface Port-channel1
 description Asymmetric EC to DHub_Y
 ip address 10.10.15.1 255.255.255.252
 ip ospf cost 100
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Port-channel2
 description Backup Link to Primary
 ip address 10.10.20.2 255.255.255.252
 ip ospf cost 200
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface GigabitEthernet1/1
 description Link for VLAN 151 standby messages
 switchport access vlan 151
 ip access-group acl_standby_link out
 speed nonegotiate
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
 description Bidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet2/2
 description Bidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
```

```
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet2/3
 load-interval 30
!
interface GigabitEthernet2/4
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet2/5
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet2/6
 description Standby interface for VOD Servers
 switchport access vlan 151
 load-interval 30
 speed nonegotiate
 unidirectional receive-only
!
interface GigabitEthernet3/1
 description Bidirectional Link of P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 channel-group 2 mode on
!
interface GigabitEthernet3/2
 description Bidirectional Link of P2
 no switchport
 no ip address
 load-interval 30
 speed nonegotiate
 channel-group 2 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
 description Bidirectional Link of P2
 no switchport
 no ip address
!
interface GigabitEthernet3/6
 description Bidirectional link of p2
 no switchport
 no ip address
!
interface Vlan1
 no ip address
!
interface Vlan150
 no ip address
!
interface Vlan151
 ip address 192.168.151.101 255.255.255.0
 no ip redirects
 standby ip 192.168.151.2
 standby timers 1 2
 standby priority 10
 standby preempt
```

```
!
router ospf 100
 log-adjacency-changes
 passive-interface Vlan151
 network 10.10.0.0 0.0.255.255 area 0
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
ip access-list extended acl_standby_link
 remark Stop video traffic on standby_link (g1/1)
 permit icmp any any
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
!
monitor session 1 source interface Gi2/5
end
```

Dhub: HSRP with Interface Tracking

This configuration is the same as that for Dhub: HSRP without UDLR, page C-7.

Layer 3 1+1 Protection

The following configurations are related in this series of examples:

- Headend: Layer 3 Protection
- Dhub: Layer 3 Protection

For topology and discussion, see Layer 3 1+1 Protection Scheme, page 4-10.

Headend: Layer 3 Protection

```
Current configuration : 4209 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname HeadEnd
!
no logging console
!
qos
ip subnet-zero
```

```
no ip domain-lookup
!
!
class-map match-all class_video
  match access-group name acl_video
!
!
policy-map setDSCP
 description Mark all video traffic with DSCP of EF
  class class_video
    set ip dscp 46
!
spanning-tree extend system-id
port-channel load-balance dst-port
!
redundancy
 mode rpr
 main-cpu
  auto-sync standard
!
process-max-time 20
!
interface Port-channel1
 description Primary Asymmetric EtherChannel to DHub_Z
 ip address 192.168.169.1 255.255.255.252
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Port-channel2
 description Secondary Asymmetric EtherChannel to Dhub_Z
 ip address 192.168.170.1 255.255.255.252
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
!
interface GigabitEthernet2/2
!
interface GigabitEthernet3/1
 description Bidirectional Link to P1 1554.13nm
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet3/2
 description Unidirectional Link to P1 1554.94nm
 no switchport
 no ip address
 load-interval 30
 unidirectional send-only
 channel-group 1 mode on
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
```

```
!
interface GigabitEthernet3/5
!
interface GigabitEthernet3/6
!
interface GigabitEthernet4/1
 description Bidirectional Link to P2 1550.12nm
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 channel-group 2 mode on
!
interface GigabitEthernet4/2
 description Unidirectional Line to P2 1550.92nm
 no switchport
 no ip address
 tx-queue 3
   priority high
 unidirectional send-only
 channel-group 2 mode on
!
interface GigabitEthernet4/3
!
interface GigabitEthernet4/4
!
interface GigabitEthernet4/5
!
interface GigabitEthernet4/6
 description VOD server input
 switchport access vlan 50
 switchport mode access
 load-interval 30
!
interface GigabitEthernet5/1
 description VOD server input
 switchport access vlan 50
 switchport mode access
 load-interval 30
 speed 1000
!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
<---output omitted--->
!
interface GigabitEthernet5/47
!
interface GigabitEthernet5/48
!
interface Vlan1
 no ip address
!
interface Vlan50
 description VOD servers
 ip address 192.168.50.150 255.255.255.0
 no ip redirects
 no ip unreachables
 service-policy input setDSCP
!
```

```
router ospf 100
 log-adjacency-changes
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
end
```

Dhub: Layer 3 Protection

```
Current configuration : 2415 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service compress-config
!
hostname DHUB_Z
!
no logging console
!
ip subnet-zero
no ip domain-lookup
!
spanning-tree extend system-id
!
process-max-time 20
!
interface Port-channel1
 description Primary Asymmetric EtherChannel from Headend
 ip address 192.168.169.2 255.255.255.252
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface Port-channel2
 description Secondary Asymmetric Etherchannel from Headend
 ip address 192.168.170.2 255.255.255.252
 ip ospf hello-interval 1
 ip ospf dead-interval 3
 load-interval 30
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet2/1
 description Bidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
```

```
 tx-queue 3
   priority high
 channel-group 1 mode on
!
interface GigabitEthernet2/2
 description Unidirectional Link of P1
 no switchport
 no ip address
 load-interval 30
 tx-queue 3
   priority high
 unidirectional receive-only
 channel-group 1 mode on
!
interface GigabitEthernet2/3
!
interface GigabitEthernet2/4
 speed nonegotiate
!
interface GigabitEthernet2/5
 description QAM1
 switchport access vlan 160
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet2/6
 description QAM2
 switchport access vlan 160
 load-interval 30
 speed nonegotiate
!
interface GigabitEthernet3/1
!
interface GigabitEthernet3/2
!
interface GigabitEthernet3/3
!
interface GigabitEthernet3/4
!
interface GigabitEthernet3/5
 description Bidirectional Link of P2
 no switchport
 no ip address
 load-interval 30
 channel-group 2 mode on
!
interface GigabitEthernet3/6
 description Unidirectional Link of P2
 no switchport
 no ip address
 load-interval 30
 unidirectional receive-only
 channel-group 2 mode on
!
interface Vlan1
 no ip address
!
interface Vlan160
 description QAM Devices
 ip address 192.168.150.1 255.255.255.0
 no ip unreachables
!
router ospf 100
 log-adjacency-changes
```

```
 network 192.168.0.0 0.0.255.255 area 0
!
ip classless
no ip http server
!
!
!
!
line con 0
 stopbits 1
line vty 0 4
 login
!
mac-address-table static 0090.f000.7268 vlan 160 interface GigabitEthernet2/5
!
monitor session 1 source interface Gi2/5
monitor session 1 destination interface Gi2/4
end
```