



Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband Solution Design and Implementation Guide, Release 1.0

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-7087-01, Rev. A1



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0601R)

Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband Solution Design and Implementation Guide, Release 1.0
Copyright © 2005, Cisco Systems, Inc.
All rights reserved.



Preface	xi
Document Version and Solution Release	xi
Document Objectives and Scope	xii
Audience	xii
Document Organization	xii
Related Documentation	xiii
Solution Documentation	xiii
Switch Documentation	xiii
Cisco Catalyst 4500 Series Switches	xiv
Cisco Catalyst 6500 Series Switches	xiv
Cisco 7600 Series Routers	xiv
Optical Component Documentation	xiv
Cisco DWDM GBICs	xiv
Document Conventions	xiv
Obtaining Documentation	xvi
Cisco.com	xvi
Product Documentation DVD	xvi
Ordering Documentation	xvi
Documentation Feedback	xvii
Cisco Product Security Overview	xvii
Reporting Security Problems in Cisco Products	xvii
Obtaining Technical Assistance	xviii
Cisco Technical Support & Documentation Website	xviii
Submitting a Service Request	xix
Definitions of Service Request Severity	xix
Obtaining Additional Publications and Information	xix

CHAPTER 1

Solution Overview	1-1
Solution Description and Scope	1-1
Generic Architecture and Scope	1-1
In Scope	1-2
Out of Scope	1-3
Solution Components	1-3
Cisco Equipment	1-3

- Third-Party Equipment 1-3
- Miscellaneous Solution Support 1-4
- Operational Support Systems 1-4
- Billing 1-4
- EMC 1-4
- Safety 1-4

CHAPTER 2

Video Application Components and Architecture 2-1

- Video Application Components 2-1
 - Broadcast Video Components 2-2
 - Real-Time Encoder 2-2
 - Electronic Program Guide 2-2
 - Broadcast Client 2-3
 - VoD Components 2-3
 - Asset Distribution System 2-3
 - Navigation Server 2-4
 - Session Manager 2-4
 - Entitlement System 2-4
 - Video Pump 2-5
 - On-Demand Resource Manager 2-5
 - On-Demand Client 2-5
 - Common Broadcast Video and VoD Components 2-5
 - Conditional Access System and Encryption Engine 2-6
 - Broadcast Video Bandwidth Enforcement 2-7
 - Set-Top-Based Video Decryption and Video Decoder 2-7
 - Set-Top Box 2-7
 - Subscriber Database 2-7
 - IPTV/VoBB Product Architecture 2-8
 - Middleware 2-8
 - VoD Server 2-9
 - Conditional Access System 2-9
 - Real-Time Encoder and Set-Top Box 2-9
- IPTV/VoBB Transport Architecture and Issues 2-9
 - Video Sites 2-10
 - Super Headend 2-10
 - Video Headend Office 2-10
 - Video Switching Office 2-10
 - Video Service Requirements 2-11
 - High Bandwidth 2-11

Asymmetric Bandwidth	2-12
Quality of Service	2-12
Service Availability	2-13
Broadcast Video Channel-Change Time	2-14
Potential Video Service Architectures	2-14
Transport-Based SLA	2-15
Managed Application-Based SLA	2-17
Service Separation in a Triple-Play Architecture	2-18
Forwarding Architectures	2-18
Service-Availability and Bandwidth Requirements	2-19
Organizational Structure	2-19
IP Infrastructure Components	2-19
Service Separation in the Release 1.0 Architecture	2-19

CHAPTER 3**Solution Transport Architecture 3-1**

Overview	3-1
Solution Components	3-3
Distribution and Aggregation Transport Architecture	3-4
Video Forwarding	3-4
Layer 3 Edge for Video Services	3-4
Video Forwarding Architecture	3-11
Multicast	3-16
Overview	3-16
Multicast Admission Control	3-17
Effect of Multicast on Channel-Change Performance	3-19
Multicast Configuration Options	3-24
IGMP Functionality in the STB	3-27
Internet Access Forwarding	3-27
Voice Forwarding	3-28
Management	3-29
Management Transport	3-29
DHCP Configuration	3-30
EMS/NMS	3-30
Redundancy	3-31
Video-Infrastructure Component Redundancy	3-31
Network Redundancy	3-31
Release 1.0 Configurations	3-32
Overview	3-32
Configuration 1: 10-GE Layer 3 Symmetric Ring	3-33

- Configuration 2: N x 1-GE Asymmetric Ring 3-34
 - Bidirectional Interface Support 3-36
 - Routing Configurations 3-37
- Edge Transport Architecture 3-39
 - Overview 3-40
 - DSLAM Functions 3-40
 - HAG Functions 3-42
 - Service Separation Functions 3-42
 - NAT/Layer 3 Functionality 3-45
- QoS Architecture 3-46
 - Overview of DiffServ Architecture 3-46
 - DiffServ Architecture in the Solution 3-47
 - Administrative Boundaries 3-49
 - DiffServ-to-Layer-2 Mapping 3-49
 - Security and Additional Boundaries of Trust 3-49
 - Triple-Play QoS Analysis 3-50
 - Internet Access 3-50
 - Voice 3-51
 - Video 3-51
 - Voice and Video Signaling 3-54
 - QoS in the Aggregation/Distribution Network 3-55
 - QoS in the Access Network 3-56
 - ATM Layer Scheduling 3-57
 - MAC/IP Layer Scheduling 3-58

CHAPTER 4

Implementing and Configuring the Solution 4-1

- Common Tasks: Configuring SSM Mapping with DNS Lookup 4-1
 - Configuring DNS Servers 4-2
 - Configuring SSM Mapping on All Switches 4-2
 - Configuring the Edge Switches for DNS Queries 4-3
- Configuring the 10-GE Symmetric Topology 4-4
 - Introduction 4-4
 - Configuring DER 4-7
 - Configuring QoS on DER 4-7
 - Establishing and Configuring Interfaces on DER 4-11
 - Configuring OSPF Routing for Video and Voice Traffic on DER 4-19
 - Configuring Spanning Tree on DER 4-20
 - Configuring AR1 4-21
 - Configuring QoS on AR1 4-21

Establishing and Configuring Interfaces on AR1	4-24
Configuring OSPF Routing for Video and Voice Traffic on AR1	4-30
Configuring Spanning Tree on AR1	4-30
Configuring AR2	4-31
Configuring QoS on AR2	4-31
Establishing and Configuring Interfaces on AR2	4-35
Configuring OSPF Routing for Video and Voice Traffic on AR2	4-40
Configuring Spanning Tree on AR2	4-41
Configuring AR3	4-42
Configuring QoS on AR3	4-42
Establishing and Configuring Interfaces on AR3	4-46
Configuring OSPF Routing for Video and Voice Traffic on AR3	4-51
Configuring Spanning Tree on AR3	4-51
Configuring the 1-GE Asymmetric Topology	4-53
Introduction	4-53
Configuring DER	4-57
Configuring QoS on DER	4-57
Establishing and Configuring Interfaces on DER	4-61
Configuring OSPF Routing for Video and Voice Traffic on DER	4-71
Configuring Spanning Tree on DER	4-73
Configuring AR1	4-74
Configuring QoS on AR1	4-74
Establishing and Configuring Interfaces on AR1	4-79
Configuring OSPF Routing for Video and Voice Traffic on AR1	4-87
Configuring Spanning Tree on AR1	4-88
Configuring AR2	4-89
Configuring QoS on AR2	4-89
Establishing and Configuring Interfaces on AR2	4-90
Configuring OSPF Routing for Video and Voice Traffic on AR2	4-98
Configuring Spanning Tree on AR2	4-99
Configuring AR3	4-100
Configuring QoS on AR3	4-100
Establishing and Configuring Interfaces on AR3	4-101
Configuring OSPF Routing for Video and Voice Traffic on AR3	4-110
Configuring Spanning Tree on AR3	4-111

CHAPTER 5**Monitoring and Troubleshooting 5-1**

Network Time Protocol (NTP) 5-1

Syslog 5-2

- Global Syslog Configuration 5-2
- Interface Syslog Configuration 5-2
- Useful Syslog Commands 5-2
 - no logging console** 5-3
 - no logging monitor 5-3
 - logging buffered 16384 5-3
 - logging trap notifications 5-3
 - logging facility local7 5-3
 - logging host 5-3
 - logging source-interface loopback 0 5-3
 - service timestamps debug datetime localtime show-timezone msec 5-3
 - logging event 5-4
- Quality of Service (QoS) 5-4
 - show class-map 5-4
 - show policy-map 5-4
 - show qos maps 5-5
 - show mls qos maps dscp-cos 5-5
 - show qos interface 5-6
 - show queueing interface 5-6
- Multicast 5-11
 - show ip mroute 5-11
 - show ip mroute ssm 5-12
 - show ip mroute active 5-13
 - show ip pim neighbor 5-13
 - show ip igmp snooping 5-13
 - show ip igmp groups 5-14
 - show ip igmp ssm-mapping 5-14
 - show ip igmp membership 5-14
 - debug ip igmp 5-15
 - debug ip pim 5-15
 - debug domain 5-16
- UDLR and Unidirectional Links 5-16
 - show interface tunnel 5-16
 - show interface 5-17
 - debug tunnel 5-17
- References 5-18

Configuration for AR1 **A-10**

Configuration for AR2 **A-16**

Configuration for AR3 **A-21**

APPENDIX B
Sample DER and AR Switch Configurations for the 1-GE Asymmetric Topology **B-1**

Configuration for DER **B-1**

Configuration for AR1 **B-15**

Configuration for AR2 **B-23**

Configuration for AR3 **B-31**

APPENDIX C
Understanding QoS as Implemented in the Solution **C-1**

Introduction **C-1**

DiffServ Classification **C-3**

Class Selector Values **C-3**

Assured Forwarding **C-4**

Express Forwarding **C-4**

Default Class **C-5**

DSCP-to-CoS Mapping **C-5**

Queueing and Thresholds **C-6**

Queueing Structures on Cisco 7600 Series and Cisco Catalyst 6000 Series Line Cards **C-6**

Queueing Structures on Cisco Catalyst 4000 Series Line Cards and Ports **C-7**

10-GE Symmetric Topology: Known and Unknown Queue Parameters **C-7**

1-GE Asymmetric Topology: Known and Unknown Queue Parameters **C-8**

1-GE Asymmetric and 10-GE Symmetric Topologies:

Known Threshold Parameters **C-9**

APPENDIX D
Configuring DSL Equipment **D-1**

Network Diagram **D-1**

Hardware and Software Versions **D-3**

Configuring Ericsson Components **D-4**

Configuring the Switch **D-4**

Configuring the DSLAM **D-4**

Configuring the HAG **D-6**

atm.conf **D-6**

bridge.conf **D-7**

Creating Line Configurations **D-8**

Creating Services and Profiles **D-9**

Creating Services and Profiles for Video	D-9
Creating Services and Profiles for Voice	D-10
Creating Services and Profiles for Data	D-11
Creating User Profiles and Adding Services	D-12
Creating Profile 1	D-12
Creating Profile 2	D-14
Creating an IP Filter	D-15
Special Issues	D-15



Preface

This preface explains the objectives, intended audience, and organization of the Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband (GOVoBB) Solution, Release 1.0. The solution supports both broadcast video and video on demand (VoD) for the IPTV/video over broadband/telco (IPTV/VoBB) market, enabling operators that use digital subscriber lines (DSL) and fiber (FTTx) to offer not only video but also voice over IP (VoIP) and data (Internet access)—collectively referred to as “triple play”—over their existing infrastructure, now intelligently optimized for video service.

The preface 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, page xi](#)
- [Document Objectives and Scope, page xii](#)
- [Audience, page xii](#)
- [Document Organization, page xii](#)
- [Related Documentation, page xiii](#)
- [Document Conventions, page xiv](#)
- [Obtaining Documentation, page xvi](#)
- [Documentation Feedback, page xvii](#)
- [Cisco Product Security Overview, page xvii](#)
- [Obtaining Technical Assistance, page xviii](#)
- [Obtaining Additional Publications and Information, page xix](#)

Document Version and Solution Release

This is the second version of this document, which covers Release 1.0 of the Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband (GOVoBB) Solution.

Document History

Document Version	Date	Notes
1	08/25/2005	This document was first released.
2	12/01/2005	Multicast Admission Control , page 17, documents a condition on the ip igmp limit command.
3	09/19/2006	Clarifies distinction between “Express forwarding” and Expedited Forwarding (EF) in Appendix C.

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 IPTV/Video over Broadband (GOVoBB) Solution, Release 1.0. The primary focus of this release is on video over broadband functionality.

**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 GOVoBB 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 series switches).

**Note**

This document addresses Cisco components only. It does not discuss how to implement third-party components typically required for a video service, such as VoD servers, encoders, headends, program guides, or DSLAMs.

Document Organization

The major sections of this document are as follows:

Section	Title	Major Topics
Chapter 1	Solution Overview	Introduces solution architecture and scope, components, and miscellaneous support topics.
Chapter 2	Video Application Components and Architecture	Discusses the segmentation of the video application architecture into logical components that are required for broadcast video and VoD services.

Chapter 3	Solution Transport Architecture	Discusses architectures for distribution, aggregation, edge transport, and quality of service. Introduces 10-GE ring and 1-GE asymmetric configurations.
Chapter 4	Implementing and Configuring the Solution	Describes the configuration and implementation of the solution, and provides example implementations.
Chapter 5	Monitoring and Troubleshooting	Provides an introduction to monitoring and troubleshooting the Cisco switches used in the solution.
Appendix A	Sample DER and AR Switch Configurations for the 10-GE Symmetric Topology	Provides example configurations for distribution edge routers and aggregation routers for this topology.
Appendix B	Sample DER and AR Switch Configurations for the 1-GE Asymmetric Topology	Provides example configurations for distribution edge routers and aggregation routers for this topology.
Appendix C	Understanding QoS as Implemented in the Solution	Presents a more detailed understanding of Quality of Service (QoS) specific to this solution. Discusses queuing structures on specific line cards and ports.
Appendix D	Configuring DSL Equipment	Provides details related to the configuration of Ericsson DSL equipment, if it is used.

Related Documentation

Solution Documentation

This document, and *Release Notes for Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband Solution, Release 1.0*, are available under the following URLs:

<http://www.cisco.com/univercd/cc/td/doc/solution/vobbsols>

http://www.cisco.com/en/US/netsol/ns524/networking_solutions_market_segment_solutions_home.html

Switch Documentation

Documentation resources for the Cisco Catalyst switches and the Cisco 7609 router are available at the following URLs:



Note

The Cisco 7609 router used in this solution functions as a switch, and is considered to be a switch in this documentation.

Cisco Catalyst 4500 Series Switches



Note

The following includes the Cisco Catalyst 4948-10GE.

For all hardware and software documentation for this series, go to the following URL:

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

Cisco Catalyst 6500 Series Switches

For all hardware and software documentation for this series, go to the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/index.htm>

Cisco 7600 Series Routers

For all hardware and software documentation for this series, go to the following URL:

<http://www.cisco.com/univercd/cc/td/doc/product/core/cis7600/index.htm>

Optical Component Documentation

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



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 .
<i>italic font</i>	Arguments for which you supply values are in <i>italics</i> .
[]	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 <i>screen font</i> .
boldface screen font	Information you must enter is in boldface screen font . ¹
<i>italic screen font</i>	Arguments for which you supply values are in <i>italic screen font</i> .
→	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

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

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 at this URL:

<http://www.cisco.com/techsupport>

You can access the Cisco website at this URL:

<http://www.cisco.com>

You can access international Cisco websites at this URL:

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

Product Documentation DVD

Cisco documentation and additional literature are available in the Product Documentation DVD package, which may have shipped with your product. The Product Documentation DVD is updated regularly and may be more current than printed documentation.

The Product Documentation DVD is a comprehensive library of technical product documentation on portable media. The DVD enables you to access multiple versions of hardware and software installation, configuration, and command guides for Cisco products and to view technical documentation in HTML. With the DVD, you have access to the same documentation that is found on the Cisco website without being connected to the Internet. Certain products also have .pdf versions of the documentation available.

The Product Documentation DVD is available as a single unit or as a subscription. Registered Cisco.com users (Cisco direct customers) can order a Product Documentation DVD (product number DOC-DOCDVD=) from Cisco Marketplace at this URL:

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

Ordering Documentation

Beginning June 30, 2005, registered Cisco.com users may order Cisco documentation at the Product Documentation Store in the Cisco Marketplace at this URL:

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

Nonregistered Cisco.com users can order technical documentation from 8:00 a.m. to 5:00 p.m. (0800 to 1700) PDT by calling 1 866 463-3487 in the United States and Canada, or elsewhere by calling 011 408 519-5055. You can also order documentation by e-mail at tech-doc-store-mkpl@external.cisco.com or by fax at 1 408 519-5001 in the United States and Canada, or elsewhere at 011 408 519-5001.

Documentation Feedback

You can rate and provide feedback about Cisco technical documents by completing the online feedback form that appears with the technical documents on Cisco.com.

You can send comments about Cisco 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.

Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

From this site, you can perform these tasks:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories and notices for Cisco products is available at this URL:

<http://www.cisco.com/go/psirt>

If you prefer to see advisories and notices as they are updated in real time, you can access a Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed from this URL:

http://www.cisco.com/en/US/products/products_psirt_rss_feed.html

Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you might have identified a vulnerability in a Cisco product, contact PSIRT:

- Emergencies—security-alert@cisco.com

An emergency is either a condition in which a system is under active attack or a condition for which a severe and urgent security vulnerability should be reported. All other conditions are considered nonemergencies.

- Nonemergencies—psirt@cisco.com

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

**Tip**

We encourage you to use Pretty Good Privacy (PGP) or a compatible product to encrypt any sensitive information that you send to Cisco. PSIRT can work from encrypted information that is compatible with PGP versions 2.x through 8.x.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one linked in the Contact Summary section of the Security Vulnerability Policy page at this URL:

http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html

The link on this page has the current PGP key ID in use.

Obtaining Technical Assistance

Cisco Technical Support provides 24-hour-a-day award-winning technical assistance. The Cisco Technical Support & Documentation website on Cisco.com features extensive online support resources. In addition, if you have a valid Cisco service contract, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not have a valid Cisco service contract, contact your reseller.

Cisco Technical Support & Documentation Website

The Cisco Technical Support & Documentation website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, at this URL:

<http://www.cisco.com/techsupport>

Access to all tools on the Cisco Technical Support & Documentation website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

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

**Note**

Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support & Documentation website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco engineer. The TAC Service Request Tool is located at this URL:

<http://www.cisco.com/techsupport/servicerequest>

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request 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 list of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/techsupport/contacts>

Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

Severity 1 (S1)—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.

Severity 2 (S2)—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.

Severity 3 (S3)—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.

Severity 4 (S4)—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, documentation, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

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

- *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 at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/packet>

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

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

or view the digital edition at this URL:

<http://ciscoiq.texterity.com/ciscoiq/sample/>

- *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>

- Networking products offered by Cisco Systems, as well as customer support services, can be obtained at this URL:

<http://www.cisco.com/en/US/products/index.html>

- Networking Professionals Connection is an interactive website for networking professionals to share questions, suggestions, and information about networking products and technologies with Cisco experts and other networking professionals. Join a discussion at this URL:

<http://www.cisco.com/discuss/networking>

- World-class networking training is available from Cisco. You can view current offerings at this URL:

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



Solution Overview

This chapter presents the following major topics:

- [Solution Description and Scope, page 1-1](#)
- [Solution Components, page 1-3](#)
- [Miscellaneous Solution Support, page 1-4](#)

Solution Description and Scope

The Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband (GOVoBB) Solution, Release 1.0, supports both broadcast video and video on demand (VoD) for the video over broadband/telco market, enabling operators that use digital subscriber lines (DSL) and fiber (FTTx) to offer not only video but also voice over IP (VoIP) and data (Internet access)—collectively referred to as “triple play”—over their existing infrastructure, now intelligently optimized for video service. (The solution assumes that Internet access is already available.)

Generic Architecture and Scope

[Figure 1-1 on page 1-2](#) presents a generic view of the Cisco GOVoBB Solution transport architecture. The solution uses a Gigabit-Ethernet (GE) transport network consisting of the following:

- A super headend (SHE), where live feeds for the broadcast video service are located
- A video headend office (VHO), where the video server complex resides
- A video switching office (VSO), where aggregation routers (ARs) that aggregate local or remotely attached GE DSLAMs are located

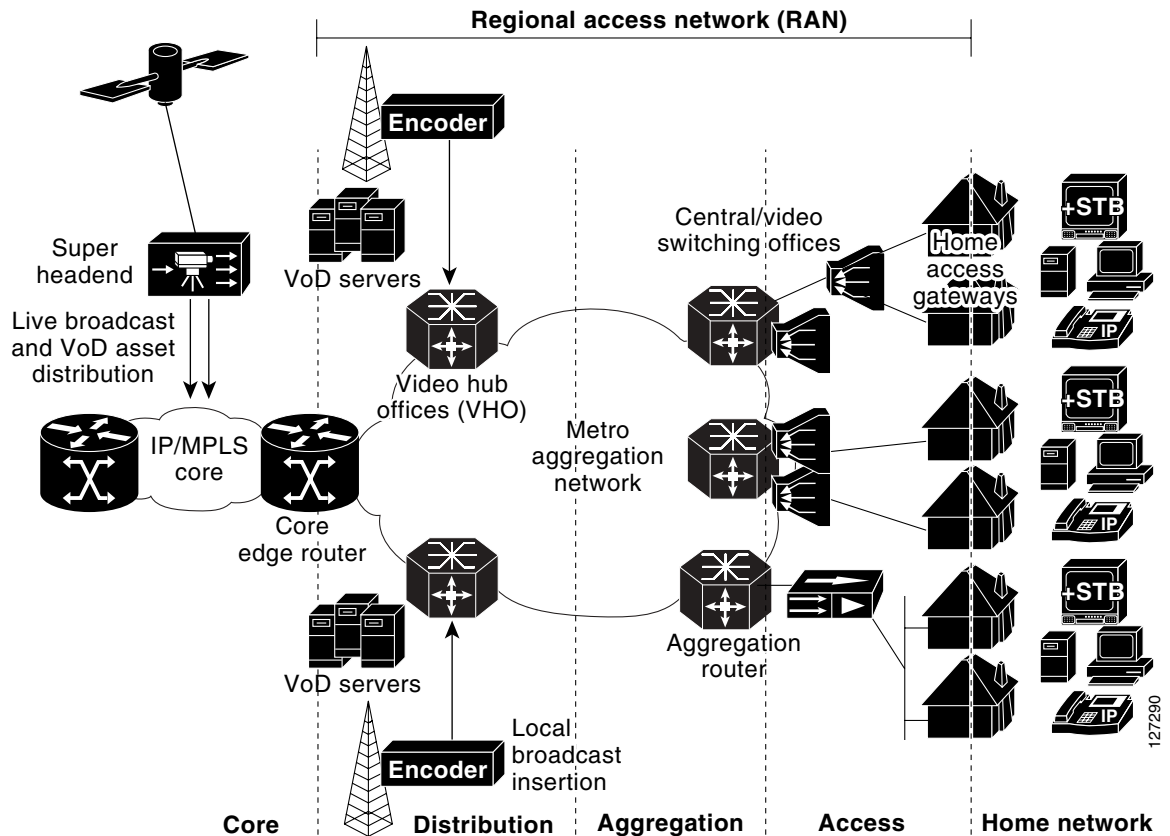
The regional access network, or RAN, consists of distribution, aggregation, and access layers. There is one SHE per region or network, and one VHE per metropolitan area. A distribution edge router (DER) provides transport for video traffic between the IP/MPLS core network and the VHO. The real-time encoder encodes and compresses analog signals. The VHO, in turn, is connected to the VSOs through one or more ARs. The customer premises equipment consists of home access gateways, or HAGs.



Note

MultiProtocol Label Switching (MPLS) was not formally tested as part of the first release of the solution.

Figure 1-1 Cisco GOVoBB Solution Transport Architecture: Generic View

**Note**

For a detailed discussion of the transport architecture, see [Chapter 2, “Video Application Components and Architecture.”](#)

In Scope

The scope of the solution comprises fully tested and supported Cisco components, as well as third-party components tested and supported by Cisco. The following aspects of the solution are fully tested and supported:

- Unidirectional optical transport network for video streams, with 1-GE drop-and-continue, asymmetric switching, and unidirectional link routing (UDLR)
- 10-GE symmetric switching
- Ethernet switching and routing at VHO and VSO interfaces
- Network and element management

**Note**

Management is provided through the Cisco IOS command line interface (CLI) only. See also [Operational Support Systems, page 1-4.](#)

- Multiservice fully converged backbone based on a ring or hub-and-spoke design

Table 1-1 summarizes the correspondence between site types and their transport network types.

Table 1-1 Site Types and Their Transport Network Types

Site Type	Super Headend	Video Headend Office	Video Switching Office	Residence
Transport Network Type	Core	Distribution	Aggregation	Home network

Out of Scope

Not included in the scope of the solution, but still required to support triple play, are items such as subscriber device authentication for one or more of the other nonvideo services. In addition, the architecture of this release places minimal requirements on the DSLAM. This allows the solution to work with as many third-party DSLAMs as possible.

This first release of the solution, which focuses on the metro/distribution-to-subscriber portion of the network, does not specify a transport architecture or test results for video transport over an MPLS core network. Instead, it uses a simplified topology in which the video components associated with the SHE and VHOs are both attached to the DER.

Solution Components

Cisco Equipment

Release 1.0 consists of core Cisco components that are tested, documented, and fully supported by Cisco. Also, third-party equipment, although not fully supported by Cisco, has been selected and tested in conjunction with the core components, to increase the number of test cases and improve the overall quality of the solution in practical networks. The following Cisco equipment has been tested in the context of the solution:

- Cisco 7609
- Cisco Catalyst 6509
- Cisco Catalyst 4948-10GE
- Cisco Catalyst 4510R
- Cisco Catalyst 4507R



Note

For the details of solution components, see [Solution Components, page 3-3](#).

Third-Party Equipment

For this release of the solution, [Table 1-2 on page 1-4](#) lists the third-party vendors and the basic functionality they provide. (For detailed descriptions of video functions, see [Video Application Components, page 2-1](#).)

Table 1-2 *Component Partners and Basic Functionality*

Vendor	Basic Functionality
Kasenna www.kasenna.com	VoD server Middleware
Amino www.aminocom.com	Set-top box
Ericsson www.ericsson.com/	DSLAM, residential gateway

For more detail, including product names and part numbers, see [Table 3-1 on page 3-3](#).

Miscellaneous Solution Support

This section clarifies the degree to which other aspects of the solution and its implementation are supported in this first release.

Operational Support Systems

Release 1.0 does not certify element management systems (EMSs) or network management systems (NMSs) operated within the context of the Cisco GOVoBB architecture. Customers continue to provide such capabilities as applicable to their particular environments. All the management information base (MIB) components for the Cisco equipment are available from Cisco, and can be incorporated into the customer's current EMS.

Billing

Billing is outside the scope of this first release of the solution.

EMC

Release 1.0, with all its platforms, accessories, and components, complies with applicable electromagnetic compliance (EMC) standards.

Safety

Release 1.0, with all its platforms, accessories, and components, complies with applicable safety standards.



Video Application Components and Architecture

This chapter discusses the segmentation of the video application architecture into logical components that are required for broadcast video and video on demand (VoD) services. The function of each component is described, as well as the basic interfaces needed between each component and other components of the system.

This chapter describes possible video architectures and components only. For the actual tested implementation, see [Chapter 4, “Implementing and Configuring the Solution.”](#)



Note

Because there are currently few standards regarding application architectures for either broadcast or on-demand IPTV/video service over a DSL infrastructure, this solution makes no specific assumptions regarding the application architectures implemented by the vendors of specific video equipment. However, although there are few standards for video application architectures, the functionality implemented is fairly consistent from vendor to vendor.

For a list of the video components that were tested in this release, including product names and part numbers, see [Table 3-1 on page 3-3](#).

This chapter presents the following major topics:

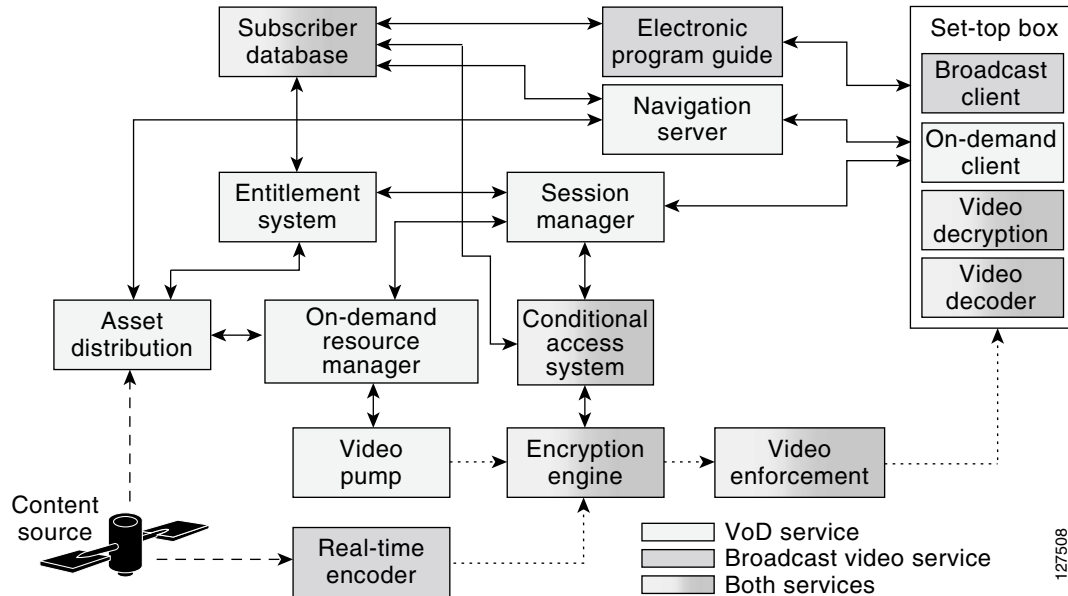
- [Video Application Components, page 2-1](#)
- [IPTV/VoBB Transport Architecture and Issues, page 2-9](#)

Video Application Components

[Figure 2-1 on page 2-2](#) illustrates the logical relationship of the application-layer video components needed to deliver broadcast video and VoD services, as well as the basic interfaces between components. Components can be categorized as follows:

- [Broadcast Video Components](#)
- [VoD Components](#)
- [Common Broadcast Video and VoD Components](#)

Figure 2-1 Video Application Component Architecture



Broadcast Video Components

Broadcast video components (see [Figure 2-1](#)) include the following:

- [Real-Time Encoder](#)
- [Electronic Program Guide](#)
- [Broadcast Client](#)

These are described below.

Real-Time Encoder

The real-time encoder takes a live feed from a broadcaster in either analog or digital format and converts it into a compressed digital stream that is encapsulated in IP packets. The input to the encoder may be in a digital format that uses digital MPEG-2 over ASI format, or it may be in an NTSC, PAL, SECAM, or other analog format. The output of the encoder is a digitally compressed stream that is encapsulated in IP headers and sent to a multicast group address. The compression method used by the encoder may be either MPEG-2, WM9, or MPEG-4/AVC, while the IP-based transport encapsulation used is MPEG-2 transport over either UDP/IP or IP/UDP/RTP. Since the real-time encoder is configured to encode a specific channel, no control interfaces are required between it and other video components.

Electronic Program Guide

The electronic program guide (EPG) provides information about available broadcast channels to the broadcast client application running on the IP set-top box (STB). The EPG is often implemented as an HTTP server and formats available channel listings as web pages. The EPG application authenticates and authorizes a subscriber for broadcast services. The EPG may also provide a customized view of channel listings that is based on the packages a particular subscriber has subscribed to. Both of these functions require an interface between the EPG application and the subscriber database. In addition to providing

a graphical listing of available channels, the EPG provides the IP multicast address to which the channel is sent in the IP network. The broadcast client uses this address in Internet Group Management Protocol (IGMP) messages that are sent during the processing of a channel change.

Broadcast Client

The broadcast client is an application process running on the STB that is responsible for providing the user and control interface for broadcast video services. The broadcast client, in conjunction with the EPG, implements a subscriber authentication interface for set-top-based services. Authentication is typically done by means of an application layer authentication protocol such as HTTP in conjunction with a shared secret such as a username/PIN pair.

The broadcast client displays available broadcast-channel information using data from the EPG and implements the control interface for channel change by means of IGMP. Since the DSL line may be capable of supporting the bandwidth of only a single broadcast channel, the IGMP process for changing channels must ensure that only a single video broadcast stream is sent to the STB at a time. The broadcast client implements this by sending an IGMP leave for the current channel and then waiting for a configurable period of time for the broadcast stream to stop. After this timer expires, the broadcast client sends an IGMP join for the new channel. The full channel-change time, documented in [Broadcast Video Channel-Change Time, page 2-14](#), includes these IGMP factors as well as other factors specific to video compression. (See also [Effect of Multicast on Channel-Change Performance, page 3-19](#).)

VoD Components

VoD components (see [Figure 2-1 on page 2-2](#)) include the following:

- [Asset Distribution System](#)
- [Navigation Server](#)
- [Session Manager](#)
- [Entitlement System](#)
- [Video Pump](#)
- [On-Demand Resource Manager](#)
- [On-Demand Client](#)

These are described below.

Asset Distribution System

The asset distribution system (ADS) takes video content from content providers and uses business rules to distribute that content to different locations in the video service provider's network. Video content may be provided to the ADS through a number of different methods. These methods include the use of pitcher/catcher systems, which receive video content from content providers over satellite links, and manual processes such as file copies from other network servers.

Standard video content objects include the actual MPEG video, images for display during content navigation, trailers, and metadata files that provide information about the files contained in the object.

The ADS may be used to modify the metadata of a video asset to add business rules such as the price of the video, the distribution window, the VoD subscription package that the video is part of, whether the content needed to be encrypted, and so on. On the basis of these business rules, the ADS replicates the video asset to the on-demand resource management component of video servers in different locations.

Navigation Server

The navigation server provides information about available VoD content to the on-demand client application running on the STB. The navigation server is often implemented as an HTTP server and formats available VoD content as web pages. The navigation server uses information provided by the asset management system to determine which VoD content to display to the subscriber. For subscription-based VoD services, the navigation server may use the information in the subscriber database to customize the view of the video content presented to the subscriber, depending on the packages the subscriber has purchased.

Session Manager

The session manager is the central point of communication for VoD session requests that originate from the on-demand client on the STB. It manages the life cycle of a video session and is responsible for arbitrating the various resources required to deliver the video stream associated with the on-demand session request. Many vendors of VoD equipment and software provide a logical “session manager” function, though this function goes by a variety of different names.

When the session manager receives an on-demand session request from an on-demand client application, it first uses the services of the entitlement system to determine whether the subscriber is authorized to view the requested video content. If the request is authorized, the entitlement server includes additional information in the authorization response, such as the encryption format to be used for the content.

When the session manager receives the authorization response, it determines the best VoD server complex to use for the session request, based on the subscriber’s IP subnet. The session manager then contacts the on-demand resource manager for that VoD server complex to request a video pump for the session. If the VoD content needs to be encrypted in real time, the session manager contacts the conditional access system (CAS) to request a real-time encryption engine for the session. The CAS responds with the decryption keys to be used by the STB to decrypt the video stream.

After all of the resources for a VoD session request are obtained, the session manager responds to the on-demand client with information about the IP/UDP/RTP transport parameters for the video stream to the STB. If the stream is to be encrypted, the session manager (or a key manager with which it coordinates) includes the decryption keys for encrypted video content in the response as well. Finally, the session manager includes the IP address of the video pump that was selected for the session. The IP address of the VoD pump is needed by the on-demand client in order to send stream control requests through Real Time Streaming Protocol (RTSP)—such as pause, fast forward, rewind—for the session.

Entitlement System

The entitlement system is responsible for determining whether the movie requested by an on-demand client is authorized for viewing by the subscriber associated with that client. The entitlement system uses information from the ADS to build a database indicating which videos are associated with different on-demand subscription packages. When the entitlement system receives an entitlement request from the session manager, it uses this database to determine with which orderable on-demand package the requested video is associated. The entitlement system then uses the services of the subscriber database to determine whether the subscriber associated with the entitlement request is entitled to view the requested video.

Video Pump

The video pump is the streaming storage component of a VoD system. The video pump contains locally or remotely connected storage that is organized in such a way that it can send any piece of stored media at a known constant rate. The streaming portion of the video pump is responsible for pulling requested files from the storage system and for pacing the output of each requested file to the network through the use of a shaper. Video pumps must be capable of implementing basic stream control functionality, such as fast-forward and rewind, to respond to requests from the on-demand client during the playout of a media file.

In addition to being able to stream media out, video pumps are also responsible for ingesting media for storage in the storage subsystem. While in general the functionality of a video pump is fairly independent of media format, the ingest function may have functionality that is specific to a particular media format. An example of this type of media-format dependence is the generation of trick files for use with stream control functionality such as fast-forward and rewind. Video pumps used in broadband environments are typically capable of storing and streaming both MPEG-2 and MPEG-4 content. (Only MPEG-2 was tested in this release.)

On-Demand Resource Manager

The on-demand resource manager (ODRM) is responsible for managing the streaming resources and storage of a group of video pumps. The ODRM is responsible for locating and replicating content, as well as for allocating video pumps for the on-demand session requests it receives from the session manager.

On the ingest side, the ODRM is responsible for taking content received from an asset management system and replicating it to one or more of the video pumps it controls. The ODRM makes decisions on when and where to replicate content on the basis of such information as asset metadata and the demand for each title (as indicated by on-demand session requests).

On the streaming side, the ODRM responds to on-demand session requests from a session manager by locating a video pump that contains the requested title, has the capacity to stream the title, and is capable of reaching over the transport network the subscriber that generated the session request.

On-Demand Client

The on-demand client (ODC), an application process running on the STB, is responsible for providing the user and control interface for on-demand services. The ODC provides the user interface for browsing on-demand content using the services of the navigation server. The browsing interface of the ODC is typically implemented by means of an embedded HTTP-based browsing application.

The ODC interfaces to the session manager to make requests to stream on-demand content. It also interfaces to video pumps to make stream-control requests for movies that are actively being streamed.

Common Broadcast Video and VoD Components

Common broadcast video and VoD components (see [Figure 2-1 on page 2-2](#)) include the following:

- [Conditional Access System and Encryption Engine](#)
- [Broadcast Video Bandwidth Enforcement](#)
- [Set-Top-Based Video Decryption and Video Decoder](#)
- [Set-Top Box](#)

- [Subscriber Database](#)

These are described below.

Conditional Access System and Encryption Engine

The conditional access system (CAS) is responsible for the key management and distribution infrastructure associated with the encryption of video content. Video encryption is used as the second tier of protection against theft of content. The first tier of protection for both broadcast and on-demand services is performed as part of the on-demand and broadcast client applications running on the STB. These applications use the services of the EPG and navigation server to authenticate the subscriber and provide a customized view of available channels and content based on the services the subscriber has purchased. For on-demand services, the entitlement system also checks whether the subscriber is authorized to view requested titles, with the result that the ODC does not allow the subscriber to view unauthorized content. While application-layer authorization protects against the theft of content from authorized STBs, it does not protect the video stream itself. Video encryption using CAS provides this second layer of protection against theft and unauthorized viewing of video content.



Note

CAS was not implemented as part of the first release of the solution.

Because conditional access adds an additional level of complexity and cost to a video delivery system, service providers typically use CAS-based encryption only on premium-tier broadcast channels and on-demand titles. For broadcast services, encryption must be done in real time as the video stream is delivered. For on-demand services, encryption may be done either in real time as the content is streamed or as part of the process of replicating content to video pumps. The process of encrypting video content as part of replication is called pre-encryption.

Video encryption may be done on either a tier or session basis. In tier-based video encryption, a single set of encryption/decryption keys is used for all of the video content associated with a particular service offering. Subscribers that are authorized to view the content associated with the service are delivered the decryption keys needed for that service ahead of time. Conditional access for broadcast video services is always implemented by means of tier-based encryption, because a single video stream may be viewed by many subscribers simultaneously. Decryption keys for broadcast video services are delivered in a secure manner to the STB through the EPG. In session-based video encryption, decryption keys for a piece of content are generated and delivered to the subscriber on a per-session basis. Session-based encryption may be used with VoD content. Because decryption keys are generated only on a per-session basis for session-based encryption, they may be used with either real-time or pre-encryption techniques.

In a typical CAS, the encryption of digital services can be achieved by using entitlement control messages (ECMs) and entitlement management messages (EMMs). In order to generate the final keys needed to decrypt a particular video stream, the STB must receive and decrypt the correct ECMs and EMMs. EMMs provide keys that can be decrypted only by a specific subscriber, while ECMs provide keys that are specific to a particular video stream. Because EMMs are specific to a subscriber, they are always generated ahead of time. Because ECMs are specific to a particular video stream, they may be generated ahead of time when pre-encryption is used, or they may be generated in real time when real-time encryption is used. ECMs are typically delivered in band as a component of the video stream.

Whether the content must be encrypted may be determined by a number of factors. For on-demand services, content providers can require content to be encrypted by enabling the “Encryption” field in the metadata file associated with the video asset. For broadcast services, the video service provider statically configures the video stream from real-time encoders to be sent either directly to a multicast group or to a real-time encryption engine, depending on whether that channel is to be encrypted.

The encryption engine takes MPEG streams in and encrypts them in real time using encryption keys received from the CAS. Encryption engines typically use a DES algorithm for encryption.

Broadcast Video Bandwidth Enforcement

Broadcast video bandwidth enforcement in Release 1.0 is implemented as part of the functionality of the aggregation router (AR). The AR enforces a maximum broadcast bandwidth limit by limiting the number of IGMP joins on the ranges of multicast addresses associated with broadcast video to a configured maximum on the aggregation links it controls. The mapping of video channels to multicast addresses can be done in such a way that the AR can associate the bandwidth for different classes of video (for example, standard definition or high definition) with different ranges of multicast addresses. IGMP join limits can then be set for each range of multicast addresses.

For more details on the network enforcement for video broadcast, refer to [Multicast Admission Control, page 3-17](#).

Set-Top-Based Video Decryption and Video Decoder

The set-top box, or STB (see below), includes two components that are responsible for turning the incoming video stream, delivered as IP packets, into an uncompressed digital stream that can be directly turned into an analog TV signal ready for display by a television set. These components are the video decoder and the video decryptor.

The video decoder is responsible for decompressing the incoming video stream. It uses a decompression algorithm that is matched to the compression algorithm used by the real-time encoders for broadcast services. The video decoder may also support additional decompression algorithms for VoD services if VoD assets are compressed by a different algorithm than broadcast channels use.

The video decryptor is responsible for performing decryption on the video stream if the stream was encrypted by the encryption engine when real-time encryption is used, or by an offline encryptor when pre-encryption is used for on-demand assets.

Set-Top Box

The set-top box (STB) is the hardware and common software infrastructure component that is used by the on-demand and broadcast clients as well as by the video decryptor and the video decoder. The STB hardware typically consists of a general-purpose processor and video subsystem that produces an analog television output. The hardware may also include a hardware-based decoder and decryption subsystem. The STB software typically includes an embedded operating system, and may also include application infrastructure components such as a web browser.

Subscriber Database

The subscriber database contains service-level information about each subscriber, for example, which services the subscriber is authorized to use, information used for billing, and so on. The subscriber database may also contain information that can be used for subscriber authentication. An example of this type of information is the username/password information that is used by the EPG to identify and authenticate a subscriber for broadcast services.

IPTV/VoBB Product Architecture

This section describes how the logical components described previously are commonly combined into products that are supplied by current video-component vendors. Figure 2-2 illustrates how the video components presented in Figure 2-1 on page 2-2 are “bundled” into application products. This bundling reduces the number of products and vendors that must be integrated to build a complete video system. It also reduces the number of interfaces that must be agreed upon by vendors.



Note

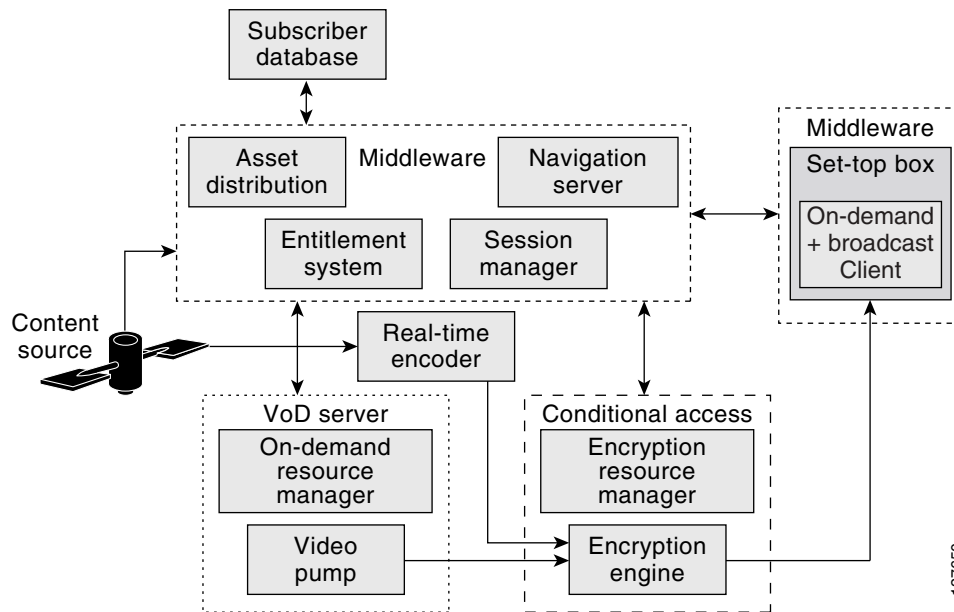
The logical components described in this section represent functional blocks that are common to most video application architectures, and do not necessarily reflect how these functions are bundled into products by video equipment vendors.

The following classes of video products are needed to build a complete broadband video solution:

- [Middleware](#)
- [VoD Server](#)
- [Conditional Access System](#)
- [Real-Time Encoder and Set-Top Box](#)

These are described below.

Figure 2-2 Common Video Components



Middleware

Middleware, as defined, has the role of gluing a number of logical components together into a more comprehensive IPTV/video software system. (Note that there are several different middleware implementations. Thus, the following description is a typical example for illustrative purposes.) Middleware implements the user interface for both broadcast and on-demand services. It is also used as

the glue software that integrates products from other vendors into an application level solution. Middleware products are often used to integrate multiple VoD servers, conditional access systems, and set-top boxes from different vendors into the same deployment.

Middleware provides the client and server functionality that implements the user interface for both broadcast and on-demand services. The components that provide the client-side functionality are the broadcast and on-demand client applications on the STB, while the components that provide the server-side functionality are the electronic program guide and the navigation server.

Middleware uses the entitlement system and session manager components to integrate the VoD servers used in an on-demand service. The entitlement system integrates the asset ingest function of a VoD server, while the session manager integrates the session plane of the VoD server into an on-demand service.

Middleware uses the session manager and on-demand client to integrate CAS into an on-demand service. These components may be used to pass decryption keys from the conditional access system to the video decryption component in the STB. These components also determine when to use the services of the CAS based on the encryption requirements of the service and each asset associated with the service. Middleware uses the EPG and the broadcast client to integrate CAS into a broadcast service. The broadcast client determines when to use the services of the CAS based on information it obtains from the EPG on each broadcast channel.

VoD Server

The VoD server (one or several) implements storage and real-time streaming functionality for on-demand services. The VoD server consists of a set of video pumps that are managed by an on-demand resource manager. The VoD server integrates with middleware and may also be integrated with the CAS when preencryption is used.

Conditional Access System

The conditional access system (CAS) provides encryption and decryption services, as well as key generation and distribution functionality, for both broadcast and on-demand services. The CAS consists of the encryption resource manager, the encryption engine, and the video decryption process in the STB.

The CAS interfaces to middleware when session-based encryption is used for on-demand services. The CAS may also interface to middleware for encryption key distribution between the encryption resource manager and the decryption process on the STB. Finally, the CAS interfaces to VoD servers where preencryption is used for on-demand content.

Real-Time Encoder and Set-Top Box

The real-time encoder and STB components described in [Real-Time Encoder, page 2-2](#), and [Set-Top Box, page 2-7](#), respectively, are identical to product classes of the same name shown in [Figure 2-2 on page 2-8](#).

IPTV/VoBB Transport Architecture and Issues

To meet the end-to-end transport requirements for broadcast VoD services, the IPTV/VoBB transport architecture provides functional requirements and configuration recommendations for each switching node in the path from the VoD servers to the STBs. This section presents the following topics:

- [Video Sites](#)

- [Video Service Requirements](#)
- [Potential Video Service Architectures](#)
- [Service Separation in a Triple-Play Architecture](#)

**Note**

Although this solution is focused on video service, it must work within the context of a triple-play solution. Because VoBB services are fairly new, vendors and service providers do not use the same terminology to describe the major sites. This section describes terminology commonly used for triple-play solutions.

Video Sites

The video sites described in this section are the super headend (SHE), the video headend office (VHO), and the video switching office (VSO). [Figure 1-1 on page 1-2](#) shows the location and roles of the sites and components in a typical IPTV/VoBB deployment.

Super Headend

The SHE is where live feeds for the broadcast video service are located. This site contains the real-time encoders used for the broadcast video service, along with the asset distribution systems for on-demand services. This site may also contain back-office systems such as the subscriber database. Most IPTV/VoBB deployments have a single SHE site; this is the source of most of the multicast streams for the broadcast video service. The SHE typically resides in the core of the transport network.

Video Headend Office

The manned VHO is where the video server complex resides (as well as where optional local/PEG content may be inserted). The VHO is where the majority of the video pumps used for on-demand services are typically located. It is also where the real-time encoders for local television stations reside. A VHO typically serves a metropolitan area of between 100,000 and 1,000,000 homes. The VHO is equivalent to (and may be contained in) the same facility as the point of presence (POP) for Internet access services. Transport for video traffic between the VHO and IP/MPLS core network is provided by a distribution edge router (DER). The DER interconnects the core network and the local video sources to a high-bandwidth distribution network that carries both broadcast and on-demand video to VSOs.

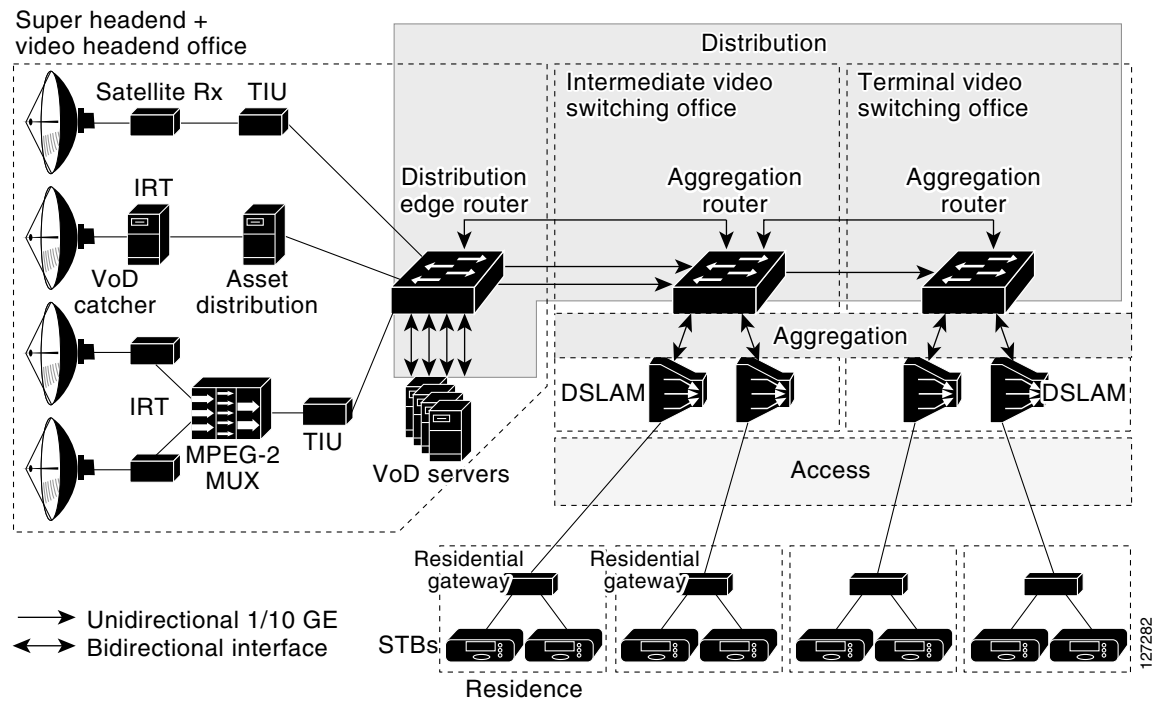
Video Switching Office

The VSOs house the aggregation routers that aggregate local or remotely attached GE DSLAMs. The VSO is typically located in the central switching office. The central switching office is the physical termination point for the majority of the copper loops for the residences it serves. Because ADSL and ADSL2+ rely on short loop lengths to obtain maximum training rates and throughput, the copper loops used for DSL service are often terminated at a location closer to the subscriber than the VSO. This means that the DSLAMs that the VSO aggregates may or may not be collocated in the VSO. The switching equipment in the VSO interconnects the aggregation and distribution networks. Traffic to and from the DSLAMs is aggregated by the aggregation router (AR). The AR resides in intermediate and terminal VSOs.

In order to minimize the bandwidth requirements between the VSO and the VHO, a VSO may include local video pumps that are used to cache popular on-demand content. While the Release 1.0 transport architecture does not preclude the use of video pumps in the VSO, this configuration is not tested as part of the solution test effort.

Figure 2-3 presents a more detailed view of the distribution, aggregation, and access layers of the IPTV/VoBB transport architecture, with video streams acquired, multiplexed, and distributed from the SHE through the VSO (not shown). Both unidirectional and bidirectional interfaces are shown.

Figure 2-3 IPTV/Video over Broadband Transport Architecture: Detailed View



Video Service Requirements

In order to understand better some of the design tradeoffs associated with the transport architecture, it is important to understand common requirements for a video service and how an IP network can be optimized to meet these requirements. This section outlines some common requirements for broadcast video and VoD services. It also describes what design factors in the transport network are relevant to these requirements.

High Bandwidth

The amount of bandwidth that a network must be capable of transporting to support video services is typically an order of magnitude more than what is required to support voice and Internet access services. A standard-definition IP video stream that is carried as an MPEG-2 SPTS stream over RTP uses about 3.75 Mbps of bandwidth. A high-definition IP video stream using the same type of compression and transport uses about 16 to 18 Mbps of bandwidth.

These bandwidth requirements mean that a DSL access infrastructure that is designed for real-time video transport must be capable of carrying significantly more bandwidth than what is needed for VoIP and Internet access services. It also means that the DSL line itself is typically constrained to carrying only one or two video streams simultaneously. The result is that video over DSL service offerings must limit the service to one or two simultaneous broadcast channels or on-demand sessions to a household.

Because the video streams associated with on-demand services are unicast while the video streams associated with broadcast services are multicast, the amount of bandwidth required in the aggregation and distribution networks to carry on-demand streams is much greater than what is required for broadcast services. Also, because broadcast video services use multicast, the amount of bandwidth required in the access and distribution networks scales with the number of channels offered. As an example, a broadcast video service that uses MPEG-2 compression and offers 300 channels of standard-definition content requires more than 1 Gbps/sec of capacity in the distribution network to handle worst-case usage patterns. On-demand services use unicast, so the amount of bandwidth required in the distribution network scales with the number of subscribers and peak on-demand utilization rates the network is designed to carry. For example, a distribution network that is designed to deliver MPEG-2 compressed standard-definition content to 50,000 on-demand subscribers at a 10% peak take rate requires about 19 Gbps of capacity.

Asymmetric Bandwidth

Video traffic is inherently asymmetric, as both broadcast video and VoD flows are unidirectional. The only traffic that is sent in the upstream direction of either service is control traffic that is used to instantiate the video flow. For on-demand services, this control traffic is the session and resource signaling that is described as part of the component descriptions in [Video Application Components, page 2-1](#). For broadcast services, the control traffic is IGMP and PIM signaling that is used to instantiate the multicast flow for the broadcast channel.

Because of this asymmetry, the cost of the distribution network can be reduced by incorporating unidirectional links in the transport path. One of the transport alternatives tested in this release of the solution includes unidirectional transport in the distribution network.

Quality of Service

When broadcast and on-demand video is carried over an IP network, there is an assumption that the video quality experienced by the subscriber is comparable to that experienced by those watching MPEG-2 video carried by cable and satellite networks today. To ensure that any degradation in video quality resulting from the IP transport network is negligible from a subscriber's point of view, most providers allow only one visible degradation in video quality roughly every two hours.

While this allowance is similar to what is allowed for VoIP services, the resulting allowed packet-drop requirement for an IP transport network designed for video services is much more stringent. There are two reasons for this:

- Video is much more highly compressed, so losing a packet may result in the loss of more-valuable encoded video information. If the network drops a single video packet, there is a visible degradation of video quality of anywhere from a single frame up to a loss of one second of video, depending on the kind of encoded information that is lost.
- The receiving decoders, such as the STBs, generally do not have loss-concealment algorithms, whereas VoIP phones and gateways typically support algorithms that conceal dropouts in the voice signal caused by lost packets. In the VoIP case, the network can drop a single voice packet without the listener perceiving any degradation in voice quality—unlike the case for video.

The DiffServ (Differentiated Services) architecture defines packet marking and scheduling behaviors that can be used to ensure that video flows meet the required 10^{-6} drop rate when links are congested. (For details on the QoS architecture for Release 1.0, see [QoS Architecture, page 3-46](#).) Packet drops due to bit errors on physical links must be addressed on a link-by-link basis.

The link-layer technologies used in video networks employ cyclic redundancy check (CRC) algorithms to ensure that packets with errors are not delivered. This means that a single bit error in a video packet results in that packet being dropped when a CRC is performed. Video over IP is typically carried in packets that are approximately 1400 bytes. If bit errors are assumed to be distributed randomly, the resulting requirement for transport links is to ensure a bit-error rate (BER) less than 10^{-10} .

The BER on optical links can be engineered to 10^{-14} or less by ensuring a high signal-to-noise ratio (SNR) on those links. Because Release 1.0 uses optical connectivity in the access and distribution networks, degradation in video quality resulting from bit errors on these links should not be an issue.

However, packet drops due to bit errors on the DSL line can have a significant effect on video quality. The SNR on a DSL line varies as a result of many factors, including loop length, proximity to noise sources, and other factors. In addition, the SNR may vary over time because of factors such as corrosion at connection points, moisture, and so on. Consequently, it may be difficult to qualify a DSL line at the time of installation to ensure a BER of less than 10^{-10} over the life of the video service.

Multiple technologies are available to deal with bit errors on the DSL line. Some common technologies are forward error correction (FEC) and real-time retransmission (RTR). While Release 1.0 does not include the testing of these technologies, future releases of the solution will include technologies to help deal with bit errors on the DSL line.

Service Availability

Service providers deploying video services often have different availability requirements for VoD and broadcast video services, as contrasted below.

Broadcast video services are inherently real time. A subscriber who experiences an outage in the broadcast service cannot come back and continue watching at that point when the outage is over. Because of this and the higher usage rates associated with broadcast services, the availability associated with broadcast services must be very high.

In contrast, the customer disruptions associated with an outage in VoD services are typically much less problematic. A subscriber who experiences an outage in a VoD service can come back at a later time and replay the content—either from the point of disruption or from the beginning. In addition, the peak usage rates associated with VoD are typically between 10 and 20% of the subscriber population. This is much lower than the peak usage rates for broadcast services.

Because of the above factors, service providers have much higher availability requirements for broadcast services than for on-demand services. Consequently, the differing availability requirements between the two services may result in differing transport requirements for each service. For example, the high-availability requirement for broadcast video typically results in the requirement that there be redundant transport paths between the DER and AR nodes of the distribution network. (See [Figure 2-3 on page 2-11](#).) Because of the higher bandwidth and lower availability requirements associated with VoD services, the topologies used for these services may not necessarily require redundant transport paths.

The test topologies for Release 1.0 include a distribution network design that provides path redundancy for both services, as well as a cost-optimized distribution design that provides path redundancy only for broadcast services. In addition, the quality of service (QoS) architecture includes DiffServ marking for broadcast and on-demand services, allowing the network to drop VoD traffic preferentially over broadcast traffic in the event of a network outage. Finally, Release 1.0 supports redundant broadcast video encoders, as well as a method to fail over in a timely manner from one encoder to another.

Broadcast Video Channel-Change Time

An important aspect of a broadcast video service is the amount of time it takes for the system to respond to a channel-change request from a subscriber. While the channel-change time for current analog broadcast services is perceived by the subscriber to be instantaneous, the channel-change time for digital broadcast services is between one and one-and-a-half seconds. The majority of this time is due to the differential encoding and decoding methods used to compress digital video streams.

To reduce the amount of bandwidth required for digital video transmission, compression methods such as MPEG compress the video frames of a digital video stream into three different types of frames. These frames are called I-frames, B-frames, and P-frames. An I-frame is a compressed version of all of the information in one frame of a video stream. An MPEG decompressor can recreate the original frame using just the information in the I-frame. A P-frame is an incrementally encoded video frame that can be decoded with the information in the preceding anchor frame (I-frame or P-frame). A B-frame is an incrementally encoded video frame that can be decoded with the information in the preceding and following anchor frames (I-frame or P-frame).

Because of incremental coding, an important factor in how long it takes to change a channel for a digital video service is the I-frame gap. The I-frame gap defines how often I-frames are included in the MPEG stream. Shorter I-frame gaps result in shorter channel-change times, while longer I-frame gaps result in longer channel-change times.

When a digital broadcast service is run over a DSL access infrastructure, the following additional factors must be added to the delay caused by the I-frame gap:

- STB performance in processing a channel-change request
- Multicast latency in terminating the IP video feed associated with the “tuned from” channel
- Multicast latency in joining the IP video feed associated with the “tuned to” channel
- Whether or not the “tuned to” channel is encrypted by means of a CAS
- Delay to the next cryptoperiod and the time needed to acquire CAS/DRM (digital rights management) decryption keys before the decryption of the “tuned to” channel begins
- Delay in refilling the jitter buffer for the decoder in the STB

The goal of this solution is to provide subscribers with a channel-change experience similar to that currently experienced for digital broadcast services. Most of the additional channel-change delay associated with a DSL access infrastructure is due to the amount of time it takes for the network to stop sending the multicast stream for the “tuned from” channel and to begin sending the multicast stream for the “tuned to” channel. [Distribution and Aggregation Transport Architecture, page 3-4](#), provides a recommendation for a scalable multicast architecture that best meets the channel-change requirements for broadcast video services.

Potential Video Service Architectures

One aspect of a transport architecture for video that must be considered initially is how the service provider sells the video service to the subscriber. This section examines how two potential video service-level agreement (SLA) models affect the requirements of a transport network implemented to deliver the service to the subscriber.

- The SLA for a video transport service is based on transport parameters. A typical transport-based SLA includes factors such as maximum bandwidth, packet-loss rate guarantees, and jitter and latency guarantees.
- The SLA for an application service is based on service-level parameters. A typical video application-based SLA includes the following:

- The number of simultaneous video channels (live or on-demand) a subscriber is authorized to view
- The broadcast channel line-up (basic or premium tier) that the subscriber has signed up for
- Any subscription VoD content that the subscriber has signed up for

The services the network provides to deliver a transport-based SLA as opposed to an application-based SLA are very different. [Table 2-1](#) provides an overview of the technologies used to deliver the basic functionality of a transport service as opposed to an application service.

Table 2-1 Service-Delivery Technologies: Transport vs. Application

Service Type	Transport Service	Managed Application Service
SLA	Transport parameters: <ul style="list-style-type: none"> • Bandwidth • Max. drop • Max. latency • Etc. 	Video application SLA: <ul style="list-style-type: none"> • Number of STBs • Basic vs. premium tier
Subscriber authentication/identification	Network based (examples): <ul style="list-style-type: none"> • PPPoE • 802.1x • Per-subscriber VLANs • DHCP option 82 	Application based: <ul style="list-style-type: none"> • Video middleware
SLA enforcement	Network based: <ul style="list-style-type: none"> • Per-subscriber shaping/policing 	Application based: <ul style="list-style-type: none"> • Based on application signaling
QoS	Per subscriber: <ul style="list-style-type: none"> • Gold, silver, bronze • Classification • Queueing 	Aggregate: <ul style="list-style-type: none"> • Single queue for video service¹

1. Assumes all network devices can support DiffServ-type congestion management

The two SLA models are examined in detail in the following sections:

- [Transport-Based SLA](#)
- [Managed Application-Based SLA](#)

Transport-Based SLA

Subscriber authentication and identification for a transport service is done at the transport layer. Subscriber authentication technologies rely on shared secrets such as passwords or private/public key pairs to establish a trust relationship between the subscriber and the network. Subscriber identification technologies use a well-known property of a subscriber (such as the DSL line to which the subscriber is attached) to identify all packets coming from or to the subscriber. Transport SLA enforcement requires a subscriber identification technology and may also include a subscriber authentication technology.

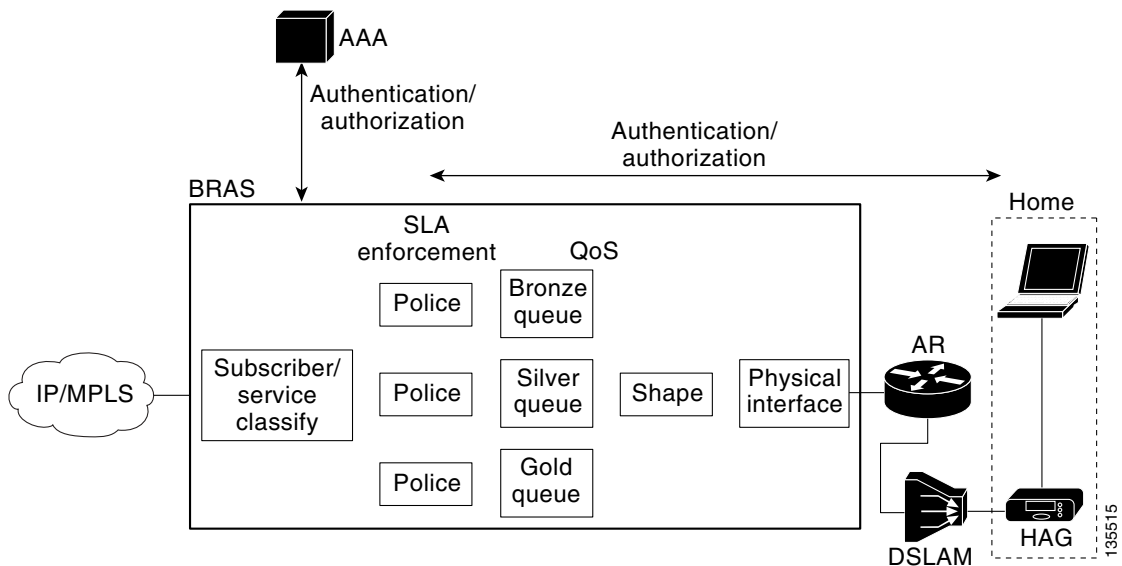
Common subscriber-authentication technologies used for a transport service include Point-to-Point Protocol over Ethernet (PPPoE) and IEEE 802.1x. These technologies are used to authenticate a subscriber transport session. To enforce a subscriber's transport SLA at the transport layer in PPPoE environments, every packet associated with a subscriber's transport session can be identified with a PPPoE session ID that is specified as part of the PPPoE tunnel encapsulation. To associate packets with an 802.1x transport session, one either incorporates SLA enforcement functionality in the switching node attached to the physical link terminating the 802.1x session, or switches the packets from the physical link terminating the 802.1x session into a VLAN that is terminated on the node that enforces the SLA. In the absence of an authentication protocol, either VLAN tags or DHCP option 82 could be used to identify the DSL line that every packet is coming from or going to. These technologies could be used to enforce a subscriber transport SLA without the use of an explicit subscriber-authentication protocol.

**Note**

PPPoE-based architectures could also use a VLAN tag as opposed to a PPPoE session ID to identify the traffic associated with a particular PPPoE.

SLA enforcement used for transport services relies on per-subscriber shaping or policing. The resulting QoS architecture relies on per-service-classification queuing and scheduling. SLA enforcement is typically implemented in the same node that terminates the transport session (PPPoE or 802.1x). Packets are classified per subscriber according to the transport session identifiers described above. The downstream traffic for each subscriber is typically shaped to a maximum rate based on the parameters of the transport SLA. If the transport SLA sold to the subscriber includes more than one class of service (gold, silver, or bronze), then additional classification, queuing, and scheduling are done to guarantee the transport parameters of the SLA associated with each class. For transport services, the node that terminates the transport session and enforces the subscriber SLA is typically the broadband remote-access server (BRAS). [Figure 2-4](#) illustrates the per subscriber control and data plane functionality used by the network to implement a transport service.

Figure 2-4 Per-Subscriber Control and Data-Plane Functionality Used to Implement a Transport Service

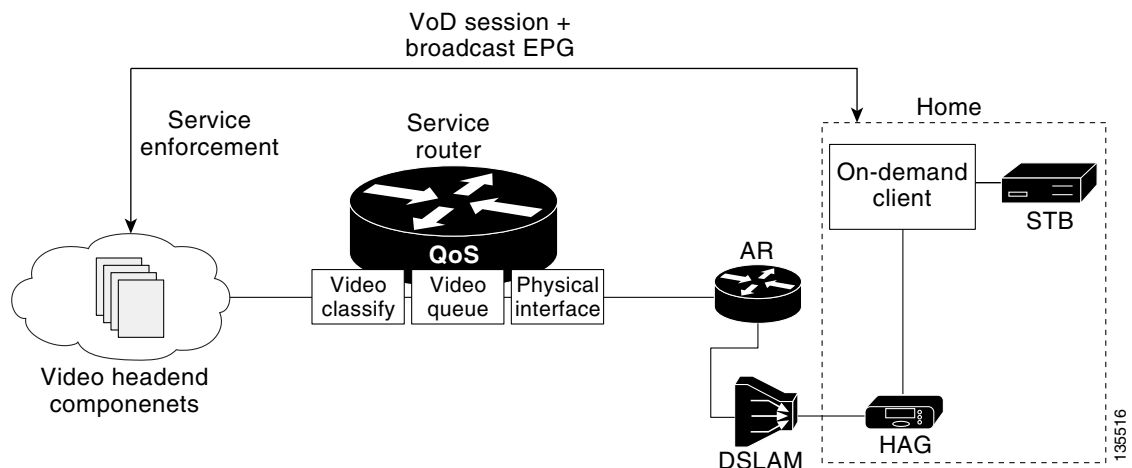


Managed Application-Based SLA

Subscriber authentication for an application service is implemented by means of application-aware components. For example, [Electronic Program Guide, page 2-2](#), describes how subscriber authentication for a video service is typically implemented as part of the electronic program guide (EPG) function. The EPG, a component of video middleware, often authenticates a subscriber's video STB by means of an application-layer challenge such as an HTTP authentication challenge. If the EPG is not able to authenticate the STB, the subscriber cannot use the STB for broadcast video services. SLA enforcement for a managed application service is also performed by application-aware components. As an example, the number of simultaneous video streams that a subscriber may have active for a video application service is limited by a combination of (1) the number of authorized STBs the subscriber has in the home, and (2) the video session limits enforced by the video middleware.

Because SLA enforcement for a managed application service is performed by application-aware components, the QoS architecture required to support an application-aware service can be greatly simplified. Instead of having shapers and queues per subscriber, QoS architectures that use class-based classification and scheduling, such as the DiffServ architecture, can be used for QoS. [Figure 2-5 on page 2-17](#) illustrates the application and transport architecture used to implement a video application service.

Figure 2-5 Application and Transport Architecture Used to Implement a Video Application Service



While an Internet access service is typically sold as a transport service, a video service may be sold to a subscriber as either a transport or an application service. From the discussion above, the transport architecture needed for an application service is significantly different from that needed for an application service.

The Release 1.0 transport architecture is optimized for service providers that sell video as a managed application service, as opposed to a transport service, with the assumption that the access network devices provide congestion management or DiffServ QoS. In the solution transport network design, architectural tradeoffs have been made with this assumption in mind.

**Note**

Alternatively, for networks where the access network devices do not provide the minimal congestion management described, distributed per-subscriber QoS control for all applications, including video, may be required. This transport SLA-like approach is not the design direction taken for the first release of the solution.

Service Separation in a Triple-Play Architecture

An important aspect of the transport network for a triple-play architecture is how much support the network provides in isolating each service.

Minimally, the network must provide the ability to meet the delay and drop requirements for each service when multiple services share the same physical link. This capability is inherent in the QoS architecture of the solution. (See [QoS Architecture, page 3-46](#).)

In addition, the network may be configured to provide separate forwarding and routing domains for each service. This level of service separation is very useful when a service provider wants to manage separately the address space, topology, and IP infrastructure associated with each service. The following section explains why a service provider may want to have different transport attributes for different sets of services.

Forwarding Architectures

The transport architecture associated with different services may require the use of different encapsulations and therefore different types of packet forwarding. If one creates separate logical topologies for different services, these services can be forwarded by means of different forwarding techniques. The paragraphs below illustrate how the different transport architectures of Internet access and video services require that there be separate logical forwarding planes for the two service categories.

As explained in [Potential Video Service Architectures, page 2-14](#), Internet access service is typically sold as a transport service. In a DSL environment, this typically results in a transport architecture that uses a PPPoE session from a CPE device to a BRAS that authenticates subscriber sessions and enforces the SLA associated with that session. Because PPPoE encapsulation requires an 802.3 header, PPPoE packets must be forwarded by means of Layer 2 switching between the PPPoE client (the CPE device) and the PPPoE server (the BRAS).

Also from [Potential Video Service Architectures, page 2-14](#), the transport architecture for Release 1.0 assumes that the SLA for video services is an application SLA. Because authentication and enforcement are application services implemented in application components, there is no need to use a Layer 2 tunneling protocol such as PPPoE or a transport-layer authentication and enforcement component such as a BRAS for video services. Instead, video services can use IP encapsulation between the STB and the video infrastructure components described in [Video Application Components, page 2-1](#). Since IP encapsulation is used, there is no need to forward packets between STBs and the video infrastructure components in the VHO using only Layer 2 switching. The solution transport architecture described in this document uses a combination of Layer 2 and Layer 3 forwarding for broadcast video and VoD services.

Note that the Internet access transport architecture described above requires that the access, aggregation, and distribution networks switch Internet access packets at Layer 2, while the video transport architecture allows these networks to switch video packets at either Layer 2 or Layer 3. To allow Layer 3 switching for video and Layer 2 switching for Internet access, the network must be configured into

separate logical topologies that are switched by means of different encapsulations and packet switching functions (Layer 2 vs. Layer 3). The transport architecture separates these logical topologies in the aggregation and distribution networks by using different 802.1q VLANs for the different services.

Service-Availability and Bandwidth Requirements

Because different services have different service-availability and bandwidth requirements, a service provider could potentially reduce the cost of the network while maintaining the requirements for each service by creating separate logical topologies for different services.

As an example of different service availability requirements, [Service Availability, page 2-13](#), describes the different availability and bandwidth requirements of broadcast video and VoD services. A service provider could optimize the network for both services by creating separate logical topologies for each service. These topologies could be created by using VRF-based technologies such as MPLS VPN or VRF-lite. [VRF stands for virtual private network (VPN) routing and forwarding, as well as a VRF instance.] In addition, the separate logical topologies could be created by populating the routing table with multiple instances of routing processes running on the different topologies and not exchanging routes between these processes. The differing availability requirements for broadcast video and VoD may lead to a transport requirement that the network must provide redundant paths for broadcast video but not for VoD. To meet this requirement cost-effectively, separate logical topologies can be created for the two services. The logical topology for broadcast video maps the address space associated with real-time encoders and STBs into a topology with redundant physical paths, while the address space associated with VoD servers and STBs maps into a VRF with nonredundant physical paths.



Note

Release 1.0 test configurations did not include the use of VRF technologies to map services to different VRFs.

Organizational Structure

A service provider may have an organizational structure in which different services are managed by different organizations. The ability to map different services to different logical topologies allows each organization to manage and debug the transport as well as the IP infrastructure components separately.

IP Infrastructure Components

When different services are managed by different organizations within a service provider, it may be operationally simpler to have separate IP infrastructure components such as Dynamic Host Configuration Protocol (DHCP) servers for different services. Using different DHCP servers for different services allows the IP address spaces for these services to be managed separately. It also allows the DHCP servers to be configured separately for different services without having to use static configuration on the DHCP server to associate different CPE devices with different services.

Service Separation in the Release 1.0 Architecture

Because of the transport architecture issues described in [Forwarding Architectures, page 2-18](#), it is likely that early IPTV/VoBB deployments will not use a unified transport architecture for all services. Because of this, Release 1.0 uses a service separation architecture in which traffic associated with each service is forwarded to or received from a separate logical access topology at the CPE device. This service-based logical topology separation is continued through the aggregation and distribution networks.

This transport architecture allows traffic associated with different services to be aggregated or terminated at different sites by means of different infrastructure components. This architecture allows traffic associated with Internet access services to be aggregated at a BRAS, while traffic associated with video services (specifically the managed video application service types) is terminated by means of the video infrastructure components described in [IPTV/VoBB Transport Architecture and Issues, page 2-9](#). [Video Forwarding Architecture, page 3-11](#), describes how service separation is implemented in the aggregation and distribution networks in Release 1.0, while [Edge Transport Architecture, page 3-39](#), describes how service separation is implemented at the CPE device and in the access network.



Solution Transport Architecture

The Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband (GOVoBB) Solution transport architecture is subdivided into recommendations for the access, aggregation, and distribution networks. While the service-separation architecture used in Release 1.0 includes requirements regarding how a home access gateway (HAG) interfaces to the home network, it does not include any recommendations for the technologies or configurations used in that network. The Release 1.0 transport architecture also does not include recommendations for the core network. Because of this, the solution test environment combines the video application components of both the super headend (SHE) and video headend office (VHO) sites into a single combined topology that connects aggregation routers (ARs) to a distribution edge router (DER).

This chapter presents the following major topics:

- [Overview, page 3-1](#)
- [Distribution and Aggregation Transport Architecture, page 3-4](#)
- [Release 1.0 Configurations, page 3-32](#)
- [Edge Transport Architecture, page 3-39](#)
- [QoS Architecture, page 3-46](#)

Overview

[Figure 2-3 on page 2-11](#) introduced the transport layers of the general IPTV/VoBB transport architecture (described in [IPTV/VoBB Transport Architecture and Issues, page 2-9](#)) that are the subject of the recommendations in this document.

While the solution transport architecture focuses on video, there is an implicit assumption that the network be able to support a full triple-play environment. Consequently, the transport architecture includes a common quality of service (QoS) architecture for video, voice, and Internet access services. Because the transport architecture is based on the service separation model described in [IPTV/VoBB Transport Architecture and Issues, page 2-9](#), the actual transport architecture used for Internet access and voice services may differ from the video transport architecture described in this document.

To ensure that the video transport architecture works in a triple-play environment, solution testing included a test bed environment in which the transport network was configured to support all three services. Because solution testing was focused on video, it included the application, control, and transport environment for video services. Testing only included enough testing of the Internet access and voice services to ensure that an example forwarding architecture for these services can coexist with video, and that the common QoS architecture specified in this document meets the jitter and packet-loss requirements for each service.

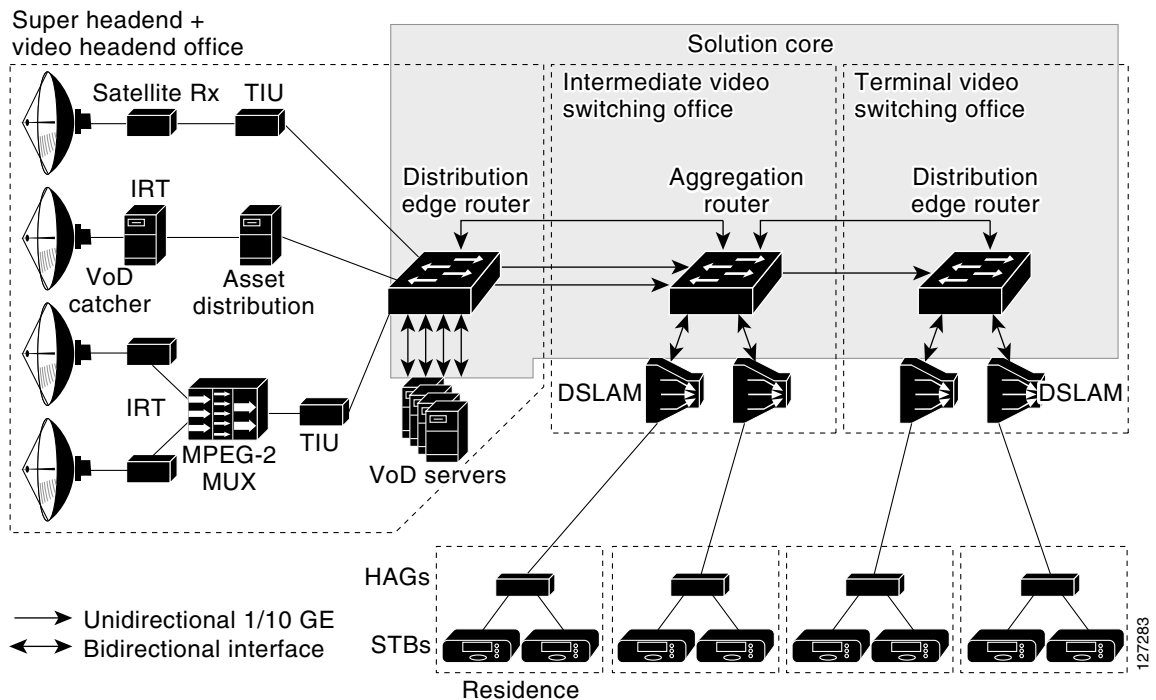
**Note**

This document specifies an example of how the transport network may be configured to support Internet access and voice services. The example configurations described in this document for those services are provided to ensure a fully specified solution test environment. However, these example configurations are not intended to constitute Cisco's recommendation for a proposed transport architecture for those services.

While the transport architecture includes configuration recommendations for all of the transport layers shown in [Figure 2-3 on page 2-11](#), this document only includes example configurations for the transport components that are implemented by means of Cisco products. These components are the DER and AR. Because the DER and ARs are the switching components that implement the distribution and aggregation networks, more detailed configuration information is provided for this portion of the network.

[Figure 3-1](#) highlights that part of the transport network that is addressed by solution testing and core configuration examples.

Figure 3-1 IPTV/Video over Broadband Transport Architecture: Solution Core



Solution Components

Table 3-1 lists the network architecture components used in Release 1.0 of the solution, with additional information. For detail regarding interfaces, see the following:

- [Table 3-5 on page 3-33](#)
- [Table 3-6 on page 3-35](#)

Table 3-1 Network Architecture Components

Network Role	Vendor	System	Product Number
DER, AR	Cisco	Catalyst switch	7609, 6509
		• Supervisor	WS-SUP720-3BXL
		• 10 GE x 4 optic	WS-X6704-10GE
		• 1 GE x 24 optic	WS-X6724-SFP
		• 1 GE x 16 DWDM optic	WS-X6816-GBIC
		• 48-port copper Ethernet	WS-X6748-GE-TX
		Catalyst switch	4507R
		• Supervisor	WS-X4515
		• 1 GE x 6 optic	WS-X4306-GB
			WS-X4448-GB-RJ45
		Catalyst switch	4510R
		• Supervisor	WS-X4516
• 1 GE x 6 optic	WS-X4306-GB		
	WS-X4448-GB-RJ45		
Catalyst switch	4948-10GE		
	WS-X4516		
DSLAM	Ericsson	Ethernet DSL Access ECN320	FAB 801 3908
		EDN312xp, version R3, revision R1A, ADSL2, ADSL2+	FAB 801 4246
HAG		HM340d, version 2, ADSL2 CPE modem	ZAT 759 94/A101
VoD server	Kasenna	GB Media Server	GB-MS-BASEA-LB
			GB-MS-GIGE-COP
Application server		Living Room Application Server	LR-VSIF-HWSW
IP STB	Amino	STB	110

Distribution and Aggregation Transport Architecture

As described in [IPTV/VoBB Transport Architecture and Issues, page 2-9](#), the transport architecture uses service separation to support the capability of having separate routing and forwarding planes for different services. This functionality is used in the aggregation and distribution networks to enable a separate logical and physical transport architecture that is optimized for the delivery of video.

This section describes how the transport architecture is optimized for video. Because an important requirement of the transport architecture is that it also must support a triple-play environment, this section also describes an example distribution and aggregation network configuration for voice and Internet access.

This section presents the following topics:

- [Video Forwarding, page 3-4](#)
- [Multicast, page 3-16](#)
- [Internet Access Forwarding, page 3-27](#)
- [Voice Forwarding, page 3-28](#)
- [Management, page 3-29](#)
- [Redundancy, page 3-31](#)

Video Forwarding

This section presents the following topics related to the delivery of video services:

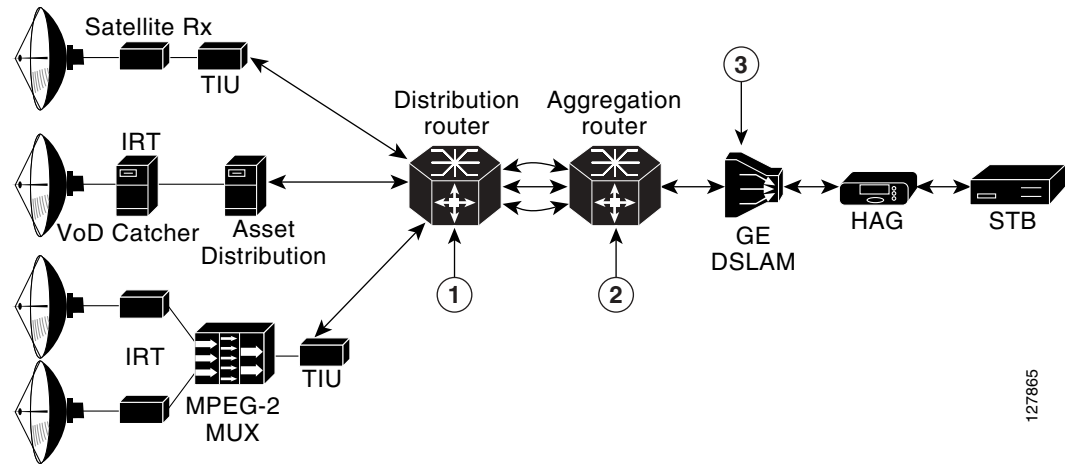
- [Layer 3 Edge for Video Services](#)
- [Video Forwarding Architecture](#)

Layer 3 Edge for Video Services

One of the primary architectural decisions that must be made in specifying a transport architecture for video services is where the Layer 3 edge of the transport network should be for video services.

[Figure 3-2 on page 3-5](#) illustrates the points in the network where the Layer 3 edge may reside, as well as the issues and benefits associated with each location. There are three points: the DSLAM, the AR, or the DER. This section describes the issues and benefits associated with each of these options, as well as the design choice that was made for Release 1.0.

Figure 3-2 Potential Layer 3 Edge Points for Video



127865

Table 3-2 summarizes issues and benefits for edge points 1, 2, and 3 in the above figure. The paragraphs that follow address issues to related to DSLAM-based, DER-based, and AR-based Layer 3 edge points.

Table 3-2 Issues and Benefits for Layer 3 Edge Points for Video

Edge point	Issue	Benefit
1	ARP/forward table scaling	Is consistent among services.
	MAC table scaling	
	Complex video VLAN topology	
	Potential problems with multicast path failover	
2	Is different for video and Internet access	Supports secure Source Specific Multicast (SSM) in distribution network.
		Supports anycast in distribution network.
		Supports multicast load balancing in distribution network.
		Supports fast failover of video encoders.
		Supports unidirectional transport in distribution network.
3	Requires IP-capable DSLAM	
	Complicates IP address management	

DSLAM-Based Layer 3 Edge

Because of their location at the edge of the network, DSLAMs have traditionally performed Layer 2 switching functions. This has kept the function of the DSLAM fairly simple and has also made DSLAMs simple to manage. However, a Layer 3-capable DSLAM is more complex to build, and therefore more complex to manage.

Issue: DSLAM Complexity

A DSLAM that supports Layer 3 functionality must be capable of a number of functions besides Layer 3 forwarding. For example, a Layer 3-capable DSLAM must be able to support a DHCP relay function. This function requires that the IP address of a DHCP server as well as the IP subnet that the DSLAM is associated with must be configured on the DSLAM. The DSLAM must also support and be configured for IP routing protocols to enable dynamic routing from the AR.

Issue: Complex Subscriber-Address Management

An IP-capable DSLAM must have an IP subnet allocated to it to allow IP packets to be routed to it. This complicates IP address management, because a separate IP subnet must be allocated for each DSLAM. This also makes IP address management for the residence more complex, as separate IP address pools must be allocated for each DSLAM.

DER-Based Layer 3 Edge

The DER may also be at the Layer 3 edge for video. With this type of design, forwarding in both the aggregation and distribution networks is performed at Layer 2. While such a design is consistent with common designs for PPPoE-based Internet access services, it creates a number of scaling issues for both the ARs and the DER. This design can also create issues for video services, because of the flooding associated with common learning-bridge architectures.

Issue: Scaling for the Layer 2 MAC Table and Layer 3 Forwarding Table

To understand the scaling issues associated with this design, it is useful to look at the number of STBs that may be aggregated by a single DER. To provide worst-case scaling numbers, we use the following numbers for a hypothetical video over broadband deployment:

- Each DSLAM serves 400 video subscribers.
- Each AR aggregates 40 DSLAMs.
- The DER aggregates 10 ARs.

Therefore, in this example, the DER is aggregating 160,000 subscribers.

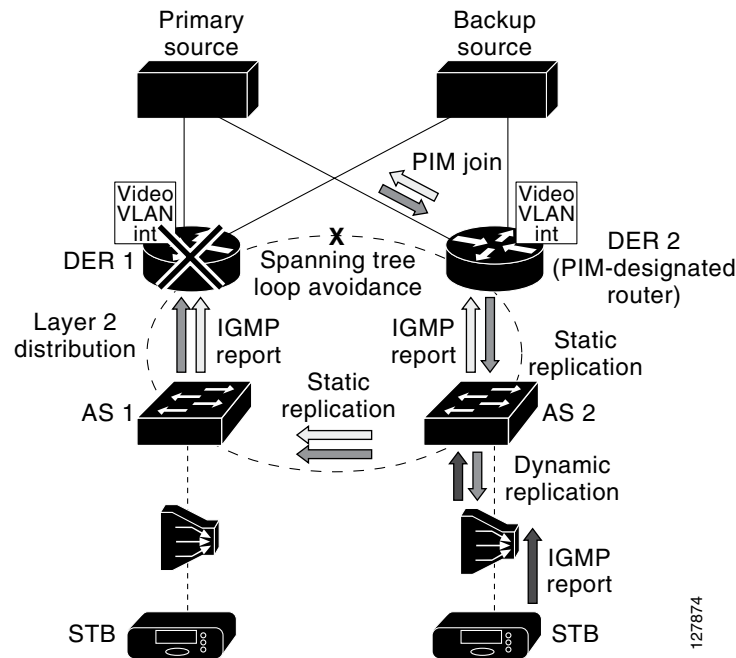
When the DER is configured as the Layer 3 edge for video services, all STBs that are connected through that router are in the same IP subnet. If the subnet is aggregated as a single Layer 2 topology, each of the ARs aggregated by the DER need to support MAC table forwarding entries for all of the STBs on that subnet. This amounts to 160,000 MAC table entries for each AR. This requirement drives up the cost of ARs, because each MAC table entry requires a hardware content-addressable memory (CAM) table entry. There are methods that use separate VLANs to divide the distribution layer topology into simpler Layer 2 topologies that are aggregated at the DER. These methods reduce the MAC table scaling requirements for the ARs, but result in a more complex Layer 2 topology to administer in the distribution network.

Another issue with configuring the DER as the Layer 3 edge for video services is that this router must maintain a separate ARP table entry and forwarding table adjacency for each STB aggregated through it. In our previous example, this amounts to 160,000 such adjacencies. Such a large number again results in higher cost for this router, because each forwarding adjacency uses a separate hardware ternary CAM (TCAM) entry. By comparison, if the Layer 3 edge in the example above were moved to the AR, this device would need to support only 16,000 ARP table entries and forwarding adjacencies.

Issue: Multicast Configuration Complexity and Transport Issues

A network design that aggregates multicast video traffic at Layer 2 results in a complex multicast configuration, as well as in significant inefficiencies in multicast traffic behavior. [Figure 3-3](#) illustrates the configuration complexity and transport inefficiencies when a Layer 2 distribution network is used for multicast video.

Figure 3-3 Multicast Traffic Flow with Layer 2 Distribution



When multicast video is aggregated at Layer 2, the resulting design typically uses more than one DER for redundancy. As a result, the PIM protocol state machine elects a designated router (DR). The DR is responsible for registering sources and sending upstream join and prunes on behalf of the members of the subnet (VLAN). In addition, the network selects an IGMP querier for the served subnet. The IGMP querier is responsible for sending IGMP queries on the subnet served by the redundant IP edge routers.

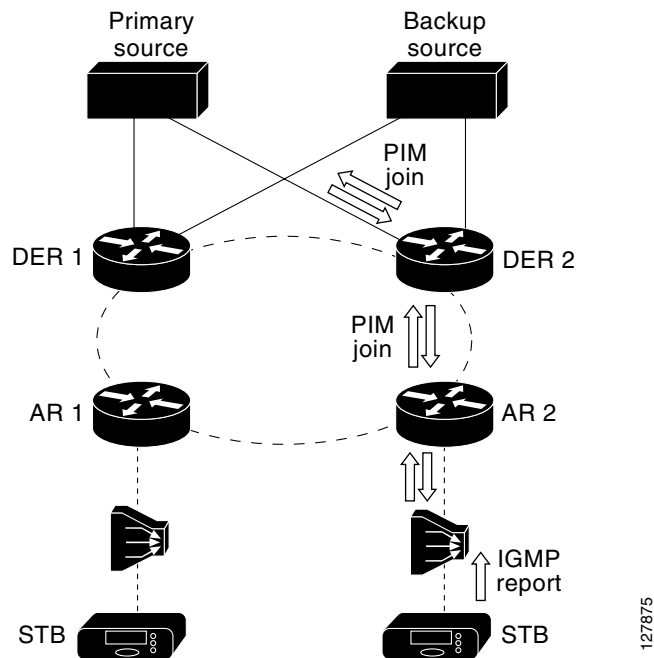
Each aggregation switch (AS) is responsible for replicating multicast streams from the distribution network to aggregation ports that have subscribers joined to them. As shown in [Figure 3-3 on page 3-7](#), there are two potential sources inserting video into the distribution network. These are DER 1 and DER 2. Because either of these two sources may be used to send multicast traffic onto the ring, each aggregation switch must send IGMP joins up both of the uplinks. This IGMP behavior makes it very difficult for the Layer 2 switches to determine when and when not to replicate multicast traffic on the distribution ring. To make multicast work properly in this type of environment, each port on each switch must be configured to replicate packets dynamically by using IGMP or statically. Ports that are configured to replicate dynamically send the traffic associated with a multicast group only if there has been an IGMP join issued for that multicast group. Ports that are configured to replicate statically send all multicast traffic all the time, independently of whether an IGMP join has been issued. In the case of [Figure 3-3 on page 3-7](#), each upstream port on each switch must be configured for static replication, because the downstream multicast traffic could potentially flow from either direction on the ring. This configuration results in additional complexity when multicast is configured on redundant topologies.

In addition to being more complex to configure in a redundant topology, multicast is less efficient. This is because multicast streams must be sent everywhere in the Layer 2 ring, independently of where an IGMP join was issued. [Figure 3-3 on page 3-7](#) illustrates an example multicast replication in a Layer 2

environment. Here the subscriber has issued a channel-change request from the STB attached to AR 2. The channel-change request results in an IGMP join message being propagated in both directions of the distribution network to both DER 1 and DER 2. DER 2 has been elected as the designated router, so it translates the IGMP join into a PIM join, while DER 1 ignores the IGMP join request. As a result of the IGMP join, DER 2 sends the multicast stream to the ring. Because AS 2 is using IGMP snooping on the downstream link, it is the only switch that replicates the stream to the DSLAM. Note, however, that the multicast traffic gets propagated all the way through the Layer 2 ring to DER 1. Each switch must replicate the traffic to other switches on the ring, because it is very difficult to determine where to send the multicast traffic on the ring based on IGMP snooping alone. DER 1 drops the multicast traffic when it receives it, because it does not have any “downstream” requestors for the stream. The result of using Layer 2 in the distribution network is that bandwidth is wasted on the distribution ring, because the multicast stream must be sent everywhere—independently of which nodes on the ring have asked for the traffic.

Figure 3-4 on page 3-8 illustrates multicast operation and traffic flow when a Layer 3 distribution network is used for video. Here all nodes in the redundant topology are in the same Layer 3 topology. This results in simpler configuration as well as a more efficient traffic flow pattern. IP multicast is inherently different from Layer 2 forms of replication, because the multicast tree is built from PIM messages that are routed from the edge of the IP network to the source by means of reverse-path routing. Reverse-path routing is essentially the same as destination-based routing, except that the path to the source is looked up on the basis of the IP source address. This figure illustrates how the PIM messages are routed to the source and how the multicast distribution tree is built more efficiently as a result.

Figure 3-4 Multicast Traffic Flow with Layer 3 Distribution



In this figure, the subscriber has again issued a channel-change request from the STB attached to AR 2. The request results in an IGMP join message being sent to DER 2. Release 1.0 uses Source Specific Multicast (SSM), along with SSM mapping, as the IP multicast technology for the broadcast video service. As a result, AR 2 can translate the IGMP join request into the IP address of the encoder that is being used to generate that stream. With the IP source address, AR 2 uses reverse-path routing to decide where to send an PIM message. In this case, the shortest path to the primary source is through DER 2.

Once PIM state is established, DER 2 replicates the multicast stream to AR 2, which in turn sends the multicast stream to the DSLAM and the STB. Note that the multicast stream is not replicated throughout the distribution ring as it was in the Layer 2 scenario. This is because reverse-path route lookup results in a multicast tree that is built from the source directly to the nodes that requested the traffic. The result of using Layer 3 in the distribution network is an IP multicast environment that is simpler to configuration and more efficient in bandwidth use than are Layer 2 environments.

AR-Based Layer 3 Edge

When the AR is configured as the Layer 3 edge for video, the network is typically configured so that the AR is located at a different point in the network than the Layer 3 edge for Internet access services. This type of configuration may be considered not as architecturally “clean” as having the Layer 3 edge for all services located at the same point in the network. However, as described below, these issues are far outweighed by the benefits of using a Layer 3 distribution network for video services. Release 1.0 uses an AR-based Layer 3 edge to take advantage of these benefits.

Note that the AR may not be the node that directly aggregates the GE uplinks from DSLAMs. The AR is defined as the first node in the physical topology that aggregates enough subscribers that either path or node redundancy is required for video services. In a ring topology, the AR is defined as the node that connects the ring to a nonredundant hub-and-spoke aggregation architecture. In a hub-and-spoke topology, the AR is defined as the first node that includes redundant uplinks to the distribution network. In topologies where the AR does not terminate the GE uplinks from DSLAMs, there may be a Layer 2 aggregation network between the DSLAMs and the AR that does not include either path or node redundancy. [Layer 2 Aggregation Alternatives, page 3-14](#), provides details on Layer 2 aggregation schemes that may be used between DSLAMs and the AR.

The sections below provide details on some of the benefits that make the AR the best choice as the Layer 3 edge in a video topology.

Benefit: Source-Specific Multicast

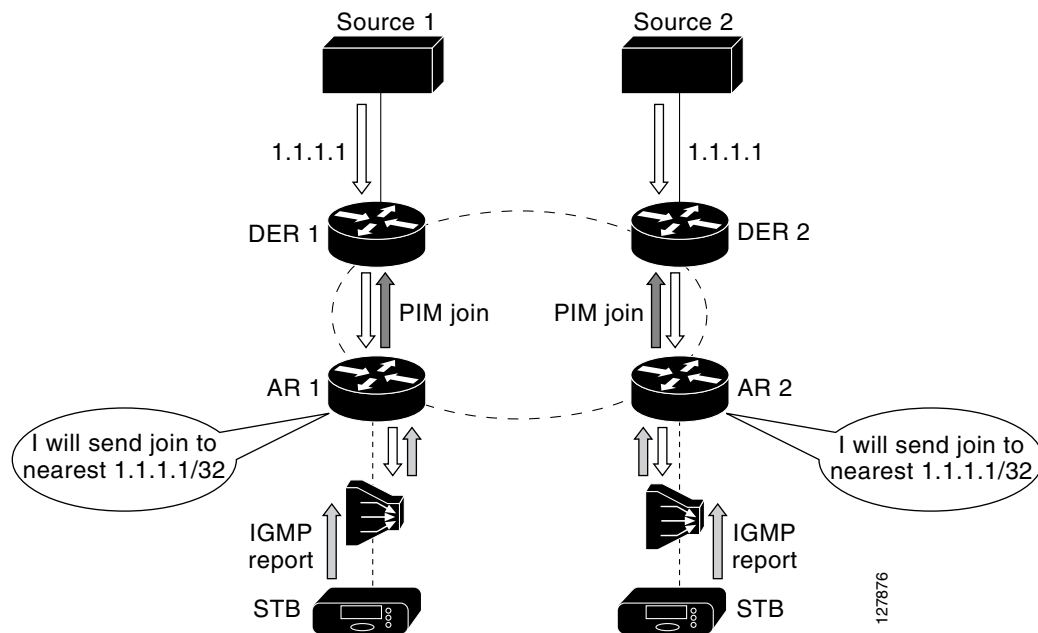
When the AR is configured as the Layer 3 edge for video services, the distribution network can take advantage of IP multicast features such as Source Specific Multicast (SSM). SSM is a technology that enables the network to build a separate distribution tree for each multicast source. SSM simplifies the operational complexity of configuring a multicast network, because it does not require the configuration of a rendezvous point (RP) to allow multicast forwarding as non-source-specific multicast technologies do. In addition, SSM only creates a multicast distribution tree to a specific multicast source address. SSM is considered more secure than non-source-specific multicast, because the multicast client must know both the multicast destination address and the multicast source address in order to join the multicast group. To create a source-specific multicast tree, SSM relies on IGMPv3 signaling from multicast hosts. IGMPv3 includes the multicast source address in the multicast join request. Because current-generation STBs do not support IGMPv3 signaling, the AR can be configured to map IGMPv2 requests received from the aggregation network to PIM SSM (S, G) (source, group) messages in the distribution network. This translation process maps the multicast destination address specified by the STBs in IGMP messages to a combination of multicast source and destination addresses in PIM messages. Release 1.0 of the solution uses SSM mapping at the AR to provide SSM support for STBs that do not support IGMPv3.

Benefit: Anycast Support

When the AR is configured as the Layer 3 edge for video, the distribution network can take advantage of “anycast” support for either the load balancing or the fast failover of video encoders. IP multicast technology natively supports the ability for “anycasting” of IP multicast sources. With anycasting, one configures two or more multicast sources that are sending to the same IP multicast group (with the same multicast destination address) and have the same IP source address. When used with PIM sparse mode, IP multicast technology uses a reverse path lookup to determine which IP source is closest to any

particular PIM edge node. The result is that the replication path for a single multicast group can consist of a separate multicast tree for each broadcast encoder. Figure 3-5 illustrates the use of anycasting for load sharing between multiple video encoders.

Figure 3-5 Anycast-Based Load Sharing between Video Encoders

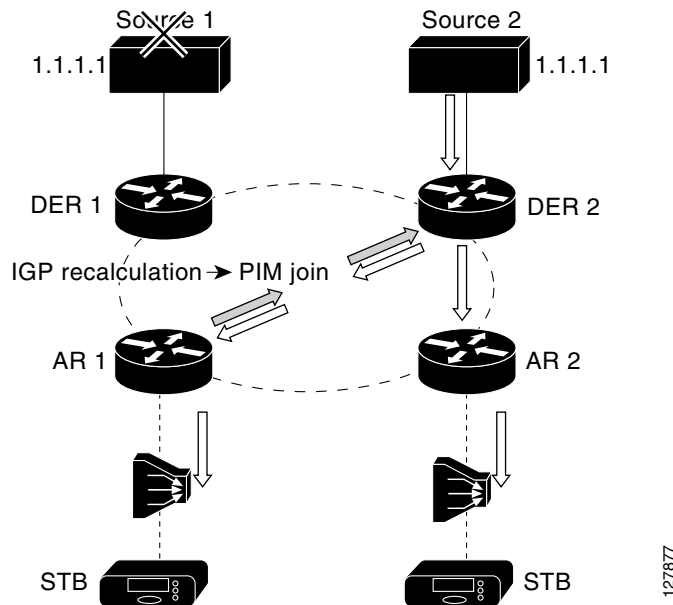


Note that the ability to instantiate multiple multicast replication trees for the same multicast destination is not possible when Layer 2 switching is used. Because each node in a Layer 2 network simply uses IGMP snooping to determine when to replicate packets, anycasting in a Layer 2 domain would result in having the stream from each multicast source replicated to all multicast destinations. Because of this, anycasting is applicable only within the context of a Layer 3 switching environment.

Benefit: Fast Failover of Video Encoders

In addition to supporting load sharing among multicast sources, anycasting can be used to support the fast failover of video encoders. When anycasting technology is combined with the ability of the network to detect the failure of an encoder, routing protocols reconverge. This reconvergence results in the reverse path from the ARs to the DER being recalculated to take into account that the location of the multicast source that has been changed. The IP reconvergence then triggers PIM to resend a join request along the path to the new multicast source. Figure 3-6 illustrates the use of anycast technology to implement the fast failover of redundant video broadcast sources. Release 1.0 used this technology to implement fast failover between redundant broadcast encoders.

Figure 3-6 Fast Multicast Source Failover Using Anycast

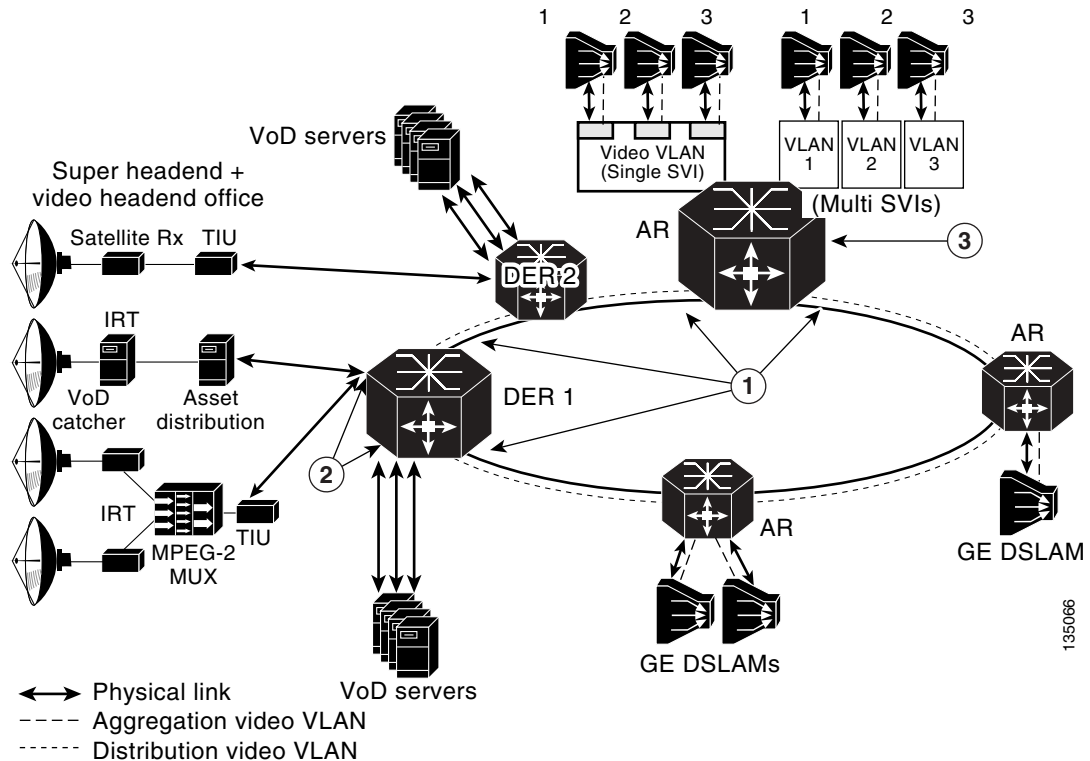
**Benefit: Asymmetric Networking**

Finally, when the AR is configured as the Layer 3 edge for video, the distribution network can be configured to support asymmetric bandwidth for video services in the distribution network. The traffic pattern associated with broadcast video and VoD services is extremely asymmetric. Each video channel or session requires multiple megabits of bandwidth in the downstream direction, while the upstream traffic is limited to control signaling for the service. Asymmetric networking allows the network to be configured for more bandwidth in the downstream direction than in the upstream direction. This reduces the cost of the transport network, because it allows the network provider to take advantage of optical components such as wavelength division multiplexing (WDM) transponders and other optical equipment that can be deployed in a unidirectional manner.

Video Forwarding Architecture

Once the choice is made to position the Layer 3 edge for video at the AR, the aggregation/distribution forwarding configuration becomes fairly straightforward. The service-separation architecture used in the solution results in the GE link from each DSLAM being configured for three separate 802.1Q VLANs to aggregate Internet access, voice, and video services. [Figure 3-7 on page 3-12](#) illustrates the overall video forwarding architecture discussed in this section.

Figure 3-7 Video Forwarding Architecture



1,	Video VLAN interfaces
3	
2	Video interfaces

AR Configuration

The AR has a set of interfaces connecting to the distribution network and a set of aggregation interfaces connecting to DSLAMs. The AR is configured to switch packets between the distribution and aggregation interfaces at Layer 3. Separate VLANs are configured for each service on each of these interfaces.

Each upstream port connected to the distribution network is configured to use 802.1q encapsulation (VLAN trunking) and contains three separate Layer 2 VLANs for the Internet access, voice, and video services. The Layer 2 video VLAN of each upstream port is terminated in a separate Layer 3 switched virtual interface (SVI). This configuration causes video coming in on any physical port to be switched at Layer 3 to any other physical port.



Note

The VLAN IDs used for video on each of the physical upstream ports must be different from the VLAN IDs configured on the downstream ports connected to the aggregation links.

There are two alternative configurations for the downstream aggregation ports: single SVI and multiple SVI configurations. In the single-SVI configuration, a single video SVI is configured for all of the GE ports connected to the DSLAMs. In the multiple SVI configuration, a separate video SVI is configured

for each of the downstream GE ports connected to the DSLAMs. Each of these two configuration options has benefits and drawbacks. The configuration option that is chosen for a particular network design depends on the benefits that the service provider finds most important in the network.

Single SVI Configuration

The single SVI configuration is the simpler of the two models to configure and maintain. With this model, there is a single IP interface and as a result a single DHCP address pool to maintain for all subscribers served by the AR. This model is also simpler in that the VLAN ID for each service can be shared across all of the DSLAMs connected to the AR.

The down side to this configuration is the potential security risk associated with Layer 2 flooding across the downstream GE interfaces. The standard bridge-learning algorithm used in Layer 2 switches floods Ethernet frames if the MAC-layer forwarding engine does not have an entry for the destination MAC address in the packet or if the destination MAC address is a broadcast address. A malicious subscriber could potentially use this flooding behavior to learn about other subscribers by snooping packets that are flooded as part of standard bridge-learning behavior. This attribute of bridge learning is often not an issue for video, because the only upstream traffic coming from other subscribers on the video VLAN is requests for on-demand content. If a service provider considers the flooding behavior associated with bridge learning to be a significant security risk, then the multiple SVI configuration option described [Multiple SVI Configuration, page 3-14](#), should be used.

With the single SVI configuration, the AR is configured to terminate the video LANs from the GE ports of the DSLAMs it aggregates into a single SVI. This results in the video VLANs from each of the GE links being switched at Layer 2 into a single Layer 3 VLAN interface on the AR.

The use of Layer 2 switching for this aggregation causes a complication for video that is dealt with in this solution. That complication is the potential flooding of VoD streams on downstream links because of issues with common learning-bridge algorithms.

Layer 2 switches that implement transparent LAN services use a learning-bridge algorithm that floods incoming unicast traffic to a particular MAC destination until a packet is received from that destination. This behavior normally is not an issue for most applications, because the MAC forwarding table is typically populated when ARP requests are sent between the routers or IP hosts attached to that LAN segment. An additional property of LAN switches is that they time out MAC table forwarding entries periodically to ensure that these entries do not become stale. This again is not an issue for most applications, because the client/server behavior of most applications means that the MAC forwarding table is repopulated when a client/server transaction occurs.

However, VoD applications do not exhibit this type of client/server behavior. While a VoD stream is being played, the traffic pattern is such that packets are being sent only from the video server to the subscriber. A separate stream-control session is used between the subscriber and the VoD server to support the ability to pause, fast forward, or rewind the video that the subscriber is playing. If the subscriber is simply playing the video, the only traffic sent on the control channel is periodic keepalives. This behavior of VoD applications can result in a learning bridge getting into a state where a MAC forwarding entry is aged and not replaced for a long period of time. The result is that the VoD stream is flooded to all downstream ports—with the ultimate result being that the video queue on downstream links such as DSL links become congested, causing all subscribers to experience poor video quality.

The solution deals with this issue by enabling unicast flood blocking on the video VLAN of the GE interfaces connected to the DSLAMs. This feature prevents the LAN switch from flooding unicast traffic when there is no bridge table entry for a destination MAC address. In addition, the Layer 3 video interface is configured to set ARP timeouts shorter than the MAC table aging timeout value. This configuration ensures that the router sends an ARP request to the downstream host before the MAC table entry times out. The resulting ARP request and response ensure that the MAC table entry gets repopulated before it is timed out.

Multiple SVI Configuration

The multiple SVI configuration is more complex to maintain than the single SVI configuration, because it requires a separate Dynamic Host Configuration Protocol (DHCP) address pool for each DSLAM connected to the AR. However, the multiple SVI configuration is considered more secure than the single SVI configuration, because there is no issue with Layer 2 flooding among downstream GE ports.



Note

The administration of a separate DHCP address pool per DSLAM could be avoided with the multiple SVI configuration if the downstream IP interfaces could be configured as IP unnumbered interfaces. IP unnumbered interfaces can be configured to take on the address of a single loopback interface on the router. Using this configuration, all IP unnumbered interfaces can share the same IP subnet. This in turn allows a single DHCP address pool to be configured across the interfaces. While it is possible to configure VLAN subinterfaces as IP unnumbered, it is currently not possible to configure SVIs as IP unnumbered. Release 1.0 solution testing was performed with SVIs instead of subinterfaces, because the Internet access service was switched at Layer 2 through the AR.

In the multiple SVI configuration, a separate video SVI is associated with the video VLAN of each downstream GE port. Each SVI is configured with a separate IP subnet, so each SVI is associated with a separate DHCP address pool.



Note

The solution test bed topologies described in [Release 1.0 Configurations, page 3-32](#), do not include multiple SVI configuration.

DER Configuration

The downstream ports of the DER are configured identically to the upstream ports of the AR. (Refer to [Figure 3-7 on page 3-12](#).) The video stream is terminated in an SVI just as it is on the AR. One difference between the AR and DER is that the different services may be aggregated by different DERs. This allows the different services to be aggregated at different sites if necessary.

In the solution, video components such as video servers and real-time encoders are connected to redundant DERs. A load-sharing scheme provides video redundancy. This means that the VoD servers and broadcast video encoders connected to each of these routers are actively sending video during normal operation. Ports connecting video components to a DER may be configured at either physical Layer 3 switched ports or Layer 2 ports terminated in an SVI. To simplify address management, ports connecting VoD servers and real-time encoders may all be configured to be in the same Layer 2 VLAN, which is terminated in a single SVI.

IP Routing

To enable dynamic routing specific to video, a routing process is configured on the ARs and DERs. This routing process is configured only on the video SVI interfaces. This enables the video topology to converge at Layer 3 independently of the topologies for the voice and Internet access services. In Release 1.0, OSPF is the routing protocol for video.

Layer 2 Aggregation Alternatives

While the AR may be directly connected to the GE uplinks of the DSLAMs it aggregates, there may be network topologies with insufficient subscriber density to warrant having DSLAMs directly connected to an aggregation router. In these types of topologies, there may be a Layer 2 aggregation network between the DSLAM and the AR.

**Note**

While this section describes an architecture that may be used for Layer 2 aggregation between DSLAMs and ARs, the solution test topologies described in [Release 1.0 Configurations, page 3-32](#), do not include Layer 2 aggregation as part of the test topology.

The solution transport architecture specifies that the AR is where the Layer 3 edge for video should be. The transport architecture also specifies that the AR is defined as the first node in the physical topology that aggregates enough subscribers to require either path or node redundancy for video services. Given these transport requirements, it is important that the Layer 2 aggregation network between DSLAMs and the AR does not include either path or node redundancy. One way to identify such an aggregation network is that it does not require spanning tree algorithms to be configured in order to avoid bridging loops.

When a Layer 2 aggregation network is used between DSLAMs and the AR, it is also important that the number of subscribers aggregated at a single AR not cause forwarding table or ARP table scalability issues for the AR. (See [Issue: Scaling for the Layer 2 MAC Table and Layer 3 Forwarding Table, page 3-6](#), for some of the issues associated with forwarding and ARP table scalability.) A general rule that can be used in network design to avoid scalability issues in the AR is that no more than 30,000 subscribers should be aggregated in a single AR.


The Layer 2 aggregation design described in this section prevents security issues in the DSL aggregation network that are associated with the flooding used in standard bridge-learning algorithms. To simplify the requirements of the aggregation switches, this design assumes that the switches support only standard bridge-learning algorithms, and do not support controlled flooding algorithms that prevent upstream packets from being flooded on down stream links. This design also assumes that aggregation switches are capable of segregating MAC broadcast domains through 802.1q VLAN tagging.

Under the above design assumptions, the Layer 2 aggregation design uses a separate VLAN ID per service per DSLAM. The use of a separate VLAN ID per service per DSLAM means that all MAC layer flooding on the aggregation switch is constrained to a single DSLAM per service. This prevents the security issues associated with MAC layer flooding, but it also means that separate copies of video broadcast channels must be sent to each VLAN—resulting in bandwidth being wasted on the link between the aggregation switch and aggregation router. To prevent multiple copies of video being sent on the link between the aggregation switch and the aggregation router, the Layer 2 aggregation design uses a separate multicast VLAN on which all multicast video traffic is sent. The multicast VLAN carries all broadcast video traffic between the aggregation router and the aggregation switch. The use of a separate multicast VLAN means that the aggregation switch that supports Layer 2 aggregation **must** be capable of performing IGMP snooping and replication between the single upstream multicast VLAN and the video VLAN on each downstream link. Cisco switches support a feature called Multicast VLAN Registration (MVR) to implement this function.

When the aggregation router is configured to use a Layer 2 aggregation network, the multi-SVI configuration described in [AR Configuration, page 3-12](#), for the downstream aggregation links must be used. In addition to this SVI configuration, the AR must have one additional SVI configured for the multicast VLAN. This VLAN has the IP multicast features described in [Multicast Configuration Options, page 3-24](#), configured on it.

[Figure 3-8 on page 3-16](#) illustrates aggregation at Layer 2.

Figure 3-8 Layer 2 Aggregation



Multicast

This section presents the following topics related to multicast:

- [Overview](#)
- [Multicast Admission Control](#)
- [Effect of Multicast on Channel-Change Performance](#)
- [Multicast Configuration Options](#)

Overview

A major component of the transport architecture is the multicast transport architecture for video. As stated previously, a Layer 3 forwarding architecture for video is used between the DER and the AR. The video topology is separated from the voice and Internet access topologies by means of a separate VLAN for video. This VLAN carries both unicast VoD streams as well as multicast broadcast-video streams.

PIM for multicast is enabled on the video VLAN interfaces on the DERs and ARs, along with OSPF. This enables a video-specific multicast topology to be built. PIM sparse mode is used for the broadcast video service.

The IGMP/PIM boundary for multicast occurs at the SVIs on the AR that are associated with the GE ports from the DSLAMs. IGMP joins are translated to PIM joins at the SVI. If the single SVI configuration described in [Multicast Configuration Options, page 3-24](#), is used, then the GE ports from all of the DSLAMs are aggregated at Layer 2 in a single SVI. In this case, multicast replication must occur as part of the Layer 2 switching process.

Source Specific Multicast (SSM) is used in the Layer 3 network. SSM simplifies the operational complexity of configuring a multicast network, because it does not require the configuration of a rendezvous point (RP) to allow multicast forwarding as non-source-specific multicast technologies do. In addition, SSM only creates a multicast distribution tree to a specific multicast source address. SSM

is considered more secure than non-source-specific multicast, because the multicast client must know not only the multicast destination address, but also the multicast source address, in order to join the multicast group.

Because SSM builds multicast replication trees that are specific to the IP address of the multicast source, there is an implicit requirement that all multicast join requests (IGMP/PIM joins) must include the address (or addresses) of the multicast source (or sources) in the request. While video STB applications could learn both the multicast source and destination address for each broadcast video channel through the electronic program guide (EPG), current-generation applications receive only the multicast destination address from the EPG. As a result, these applications send IGMPv2 join requests that contain only the destination multicast address in the request. The solution works around this by translating IGMPv2 requests that contain only the destination multicast address into SSM PIM join requests that contain both the multicast source and destination address at the AR. The ability to map the multicast destination address contained in IGMPv2 requests to a source/destination pair is called SSM mapping. To map between multicast destination addresses and source/destination pairs, SSM mapping can be configured to use either statically configured maps on each AR, or the services of a Domain Name System (DNS) server that contains a single map for all ARs. The solution uses the DNS-based approach to simplify the administration of this map.

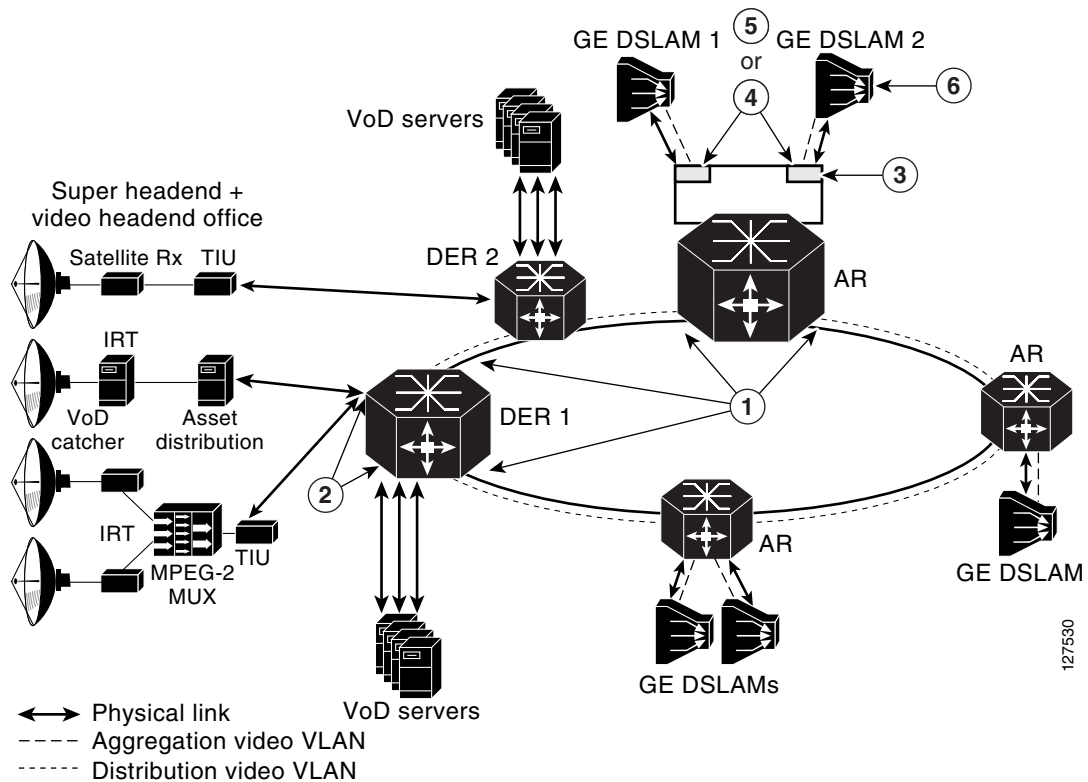
[Figure 3-9 on page 3-18](#) (similar to [Figure 3-7 on page 3-12](#)) illustrates the multicast features used in the solution with the aggregation and distribution networks.

Multicast Admission Control

The Release 1.0 architectural design does support the ability to perform a network-based connection admission control (CAC) function for the broadcast video service. In some broadcast video deployments, it may not be reasonable to support the transmission of all of the broadcast channels offered by the video service on the links between the AR and DSLAMs at the same time. For example, a broadcast video service may offer 150 standard-definition channels and 20 channels of high-definition television. If the channels are encoded by means of MPEG-2, the bandwidth required to support the transmission of all channels simultaneously is 862 Mbps. To ensure that there is no congestion in the queue used for the broadcast video service, bandwidth must be reserved on the GE aggregation links to the DSLAMs. This can be done by simply subtracting the bandwidth used for broadcast video from the bandwidth pools used by the application components of the other services (such as voice and VoD) that require guaranteed bandwidth. In the example above, if the amount of bandwidth that was reserved for broadcast video was based on supporting all channels simultaneously, only 138 Mbps of bandwidth would be available for voice and VoD. This is not enough bandwidth to implement a reasonable VoD service.

The amount of bandwidth reserved for broadcast video can be controlled by implementing an admission control function for that service. This can be implemented by limiting the number of broadcast streams that are replicated on a particular link. Because the GE aggregation links between the AR and the DSLAM are typically the most likely links to be oversubscribed, they are the best place to enforce a stream limit. When stream limits are used for broadcast video, there is a probability that an IGMP join sent by the broadcast video client application as a result of a channel-change request will fail. Because IGMP signaling has no acknowledgement associated with it, there is no explicit failure indication associated with a failed IGMP join request. Instead, a failed IGMP join request simply results in the requested MPEG stream not being delivered to the STB. The subscriber sees a blank picture as the result of a failed channel-change request. While this user interface is nonoptimal, it is consistent with what video subscribers currently experience when a broadcast channel is not available for some reason.

Figure 3-9 Multicast Forwarding Architecture



1	Video subinterface with SSM multicast forwarding, PIM sparse mode
2	Video interface with SSM multicast forwarding, PIM sparse mode
3	DNS-based SSM mapping, static multicast group
4	IGMP snooping or static IGMP join
5	IGMP snooping with report suppression
6	IGMP fast-leave processing

When a CAC function is used for broadcast video, it is important that the service provider sets the stream limit high enough that subscribers very seldom experience failures as a result of a channel-change request. This can be done by using statistical analysis methods such as Erlang analysis. The statistical analysis described in [Static IP Multicast Joins on the AR, page 3-26](#), is an example of the type of analysis that can be used to determine what the stream limit should be set to in order to ensure a low blocking factor for a group of broadcast channels.

When the `ip igmp limit` command is configured on an AR, that router can enforce a maximum broadcast-bandwidth limit by limiting the number of IGMP joins on the ranges of multicast addresses associated with broadcast video to a configured maximum on the aggregation links that the router controls. The mapping of video channels to multicast addresses can be done in such a way that the AR can associate the bandwidth for different classes of video (standard definition, high definition, and so

on) with different ranges of multicast addresses. IGMP join limits can then be set for each range of multicast addresses. For example, a service provider may choose to exclude some video channels from the video CAC function and instead reserve bandwidth for all of the channels that are excluded from that function. This configuration may be useful for managing popular channels that the service provider wants to ensure are never blocked. These channels can be excluded from the CAC function by simply not associating an IGMP limit with their multicast addresses.



Caution

The **ip igmp limit** command on an AR can be used only when that AR is not performing SSM mapping. For details, see the most current version of the *Release Notes for Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband Solution, Release 1.0*.

Effect of Multicast on Channel-Change Performance

One of the important aspects of a broadcast video service that this solution characterizes is the effect of multicast join and leave latency on channel-change performance. This section documents the multicast configurations that testing has evaluated, and makes recommendations that achieve the following design goals:

- Efficiency in bandwidth use
- Scalability to large numbers of subscribers
- Minimal impact on channel-change performance

Table 3-3 illustrates the major components of channel-change latency. Note that the largest factor in the channel-change delay is the I-frame delay associated with the video decoder. (The I-frame is a keyframe used in MPEG video compression.) As the table indicates, multicast performance should not have a significant effect on channel-change delay.

Table 3-3 Major Components of Channel-Change Latency

Channel-Change Latency Factor	Typical Latency, msec
Multicast leave for old channel	50
Delay for multicast stream to stop	150 ¹
Multicast join for new channel	50–200
Jitter buffer fill	200
Conditional access delay ²	0–2000
I-frame delay	500–1000

1. Assumes that the DSLAM implements IGMP fast-leave processing.
2. The conditional access delay is applicable to broadcast channels that are encrypted by means of a conditional access system (CAS) that modifies decryption keys periodically and carries updated decryption keys in-band in the video stream. The STB must wait for the latest set of decryption keys to be delivered in the video stream before it can perform any decoding. The amount of time associated with this delay depends on how often the CAS sends updated decryption information in the video stream.

Analysis of Multicast Bandwidth vs. Delay

The best approach to use for an IGMP/multicast configuration is based on a tradeoff between bandwidth and delay. IP multicast natively supports the ability to perform replication on a stream only when that stream is requested by a downstream device. While IP multicast and IGMP natively support dynamic replication, each can be configured always to replicate multicast data for a particular channel or channel group to any node in the network. When a channel or channel group is always replicated from the source to a particular node, that node is said to be configured for static joins of the channel or channel group. The benefit of configuring static joins at a particular node is that no channel-change latency is associated with dynamic signaling and replication from the source to the node on which static joins are configured. The down side of configuring static joins at a node is that the video streams for the channels that are statically joined are always sent whether a subscriber is watching them or not.

Statistical analysis can be used to determine when the benefits of static joins (less channel-change latency) outweigh the costs (additional bandwidth usage). This section describes the statistical analysis that was done as part of the solution to determine the recommendations for where in the network static joins should and should not be configured.

The behavior of a population of subscribers can be modeled statistically to determine, for a population of subscribers, the probability of at least one subscriber in the group being tuned to a set of television channels. If the probability of at least one subscriber being tuned to each of the channels in a broadcast channel group is fairly high, then the amount of bandwidth that is saved by performing dynamic joins on that group of channels is statistically insignificant. When statistical analysis shows insignificant bandwidth savings for a group of channels, static joins can be used on those channels without having a significant impact on the amount of bandwidth on the GE aggregation links.

The factors used in this analysis included the following:

- The number of subscribers in a video broadcast population
- The number of channels in the broadcast channel group
- The popularity of each channel in this broadcast group

The number of video subscribers served by a particular node depends on where that node is located in the network. Based on common expected video service take rates, the number of subscribers served by a DSLAM is typically about 500 while the number of subscribers served by an AR is typically about 5000.

The following is a statistical analysis model that is helpful in determining when to use dynamic joins, and when to use static joins.

Analysis of Dynamic Joins in a Video over IP Environment

Each subscriber is modeled as a random process selecting a channel to watch according to a given probability distribution across all possible channels. Given a group of channels, we would like to calculate the average number of channels in use, given the “popularity” probabilities of the channels. Because we are interested in determining the average number of channels in use, we can consider the channels to be probabilistically independent of each other and consider the channels one at a time.

For a single channel, the probability that this channel is idle is calculated as follows:

Let

$p = P\{\text{a subscriber tunes to this channel}\}$, and

$N = \text{number of subscribers subtended by the given AR or DSLAM}$

so that

$$P\{\text{channel is idle}\} = (1-p)^N$$

For multiple channels, we sum the above expression.

Let

C = number of channels, and

$p_k = P\{\text{a subscriber tunes to } k^{\text{th}} \text{ channel}\}$

so that the average number of channels in use, C_{IU} , is

$$C_{IU} = \sum_{k=1}^C [1 - (1 - p_k)^N]$$

Channel-Change Latency Probabilities

When subscribers change channels, if they change to a channel that is not part of a static join and that no one else is watching, they experience some latency while the dynamic join is established before they can view the channel's content.

We assume that when there is a channel-change event, the probability a particular channel is changed to is proportional to that channel's popularity. This assumption can be combined with the above calculated $P\{\text{channel is idle}\}$ and the knowledge of which channels are associated with static or dynamic joins to determine the probability that a given channel change results in the latency associated with establishing a dynamic join.

Let

D = the set of all channels involved in dynamic joins, and

$P_L = P\{\text{a channel change experiences latency due to a newly created dynamic join}\}$

so that

$$P_L = \frac{\sum_{k \in D} p_k (1 - p_k)^N}{\sum_{k=1}^C p_k}$$

Analysis

Given a set of channels with probabilities p_k , they can be ranked from highest to lowest p_k . Then, once they are ranked, we can have a cutoff value so that channels with higher p_k get a static join and those with lower p_k get a dynamic join. The questions then are, for a given cutoff,

- What is the bandwidth use (relative to the all static-join case)?
- What is the probability of channel-change latency?

As an example, consider a 150-channel system with an exponential decay function for

$P\{\text{subscriber tunes to } k^{\text{th}} \text{ channel}\}$

[Figure 3-10](#) graphs the channel popularity for this example.

Figure 3-10 Channel Popularity for a 150-Channel System





Figure 3-11 add curves showing the probability that a given channel is busy for subscriber bases of 200 (at a DSLAM) or 5000 (at an AR).

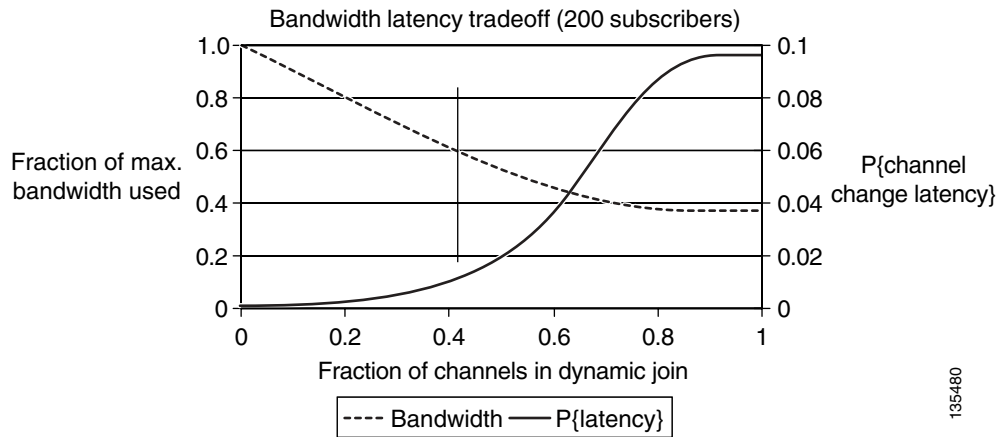
Figure 3-11 Probability a Given Channel is Busy for Subscriber Bases of 200 (at a DSLAM) or 5000 (at an AR)



The important thing to note here is that the $P\{\text{busy}\}$ curve shifts dramatically to the right when the number of subscribers is increased from 200 to 5000.


Figure 3-12 and Figure 3-13 show the tradeoff between average bandwidth requirements and channel-change latency probability for a DSLAM and an AR, respectively. The horizontal axis is the fraction of channels moved from a static join to a dynamic join. The two curves show the bandwidth required (as a percentage of bandwidth required in the static join case) and the probability of channel-change latency. (The two curves in each figure are shown on different scales to make them both visible.)

Figure 3-12 Tradeoff Between Average Bandwidth Requirements and Channel-Change Latency Probability for a DSLAM (200 Subscribers)



135480

Figure 3-13 Tradeoff Between Average Bandwidth Requirements and Channel-Change Latency Probability for an AR (5000 Subscribers)



135481

Figure 3-12 shows the tradeoff for the DSLAM (200 subscribers). There seems to be substantial opportunity in using dynamic joins where about half the bandwidth can be saved with a channel-change latency probability of about 1 in 50 (0.02). (See the vertical black line in the graph.)

Figure 3-13 shows the tradeoff for the AR (5000 subscribers). In this case, the best possible bandwidth savings is about 20%, even with all channels in dynamic joins. Here the channel-change latency probability is uniformly low, with a maximum value of about 1 in 300.

From the statistical analysis results described above, you can see that there is a typically a significant bandwidth savings to be gained (~60%) by using dynamic joins at the DSLAM. Because of this, we recommend that the multicast configuration models used on the links between the AR and the DSLAM take advantage of the dynamic replication capabilities native to IP multicast. Also from these results, it can be seen that the benefit of using dynamic vs. static joins at the AR depends heavily on the popularity of a channel or channel group. It may be best to join popular channels statically, and join less-popular

channels dynamically. [Static IP Multicast Joins on the AR, page 3-26](#), describes additional analysis that was performed to determine when it is best to perform static vs. dynamic joins of a channel group on the AR.

Multicast Configuration Options

From the above analysis, the solution architecture assumes that multicast traffic is replicated by means of dynamic Internet Group Management Protocol (IGMP) signaling on the GE aggregation links between ARs and DSLAMs, and also on the DSL access links between the DSLAM and HAG. The following sections detail the multicast configuration options included in the solution.

IGMP-Based Replication in the DSLAM

Because the DSLAM performs packet switching at Layer 2, it must use a Layer 2 method of implementing multicast replication based on dynamic signaling. In the transport architecture, the DSLAM performs multicast replication by means of IGMP snooping. A Layer 2 switching node that implements IGMP snooping uses the IGMP state machine to determine when to perform multicast replication to a particular link.

IGMP Snooping vs. IGMP Proxy Functionality

Note that the recommendation for the transport architecture is to use IGMP snooping and not an IGMP proxy function. IGMP snooping is defined as a function whereby the DSLAM uses IGMP messages and the associated IGMP state machine to determine when to perform replication of an incoming multicast stream on outgoing DSL lines. When IGMP snooping is used, the DSLAM appears totally transparent to the IGMP signaling path. It does not modify IGMP messages in either the upstream or downstream directions. With an IGMP proxy function, the DSLAM acts as an IGMP server to video STBs and as an IGMP client to upstream routers. With an IGMP proxy function, the DSLAM can statically join multicast streams coming from the AR and replicate them on demand, based on IGMP messages coming from the STBs.

IGMP proxy functionality is not recommended on the DSLAM for a couple of reasons. First, the IGMP proxy function complicates both the operation and configuration of IGMP signaling. This is because the signaling path is now split into two separate IGMP sessions between the STB and the AR. Second, the main benefit of an IGMP proxy function is to allow the DSLAM to join multicast groups statically from the AR and perform dynamic replication to the DSL line. As shown from the analysis in [Analysis of Multicast Bandwidth vs. Delay, page 3-20](#), the benefits of statically joining broadcast channels at the DSLAM (decreased channel change latency) are far outweighed by the cost (additional bandwidth on the GE aggregation links).



Note

Some, and only some, DSLAMs support IGMP snooping with report suppression. When IGMP snooping with report suppression is configured on a DSLAM, the DSLAM forwards only the first IGMP join request for a particular multicast address on the upstream GE link. In addition, the DSLAM sends an IGMP leave request only when it sees a single DSL line currently joined to the multicast stream. This behavior reduces the number of IGMP joins and leaves that the AR must process, and some have recommend its use to provide a more scalable IGMP snooping configuration.

Cisco has load tested IGMP signaling on the Cisco 7600 series, for example, and no join or leave performance degradation was experienced with over 10,000 IGMP messages (join/leaves) per second. Thus, although the Release 1.0 multicast architecture does not require IGMP report suppression on the DSLAM, using this report suppression feature does not cause any issues with the multicast architecture.

IGMP Fast Leave Processing

To meet the channel-change time requirements, the DSLAM must perform IGMP snooping with fast leave processing. Fast leave processing is a modification of the normal IGMP Version 2 host state machine. In IGMPv2, when a router (IGMP server) receives an IGMP leave request from a host (IGMP client), it must first send an IGMP group-specific query to learn whether other hosts on the same multi-access network are still requesting to receive traffic. If after a specific time no host replies to the query, the router stops forwarding the traffic. This query process is required because, in IGMP Versions 1 and 2, IGMP membership reports are suppressed if the same report has already been sent by another host in the network. Therefore, it is impossible for the router to know reliably how many hosts on a multi-access network are requesting to receive traffic.

The requirement of making IGMP queries and waiting for a response can be removed if there is only a single video STB per DSL line that is making IGMP requests. In this case, when an STB sends an IGMP leave request, the DSLAM can safely and immediately stop sending the multicast stream down the DSL line from which the request came. The ability for a node that supports IGMP snooping to stop sending a multicast stream immediately on the receipt of an IGMPv2 leave request is called fast leave processing. The solution requires that DSLAMs support IGMP snooping with fast leave processing.

However, IGMP snooping with fast leave processing does not work when more than one STB is connected to a DSL line. The problem with fast leave processing is that if two STBs attached to the same DSL line are tuned to the same channel, the first STB that tunes off that channel causes the DSLAM to stop sending the multicast stream for that channel. This in turn causes the second STB to stop receiving video. The workaround for this problem requires additional functionality in both the STBs and the DSLAM. STBs **must** always send IGMPv2 join and leave requests during a channel-change operation, independently of whether other STBs on the same network segment are currently joined to the same multicast group. The DSLAM **must** keep track of the IP source address associated with each IGMP join and leave request. The DSLAM stops sending a multicast stream to a particular DSL line when all of the IGMP hosts (as specified by the IP source address in each IGMP message) have issued IGMP leave requests. (In fact, these modifications to the IGMPv2 state machine are required in order to make IGMP hosts compliant with IGMPv3.)

IGMP-Based Replication in the AR

To support dynamic replication to aggregation links, the AR is configured in one of two ways, depending on which configuration from [AR Configuration, page 3-12](#), is used. If the AR is configured by means of the single SVI configuration described in [Single SVI Configuration, page 3-13](#), it is configured to perform multicast replication at Layer 2 by means of IGMP snooping. If the AR is configured to use the multiple SVI configuration described in [Multiple SVI Configuration, page 3-14](#), multicast replication is performed at the IGMP/PIM boundary. Both of these methods depend on the AR's ability to process IGMP messages in order to determine when to replicate multicast traffic to the GE aggregation links.

Because the AR potentially aggregates many subscribers, it must be capable of processing a high volume of IGMP join and leave requests if many subscribers are changing channels at the same time.

The solution testing effort characterized the performance of IGMP on the AR by flooding the AR with a constant rate of IGMP join and leave requests, in order to determine the effect on CPU performance in the AR, as well as on the network multicast join delay that contributes to the channel-change performance experienced by an STB. To determine that, an IGMP host makes an IGMP join request for a multicast address that is currently not being sent on the GE aggregation link while the AR is being flooded with IGMP join and leave requests for a different multicast address. The test measures the amount of time it takes from the time the join is sent until the time the stream is delivered, both when the AR is not busy and when it is under various IGMP load conditions.

Static IP Multicast Joins on the AR

If the AR is configured to use static IP multicast joins, all of the multicast streams that are configured with static joins are sent through the distribution network to the AR independently of whether or not IGMP requests have been made by STBs.

Statistical analysis can be used to determine when the use of static joins in the AR does not result in a significant amount of additional bandwidth on the GE aggregation links. The results of this statistical analysis are shown below.

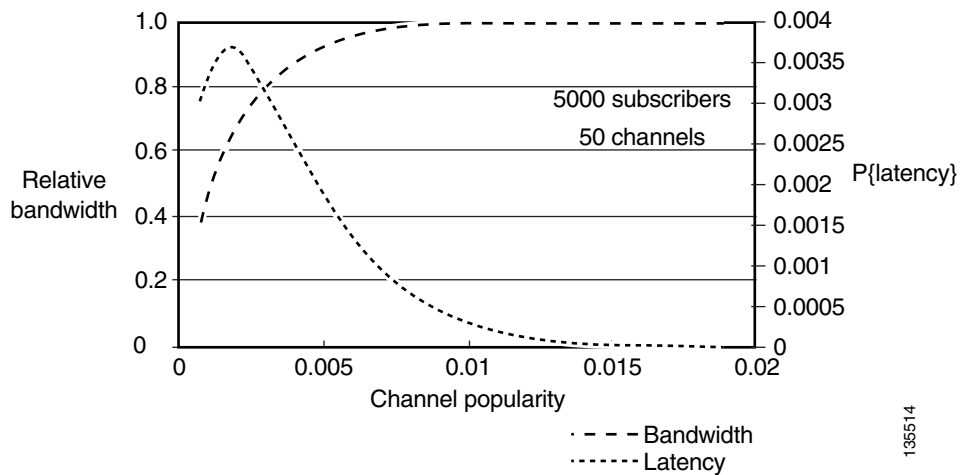
Each service provider must decide, for each channel group, whether that channel group should be a static join or a dynamic join, based on a balance of configuration overhead vs. delay probabilities. Table 3-4 summarizes the factors and formulas used in this analysis.

Table 3-4 Summary of Statistical Analysis

Inputs	Outputs	Formulas
Number of subscribers, N	Bandwidth use of dynamic join vs. static join, B	$B = 1 - (1 - p)^N$ = dynamic-join bandwidth relative to static-join bandwidth
Number of channels, C	Probability of channel-change latency, L	$L = Cp(1 - p)^N$ = contribution to total channel-change latency by this channel group, if joined dynamically
Average channel popularity, p		

Figure 3-14 illustrates the results of the statistical analysis model for bandwidth/latency tradeoff at the AR. Here fixed values are used for the number of channels in a channel group ($C = 50$) and the number of subscribers served by the node ($N = 5000$). Note how bandwidth and latency vary with average channel popularity.

Figure 3-14 Bandwidth and Latency vs. Channel Popularity: 5000 Subscribers at the AR



135514

Based on the above, we can make a general recommendation that channel groups with an average per-channel popularity of 0.05% or less should be joined dynamically at the AR, while channel groups with an average per-channel popularity of greater than 0.05% could be joined statically.

**Note**

From [Figure 3-14 on page 3-26](#), the probability of any additional latency being caused by dynamic multicast joins is at most 0.37%—and typically much less. Because of this, the additional configuration effort required to set up static groups may not result in much benefit other than 100% consistent delay, because it is very rare for a subscriber to experience the additional delay associated with the IGMP join time.

IGMP Functionality in the STB

As described in [Broadcast Client, page 2-3](#), the broadcast client in the video STB is responsible for implementing channel-change requests from a subscriber by issuing an IGMP leave followed by an IGMP join.

Because the bandwidth on the DSL line is often limited, the broadcast client on the STB typically implements the channel-change function by sending an IGMP leave, waiting for the video stream from the channel that is being tuned away to stop, and then an IGMP join. The broadcast client **must** support IGMPv2, because version 2 is the first release of IGMP that provides the ability for a client to signal explicitly when it wants to leave a multicast group. Broadcast clients that support IGMPv2 **should** also send IGMP joins during a channel change, independently of whether other STBs have also sent IGMP joins for the same channel.

**Note**

This behavior is in fact consistent with the IGMP state machine required to support the IGMPv3 state machine documented in RFC 3376. This modified IGMP behavior is needed in order to support fast leave processing in the DSLAM with multiple STBs in the home.

Broadcast clients **should** also support IGMPv3. In addition to IGMP state machine enhancements, the support of IGMPv3 by the broadcast client enables the client to specify one or more IP source addresses of broadcast encoders from which it wishes to receive the broadcast channel. To support this function, the electronic program guide (EPG) must be updated to send both the multicast group address as well as a list of the IP addresses of real-time encoders that may be used for each broadcast channel. When the broadcast client as well as the EPG are updated to support IGMPv3, the multicast solution is significantly simplified, because Source Specific Multicast (SSM) is supported from the STB all the way to the real-time encoder. As a result, there is no need to turn on SSM mapping in the AR.

Internet Access Forwarding

Because different services in the transport architecture use separate VLANs, the forwarding architecture for Internet access may be different from that for video. The Internet access forwarding architecture used in the solution provides an example of how Internet access can be implemented alongside a video service.

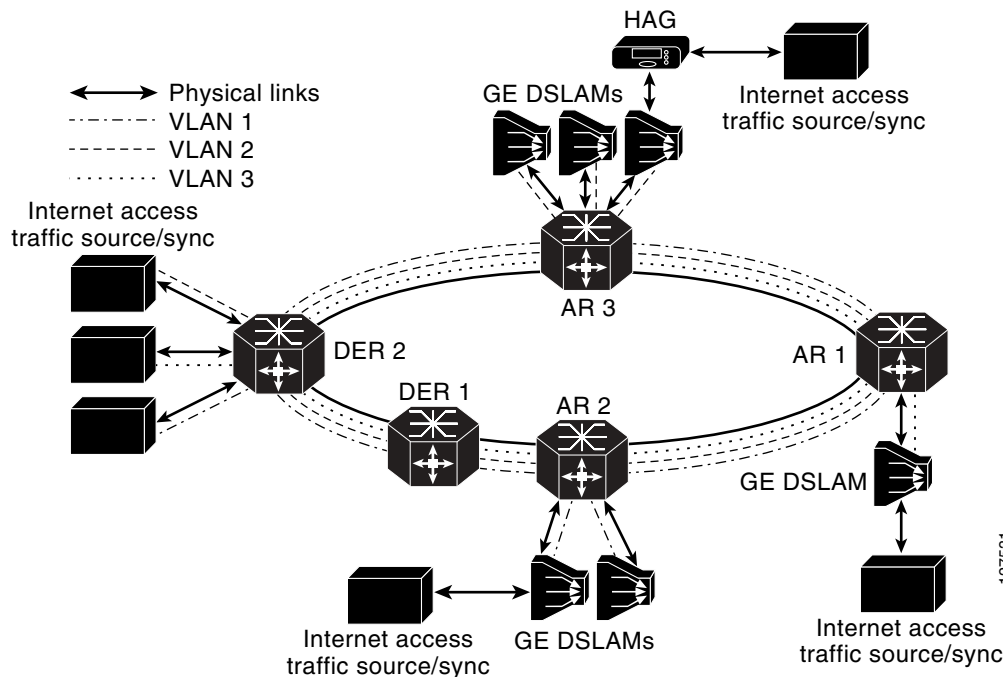
The solution uses Layer 2 forwarding in the aggregation and distribution networks for Internet access. An example Internet access service that could be implemented by this type of architecture would be PPPoE aggregation to a broadband remote-access server (BRAS) that is connected to the DER. [Figure 3-15 on page 3-28](#) illustrates the Layer 2 forwarding architecture used for the Internet access service.

To conserve MAC forwarding entries in both the ARs and DERs, a separate VLAN is used for Internet access for each AR. This configuration makes the VLAN topology look like a hub-and-spoke topology with a separate logical network between each AR and the two DERs. Spanning tree is configured to break the link between the DERs to avoid a forwarding loop. After the spanning tree converges, the VLAN topology looks like separate point-to-point connections between each AR and the DERs. This logical topology conserves MAC address forwarding entries on both the ARs and DERs, because each VLAN now connects only two physical ports. MAC learning algorithms are not needed when a logical topology consists of only two physical ports, because each MAC frame that arrives at one port is always sent on the other port.

**Note**

Solution testing provides only enough testing of the Internet access service to ensure that the transport network forwards frames correctly, and that the Quality of Service (QoS) configuration provides the guarantees required for each service. Because of this, solution testing includes only traffic sources and syncs that emulate Internet access traffic patterns.

Figure 3-15 Configuration for Internet Access Forwarding



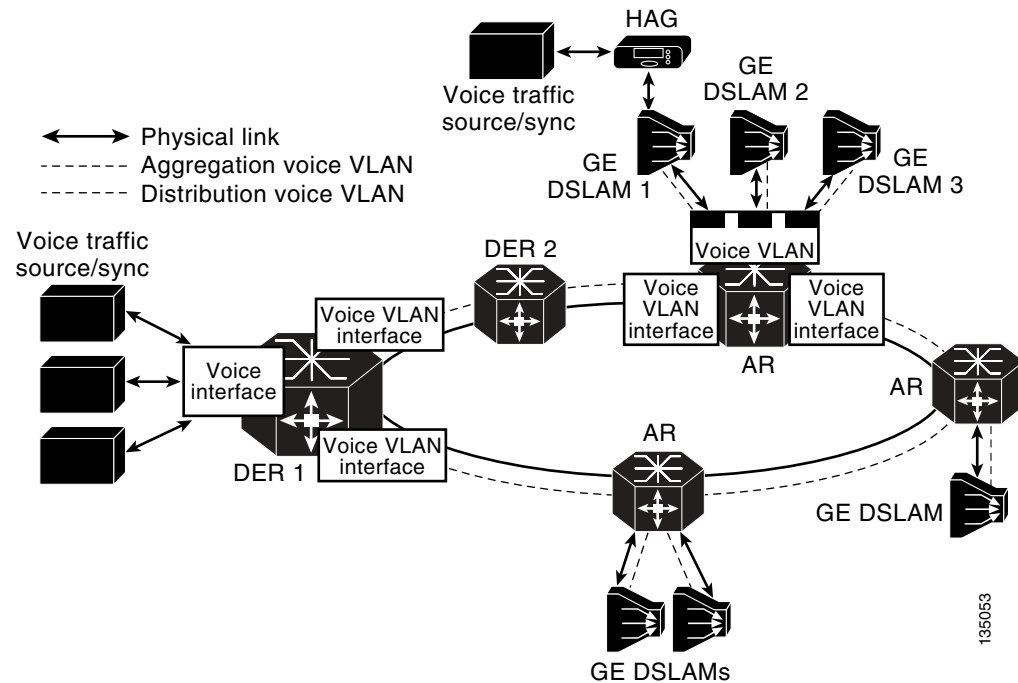
Voice Forwarding

Because the transport architecture uses separate VLANs for each service, the forwarding architecture for voice services may be different from that for video and Internet access. The voice forwarding architecture provides an example of how a voice service may be implemented alongside video and Internet access services.

The transport configuration for the voice service uses a transport architecture similar to that for video. The AR is the Layer 3 edge for voice services. Voice packets are forwarded through the distribution network on a separate VLAN that is terminated in each AR and in the DERs. Voice traffic sources and

syncs are attached to the DERs through separate physical or logical interfaces. A separate routing process is configured for voice and includes all of the voice interfaces on the ARs and DERs. [Figure 3-16 on page 3-29](#) illustrates the voice forwarding configuration used in Release 1.0.

Figure 3-16 Voice Forwarding Configuration



Management

Aspects of management include the separation of services, the management of address spaces, element and network management systems, and service monitoring. These topics are addressed below:

- [Management Transport](#)
- [DHCP Configuration](#)
- [EMS/NMS](#)

Management Transport

The service separation architecture supported in this release of the solution provides the flexibility to allow service providers either to manage each service independently or use a common infrastructure for all services. The level of sharing among services can be controlled by configuring a subset of the IP address to be shared, and deploying common infrastructure components within the shared IP address space. A service provider could thereby share some components such as DHCP and DNS servers, while making other components such as VoD servers specific to the video service.

Solution testing included the configuration and testing of the scenarios where DNS and DHCP servers are shared across services. In Release 1.0, a separate management subnet is configured for components that may be shared across services such as DHCP and DNS servers as well as management hosts and as element management systems (EMS) and network management systems (NMS). These components are

connected to the DER through either a separate physical port or a separate VLAN than are devices associated with video or voice services. The address spaces associated with different services such as voice and video are separated by configuring a separate routing process per service. The management subnetwork can be shared across services by including the interface associated with that subnetwork in the routing process associated with each service.

DHCP Configuration

To enable dynamic address allocation for the devices in the home, the network is configured to support Dynamic Host Configuration Protocol (DHCP). Because the AR is the Layer 3 edge device, DHCP relay functionality is configured on the downstream video VLAN interface of that router. The helper address used with DHCP relay points to a DHCP server located in the management network.

Release 1.0 supports a segmented address allocation scheme that uses a separate DHCP address pool per service. With segmented address allocation, the downstream voice and video SVIs on the AR are associated with separate DHCP address pools. In this environment, the home network is actually divided into three separate IP subnets, where each subnet is associated with a different service topology. The issue with this form of address assignment is that it results in a home network environment where devices within the home network that are associated with different services will be able to communicate with each other only by using a Layer 3 capable home access gateway. (NAT/Layer 3 Functionality, page 3-45, describes an example of functionality that a HAG could implement to enable devices associated with different services to communicate with each other.)

In some environments, the service provider may choose to identify video subscribers by identifying the DSL port that connects the subscriber to the network. In these environments the DSLAM must be capable of snooping DHCP requests from devices in the home network and inserting a DSL port ID in the DHCP request by using DHCP option 82. The DHCP server can then extract this port ID from the DHCP request and use it to identify the subscriber. (DHCP option 82 is described in RFC 3046.)

Note that because the AR is acting as a DHCP relay agent, a DSLAM that supports DHCP option 82 **must** support the ability to appear as a trusted downstream (closer to client) network element (bridge) between the relay agent (AR) and the client (STB). In this mode, the DSLAM inserts DHCP option 82 information but does not set the “giaddr” field in the DHCP request. In addition, because the DSLAM is not acting as a DHCP relay agent, it does not modify the destination MAC address of the DHCP request, and just forwards it using Layer 2 forwarding.



Note

DSLAMs that support option 82 **must** support the relay agent information option of RFC 3046. To enable the DHCP server to identify both the DSLAM and DSL line with which a DHCP request is associated, it is recommended that DSLAMs insert both the management IP address of the DSLAM and the ATM virtual circuit number identifier (VPI/VCI) into the circuit ID suboption field. For consistency, it is recommended that the upper 48 bits of the circuit ID suboption field be the management IP address of the DSLAM, the middle 8 bits be the VPI value of the ATM VC from which the subscriber request originated, and the lower 16 bits be the VCI value of the ATM VC from which the subscriber request originated.

EMS/NMS

The solution does not yet include the integration of element management or network management systems into a video transport solution. The Cisco command line interface (CLI) is the method of configuring the Cisco platforms included in the solution.

Redundancy

The solution addresses fast recovery from the failure of video infrastructure components, as well as of network components in the distribution network, such as physical links or network switching components. Solution testing looked at the recovery associated with failures of video components such as the VoD servers used for on-demand services and the real-time encoders used for broadcast services.

**Note**

Testing focused on Cisco equipment, with generic failures tested on ingress ports for video services. Only multicast reconvergence was tested.

In addition, solution testing has determined how to optimize the network reconvergence time associated with the failures of links in the distribution network, as well as the failure of a DER.

This section discusses two types of redundancy:

- [Video-Infrastructure Component Redundancy](#)
- [Network Redundancy](#)

Video-Infrastructure Component Redundancy

[Figure 3-7 on page 3-12](#) illustrates how the transport architecture supports the redundancy of video infrastructure components such as VoD servers and real-time encoders. The solution test bed included redundant video pumps and real-time encoders attached to redundant DERs.

The solution relies on application-layer failover between the redundant video pumps attached to the DERs in one or more video headends. The video server must support the ability to load-balance VoD sessions between the video pumps attached to the redundant DERs. In addition, a video server must be capable of detecting the failure of a video pump and routing new VoD requests from STBs to still-active video servers in the event of the failure of a video pump.

Solution testing has also characterized the recovery time associated with the failure of real-time encoders by using anycast services. As discussed in [Benefit: Fast Failover of Video Encoders, page 3-10](#), anycast technology can be used to support the ability to detect and recover from the failure of a real-time encoder in the time it takes for the network to reconverge. Release 1.0 testing used redundant real-time encoders configured with the same IP source address attached to the redundant DERs to implement the failover of encoders by using anycast.

Testing simulated the signaling of an encoder (broadcast source) failure, which effectively removes the host route for the failed encoder from the DER. The multicast network between the DERs and the ARs then reconverge. The result is that all that the IP multicast trees for the affected broadcast channel consist of sources from the encoder that is still available.

Network Redundancy

The transport architecture uses dynamic IP routing in the distribution network. This means that the failure of either a physical link or a DER should cause both unicast and multicast routing in the IP transport network to reconverge.

Solution testing has characterized the average and maximum reconvergence times for both unicast and multicast in the event of a link failure or the failure of a DER in the distribution network. The reconvergence trigger events that have been characterized by testing include the following:

- Both an interface and DWDM loss of signal (LOS) caused by a fiber cut
- The failure of a line card within a switching platform

- The loss of an entire DER

Average and worst-case reconvergence times were measured by measuring how long video streams are disrupted at the STB. Testing has also characterized the effect on video quality of the loss of IP video to the STB. During testing, the IP video stream was disrupted for different periods of time (50, 100, 200, 500, and 1000 msec) in order to determine quantitatively the effect of this on video quality. Using this reconvergence and video quality information, the service provider should be able to determine accurately the effect of various network outages in various locations in the video transport network.

Solution testing has also determined the optimal configuration for IP unicast and multicast parameters to optimize reconvergence time for video. Finally, testing has determined the ability of the Quality of Service configuration described in [QoS Architecture, page 3-46](#) to enable the service provider to degrade on-demand services without affecting video broadcast services in the event of a failure in the distribution network.

**Note**

The solution does not include the use of redundant ARs to provide physical-link redundancy to GE DSLAMs. This functionality is planned for a future release of the solution.

Release 1.0 Configurations

Two physical distribution-network topologies based on the transport architecture described in [Distribution and Aggregation Transport Architecture, page 3-4](#) were tested. Both distribution topologies are based on GE rings between the video headend office and video switching offices.

This section presents the following topics:

- [Overview, page 3-32](#)
- [Configuration 1: 10-GE Layer 3 Symmetric Ring, page 3-33](#)
- [Configuration 2: N x 1-GE Asymmetric Ring, page 3-34](#)

Overview

One topology (referred to as Configuration 1) uses a 10-GE ring between the VHO and VSOs. This configuration uses symmetric bandwidth around the ring to provide physical link redundancy for all services.

The other topology uses asymmetric transport links between the VHOs and VSOs to provide a cost-reduced 1-GE transport solution that is optimized for video. This topology reduces the cost of the distribution network by combining bidirectional and unidirectional 1-GE links between the VHO and VSOs. This topology provides asymmetric bandwidth pipes that are optimized for the traffic pattern associated with video. It reduces cost in 1-GE deployments by eliminating bidirectional optics such as lasers when they are not needed. This topology provides physical link redundancy for the Internet access, voice, and broadcast video services, but it does not provide full redundancy for VoD services. In the event of a link failure, VoD services are degraded without affecting any of the other services through the use of the QoS architecture described in [QoS Architecture, page 3-46](#).

Configuration 1: 10-GE Layer 3 Symmetric Ring

Figure 3-17 on page 3-33 illustrates the 10-GE-based symmetric ring topology used in Release 1.0. This topology uses the aggregation/distribution transport architecture described in [Distribution and Aggregation Transport Architecture, page 3-4](#), with the Layer 3 edge at the AR.

Figure 3-17 Configuration 1: 10-GE Symmetric Ring

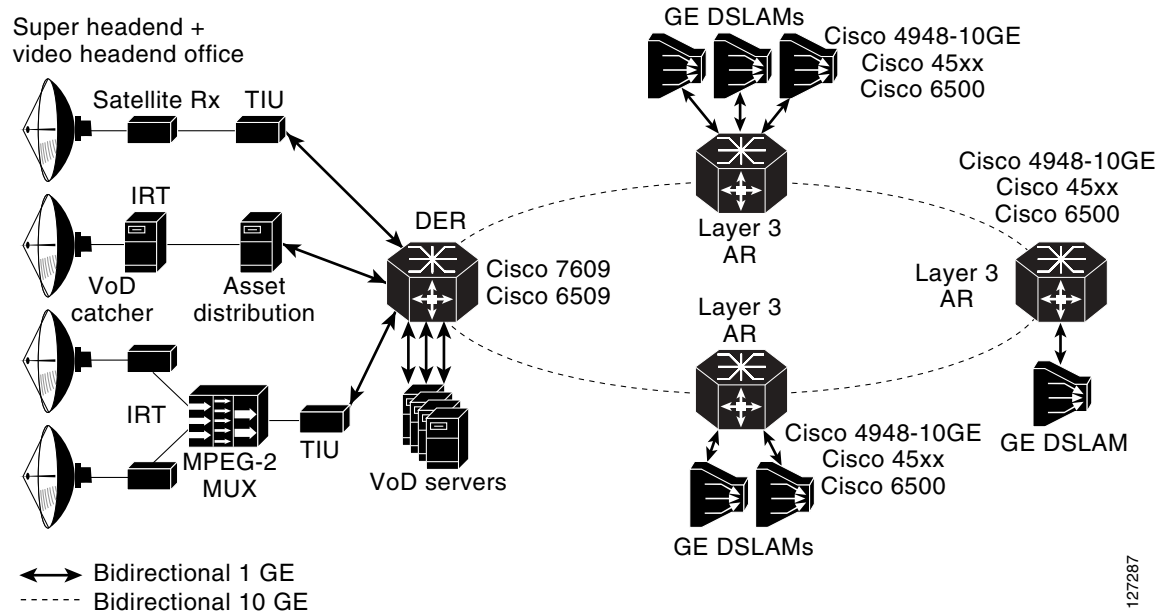


Table 3-5 on page 3-33 lists the transport components tested for Configuration 1.

Table 3-5 Transport Components Tested for Configuration 1

Network Role	Line Card Role	System	Product Number	Interface Type
DER		Cisco Catalyst switch	7609, 6509	
		Supervisor	WS-SUP720-3BXL	N/A
	DER <--> AR	10 GE x 4 optic	WS-X6704-10GE	XENPAK-10GB-LR
	DER <--> VoD servers	1 GE x 24 optic	WS-X6724-SFP	1000BASE-SX, -LX/LH
48-port copper Ethernet		WS-X6748-GE-TX	N/A	
AR		Cisco Catalyst switch	7609, 6509	
		Supervisor	WS-SUP720-3BXL	N/A
	DER <--> AR, AR <--> AR	10 GE x 4 optic	WS-X6704-10GE	XENPAK-10GB-LR
	AR <--> DSLAM	1 GE x 24 optic	WS-X6724-SFP	1000BASE-SX SFP; 1000BASE-LX/LH SFP
		1 GE x 16 optic	WS-X6816-GBIC	1000BASE-SX GBIC; 1000BASE-LX/LH GBIC

Table 3-5 Transport Components Tested for Configuration 1 (continued)

Network Role	Line Card Role	System	Product Number	Interface Type
AR		Cisco Catalyst switch	4510R	N/A
	DER <--> AR, AR <--> AR	Supervisor	WS-X4516	X2-10GB-LR
		1 GE x 6 optic	WS-X4306-GB	1000BASE-SX GBIC; 1000BASE-LX/LH GBIC
			WS-X4448-GB-RJ45	N/A
AR		Cisco switch	4948-10GE	N/A
		Supervisor	WS-X4516	X2-10GB-LR

The 10-GE topology shown in [Figure 3-17 on page 3-33](#) provides fiber redundancy for all services. A link cut anywhere on the ring results in traffic from all services being rerouted in the other direction around the ring. Solution testing includes a test scenario where the failure of a link in the 10-GE ring and the resulting rerouting of video traffic results in steady-state congestion of video traffic on the remaining 10-GE links. In this scenario, the QoS configuration described in [QoS Architecture, page 3-46](#), causes the VoD flows to be affected, while not affecting the broadcast video service at all.

Configuration 2: N x 1-GE Asymmetric Ring

[Figure 3-18 on page 3-35](#) illustrates the N x 1-GE-based topology used in Release 1.0, where N represents any multiple of 1-GE rings. As in Configuration 1, this topology uses the aggregation/distribution transport architecture described in [Distribution and Aggregation Transport Architecture, page 3-4](#), with the addition of unidirectional transport links between the DERs and the AR. Also as in Configuration 1, the Layer 3 edge is at the AR.

Figure 3-18 Configuration 2: N x 1-GE Asymmetric Ring

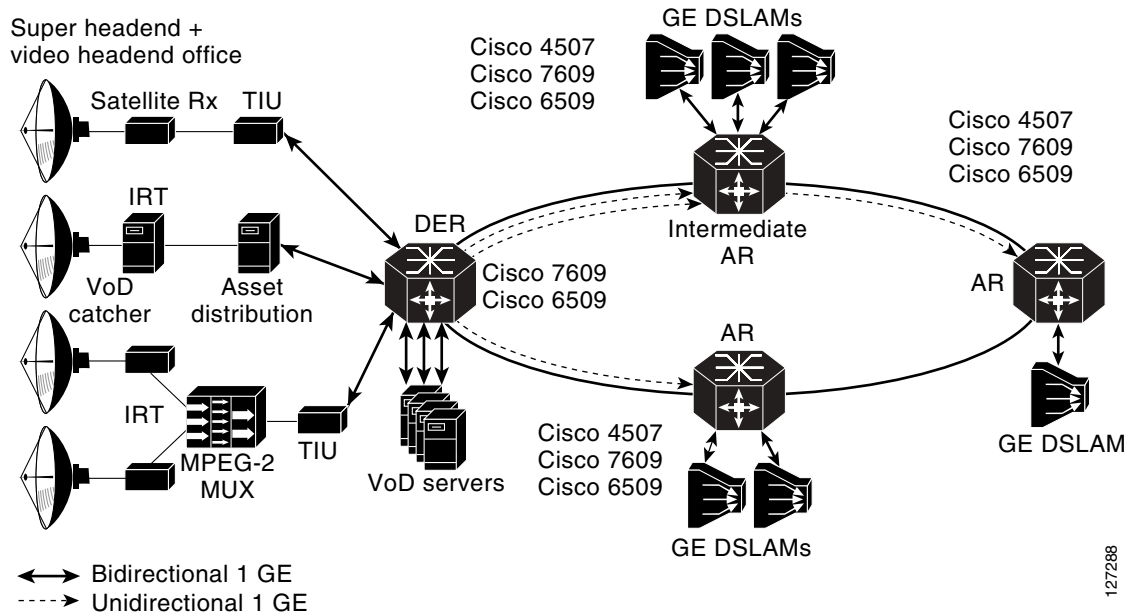


Table 3-5 on page 3-33 lists the transport components tested for Configuration 2.

Table 3-6 Transport Components Tested for Configuration 2

Network Role	Line Card Role	System	Product Number	Interface Type
DER		Cisco Catalyst switch	7609, 6509	
		Supervisor	WS-SUP720-3BXL	N/A
	DER <--> AR	1 GE x 24 optic	WS-X6724-SFP	1000BASE-SX, -LX/LH; 1000BASE DWDM
		1 GE x 16 optic	WS-X6816-GBIC	1000BASE-SX GBIC; 1000BASE-LX/LH GBIC
	DER <--> VoD servers	1 GE x 24 optic	WS-X6724-SFP	1000BASE-SX SFP; 1000BASE-LX/LH SFP
48-port copper Ethernet		WS-X6748-GE-TX	N/A	
AR		Cisco Catalyst switch	7609, 6509	
		Supervisor	WS-SUP720-3BXL	N/A
	DER <--> AR, AR <--> AR	1 GE x 24 optic	WS-X6724-SFP	1000BASE-SX SFP; 1000BASE-LX/LH SFP
		1 GE x 16 optic	WS-X6816-GBIC	1000BASE-SX GBIC; 1000BASE-LX/LH GBIC
	AR <--> DSLAM	1 GE x 24 optic	WS-X6724-SFP	1000BASE-SX SFP; 1000BASE-LX/LH SFP
		1 GE x 16 optic	WS-X6816-GBIC	1000BASE-SX GBIC; 1000BASE-LX/LH GBIC

Table 3-6 Transport Components Tested for Configuration 2 (continued)

Network Role	Line Card Role	System	Product Number	Interface Type
AR		Cisco Catalyst switch	4507R	N/A
		Supervisor	WS-X4515	X2-10GB-LR
	DER <--> AR, AR <--> AR	1 GE x 6 optic	WS-X4306-GB	1000BASE-SX GBIC; 1000BASE-LX/LH GBIC
			WS-X4448-GB-RJ45	N/A

The 1-GE configuration could be used in networks where the amount of traffic in the distribution network does not initially support the deployment of 10-GE links. The transport architecture supports a “pay as you grow” model of deployment, where additional bandwidth could be deployed by adding 1-GE links in the ring as the amount of traffic grows. Because both VoD and broadcast video essentially generate traffic in only one direction, the asymmetric transport architecture allows bandwidth to be added for video in a more cost-effective manner than is provided by fully symmetric bandwidth. Video bandwidth can be added to the distribution network by deploying additional unidirectional 1-GE links between the DER and the ARs. The use of unidirectional links provides a significant savings in the transport network, because most of the cost associated with a DWDM GE interface (approximately 80 percent) is the cost of the transmit laser. The transmit laser may be integrated into the line card by means of pluggable DWDM optics, or it may be implemented externally from the line card by means of a DWDM transponder. The 1-GE line cards tested in this configuration support receive-only pluggable DWDM optics that can be used on the receive side of a unidirectional link. In addition, Cisco supports unidirectional 1-GE DWDM transponders.

The topology illustrated in [Figure 3-18 on page 3-35](#) is implemented by means of a 1-GE bidirectional ring between the DER and the ARs. Additional unidirectional 1-GE links are then used from the DER to the first-hop ARs, to provide additional bandwidth for VoD. Unidirectional 1-GE links may also be used between ARs for this purpose as well.

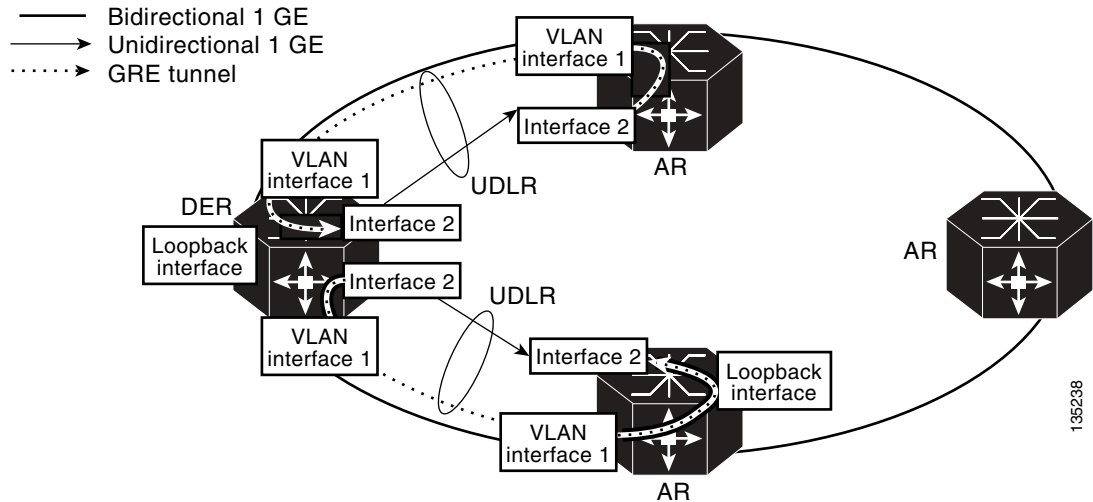
Bidirectional Interface Support

Some protocols that are required to support dynamic routing such as ARP, OSPF, and other routing protocols make an implicit assumption that they are running on bidirectional interfaces. When these protocols are enabled on an interface they respond to protocol requests on the same interface that the request was received on. Because the solution transport architecture is intended to support dynamic routing, it is important that the unidirectional links appear as if they are bidirectional to these protocols. Consequently, Unidirectional Link Routing (UDLR) is used on unidirectional interfaces to make them appear bidirectional at the interface layer. UDLR combines a unidirectional link with a GRE tunnel that is provisioned between the same two nodes that the unidirectional link runs between. Both the unidirectional link and the GRE tunnel are combined at the interface layer to make the unidirectional link appear to be bidirectional.

In Release 1.0, the upstream GRE tunnel is configured between loopback interfaces on the downstream and upstream routers. The tunnel is routed back to the upstream router through the bidirectional 1-GE port. OSPF cost metrics are configured on the downstream interface to direct upstream IP packets through the bidirectional IP interface, as opposed to the UDLR back-channel interface. This is needed because the UDLR back channel is process switched and will have very low throughput capabilities as a result.

[Figure 3-19](#) illustrates how UDLR is configured on the unidirectional interfaces in the asymmetric Ethernet topology.

Figure 3-19 Interface Configuration for Bidirectional Interface Support



Routing Configurations

Note that the unidirectional links do not provide path redundancy around the ring. Because of this, if a unidirectional link fails, there will be no alternate path to carry the traffic on that link around the ring in the other direction. The unidirectional links provide cost-optimized transport for VoD traffic. Because of the relaxed availability requirements for VoD services, it may not be necessary to provide path redundancy on links used to carry VoD traffic. Note that this is typically not the case for other triple-play services that include Internet access, voice, and broadcast video. Because of this, it is important that packets associated with the Internet access and voice services be routed only to the bidirectional links that provide path redundancy, while packets associated with video services may be routed to either the bidirectional or unidirectional links.

As discussed previously, the availability requirements associated with broadcast video are typically more stringent than those for VoD. In the event of a link failure, the QoS architecture described in [QoS Architecture, page 3-46](#), ensures that VoD packets are dropped before broadcast video packets are. While this QoS configuration should be sufficient to ensure that broadcast video and other services are unaffected in the event of a link failure, some providers may want to ensure that a failure that causes a disruption of the VoD service will not cause any degradation to the other services. This can be achieved by ensuring that VoD packets are routed only to unidirectional links, while broadcast video packets are routed only to bidirectional links.

Two routing configurations tested as part of the solution can be used as part of the asymmetric topology shown in [Figure 3-18 on page 3-35](#). These support broadcast and on-demand video routed together, as well as routed separately. Both configurations support Internet access and voice services (which are configured identically for both of the routing options). These two configurations implement the methods of routing the broadcast video and VoD services described above, and route Internet access and voice packets to the bidirectional links that provide path redundancy.

However, the two configurations differ in how they forward broadcast video and VoD packets. The first configuration routes broadcast video and VoD packets in the same logical topology, which uses both the bidirectional and unidirectional links of the asymmetric topology. The second configuration uses separate logical topologies for broadcast video and VoD. Broadcast video shares the same logical

topology as voice in the distribution network and is routed only through the bidirectional links. On the other hand, VoD is carried in its own logical topology, which is routed only through the unidirectional links.

Broadcast and On-Demand Video Routed Together

In this configuration, both broadcast video and VoD traffic are carried in a single logical pipe through the distribution network. This logical pipe consists of a video SVI on the bidirectional GE links in the distribution network and a Layer 3 video interface on each unidirectional link. A video routing process (OSPF) is configured across the video SVIs on the bidirectional links as well as on Layer 3 video interfaces on the unidirectional links.

Because each of these interfaces is configured as a Layer 3 interface, CEF switching is used to switch packets among the interfaces. An implicit property of CEF switching is that it performs load balancing between IP interfaces when the CEF forwarding table shows the interfaces have the same cost to reach a particular destination. When IP interfaces are configured on two or more physical links between a pair of forwarding nodes and the interfaces are configured to have the same cost, all of the destinations reachable by these interfaces have the same cost in the CEF forwarding table. Because of this, all of the IP flows that are routed over the video interfaces configured on the unidirectional and bidirectional links are load balanced. On the AR and DER platforms supported in Release 1.0, the IP CEF load-balancing function is implemented as part of the hardware-accelerated forwarding function. The load-balancing function ensures that all packets associated with each flow are forwarded over the same link within the load-balanced group. The load-balancing function is implemented by feeding specific Layer 3 fields (IP Src, IP Dest) and Layer 4 headers (UDP/TCP src, UDP/TCP dest) of each packet into a polynomial hash function. The result of this hash is then used to select the interface through which the IP packet is sent.

Because of CEF load balancing, the video interfaces configured on the unidirectional and bidirectional links can be treated as one large pipe in the transport network design. Consequently, when a unidirectional link in this design fails, IP routing reconverges with the result that broadcast video and VoD traffic are load balanced across the remaining links. This reconvergence may result in a condition where the remaining links in the bundle are in a steady state of congestion resulting from the broadcast video and VoD traffic. The QoS configuration described in [QoS Architecture, page 3-46](#), ensures that the broadcast video service is not affected when this condition occurs. Unfortunately, the QoS configuration cannot prevent the Internet access service from being affected, because this is essentially a best-effort service. Service providers that want to ensure that only the VoD service is affected in the event of a unidirectional link failure should use the routing configuration described in [Broadcast and On-Demand Video Routed Separately](#), below.

Broadcast and On-Demand Video Routed Separately

In this configuration, broadcast video and VoD traffic are carried in separate logical topologies through the distribution network. The real-time encoders used for broadcast video are connected to the DER through different IP interfaces than are the VoD components. As with the previous configuration, a video SVI is configured on each bidirectional GE port, while a Layer 3 video interface is configured on each unidirectional GE port.

To ensure that broadcast video traffic is constrained to the bidirectional links and VoD traffic is constrained to the unidirectional links, separate routing processes are configured for respective interfaces of the DER and ARs. The broadcast video routing process includes the interfaces that connect the real-time encoders as well as the video SVIs configured on the bidirectional GE links. The VoD routing process includes the interfaces that connect the VoD components and the Layer 3 video interfaces configured on the unidirectional links. When there is more than one unidirectional link between a pair of nodes, CEF forwarding ensures that VoD traffic is load balanced across all of the unidirectional links.

Both broadcast video and VoD are carried on the same VLAN on the GE aggregation links between the AR and the DSLAMs. This means that the AR must merge the broadcast video and VoD traffic, which are carried on separate Layer 3 interfaces through the distribution network, onto a single VLAN on the GE aggregation link to the DSLAM. This is implemented on the AR by configuring a static route to the interface associated with the video VLAN of the GE aggregation links to the DSLAM. This static route is injected into the VoD routing process configured on the unidirectional upstream GE links, but not into the broadcast routing process that is configured on the broadcast VSIs associated with the broadcast VLAN on the bidirectional upstream GE links. This configuration works because VoD forwarding relies on destination-based routing; on the other hand, broadcast video forwarding is based on multicast forwarding, which uses reverse-path forwarding to build the multicast tree. Reverse-path forwarding for multicast works because the broadcast routing process installs upstream routes for the broadcast network in the AR's routing table. Destination-based routing for VoD works because the static route configured for the downstream VLAN interface is redistributed to the VoD routing process, which in turn distributes those routes through the upstream VoD routing domain.

**Note**

There are various ways to configure the AR to merge the broadcast video and VoD traffic on the downstream VLAN interface. For the approach tested by the solution, see [Chapter 4, "Implementing and Configuring the Solution."](#)

Internet Access and Voice

Internet access and voice services are configured identically for both of the routing options described above. As described in [Internet Access Forwarding, page 3-27](#), Internet access is aggregated at Layer 2 in the distribution network by means of a separate VLAN for each AR. To ensure that Internet access is constrained to the bidirectional link, the VLAN used to forward Internet access is assigned only to the upstream and downstream bidirectional links. The VLANs used for the voice service are configured in a similar manner, except that the voice VLAN associated with each bidirectional link is terminated in an SVI, so that it is switched at Layer 3 in the distribution network.

To enable the dynamic routing of voice in the distribution network, an OSPF routing process is configured on the DERs as well as on the ARs on each of the voice SVIs. This routing process is configured only on the voice SVIs, to ensure that the Layer 3 topology for voice converges independently of the video topology.

Edge Transport Architecture

The edge transport architecture specifies how traffic from the voice, Internet access, and video services are aggregated in separate logical topologies in the aggregation and access networks. The edge network consists of the GE aggregation links between the ARs, the DSLAMs, the DSL links, and the HAG.

**Note**

While the solution specifies the interfaces between the home network and the HAG that are needed to support service separation, it does not specify either the transport technology or architecture that is used in the home to support service separation.


This section presents the following topics:

- [Overview, page 3-40](#)
- [DSLAM Functions, page 3-40](#)
- [HAG Functions, page 3-42](#)

Overview

Figure 3-20 illustrates the nodes and links in the edge network.

Figure 3-20 Edge Transport Network



As described previously in [AR Configuration, page 3-12](#), voice, video, and Internet access services are separated on the aggregation GE links by assigning each service to a separate 802.1q VLAN. On the DSL link, services are separated by using a separate ATM PVC or a separate 802.1q VLAN tag per service. The DSLAM and HAG each include the service with which a packet is associated as part of the algorithms for switching packets.

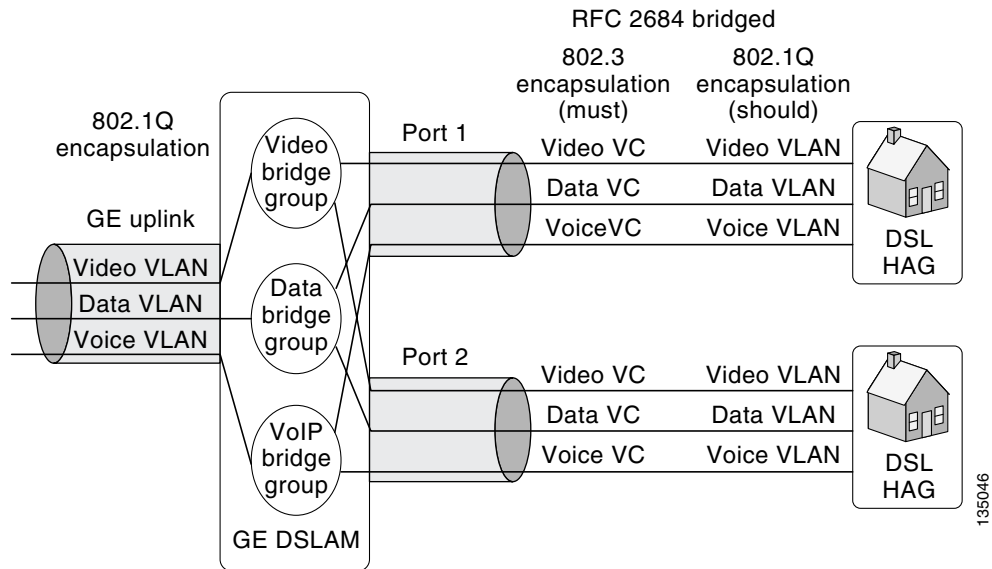
DSLAM Functions

The solution uses an Ethernet DSLAM that performs MAC layer switching between the ATM VCs on the DSL links and the GE uplink. Mac layer bridging in the DSLAM is enabled through the use of RFC 2684 bridged encapsulation on each VC of the DSL link.

The DSLAM maintains a separate Ethernet bridge group per service. A separate instance of the MAC layer learning and forwarding function is maintained per bridge group. Incoming packets from the GE link are mapped to a bridge group according to the 802.1q VLAN tag in the packet. Incoming packets from each DSL link are mapped to a bridge group by means of one of two methods.

The first method is to use the ATM VPI/VCI value associated with each AAL-5 frame. The second method is to have the DSLAM and HAG support the use of 802.1q VLAN tags over the DSL link. With this form of encapsulation, the DSLAM can use the 802.1q VLAN tag in packets from the DSL link to determine the bridge group to which each packet should be mapped. To simplify DSLAM and HAG requirements, Release 1.0 requires that DSLAMs support the ability to map incoming frames from the DSL link to a bridge group by using the ATM VPI/VCI, and recommends that DSLAMs support this ability by using an 802.1q VLAN tag. Both of these methods are illustrated in [Figure 3-21 on page 3-41](#).

Figure 3-21 DSLAM Bridge-Group Mapping Options



Because ATM encapsulation is used only on the DSL link between the DSLAM and the HAG, no ATM-layer switching is performed in the edge network. Also, because the ATM VPI/VCI value for each virtual circuit is meaningful only in the context of the DSL link, the DSLAM can assign the same ATM VPI/VCI value to use on every DSL link for each service. This greatly simplifies the configuration of the DSLAM, because every DSL link can be configured identically.

To ensure that subscribers connected to the DSL links are not able to snoop each other's packets, Ethernet frames (including broadcast frames) that arrive in a bridge group from an ATM VC are always transmitted on the upstream GE link, independently of the state of the MAC forwarding table. Ethernet frames that arrive from the GE uplink are forwarded to a DSL link by means of standard MAC layer learning and forwarding algorithms. Also, there is no need for the DSLAM to participate in the spanning-tree loop-detection algorithms, because the links connected to the DSLAM cannot form a loop.

As previously discussed in [AR Configuration, page 3-12](#), when MAC layer forwarding is used for unicast video applications such as VoD, it is possible that the downstream MAC table entry for a video flow may time out if no packets are sent from the STB during the DSLAM's MAC aging period. To prevent this problem from occurring, it is recommended that DSLAMs implement functionality similar to that described in [AR Configuration, page 3-12](#), as part of the downstream bridge-learning algorithm on the video bridge group. Specifically, it is recommended that DSLAMs support the ability to enable unicast flood blocking on the video bridge group. This feature prevents the DSLAM from flooding unicast traffic when there is no bridge table entry for a destination MAC address. In addition, it should be possible to configure the MAC-table aging timeout value on the video bridge group. In this solution, the MAC-table aging time on the DSLAM should be set to a period longer than the ARP timeout configured on the downstream video VLAN interface of the AR. This configuration ensures that the AR sends an ARP request through the DSLAM to the downstream host before the MAC table entry times out. The resulting ARP request and response ensure that the MAC table entry gets repopulated before it is timed out.

HAG Functions

The home access gateway, or HAG, performs physical adaptation as well as Layer 2 bridging between one or more physical media in the home and the upstream DSL link that uses RFC 2684 bridged encapsulation. The transport architecture does not make any assumptions regarding the physical media used for triple-play services within the home.

The transport architecture also assumes that the home devices that terminate the IP streams for video and Internet access services are typically not integrated into the HAG. Because of this, the architecture assumes that the physical media within the home are capable of transporting IP packets, and use a Layer 2 encapsulation method that can be translated to an 802.3 transport header in a straightforward manner.

For the voice service, the HAG may include an integrated voice gateway that translates VoIP into one or more FXS ports that connect to telephone wiring in the home via one or more RJ-11 ports. In this case, there is no need to carry VoIP traffic within the home network.

[Table 3-7](#) specifies the potential physical media and the associated Layer 2 encapsulations that a HAG may have to translate to the upstream DSL link with RFC 2684 bridged encapsulation.

Table 3-7 Potential Home Wiring Technologies Requiring HAG Support

Physical media	Layer 2 Encapsulation
Air	802.11
Category 5 cable	802.3
Coaxial cable	Media over Coax Alliance (MoCA)
Power line	HomePlug Alliance
Phone line	HomePhoneNetwork Alliance (HomePNA v3)

In Release 1.0 the HAG is responsible for identifying the service topology with which each device in the home network should be associated. This is a very important aspect of the service separation architecture, because it has implications for how services are delivered to the home network.

Service Separation Functions

The basic premise of the service separation architecture is that each device in the home network is associated with a primary service—voice, video, or Internet access. The primary service with which each device is associated determines which service-specific topology that device will use to communicate with the network. On the DSL link the service-specific topology is represented by either the ATM PVC on which the packet is forwarded or the 802.1q VLAN tag that is given to the packet. This form of service mapping and the resulting upstream forwarding algorithms are very different from typical Layer 2 or Layer 3 forwarding algorithms, because the HAG forwards packets that originate from the home network to a VLAN or ATM VC on the DSL uplink according to where the packets came from, as opposed to where they are going.

Release 1.0 includes two methods for how the HAG determines which device a packet came from on the home network. The physical port method uses the physical port on which a packet arrives to associate the packet with a service topology. The device MAC address method uses the MAC address of the packet to associate the packet with a service topology. Each method is described below.

**Note**

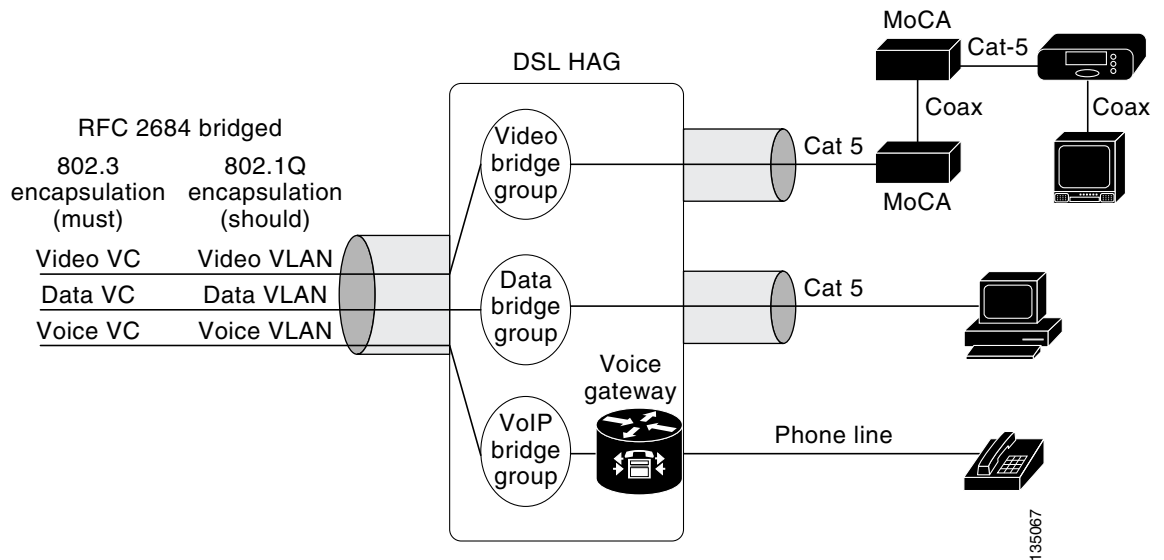
While the solution requires that compliant STBs support one of these two methods of traffic separation, a HAG vendor may implement either method and be compliant with the transport architecture.

Traffic Separation Based on Physical Ports

The physical port method of traffic separation takes advantage of the fact that most existing home communications wiring uses a separate physical medium for each of the triple-play services. The physical wiring of most homes today includes telephone wiring to telephones for telephony services, coax wiring to television sets or STBs for video services, and may include Category 5 wiring for Internet access services (see [Table 3-7 on page 3-42](#)).

A HAG could take advantage of this existing wiring to provide service separation by including an integrated VoIP gateway for the voice service, terminating Media over Coax Alliance (MoCA) wiring for the video service, and terminating either 802.11 or 802.3 for the Internet access service. Because each of these services is terminated in the HAG by means of different physical media, the HAG can determine which upstream VLAN or ATM VC to associate with each packet by determining the physical port on which the packet arrived. [Figure 3-22](#) illustrates traffic separation based on physical ports in the HAG.

Figure 3-22 Traffic Separation Based on Physical Ports



When this algorithm is used without MAC-layer bridging algorithms, it is used for both upstream and downstream traffic. In other words, the mapping of the physical port to the VLAN or the VC is used for packets that arrive at the HAG from the home network and for packets that arrive at the HAG from a VLAN or ATM VC on the DSL link.

**Note**

Because the MoCA technology has not yet been widely deployed, most currently available HAGs and IP-capable STBs do not yet support integrated MoCA and conventional coax capabilities. In this case, the HAG includes a separate Ethernet port for the video service. This port can be bridged to an external MoCA converter that feeds the subscriber's home coax network. A MoCA converter can also be used in front of an Ethernet-based IP STB.

One issue with physical port-based traffic separation is that it enforces a rather rigid mapping of home devices to services. As home network technologies become more ubiquitous, it may not be practical to use the physical media connecting a device to associate a medium with a particular service. In these environments, it may be more practical to use traffic separation based on source MAC address.

Traffic Separation Based on Device MAC Address

A HAG may use the source MAC address of packets coming from the home network to determine the service with which each packet should be associated. In this case, the HAG maintains a mapping of device MAC address to service for one or more services. The mapping of device MAC address to service may be provisioned statically in the HAG or it may be learned dynamically. (The methods by which the HAGs may dynamically learn this mapping are outside the scope of this document.)

The device MAC address table is used for both upstream and downstream traffic. This table maintains a mapping among the device MAC address, the home network port, and the service VC and VLAN on the DSL link. Entries in the table are looked up according to the source MAC address for packets received on a port connected to the home network, and according to the destination MAC address for packets received on the DSL link. Broadcast or multicast packets received from the DSL link are broadcast to all ports connected to the home network. If the table does not contain a match for a packet received on the DSL link, the packet is flooded to all ports connected to the home network. If the table does not contain a match for a packet received on a home network port, the packet is flooded to all VCs and VLANs on the DSL link.

The following algorithm provides an example of how a HAG could implement traffic separation based on Device MAC address:

- Including an integrated VoIP gateway for the voice service
- Having a provisioned range of MAC addresses for the video STBs that may be deployed in the home network
- Associating any unrecognized MAC address with the Internet access service

QoS Tagging

The Quality of Service architecture described in [QoS Architecture, page 3-46](#), is based on the IP Differentiated Services (DiffServ) architecture. Because the HAG in a triple-play environment is managed by the service provider (SP), it is considered to be at the edge of the SP's DiffServ domain. In the DiffServ architecture, one of the important functions of a device at the edge of a DiffServ domain is to mark traffic coming into the domain with the appropriate DiffServ marking. When this function is applied to a HAG, it means that the HAG is responsible for marking all upstream traffic with the correct DiffServ marking for use within the SP's network.

In the QoS architecture described in [QoS Architecture, page 3-46](#), upstream packets associated with the video and Internet access services can be mapped to a single DiffServ code point (DSCP), while upstream packets associated with the voice service can be mapped to one of two different DSCPs reserved for voice bearer and voice signaling traffic.

In addition to associating packets with a logical topology, the traffic classification methods based on physical port and MAC address, described above, are used to mark the correct DSCP in each upstream packet according to the service with which the packet is associated. Because upstream traffic associated with the video and Internet access services is associated with a single DSCP, no additional logic is required to determine the correct marking for these packets. However, upstream traffic associated with the voice service requires additional classification logic to distinguish voice bearer packets from voice signaling packets. [QoS in the Access Network, page 3-56](#), describes how the HAG uses these markings to schedule packets on the DSL link.

NAT/Layer 3 Functionality

To limit the number of IP addresses the service provider must allocate in the home network, and to allow home devices associated with different services to communicate with each other, the HAG may include NAT/Layer 3 functionality instead of MAC learning and forwarding algorithms.

A HAG that includes NAT/Layer 3 functionality implements a separate DHCP (or PPPoE for Internet access) client instance for each service. Each DHCP/PPPoE client communicates with its associated server by using the VC or VLAN associated with that service. As the result of the client/server exchange, the HAG maintains a separate external IP address for each service.

The HAG also implements a local DHCP server that is used to allocate local IP addresses to devices within the home network. All devices within the home network are assigned to a single IP subnetwork whose address is configured in the HAG. When the HAG receives a DHCP request from a home device, it uses one of the service mapping functions described in [Service Separation Functions, page 3-42](#), to determine the service or external address with which that device is associated. When the DHCP allocation algorithm is combined with the service mapping function, DHCP transactions can be used to populate a combined NAT/service mapping table. Each NAT/service mapping table entry includes a device's dynamically allocated internal IP address, the external IP address that is associated with it, and the resulting VC or VLAN on which upstream packets should be sent.

Because all of the addresses in the home network are assigned by the HAG to be in a single IP subnet, devices in the home that are associated with different services can communicate by means of standard IP host functionality. A HAG that implements NAT/Layer 3 functionality should also implement standard MAC-layer learning and forwarding functionality between the physical ports attached to the home network. This functionality enables all devices in the home network to communicate independently of the physical port to which they are attached. Because the home network is typically capable of carrying much more bandwidth than is generated as part of the video broadcast service, the MAC-layer forwarding function in the HAG may broadcast Ethernet frames received from the DSL port with a destination MAC address in the multicast address range to all downstream ports.

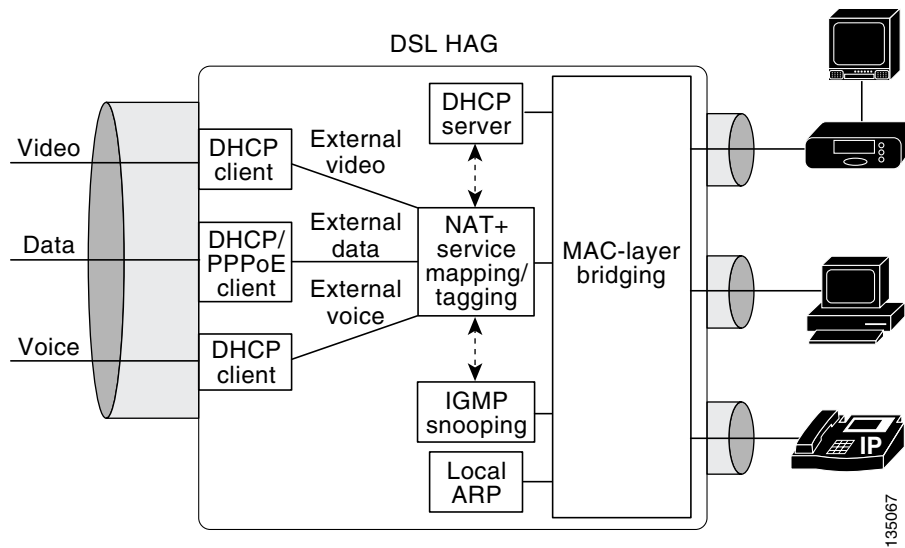
A HAG that implements NAT/Layer 3 functionality **must** implement an IGMP snooping function, in order to enable multicast streams to bypass the NAT logic and be sent directly to the home network without performing address translation. The IGMP snooping function on the HAG **must not** repress any IGMP report messages from home devices. This is needed to enable the DSLAM to implement a fast-leave algorithm that tracks IGMP requests from each home network device. The home network is typically capable of carrying much more bandwidth than is generated as part of the broadcast video service. Because of this, the MAC-layer forwarding function in the HAG may broadcast Ethernet frames received from the DSL port with a destination MAC/IP address in the multicast address range and an active IGMP session to all downstream ports.

A HAG that implements NAT/Layer 3 functionality **must** support a local ARP function for devices on the home network. The ARP function responds to ARP requests for nonlocal IP addresses (IP addresses that have not been locally allocated by the HAG's DHCP server) with its own MAC address.

A HAG that implements NAT functionality **must** implement stateful inspection for Real Time Streaming Protocol (RTSP, RFC 2326) and Session Initiation Protocol (SIP, RFC 3261) as part of the NAT function. RTSP is the most common session-signaling protocol for a VoD session initiated by video STBs, while SIP is the most common session-signaling protocol for IP telephony applications. Stateful inspection is required by NAT functions that support RTSP and SIP because these protocols specify the IP address and Layer 4 port values for video and voice media streams in the payload of signaling messages.

[Figure 3-23 on page 3-46](#) illustrates how Layer 3 functionality in the HAG can be used to implement service separation.

Figure 3-23 NAT/Layer 3 Functionality in the HAG



QoS Architecture

The Quality of Service (QoS) architecture in the solution is based on the IETF Differentiated Services (DiffServ) Architecture described in RFC 2475. The DiffServ architecture assumes that each node in a transport network that is connected to physical links where congestion can occur **must** be capable of scheduling packets from different services separately. In an environment where all services are not aggregated at the BRAS, the DSL access links between home access gateways (HAGs) and DSLAMs, as well as the aggregation links between DSLAMs and aggregation routers, can become congested. This means that aggregation routers, DSLAMs, and HAGs **must** be capable of basic DiffServ functionality.

This section presents the following topics:

- [Overview of DiffServ Architecture](#)
- [DiffServ Architecture in the Solution](#)
- [Triple-Play QoS Analysis](#)
- [QoS in the Aggregation/Distribution Network](#)
- [QoS in the Access Network](#)

Overview of DiffServ Architecture

The DiffServ architecture specifies different requirements for nodes at administrative boundaries than for nodes in the interior of a DiffServ domain. A DiffServ domain is defined as an area where all nodes are configured with the same DiffServ policies for QoS. The edge of a DiffServ domain is the administrative boundary of that domain.

In the DiffServ architecture, nodes at administrative boundaries must implement a superset of the functionality that nodes in the interior of a DiffServ domain implement. Nodes at administrative boundaries must be capable of rate-limiting traffic coming into a DiffServ domain by using a rate-limiting technology such as policing or shaping. Nodes at administrative boundaries must also be

capable of marking traffic that is supposed to have different per-hop behaviors, by using separate DSCP code points. An example of a node at the edge of a DiffServ domain in a residential triple-play architecture is the HAG. While the HAG is managed by the service provider, the home network typically is not. Because of this, the HAG must associate packets that arrive from ports attached to the home network with a service and its associated QoS. The functionality that the HAG implements to classify and mark traffic from the home network is an example of the DiffServ functionality required at administrative boundaries.

All nodes in a DiffServ domain that may experience packet congestion must be capable of classifying packets by means of a DiffServ code point (DSCP) and implementing the specified DiffServ per-hop behavior (PHB) accordingly. This functionality must be implemented on nodes in the interior of a DiffServ network as well as on nodes at an administrative boundary.


The DiffServ architecture described in RFC 2475 assumes that all nodes where congestion can occur are capable of implementing QoS functionality at the IP layer. One can extend this basic architecture to nodes that implement QoS functionality at Layer 2 by mapping the DiffServ PHBs specified by DiffServ code points to Layer 2 functionality at the edges of the Layer 2 network. An example of a Layer 2 technology that implements QoS is ATM. The ATM specification defines its own methods of obtaining QoS by using functionality that is part of ATM switching. The ATM traffic-management specification defines classes of service that must be implemented by ATM switching nodes as well as by the nodes at the edge of an ATM network that implement the Segmentation and Reassembly (SAR) function. Examples of services classes defined by the ATM traffic-management specification are Constant Bit Rate (CBR), Variable Bit Rate (VBR), and Unspecified Bit Rate (UBR). In a DSL environment, each of the ATM CoS values can be mapped to a DiffServ PHB without sacrificing the overall QoS requirements of the network.

While the edge of a DiffServ domain represents one level of boundary of trust, service providers (SPs) may choose to implement a second, more secure boundary of trust within the interior of the DiffServ domain. For example, while the functions the HAG are considered functions of a boundary of trust, the HAG may be compromised because it is not located within the SP's premises. Because of this, an SP may choose to implement additional enforcement functions such as policing at a location in the network that is considered more secure.

DiffServ Architecture in the Solution

This section describes how the Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband (GOVoBB) Solution uses the DiffServ architecture to implement QoS in support of triple play. [Figure 3-24](#) illustrates the upstream QoS and security functionality used in the solution, while [Figure 3-25 on page 3-48](#) illustrates the downstream QoS and security functionality used.


Figure 3-24 Upstream QoS and Security



135270

1	Internet access service enforcement (per-subscriber policing)
2	Administrative boundary (service-based marking); DiffServ-to-Layer 2 mapping (802.1p, ATM CoS)
3	VoD boundary of trust (per-flow policing of signaling)
4	Broadcast video boundary of trust (multicast access lists, IGMP policing)

Figure 3-25 Downstream QoS and Security



135271

1	Internet access service enforcement (per-subscriber policing, shaping, queueing)
2	DiffServ-to-Layer 2 mapping (ATM CoS)
3	DiffServ-to-Layer 2 mapping (802.1p)
4	Video administrative boundary (service-based marking)

Administrative Boundaries

When the DiffServ architecture is mapped to the solution transport architecture, the DiffServ administrative boundaries can be mapped to specific nodes, as discussed below.

In the upstream direction, the administrative boundary is the HAG. While the HAG is managed by the SP, the home network typically is not. Because of this, the HAG must associate packets that arrive from ports attached to the home network with a service and its associated QoS. Since the HAG is at the edge of the SP's DiffServ domain, it **should** be capable of writing a configurable DSCP to each upstream packet, based on the service associated with that packet, by means of the service classification rules described in [Service Separation Functions, page 3-42](#).

The service separation architecture ensures that the video headend infrastructure resides in a separate logical topology from other services such as Internet access. Because of this, the video headend topology is managed by the SP and can be contained within a single DiffServ administrative domain. If VoD servers and real-time encoders are capable of marking the video streams, as well as the control traffic, with the appropriate DSCP values, then there is no need for the video topology to implement DiffServ edge functionality in the downstream direction. If the VoD servers or real-time encoders are not capable of implementing the DiffServ marking functionality, then the DER should perform this function on their behalf.

DiffServ-to-Layer-2 Mapping

In most DSL/Ethernet aggregation architectures, the access and aggregation networks include nodes that are not capable of supporting IP-layer QoS. In such architectures, the DSLAM and the HAG typically implement both packet forwarding and QoS algorithms at Layer 2. [QoS in the Access Network, page 3-56](#), provides the details of how DiffServ-to-Layer-2 Mapping is used to provide proper scheduling behavior in the DSL access network.

Security and Additional Boundaries of Trust

While the edge of a DiffServ domain represents one level of a boundary of trust, SPs may choose to implement a second, more secure boundary of trust within the interior of the DiffServ domain. While the HAG limits upstream traffic by means of its ATM-based scheduling functionality, the HAG is not located within the SP's premises and may therefore be compromised.

The solution architecture specifies additional functionality that is used to provide additional security for both VoD and broadcast video services, as discussed below.

For VoD services a second upstream policing function is implemented on the aggregation platforms located in either the video switching office or the video headend office. The upstream policing function uses the per-flow policing functionality of the Cisco Catalyst 6000 and Cisco Catalyst 7600 series

switches used in the solution to limit the amount of upstream video signaling traffic for VoD to a specified upper limit. With per-flow policing, each upstream flow is recognized dynamically in hardware, and for each new flow a separate hardware-based policer is instantiated. Each policer limits the amount of traffic that is passed upstream to a configured maximum bandwidth and burst size. The bandwidth and burst rate of the upstream policer are determined by the expected maximum bandwidth and burst that video signaling is expected to generate. Any traffic that exceeds the per-flow rate or burst size is dropped.

To limit control-plane-based denial of service (DoS) attacks on the broadcast video service, IGMP access lists can be used on DSLAMs and ARs to restrict multicast join requests to the multicast address range that is known to be valid for the video broadcast service. Any IGMP join requests that fall outside of this address range are dropped. Note that this function is not intended as an enforcement method to limit subscribers to the set of broadcast channels they are authorized to view. The Conditional Access System (CAS) described in [Conditional Access System and Encryption Engine, page 2-6](#) is typically used to implement this function.

In addition to IGMP access lists, DSLAMs and ARs can also restrict the rate at which multicast join requests are accepted by the network through the upstream policing of IGMP traffic. In the AR, IGMP policing can be configured by policing all traffic that matches the IP protocol ID for IGMP to a rate that is less than the maximum performance determined by the IGMP performance testing described in [IGMP-Based Replication in the AR, page 3-25](#).

Triple-Play QoS Analysis

This section describes the analysis behind the DiffServ PHBs and the resulting scheduling QoS configuration recommendations used in Release 1.0 of the solution. This discussion assumes a residential triple-play service with Internet access, voice, and video services. Internet access is assumed to be a best-effort service, with the customer's service-level agreement (SLA) specifying only a maximum (but not a guaranteed minimum) rate. Voice and video are assumed to be managed application services, where the SP provides the subscriber with a video STB and sells the subscriber a video or voice SLA.

An example of a video SLA that an SP may offer is a single channel of VoD or broadcast video delivered to each STB for which the subscriber signs up. The subscriber may sign up for basic or premium-tier broadcast services, and may also sign up for a set of VoD services offered by the provider. The maximum number of STBs that the subscriber may sign up for is limited by the following:

- The type of video service the subscriber requests (for example, standard vs. high definition)
- The video encoding technology used by the SP (for example, MPEG-2 vs. MPEG-4)
- The total amount of DSL bandwidth available to the subscriber

Internet Access

If Internet access is sold as a best-effort service, the DiffServ Default PHB can be used to schedule packets classified as belonging to the Internet access service. The DiffServ Default PHB is described in RFC 2474. The DiffServ PHB provides a best-effort packet-scheduling behavior.

On the aggregation and distribution edge routers, the Default PHB is implemented by using a weighted scheduler that is configured for a minimum bandwidth guarantee. This configuration ensures that Internet access traffic does not significantly affect jitter, latency, or drop for packets associated with the voice or video services.

Voice

End-to-end latency and jitter are very important for a VoIP service. A typical end-to-end jitter requirement for a carrier-class VoIP service is 60 msec. Low jitter and latency are essential to a voice service because the additional delay that results from both factors makes conversations less of an interactive experience, degrading the telephone user's experience.

While delay is an extremely important factor for a successful voice service, the drop requirements for a voice service are not as stringent as they are for video. The reasons that packet drop requirements are not as stringent for a voice service as for a video service are due to digital-to-analog translation algorithms available in current VoIP endpoint implementations. These implementations include a concealment algorithm that can conceal the effects of the loss of a 30-msec voice sample. This means that a packet loss that causes less than a 30-msec loss of digital audio results in an analog signal with no noticeable impairment to the user. With voice concealment algorithms it takes a loss of two or more consecutive 20-msec voice samples to result in a perceptible loss of voice quality. A drop rate of 1% in a voice stream results in a loss that could not be concealed every three minutes when concealment algorithms are taken into account. A 0.25% drop rate results in a loss that could not be concealed once every 53 minutes on average.

Because of this stringent latency requirement, voice services use the DiffServ EF (Express Forwarding) PHB. The DiffServ EF PHB is described in RFC 3246. The EF PHB defines a scheduling behavior that guarantees an upper bound on per-hop jitter that can be caused by packets from non-EF services.

On the aggregation and distribution edge routers, the EF PHB is implemented by means of a priority scheduling algorithm. This algorithm ensures that voice packets can only be delayed by at most one packet serialization time by nonvoice packets per network hop. This delay amounts to a maximum of 12 microsec per hop on 1-GE links configured for a 1500-byte MTU.

Video

While voice has stringent jitter and latency requirements and a relaxed loss requirement, video has a very stringent loss requirement and a relatively relaxed jitter requirement.

Current video-encryption technologies are not resilient to a loss of information in the compressed video stream. As a result, the loss of a single IP packet in a video stream typically causes a noticeable degradation of video quality. The hit to video quality can vary from pixelization across a few frames to a video stream that is frozen for up to 1 second depending on which information in the video stream is lost. The result is that the packet loss requirements for video are extremely stringent. Because of the lack of a concealment algorithm for video, the allowed drop rate for a video service with at most one visible defect per hour is 10^{-6} .

The maximum jitter requirement for video can be determined by examining the maximum channel change delay for broadcast video. [Broadcast Video Channel-Change Time, page 2-14](#), describes the components of channel change delay for a broadcast video over IP service. From [Table 3-3 on page 3-19](#), the component of channel-change delay associated with the jitter buffer on the STB is typically around 200 msec; the size of this buffer determines the maximum allowed jitter—200 msec—for a VoIP service.

Because traffic associated with the Internet access service is carried in a best-effort queue, it does not have a significant impact on jitter for video. However, because voice is carried in a priority queue, it does have an impact on jitter for video. The impact of voice on video jitter is minimized by the fact that in most triple-play deployments, the link utilization of voice traffic is not a significant amount of the total link bandwidth. When the relatively loose jitter requirement for video (200 msec) is taken into account, the relatively low link utilization does not result in video jitter above the maximum limit.

Because of the above factors, video flows are scheduled by means of the DiffServ AF PHB. (The DiffServ AF PHB is described in RFC 2597.) While the AF PHB does not provide as stringent a jitter guarantee as the EF PHB used for voice, it can be used to guarantee a maximum jitter/latency and drop rate for a class of packets. RFC 2597 defines four different classes of the AF PHB. To maintain consistency with current IETF DiffServ marking conventions, this document uses the AF4 PHB. This means that all video flows are marked with a DiffServ code point in the AF 4X range. In addition to four scheduling classes, the AF PHB makes it possible to provide four different drop characteristics for each scheduling class. These different drop characteristics are called drop precedence values. [Broadcast Video vs. Video on Demand, page 3-53](#), provides details on how the different availability requirements of broadcast video and VoD traffic can be implemented by using different drop precedence values within the AF PHB.

On the aggregation and distribution routers, the AF PHB is implemented by means of a weighted scheduling algorithm. To ensure the packet loss and drop requirements for video, the weight configured on the video queue should be greater than the combined bandwidth of the traffic associated with both services under normal operating conditions. [Broadcast Video vs. Video on Demand, page 3-53](#), provides specific recommendations for the weight that should be applied to this queue.

Burst Accumulation


The last factor that must be taken into account in determining QoS requirements for video is burst accumulation. Burst accumulation is an instantaneous burst of traffic that is caused by multiple sources of video transmitted through an IP network. If two or more sources are unsynchronized, there is a probability that the packets they generate are transmitted at exactly the same time. If this traffic converges at an intermediate physical link, the link experiences an instantaneous build up of packets. As these bursts of packets are transmitted to downstream routers, the bursts becomes even larger. This is called burst accumulation. Burst accumulation is influenced by a number of factors, including the following:

- The number of sources in the network
- The number of hops between the sources and the receivers
- The amount of video traffic carried on network links

Burst accumulation may be characterized by means of probability analysis. A typical burst-accumulation analysis shows the probability of a burst of a particular size based on the number of sources, the number of hops between the sources and the receivers, and the amount of video traffic carried on network links. A burst results in either (1) network jitter, if the router has enough buffering for the burst, or (2) a packet drop, if there is not enough buffering to handle the burst.

[Figure 3-26 on page 3-53](#) provides an example of what a set of burst-accumulation curves look like. (This is an example only and does not represent the data from an actual simulation.) In this example, the number of hops and the number of video sources have been fixed, with separate curves for different values of video link utilization. Note that as the probability decreases, the maximum delay increases. When probability curves such as these are mapped to network design constraints, the probabilities can be mapped to the maximum allowed end-to-end drop rate for a class of traffic, and the delay can be mapped to the maximum amount of jitter that can be expected for one or more flows within a class of traffic. For video, the maximum jitter number has two implications for the transport network and for the video STB:

- The network must have enough buffering to buffer worst-case flows. If there is not enough buffering, the large bursts result in packet loss instead of delay.
- The video STBs must have a large enough jitter buffer to hold the maximum burst size.

Figure 3-26 Example Burst-Accumulation Curves

When burst accumulation is applied to video, the low allowed drop rate for video (10^{-6}) should be mapped to the same probability value on a burst-accumulation probability curve. From Figure 3-26, this low probability is associated with a relatively high maximum delay. If the probability curves were based on real data, the STBs would need 250 msec of buffering to make up for network jitter to a probability of 10^{-5} . It would also mean that the network would need to provide 250 msec of buffering for these worst-case flows.

In current video deployments, the effects of burst accumulation are dramatically reduced by the fact that there are very few sources of video in the network. In current deployments, the sources of video for VoD services are video servers that are located in video headends. The video servers in each headend stream video only to the STBs served by that headend. In addition, the sources of video for broadcast video services are typically located in a single super headend as well as in local headends. The result is that each STB receives broadcast video and VoD streams from at most two locations in the network.

However, burst accumulation may become more of a factor for video services in the future, as diverse VoD and broadcast video content is distributed to VoD servers and broadcast encoders located at multiple points in the network. When video streams destined to a group of STBs can originate from many different locations in the network, video streams from many different sources may converge at multiple locations in the network. In such an environment, both the jitter buffers on STBs, as well as the amount of buffering available in routers in the network, will need to increase.

Broadcast Video vs. Video on Demand

The QoS requirements for video do not change depending on the service with which the video is associated. For example, the allowed drop rate (10^{-6}) and maximum jitter (200 msec) allowed for both broadcast video and VoD services are the same.

Even though the QoS requirements for these two services are the same, the availability requirements typically are not. As described in [Service Availability, page 2-13](#), the availability requirements for a broadcast video service are typically higher than those of a VoD service. In addition, [High Bandwidth, page 2-11](#), explains why the amount of bandwidth consumed by VoD services is typically much higher than that of broadcast video services in aggregation and distribution networks. Because of the different availability and bandwidth requirements associated with both video services, service providers may decide to reduce the cost of the video transport network by not providing as much backup bandwidth for VoD services as for broadcast video services. When transport networks are designed in this way, a network failure should result in reduced capacity of the VoD service, while not affecting the broadcast video service.

A future release of the solution transport architecture will include support for a video admission control (VAC) function. When this functionality is available in the network, a network failure that reduces the amount of transport capacity for video results in the admission control function accepting fewer VoD requests than under normal circumstances. This functionality also results in the failure of existing VoD sessions if the number of existing sessions is greater than the capacity the network can support in its degraded state.

Until video admission control is available, however, a network failure that results in reduced video capacity should result in only the VoD service being affected. Flows associated with the broadcast video service should not be affected even in the event of a network failure. Unfortunately, without admission control all VoD flows that are sent through a link that is congested because of a link failure are affected. This is because there is no way to distinguish one VoD flow from another.

The solution has implemented a VoD priority queueing mechanism that may help users until VAC is available. (For details, see [Configuring QoS on DER, page 4-7.](#)) QoS can be configured to ensure that in the event of a link failure that causes reduced video capacity, only VoD flows are affected. The drop precedence characteristic of the AF PHB can be used to ensure this behavior in the event of a network failure.

Drop precedence associated with the DiffServ AF PHB is implemented in the solution architecture by marking all broadcast video traffic with DSCP value AF41, and marking VoD traffic with DSCP values AF42 and AF43 (high- and low-priority packets, respectively). This marking can be done either by the VoD servers and real-time encoders, or by the service router for VoD servers and encoders that do not support DiffServ marking capabilities.

The two DiffServ drop-precedence behaviors are configured by configuring separate queue thresholds for VoD and broadcast traffic on the weighted queue configured for the AF PHB. Queue thresholds set a limit on the effective queue length for a particular class of traffic that is less than the size of the physical queue. Queue-threshold algorithms are run when packets are received at an egress interface before the packets are entered into the output queue. Threshold algorithms can provide simple algorithms such as tail drop or more complex algorithms such as weighted random early discard (WRED). A tail-drop algorithm compares the current queue length to the length of the threshold configured for a particular class of traffic and drops the packet if the current queue length is greater than the configured threshold. Video packets are not transmitted by means of reliable transport protocols that implement congestion-avoidance mechanisms such as TCP/IP. Consequently, simple threshold algorithms such as tail drop can be used for video.

To ensure that VoD packets are dropped before broadcast video packets in the event of link congestion, a queue threshold is configured for packets marked with AF42 and AF43. The size of this threshold should be the expected ratio of VoD traffic to all video traffic (VoD + broadcast) on the egress link multiplied by the configured queue length. In the distribution network, the ratio of VoD traffic to all video traffic is often between 50% and 80%. If a link ever gets congested with video traffic because of an unexpected network failure, the queue threshold configured for VoD ensures that VoD packets are dropped before they are entered into the output queue.

Voice and Video Signaling

Both voice and video services use IP-based signaling to set up and tear down subscriber-initiated sessions. For voice services, Session Initiation Protocol (SIP) is often used to set up and tear down telephone calls between subscribers. For VoD services, Real Time Streaming Protocol (RTSP) is often used to set up sessions to a VoD server that streams on-demand content. The subscriber-initiated signaling protocol used to change channels for broadcast video services is IGMP.

Both voice and video signaling require better than best-effort treatment in order to have an effective service. Drops of signaling packets delay session setup. Since RTSP and SIP normally use TCP as a reliable transport protocol, the additional delay is caused by dropped packets within the period of a TCP

retransmission window. The service that is most affected by drops of signaling packets is broadcast video. This is because channel-change latency has the most stringent requirements of the three services described, and IGMP does not include a reliable transport method.

To ensure that voice and video session-setup latency is not adversely affected by interface congestion, voice and video signaling are scheduled by means of a class selector PHB. The recommended DiffServ code point (DSCP) to use for voice and video signaling is CS3. On the AR and DER, the class selector PHB is implemented by means of a weighted scheduling algorithm. To ensure that signaling packets are not dropped in the event of congestion, the weight configured on the queue should be greater than the maximum bandwidth expected for voice and video signaling traffic under normal operating conditions.

QoS in the Aggregation/Distribution Network

In the solution architecture, downstream packet scheduling in the aggregation and distribution networks is implemented by means of line-card-based scheduling algorithms to implement the DiffServ PHBs recommended for the voice, video, and Internet access services. (See [Internet Access, page 3-50](#); [Voice, page 3-51](#); [Video, page 3-51](#); and [Voice and Video Signaling, page 3-54](#).)

Based on the DiffServ DSCP values and associated PHBs for each of the four traffic classes listed above, there is an implied assumption that line cards in the aggregation and distribution networks should support four queues (one priority, three weighted) as well as one threshold to differentiate broadcast video from VoD traffic. Although some of the line cards used in the solution transport architecture support only three queues, this in fact does not turn out to be a problem, because the QoS architecture carries video traffic as well as voice plus video signaling in two weighted queues. Both video and voice plus video signaling use the AF PHB, both classes of traffic can be scheduled by using the same queue on the line card without adversely affecting either class of traffic. Combining both classes in the same queue also simplifies configuration, because queue weights need to be configured for only two weighted queues.

[Table 3-8 on page 3-56](#) provides the recommendations for line card configuration using the DiffServ recommendations described in [Internet Access, page 3-50](#); [Voice, page 3-51](#); [Video, page 3-51](#); and [Voice and Video Signaling, page 3-54](#).



Note

Some of the broadcast video and video on demand traffic classes shown in the table are relevant only in the downstream direction. Consequently, the queue weight recommendation shows different values for the downstream and upstream directions.

In addition to DiffServ-based scheduling, the aggregation router sets the 802.1p value of packets being sent on aggregation links according to the DSCP value in each packet. [Table 3-8 on page 3-56](#) also provides the recommendations for marking the 802.1p value in the Ethernet header based on an IP packet's DSCP value.

Table 3-8 Recommendations for Configuring Line Cards for Access/Aggregation Networks

Service	DiffServ PHB	DiffServ DSCP Value	Line Card Queue	Queue Weight	Queue Threshold
Broadcast video	Assured Forwarding (AF)	AF41	Weighted (1)	80% downstream, ¹ 20% upstream ²	N/A
VoD		AF42, AF43			$VoD/(VoD + Broadcast) * Queue_Length$
Voice + video signaling	Class Selector (CS)	CS3			N/A
Voice	Expedited Forwarding (EF)	EF	Priority	N/A	
Internet access	Default	0	Weighted (2)	UBR	

1. The downstream queue weight for video is a recommendation that assumes all video traffic consumes no more than 70% of the physical link bandwidth for the link being configured. If the expected ratio of video traffic to total link bandwidth is significantly less, then a lower queue weight may be used.
2. The upstream queue weight for video takes into account only voice plus video signaling, because broadcast video and VoD traffic is unidirectional. The actual value used for the queue weight may vary, depending on the expected ratio of signaling traffic compared to total link bandwidth.

If the density of the DSLAM is such that the Ethernet uplink can become congested, the DSLAM **must** include upstream scheduling functionality. Since Ethernet-capable DSLAMs forward Ethernet frames by means of MAC-layer switching, they typically implement QoS on the Ethernet uplink by using MAC-layer classification techniques. This is done either by (1) associating the incoming ATM VC of an upstream packet with a service and an associated QoS scheduling class, or (2) by using the 802.1p marking on the packet to associate the packet with a specific scheduling class. DSLAMs compliant with the solution's QoS architecture **must** be capable of associating the incoming ATM VC of an upstream packet with a service and an associated QoS scheduling class. DSLAMs compliant with the solution's QoS architecture **should** be capable of using the 802.1p marking on upstream packets, to associate them with a specific scheduling class. Note that this second form of classification on the DSLAM implies that the HAG **should** be capable of marking the 802.1p value of upstream Ethernet frames according to the service with which the HAG associates a packet.

QoS in the Access Network

The HAG is located at the DiffServ administrative boundary in the upstream direction. While the HAG is managed by the service provider, the home network typically is not. Consequently, the HAG must associate packets that arrive from ports attached to the home network with a service and its associated QoS. Because the HAG is at the edge of the SP's DiffServ domain, it **should** be capable of writing a configurable DiffServ code point to each upstream packet according to the service with which it has associated that packet, using the service classification rules described in [Service Separation Functions, page 3-42](#). For HAGs that implement DSCP marking, the DSCP value with which the HAG marks each packet based on service classification follows the conventions illustrated in [Table 3-9 on page 3-57](#).

The solution provides two potential methods for implementing packet scheduling in the access network. The method used depends on the capabilities of the DSLAM and the HAG. [ATM Layer Scheduling](#), below, describes the required method of packet scheduling based on the ATM layer scheduling methods. [MAC/IP Layer Scheduling, page 3-58](#), describes an additional optional method of packet scheduling, based on MAC/IP-layer scheduling, that DSLAMs and HAGs may use.

ATM Layer Scheduling

Both the DSLAM and the HAG include ATM Segmentation and Reassembly (SAR) functionality. The ATM SAR function encapsulates IP packets in ATM AAL-5 frames and segments each frame into ATM cells. The ATM SAR function that is incorporated in HAGs and DSLAMs is typically capable of implementing the cell-scheduling algorithms required for most of the ATM classes of service defined in the ATM traffic management specification. Consequently, the solution QoS architecture requires support for ATM-layer Quality of Service (QoS) for scheduling across the DSL link.

The use of ATM-layer QoS in the HAG means that the HAG **must** be capable of mapping the service with which it associated with each upstream packet to the appropriate ATM Class of Service (CoS). The use of ATM-layer QoS in the DSLAM means that the DSLAM **must** be capable of mapping the incoming VLAN of downstream packets and the service with which that VLAN is associated to the appropriate ATM CoS on the DSL line. In addition, the DSLAM **should** be capable of mapping the incoming 802.1p value of downstream packets to the appropriate ATM CoS. Note that this second form of classification on the DSLAM relies on the AR to set the 802.1p value in Ethernet frames before they are sent on the aggregation links to the DSLAM.

When ATM scheduling is provided on the DSL line by means of multiple VCs with an ATM SAR function, it provides both scheduling and a link fragmentation and interleaving (LFI) functionality. LFI may be needed for voice services when the amount of upstream bandwidth available on the DSL line is below 400 kbps. LFI is needed in this case because the serialization delay for a single 1500-byte packet could exceed 30 msec on the DSL line (30 msec is about half of the end-to-end jitter budget for a voice service). When multiple ATM VCs are configured on the DSL line, the ATM SAR function breaks each IP packet into a sequence of 53-byte cells and then schedules each cell by means of ATM-based scheduling algorithms. This process ensures that the maximum delay that may be experienced by a voice packet because of the serialization of video or data packets is 1 msec on a 400-kbps DSL link.

[Table 3-9](#) shows the mapping between traffic classes, DiffServ Per Hop Behavior (PHB), and ATM-based scheduling (CoS). [Triple-Play QoS Analysis, page 3-50](#), provides details on the analysis used to determine the mapping used in the solution between services and DiffServ PHBs.

Table 3-9 DiffServ-to-ATM CoS Mapping

Traffic Class	DiffServ PHB	DiffServ DSCP Value	ATM CoS	SCR Value
Broadcast video	Assured Forwarding (AF)	AF41	VBR	Expected bandwidth (bps) * 1.25424
VoD		AF42, AF43		
Video signaling		CS3		
Voice	Expedited Forwarding (EF)	EF	CBR	
Voice signaling	Class Selector (CS)	CS3		
Internet access	Default	0	UBR	—

This table also provides a recommendation for configuring the sustained cell rate (SCR) for VBR and CBR virtual circuits (VCs). The expected bandwidth that is used to calculate the SCR for the voice and video VCs can be determined by multiplying the maximum number of voice and video streams by the maximum bandwidth per stream that is expected on the DSL line. Additional bandwidth may also need

to be added to take into account voice and video signaling through the VC. Note that the recommended SCR values shown in the table provide about 25% extra bandwidth over what is required to support the expected bandwidth value. This is because the SCR value for CBR and VBR VCs is used to guarantee a minimum rate and also enforce a maximum rate for traffic through that VC. The additional 25% value ensures that the maximum rate enforcement does not degrade the voice and video services during normal operation.

MAC/IP Layer Scheduling

Some DSLAMs and HAGs support the ability to schedule packets by means of DiffServ-based scheduling algorithms. [Edge Transport Architecture, page 3-39](#), describes an optional method that uses 802.1q VLAN tags that the DSLAM and HAG can also use to identify the service topology with which each packet on the DSL line is associated. In environments where both the DSLAM and HAG support DiffServ-based packet scheduling and 802.1q encapsulations on the DSL line to identify service topology, a single ATM VC can be used between the HAG and the DSLAM.

When a single ATM VC is configured between the HAG and the DSLAM, the 802.1q marking is used to identify the service topology with which the VC is associated, while either the DSCP value or the 802.1p value is used to classify packets for scheduling. [Table 3-10 on page 3-58](#) shows the mapping between traffic classes, DiffServ PHBs, and DiffServ-based scheduling algorithms on the DSL line. [Triple-Play QoS Analysis, page 3-50](#), provides details on the analysis used to determine the mapping used in the solution between services and DiffServ PHBs.

Table 3-10 DiffServ-to-MAC-Based Scheduling

Service	DiffServ PHB	DiffServ DSCP Value	802.1p Value	Queue	Queue Weight
Multicast (broadcast) video	Assured Forwarding (AF)	AF41	4	Weighted (1)	80% downstream, ¹ 20% upstream ²
Unicast video (VoD) (50%)		AF42	2		
		AF43	1		
Voice + video signaling	Class Selector (CS)	CS3	3		
Voice	Expedited Forwarding (EF)	EF	5	Priority	N/A
Internet access	Default	0	0	Weighted (2)	UBR

1. The downstream queue weight for video is a recommendation that assumes all video traffic consumes no more than 70% of the physical link bandwidth for the link being configured. If the expected ratio of video traffic to total link bandwidth is significantly less, then a lower queue weight may be used.
2. The upstream queue weight for video takes into account only voice plus video signaling, because broadcast video and VoD traffic is unidirectional. The actual value used for the queue weight may vary, depending on the expected ratio of signaling traffic compared to total link bandwidth.

When a single VC is used between the DSLAM and the HAG, IP-based link fragmentation and interleaving functionality may be needed on the HAG and DSLAM. LFI may be needed for voice services when the amount of upstream bandwidth available on the DSL line is below 400 kbps. LFI is needed in this case because the serialization delay for a single 1500-byte packet could exceed 30 msec on the DSL line (30 msec is about half of the end-to-end jitter budget for a voice service). DSLAMs and

HAGs that support the DiffServ-based scheduling and 802.1q-based service-mapping functionality required in single-VC environments **should** also support the ability to implement LFI across the DSL line by means of Multilink Point-to-Point Protocol (MLPPP).



Implementing and Configuring the Solution

This chapter begins with tasks common to the 10-GE symmetric and 1-GE asymmetric topologies used in the Cisco GOVoBB solution:

- [Common Tasks: Configuring SSM Mapping with DNS Lookup, page 4-1](#)

It then presents the details of configuring each topology:

- [Configuring the 10-GE Symmetric Topology, page 4-4](#)
- [Configuring the 1-GE Asymmetric Topology, page 4-53](#)



Note

For command references and best practices for the switches used, see the following:

— Cisco Catalyst 6500 Series Switches:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/index.htm>

— Cisco 7600 Series Router:

<http://www.cisco.com/univercd/cc/td/doc/product/core/cis7600/index.htm>

—Cisco Catalyst 4500 Series Switches:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/>

Common Tasks: Configuring SSM Mapping with DNS Lookup

As discussed in [Multicast, page 3-16](#), Source Specific Multicast (SSM) is used simplify the configuration of a multicast network, and is common to both topologies. The solution uses edge devices that do not support IGMPv3. The switches accept IGMPv2 messages and convert these to IGMPv3 by resolving the source IP address of the multicast group by means of either a static mapping or a DNS resource record. This solution uses a DNS lookup method.



Note

For the details and an extended discussion of SSM mapping, see “Source Specific Multicast (SSM) Mapping” at the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_2/gtssmma.htm

The following tasks are presented:

- [Configuring DNS Servers](#)
- [Configuring SSM Mapping on All Switches](#)
- [Configuring the Edge Switches for DNS Queries](#)

Configuring DNS Servers

The following steps are general. Refer to your DNS server documentation for details.

-
- Step 1** For background, refer to “DNS-Based SSM Mapping” in “Source Specific Multicast (SSM) Mapping,” referenced above.
- Step 2** Configure the following parameters, as appropriate:
- a. Resource records for the first multicast IP address associated with a source
 - b. All other multicast IP addresses from the same source
 - c. The multicast domain
 - d. The timeout (optional)
-

Configuring SSM Mapping on All Switches

Configure the following on all switches (the DER and the ARs) in both topologies.

-
- Step 1** Enable multicast routing.
- ```
ip multicast routing
```
- Step 2** Enable SSM mapping.
- ```
ip igmp ssm-map enable
```



Note

Although the document Source Specific Multicast (SSM) Mapping, referenced above, states that the **ip igmp ssm-map enable** command needs to be configured only on switches that are connected to IGMP clients, it was found that this led to inconsistent recovery times during solution network failure and recovery tests. A majority of the time, recovery was fast, but occasionally recovery times were poor. It was found that configuring this command on the headend switch, recovery times were more consistent, although slightly slower than the best recovery times when SSM mapping was not configured on the headend switch.

- Step 3** Enable SSM on the edge switches. The default IP address range for SSM is 232.0.0.0 to 232.255.255.255.

**Note**

The above command also enables the **ip igmp ssm-map query dns** command. By default, IGMPv2 is configured on the Layer 3 interfaces, so no commands are required to enable SSM mapping with DNS query on the interfaces connected to the device that receives multicast. Also, no special commands are required to enable SSM mapping with DNS query on the Cisco 7609 interfaces that connect to the DNS servers.

Configuring the Edge Switches for DNS Queries

On the edge switches that perform the DNS queries, you must configure the domain and IP addresses of the domain name servers. The domain for the multicast video in the following example is coronado.net. (Domain names will vary.) The switches send queries to the first DNS listed in the running configuration. If the first query fails, the next query is sent to the second DNS.

Step 1 Configure the domain for multicast video.

```
ip domain multicast coronado.net
```

Step 2 Configure the IP address of the first DNS.

```
ip name-server 192.168.10.101
```

Step 3 Configure the IP address of the second DNS.

```
ip name-server 192.168.11.101
```

Configuring the 10-GE Symmetric Topology

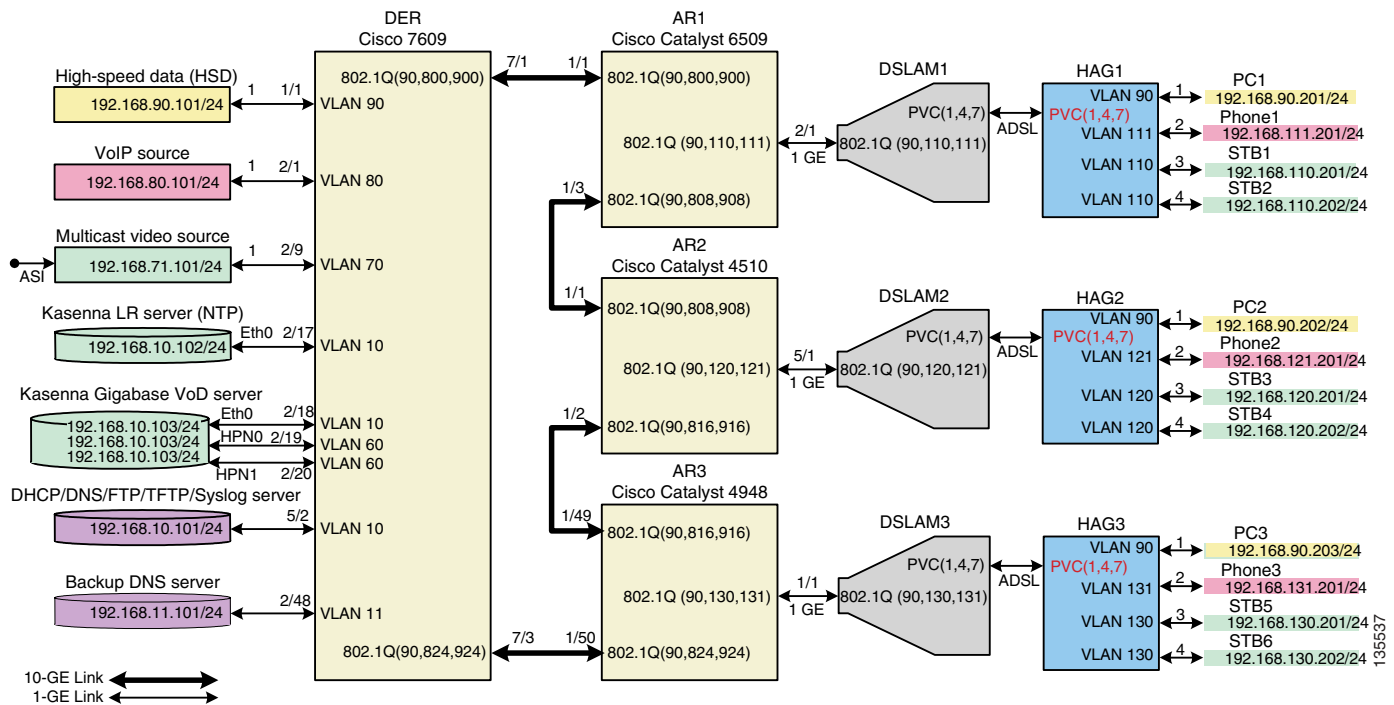
This section presents the following major topics:

- [Introduction, page 4-4](#)
- [Configuring DER, page 4-7](#)
- [Configuring AR1, page 4-21](#)
- [Configuring AR2, page 4-31](#)
- [Configuring AR3, page 4-42](#)

Introduction

Figure 4-1 illustrates the 10-GE symmetric topology used in the solution. (See [Configuration 1: 10-GE Layer 3 Symmetric Ring, page 3-33](#).) All video traffic sources are on DER. Policy maps are applied to the ingress ports on DER in order to mark the DSCP values of the different service types. Traffic is routed through 10-GE bidirectional links, configured as IEEE 802.1q trunks that carry three VLANs: one for video, one for VoIP, and one for high-speed data (HSD). Two OSPF processes are used for the routing protocol. The first advertises routes for the video-related interfaces, and second advertises routes for the VoIP-related interfaces. HSD is carried around the ring on Layer 2. The HAG used in the test bed used service separation based on physical ports, as described in [Traffic Separation Based on Physical Ports, page 3-43](#).

Figure 4-1 10-GE Symmetric Topology



The switches in Figure 4-1 use the line cards, hardware versions, and IOS versions listed in [Table 4-1 on page 4-5](#).

Table 4-1 Hardware and IOS Versions for the 10-GE Symmetric Topology

Switch	Module	Line Card	Hardware Version	IOS Release	Submodule	Hardware Version
DER	1	WS-X6816-GBIC	1.7	12.2(18)SXE1	WS-F6K-DFC3BXL	2.2
	2	WS-X6748-GE-TX	1.4		WS-F6700-DFC3BXL	4.0
	5	WS-SUP720-BASE	3.1		WS-F6K-PFC3BXL	1.2
	7	WS-X6704-10-GE	1.2		WS-SUP720 (MFSC)	2.1
AR1	1	WS-X6704-10GE	1.2		WS-F6700-DFC3BXL	4.0
	2	WS-X6816-GBIC	1.7		WS-F6K-DFC3BXL	2.2
	5	WS-SUP720-BASE	3.1		WS-F6K-PFC3BXL	1.2
AR2	1	WS-X4516-10GE	2.0		12.2(25)EWA	—
	5	WS-X4448-GB-RJ45	1.1			
AR3	1	WS-C4948-10GE	1.0			

Table 4-2 lists VLANs, their descriptions (service types), and IP addresses, for the DER and ARs in Figure 4-1 on page 4-4.

Table 4-2 VLANs, Descriptions, and IP Addresses for the 10-GE Symmetric Topology

Node	VLAN	Description	IP Address
DER	10	Management (VoD signaling, DHCP, DNS, FTP, TFTP, Syslog servers)	192.168.10.1/24
	11	Management (backup DNS server)	192.168.11.1/24
	60	Unicast video aggregation	192.168.60.1/24
	70	Multicast video aggregation	192.168.70.1/24
	80	VoIP	192.168.80.1/24
	90	HSD	Layer 2
	800	VoIP to/from AR1	192.168.252.1/30
	824	VoIP to/from AR1	192.168.252.25/30
	900	Video transport to/from AR1	192.168.254.1/30
	924	Video transport to/from AR3	192.168.254.25/30
AR1	90	HSD	Layer 2
	110	Video edge	192.168.110.1/24
	111	VoIP edge	192.168.111.1/24
	800	VoIP to/from DER	192.168.252.2/30
	808	VoIP to/from AR2	192.168.252.9/30
	900	Video transport to/from DER	192.168.254.2/30
	908	Video transport to/from AR2	192.168.254.9/30

Table 4-2 VLANs, Descriptions, and IP Addresses for the 10-GE Symmetric Topology (continued)

Node	VLAN	Description	IP Address
AR2	90	HSD	Layer 2
	120	Video edge	192.168.120.1/24
	121	VoIP edge	192.168.121.1/24
	808	VoIP to/from AR1	192.168.254.10/30
	816	VoIP to/from AR3	192.168.254.17/30
	908	Video transport to/from AR2	192.168.254.10/30
	916	Video transport to/from AR3	192.168.254.17/30
AR3	90	HSD	Layer 2
	130	Video edge	192.168.130.1/24
	131	VoIP edge	192.168.131.1/24
	816	VoIP to/from AR2	192.168.254.18/30
	824	VoIP to/from DER	192.168.254.26/30
	916	Video transport to/from AR2	192.168.254.18/30
	924	Video transport to/from DER	192.168.254.26/30

Table 4-3 on page 4-6 lists the parameters used to configure the home access gateway (HAG).

**Note**

See [HAG Functions, page 3-42](#), and [Appendix D, “Configuring DSL Equipment.”](#)

Table 4-3 HAG Configuration Parameters

Traffic	VLAN	HAG Ports	PVC ¹	VPI ²	VCI ³	Encapsulation	Service Class	PCR ⁴	SCR ⁵	MBS ⁶
HSD	90	0	1	8	35	LLC	UBR	—	—	—
VoIP	1x0 ⁷	1	4	0	51		CBR	—	300	—
Video	1x1	2, 3	7	8	59		VBR-RT	1200	600	10

1. Permanent virtual connection
2. Virtual path identifier
3. Virtual connection identifier
4. Peak cell rate
5. Sustained cell rate
6. Maximum burst size
7. The *x* corresponds to the AR number 1, 2, or 3 in the corresponding VLAN

Configuring DER

This section addresses the configuration required on the switch labeled DER in [Figure 4-1 on page 4-4](#), to route multiple services from that switch to the ARs.

See [Configuring DNS Servers, page 4-2](#).

**Note**

A Cisco Catalyst 6509 can also be used, as it uses the same supervisor engine, line cards, and Cisco IOS code as the Cisco 7609 router.

This section addresses the following:

- [Configuring QoS on DER](#)
- [Establishing and Configuring Interfaces on DER](#)
- [Configuring OSPF Routing for Video and Voice Traffic on DER](#)
- [Configuring Spanning Tree on DER](#)

**Note**

For a complete configuration example, see [Appendix A, “Sample DER and AR Switch Configurations for the 10-GE Symmetric Topology.”](#)

Configuring QoS on DER

This section presents the following topics:

- [Overview of QoS on a Cisco 7600 Series and Cisco Catalyst 6500 Series](#)
- [Configuring Marking and Classification on DER](#)
- [Configuring Mapping on DER](#)

**Note**

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Overview of QoS on a Cisco 7600 Series and Cisco Catalyst 6500 Series

This section addresses the configuration of quality of service (QoS) on the DER, through marking, classification, mapping, and queuing, to provide different degrees of quality of service for the different types of services supported by the solution architecture. For example, it is important to ensure the expeditious delivery of video and VoIP traffic, while providing only best-effort delivery for high-speed data (HSD).

By default, the Cisco 7600 series router and Cisco Catalyst 6500 series switch do not trust the incoming QoS markings, and therefore rewrite these bits with zeros. In this solution, packets at the network ingress ports are identified, classified, and marked according to type of traffic. The packets are marked with one of 64 possible Differentiated Services Code Point (DSCP) values at the ingress ports. These in turn are internally mapped to one of eight possible Class of Service (CoS) values, because CoS is used to determine the appropriate transmit queue for each packet. Queuing is configured on the individual 10-GE interfaces.

**Note**

For more information on class of service, see “White Paper: Cisco IOS Software Features for Differentiated Class of Service for Internetworks,” at the following URL:

http://www.cisco.com/warp/public/cc/pd/iosw/iore/tech/osfea_wp.htm

Configuring Marking and Classification on DER

Do the following to enable marking and classification on DER.

Step 1 Enable QoS in global configuration mode.

```
mls qos
```

Step 2 Create access lists to identify the different service types in the network.

```
ip access-list extended acl_HSD
  remark Identify HSD traffic
  permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
  remark Identify VoD signaling traffic
  permit ip host 192.168.10.102 any
  permit ip host 192.168.10.103 any
ip access-list extended acl_VoIP
  remark Identify VoIP traffic
  permit ip 192.168.80.0 0.0.0.255 any
ip access-list extended acl_video_VoD_high
  remark Identify high priority VoD traffic
  permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 5000 9000
  permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 5000 9000
  permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 5000 9000
ip access-list extended acl_video_VoD_low
  remark Identify low priority VoD traffic
  permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 1000 4999
  permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 1000 4999
  permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 1000 4999
ip access-list extended acl_video_broadcast
  remark Identify broadcast video traffic (multicast)
  permit ip 192.168.70.0 0.0.0.255 232.0.0.0 0.255.255.255
```

Step 3 Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_video_VoD_high
  match access-group name acl_video_VoD_high
class-map match-all class_video_VoD_low
  match access-group name acl_video_VoD_low
class-map match-all class_video_broadcast
  match access-group name acl_video_broadcast
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
```

Step 4 Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
```



```

class class_HSD
  set dscp default
class class_VoD_signaling
  set dscp cs3
class class_video_broadcast
  set dscp af41
class class_video_VoD_high
  set dscp af42
class class_video_VoD_low
  set dscp af43

```

Step 5 Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note

Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

Step 6 To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
mls qos trust dscp
```

Configuring Mapping on DER

Do the following to configure mapping on DER.

Step 1 View the Cisco 7600 and Cisco Catalyst 6500 default DSCP-to-CoS mapping for the different services. Use the **show mls qos maps dscp-cos** command.



Note

At the beginning of this section, we mentioned that there are 64 possible DSCP values and only 8 CoS values. This means that there could be more than one DSCP value for one CoS value. The following command shows the default DSCP-to-CoS mapping on the Cisco 7600 and Catalyst 6500.



Note

In the map, d1 corresponds to the y-axis value of the table, and d2 to the x-axis value.

```

DER# show mls qos maps dscp-cos

Dscp-cos map:                               (dscp= d1d2)
d1 : d2 0 1 2 3 4 5 6 7 8 9
-----
0 : 00 00 00 00 00 00 00 00 01 01
1 : 01 01 01 01 01 01 02 02 02 02
2 : 02 02 02 02 03 03 03 03 03 03
3 : 03 03 04 04 04 04 04 04 04 04
4 : 05 05 05 05 05 05 05 05 06 06
5 : 06 06 06 06 06 06 07 07 07 07
6 : 07 07 07 07

```

This table shows the following mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	34	4
VoD high priority	36	4
VoD OOB	24	3
Broadcast video	38	4
VoIP	46	5

- Step 2** Change the Cisco 7600 and Cisco Catalyst 6500 DSCP-to-CoS mapping for the different services to match the specifications of the solution.

The solution specifies the following DSCP-to CoS-mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	38	1
VoD high priority	36	2
VoD OOB	24	3
Broadcast video	34	4
VoIP	46	5

- a. Execute the following command on the Cisco 7600 and Cisco Catalyst 6500 to modify the DSCP-to-CoS mapping.

```
mls qos map dscp-cos 36 to 2
mls qos map dscp-cos 38 to 1
```

- b. Verify the changes to the DSCP-to-CoS mappings.

```
DER# show mls qos maps dscp-cos
DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 02 04 01 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07
```

Establishing and Configuring Interfaces on DER

Refer to [Figure 4-1 on page 4-4](#).

This section addresses the following:

- [Establishing VLANs for Services on DER](#)
- [Establishing 1-GE Interfaces for Servers, HSD, and Management on DER](#)
- [Establishing 10-GE Interfaces for Transport on DER](#)

Establishing VLANs for Services on DER

Before 1-GE interfaces and 10-GE trunks can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-2 on page 4-5](#).)

The following is configured on DER.



Tip

For convenience in establishing these VLANs and others, you can establish all VLANs in global configuration mode first, then configure all the interfaces in interface configuration mode.

Step 1 Establish VLANs and VLAN interfaces for management (including VoD signaling, connectivity with DHCP, DNS, FTP, TFTP, and Syslog servers).

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 10
name VLAN_10_Management
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan10
description Management VLAN (VoD signaling, DNS, DHCP, etc)
ip address 192.168.10.1 255.255.255.0
no ip redirects
no ip unreachablees
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```

- d. Repeat Step 1a through Step 1c, as appropriate, for the remaining management and video aggregation VLANs and interfaces. The abbreviated configurations are shown below.

Backup DNS server

```
vlan 11
name VLAN_11_Management
```

```
interface Vlan11
description Management VLAN (Backup DNS)
ip address 192.168.11.1 255.255.255.0
no ip redirects
no ip unreachablees
load-interval 30
```

Unicast video aggregation

```

vlan 60
name VLAN_60_Unicast_Video

interface Vlan60
description VoD server VLAN (Unicast Video)
ip address 192.168.60.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30

```

VoIP

```

vlan 80
name VLAN_80_VoIP

interface Vlan80
description VoIP gateway VLAN
ip address 192.168.80.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30

```

Step 2 Establish a VLAN for multicast video aggregation.

- a.** In global configuration mode, add the VLAN to the VLAN database.

```

vlan 70
name VLAN_70_Multicast_Video

```

- b.** In interface configuration mode, create and configure the VLAN interface.

```

interface Vlan70
description Broadcast video source VLAN (Multicast Video)
ip address 192.168.70.1 255.255.255.0
no ip redirects
no ip unreachable

```

- c.** Enable PIM sparse mode. This is the ingress port for broadcast video traffic, which is multicast addressed.

```

ip pim sparse-mode

```

- d.** Change the load interval from the default of 300.

```

load-interval 30

```

Step 3 In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```

vlan 90
name VLAN_90_HSD

```

Step 4 Establish VLANs for VoIP transport. The first is for transport to and from AR1.

- a.** In global configuration mode, add the VLAN to the VLAN database.

```

vlan 800
name VLAN_800_VoIP_to/from_AR1

```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan800
  description VoIP transport to/from AR1
  ip address 192.168.252.1 255.255.255.252
```

- c. Configure Open Shortest Path First (OSPF) on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```



Note To avoid the election of the designated router (DR) and backup designated router (BDR), and prevent the origination of an unnecessary network link state advertisement (LSA), configure the transport VLAN as a point-to-point network. In addition, reduce the interval between OSPF hello messages from 10 seconds to 1 second. This improves reconvergence in the event of failure in the transport or in a neighboring switch.

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Repeat Step 4a through Step 4d for VoIP transport to and from AR3.

```
vlan 824
name VoIP transport to/from AR3

interface Vlan824
  description VoIP transport to/from AR3
  ip address 192.168.252.25 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
```

Step 5 Establish VLANs for video transport. The first is for transport to and from AR1.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 900
name VLAN_900_Video_to/from_AR1
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan900
  description Video transport VLAN to/from AR1
  ip address 192.168.254.1 255.255.255.252
```

- c. Enable PIM sparse mode. This is the ingress port for broadcast video traffic, which is multicast addressed.

```
ip pim sparse-mode
```

- d. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

- f. Repeat Step 5a through Step 5e to establish a VLAN for video transport to and from AR3.

```
vlan 924
name VLAN_924_Video_to/from_AR3

interface Vlan924
description Video transport VLAN to/from AR3
ip address 192.168.254.25 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Establishing 1-GE Interfaces for Servers, HSD, and Management on DER

VoD servers, high-speed data sources, and management resources connect to Layer 2 interfaces on DER, and their traffic is aggregated into the appropriate service VLANs.

The following is configured on DER.

Step 1 Establish an interface.

- a. Establish an interface for high-speed data.

```
interface GigabitEthernet1/1
description High speed data ingress/egress port
no ip address
```

- b. Configure the interface as a Layer 2 access port and assign it to VLAN 90.

```
switchport
switchport access vlan 90
switchport mode access
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```

- d. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- e. Enable PortFast on the interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast
```

- f. Configure the switch to disable any interface that is configured for PortFast and receives a Bridge Protocol Data Unit (BPDU).

```
spanning-tree bpduguard enable
```



Note This guards against a user accidentally connecting a switch to a switchport that is intended for a VoD server or other host. The switchport is disabled and the user must investigate why the port is down. If this command is not used and such an accidental connection were to happen, STP could reconverge and block other connections in the switch.

- g. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

- Step 2** Repeat Step 1a through Step 1g for the remaining server, HSD, and management 1-GE interfaces and their associated VLANs, changing the value in **switchport access vlan xxx** as appropriate. Those configurations are shown abbreviated below.

VoIP traffic

```
interface GigabitEthernet2/1
description VoIP traffic ingress/egress

switchport access vlan 80
```

Ingress multicast broadcast video

```
interface GigabitEthernet2/9
description Broadcast video source (multicast 232.1.1.1 - 232.1.1.10)

switchport access vlan 70
```

Management for the Kasenna LR server

```
interface GigabitEthernet2/17
description Management port from Kasenna LR Server (Eth0)

switchport access vlan 10
```

Management for the Kasenna VoD pump

```
interface GigabitEthernet2/18
description Kasenna VoD Pump Management

switchport access vlan 10
```

Ingress unicast video from the Kasenna VoD pump (1)

```
interface GigabitEthernet2/19
description Unicast video from Kasenna VoD Pump (HPN0)

switchport access vlan 60
```



Note In Kasenna’s terminology, HPN0 stands for High-Performance Network interface 0.

Ingress unicast video from the Kasenna VoD pump (2)

```
interface GigabitEthernet2/20
description Unicast video from Kasenna VoD Pump (HPN1)

switchport access vlan 60
```

Backup DNS server

```
interface GigabitEthernet2/48
description Backup DNS server

switchport access vlan 11
```

Primary DNS, DHCP, NTP, TFTP, and Syslog servers

```
interface GigabitEthernet5/2
  description Primary DNS/DHCP/NTP/TFTP/Syslog servers

  switchport access vlan 10
```



Note In this case, specify the physical connection on a Gigabit Ethernet interface as RJ-45.

```
media-type rj45
```

Establishing 10-GE Interfaces for Transport on DER

The 10-GE trunk interfaces create the ring topology from the DER through the ARs and back to the DER. The following is configured on DER.

Step 1 Establish an interface to and from AR1.

a. Establish the interface.

```
interface TenGigabitEthernet7/1
  description Transport to/from AR1 (TenGig1/1)
  switchport

  switchport mode trunk
  dampening
  no ip address

  carrier-delay msec 0
```

b. Configure the trunk for 802.1q encapsulation.

```
switchport trunk encapsulation dot1q
```

c. Assign the trunk to VLANs 90, 800, and 900. (See [Table 4-2 on page 4-5](#).)

```
switchport trunk allowed vlan 90,800,900
```

d. Change the load interval from the default of 300.

```
load-interval 30
```

Step 2 Configure QoS on the interface.



Note The 10-GE transport links from the DER to the ARs require modifications to the transmit queues. There are eight transmit queues, but this solution uses only three.

- a. View the default CoS to Tx-Queue mapping. The following information was extracted from the **show queueing interface** command.

TxQueue	CoS
1	0, 1
2	2, 3, 4
3	6, 7
4	—
5	—
6	—
7	—
8	5

- b. Configure the CoS-to TxQueue mapping on the 10-GE transport interfaces. HSD (CoS = 0) remains in TxQueue1 and VoIP (CoS = 5) remains in TxQueue8. The other six CoS values are associated with TxQueue2.

```
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2
wrr-queue cos-map 2 3 3 4 6 7
```



Note TxQueue1 and TxQueue8 use the default mappings. TxQueue2 has three thresholds: Threshold 1 = CoS 1, Threshold 2 = CoS 2, and Threshold 3 = CoS 3, 4, 6, and 7. For details, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

- c. Verify the modified CoS-to-Tx-Queue mapping. The following information was extracted from the **show queueing interface** command.

TxQueue	CoS
1	0
2	1, 2, 3, 4, 6, 7
3	6, 7
4	—
5	—
6	—
7	—
8	5

- d. Configure the TxQueue thresholds.

TxQueue1 uses Weighted Random Early Drop (WRED) for queue-congestion management. Only HSD is queued in this queue, and when the amount of HSD in the queue reaches 75%, random packets are dropped in an attempt to keep the queue from reaching 100% utilization.

```
wrr-queue threshold 1 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100
```

TxQueue2 uses tail drop for queue congestion management. Low-priority VoD is assigned to the first threshold and is dropped once the queue reaches 45% utilization. High-priority VoD is assigned to the second threshold and is dropped once the queue reaches 85% utilization. VoD signaling, network signaling, and broadcast video are assigned to the third threshold and are dropped once the queue reaches 100% utilization.

```
wrr-queue threshold 2 45 85 100 100 100 100 100 100
no wrr-queue random-detect 2
```

- e. Configure the bandwidth of the weighted queues.

The weighted queues need to be modified to handle our modified TxQueue mappings. The ratio between TxQueue2 and TxQueue1 is $255/64 = 4$, so TxQueue2 needs four times as much bandwidth as TxQueue1. Therefore, TxQueue1 is allocated 20% of the bandwidth on the interface, and TxQueue2 is allocated 80% of the bandwidth.

```
wrr-queue bandwidth 64 255 0 0 0 0 0
```

- f. Configure the size of the weighted queues.

Each line card has a limited amount of buffer for the transmit queues. For this interface, 40% of the buffer is allocated for TxQueue1, and 50% of the buffer is allocated for TxQueue2.

```
wrr-queue queue-limit 40 50 0 0 0 0 0
```

- g. Configure this interface (and all noningress transport interfaces) to trust the incoming DSCP markings. (This maintains the DSCP marking applied at the network ingress interface.)

```
mls qos trust dscp
```

Step 3 Establish an interface to and from AR3.

- a. Establish the interface, configure the trunk for 802.1q encapsulation, and assign it to VLANs 90, 800, and 900. (See [Table 4-2 on page 4-5](#).)

```
interface TenGigabitEthernet7/1
description Transport to/from AR1 (TenGig1/1)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,800,900
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
```

- b. Proceed as in Step 1b through Step 2 of this task.
-

Configuring OSPF Routing for Video and Voice Traffic on DER

Two OSPF routing processes must be established:

- One to route the video traffic over the transport VLANs for video
- One to route VoIP traffic over the transport VLANs for VoIP

The first OSPF process (100) associates the management VLANs, the VoD VLAN, and the broadcast VLAN with the two transport VLANs that carry video. The second OSPF process (101) associates the VoIP VLAN with the two transport VLANs that carry VoIP. Routing advertisements are enabled on the transport VLANs, but are turned off on the aggregation VLANs by means of the **passive-interface** command.

The following is configured on DER.

Step 1 Define an OSPF routing process to route video traffic.

```
router ospf 100
router-id 1.1.1.1
log-adjacency-changes
```

- a. The OSPF timers are modified to provide fast convergence. The following command enables OSPF SPF throttling: **timers throttle spf** *spf-start spf-hold spf-max-wait*

```
timers throttle spf 10 100 1000
```

- b. The following command sets the rate-limiting values for OSPF link-state advertisement (LSA) generation: **timers throttle lsa all** *start-interval hold-interval max-interval*

```
timers throttle lsa all 1 10 1000
```

- c. The following command controls the minimum interval for accepting the same LSA: **timers lsa arrival** *milliseconds*

```
timers lsa arrival 100
```

If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped.

- d. Apply the **passive-interface** statements to the aggregation VLANs.

```
passive-interface Vlan10
passive-interface Vlan11
passive-interface Vlan60
passive-interface Vlan70
```

- e. Advertise the networks in the first OSPF routing process.

```
network 192.168.10.0 0.0.1.255 area 0
network 192.168.60.0 0.0.0.255 area 0
network 192.168.70.0 0.0.0.255 area 0
network 192.168.254.1 0.0.0.0 area 0
network 192.168.254.25 0.0.0.0 area 0
```

Step 2 Define a second OSPF process to route voice traffic. For details, refer to Step 1.

```
router ospf 101
router-id 1.1.1.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan80
```

```
network 192.168.80.0 0.0.0.255 area 0
network 192.168.252.1 0.0.0.0 area 0
network 192.168.252.25 0.0.0.0 area 0
```

Configuring Spanning Tree on DER

Because VLAN 90 is at Layer 2 around the 1-GE ring, Spanning Tree Protocol (STP) is needed to guard against loops. To improve convergence time, the four switches are configured for IEEE 802.1w Rapid Spanning Tree Protocol (RTSP), with the root at DER.

Do the following in global configuration mode to configure spanning tree parameters on DER.

Step 1 Configure DER as the root node of the spanning tree for VLAN 90. There are two ways to do this.

a. Use the **root primary** option.

```
spanning-tree vlan 90 root primary
```

or

b. Set the priority to 24576.

```
spanning-tree vlan 90 priority 24576
```

Step 2 Configure RTSP.

```
spanning-tree mode rapid-pvst
```

Step 3 Because the transport VLANs in the 10-GE ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 800, 808, 900, 908
```

Configuring AR1

This section addresses the configuration required on the switch labeled AR1 in [Figure 4-1 on page 4-4](#), to route multiple services from AR1 to DER and AR2.

See [Configuring DNS Servers, page 4-2](#).

**Note**

A Cisco Catalyst 6509 can also be used, as it uses the same supervisor engine, line cards, and Cisco IOS code as the Cisco 7609 router.

This section addresses the following:

- [Configuring QoS on AR1](#)
- [Establishing and Configuring Interfaces on AR1](#)
- [Configuring OSPF Routing for Video and Voice Traffic on AR1](#)
- [Configuring Spanning Tree on AR1](#)

**Note**

For a complete configuration example, see [Appendix A, “Sample DER and AR Switch Configurations for the 10-GE Symmetric Topology.”](#)

Configuring QoS on AR1

See [Overview of QoS on a Cisco 7600 Series and Cisco Catalyst 6500 Series, page 4-7](#).

This section presents the following topics:

- [Configuring Marking and Classification on AR1](#)
- [Configuring Mapping on AR1](#)

**Note**

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Configuring Marking and Classification on AR1

Do the following to enable marking and classification on AR1.

Step 1 Enable QoS in global configuration mode.

```
mls qos
```

Step 2 Create access lists to identify the different service types in the network.

```
ip access-list extended acl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip 192.168.110.0 0.0.0.255 192.168.10.102
 permit ip 192.168.110.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.111.0 0.0.0.255 any
```

Step 3 Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
```

Step 4 Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
```

Step 5 Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note

Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

Step 6 To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
mls qos trust dscp
```

Configuring Mapping on AR1

Do the following to configure mapping on AR1.

Step 1 View the Cisco 7600 and Cisco Catalyst 6500 default DSCP-to-CoS mapping for the different services. Use the **show mls qos maps dscp-cos** command.



Note

At the beginning of this section, we mentioned that there are 64 possible DSCP values and only 8 CoS values. This means that there could be more than one DSCP value for one CoS value. The following command shows the default DSCP-to-CoS mapping on the Cisco 7600 and Cisco Catalyst 6500.

**Note**

In the map, d1 corresponds to the y-axis value of the table, and d2 to the x-axis value.

```
AR1# show mls qos maps dscp-cos
Dscp-cos map: (dscp= d1d2)
d1 : d2 0 1 2 3 4 5 6 7 8 9
-----
0 : 00 00 00 00 00 00 00 00 01 01
1 : 01 01 01 01 01 01 02 02 02 02
2 : 02 02 02 02 03 03 03 03 03 03
3 : 03 03 04 04 04 04 04 04 04 04
4 : 05 05 05 05 05 05 05 05 06 06
5 : 06 06 06 06 06 06 07 07 07 07
6 : 07 07 07 07
```

This table shows the following mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	34	4
VoD high priority	36	4
VoD OOB	24	3
Broadcast video	38	4
VoIP	46	5

- Step 2** Change the Cisco 7600 and Cisco Catalyst 6500 DSCP-to-CoS mapping for the different services to match the specifications of the solution.

The solution specifies the following DSCP-to CoS-mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	38	1
VoD high priority	36	2
VoD OOB	24	3
Broadcast video	34	4
VoIP	46	5

- a. Execute the following command on the Cisco 7600 and Cisco Catalyst 6500 to modify the DSCP-to-CoS mapping.

```
mls qos map dscp-cos 36 to 2
mls qos map dscp-cos 38 to 1
```

- b. Verify the changes to the DSCP-to-CoS mappings.

```
AR1# show mls qos maps dscp-cos
DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0  :    00 00 00 00 00 00 00 00 01 01
1  :    01 01 01 01 01 01 02 02 02 02
2  :    02 02 02 02 03 03 03 03 03 03
3  :    03 03 04 04 04 04 02 04 01 04
4  :    05 05 05 05 05 05 05 05 06 06
5  :    06 06 06 06 06 06 07 07 07 07
6  :    07 07 07 07
```

Establishing and Configuring Interfaces on AR1

Refer to [Figure 4-1 on page 4-4](#).

This section addresses the following:

- [Establishing VLANs for Services on AR1](#)
- [Establishing 10-GE Interfaces for Transport on AR1](#)
- [Establishing 1-GE Interfaces to a DSLAM on AR1](#)

Establishing VLANs for Services on AR1

Before 1-GE interfaces and 10-GE trunks can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-2 on page 4-5](#).)



Note

For additional details, see [Establishing VLANs for Services on DER, page 4-11](#).

The following is configured on AR1.

- Step 1** In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```
vlan 90
name VLAN_90_HSD
```

- Step 2** Establish a VLAN for video at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 110
name VLAN_110_Video
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan110
description Video edge VLAN
ip address 192.168.110.1 255.255.255.0
no ip redirects
no ip unreachable
```


- c. Enable PIM sparse mode. This is the aggregation VLAN for video traffic to the DSLAMs. Broadcast video is multicast addressed.

```
ip pim sparse-mode
```

- d. To ensure consistently fast PIM convergence times, statically join the aggregation VLAN for video at the AR to the multicast groups.

```
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

- f. Change the ARP timeout from the default.

```
arp timeout 250
```



Note

The default timeout for an entry in the ARP cache is 4 hours. The default timeout for an entry in the MAC address table is only 5 minutes. Because video traffic is mostly unidirectional, the MAC address table may not be refreshed within the 5-minute timeout. This causes video traffic to be flooded until the destination MAC address is found. To prevent this, reduce the ARP cache timeout to 250 seconds. This forces the switch to re-ARP for the entries in the ARP cache before the entries in the MAC address table time out, avoiding the disruptive behavior.

Step 3 Establish a VLAN for VoIP at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 111
name VLAN_111_VoIP
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan111
description VoIP edge VLAN
ip address 192.168.111.1 255.255.255.0
no ip redirects
no ip unreachablees
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```

Step 4 Establish VLANs for VoIP transport. The first is to and from DER.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 800
name VLAN_800_VoIP_to/from_DER
```

- b.** In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan800
  description VoIP transport VLAN to/from DER
  ip address 192.168.252.2 255.255.255.252
```

- c.** Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- d.** Change the load interval from the default of 300.

```
load-interval 30
```

- e.** Repeat Step 4a through Step 4d, as appropriate, to establish a VLAN for VoIP transport to and from AR2.

```
vlan 808
name VLAN_808_VoIP_to/from_DER
```

```
interface Vlan808
  description VoIP transport VLAN to/from AR2
  ip address 192.168.252.9 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
```

- Step 5** Establish VLANs for video transport. The first is to and from DER.

- a.** In global configuration mode, add the VLAN to the VLAN database.

```
vlan 900
name VLAN_900_Video_to/from_DER
```

- b.** In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan900
  description Video transport VLAN to/from DER
  ip address 192.168.254.2 255.255.255.252
```

- c.** Enable Protocol Independent Multicast (PIM) sparse mode. Broadcast video is multicast addressed.

```
ip pim sparse-mode
```

- d.** Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- e.** Change the load interval from the default of 300.

```
load-interval 30
```

- f. Repeat Step 5a through Step 5e, as appropriate, to establish a VLAN for video transport to and from AR2.

```
vlan 908
name VLAN_908_Video_to/from_AR2

interface Vlan908
description Video transport VLAN to/from AR2
ip address 192.168.254.9 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Establishing 10-GE Interfaces for Transport on AR1

The 10-GE trunk interfaces provide the transport between AR1 and DER and AR2.



Note

For additional details, see [Establishing 10-GE Interfaces for Transport on DER, page 4-16](#).

The following is configured on AR1.

Step 1 Establish an interface. The first is to and from DER.

- a. Establish the interface to and from DER, configure the trunk for 802.1q encapsulation, and assign it to VLANs 90, 800, and 900. (See [Table 4-2 on page 4-5](#).)

```
interface TenGigabitEthernet1/1
description Transport to/from DER (TenGig7/1)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,800,900
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
```

- b. Proceed as in Step 2 of [Establishing 10-GE Interfaces for Transport on DER, page 4-16](#).

```
wrr-queue bandwidth 64 255 0 0 0 0 0
wrr-queue queue-limit 40 50 0 0 0 0 0
wrr-queue threshold 1 100 100 100 100 100 100 100 100
wrr-queue threshold 2 85 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
no wrr-queue random-detect 2
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2
wrr-queue cos-map 2 3 3 4 6 7
mls qos trust dscp
```

- Step 2** Repeat Step 1, as appropriate, to establish an interface to and from AR3 and assign it to VLANs 90, 824, and 924.

```
interface TenGigabitEthernet1/3
description Transport to/from AR2 (TenGig1/1)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,808,908
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
wrr-queue bandwidth 64 255 0 0 0 0 0
wrr-queue queue-limit 40 50 0 0 0 0 0
wrr-queue threshold 1 100 100 100 100 100 100 100 100
wrr-queue threshold 2 45 85 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100
no wrr-queue random-detect 2
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2
wrr-queue cos-map 2 3 3 4 6 7
mls qos trust dscp
```

Establishing 1-GE Interfaces to a DSLAM on AR1

The only 1-GE interface is to and from DSLAM1.

The following is configured on AR1.

- Step 1** Establish an interface to DSLAM1.

- a. Establish the interface and assign it to VLANs 90, 110, and 111.

```
interface GigabitEthernet2/1
description GigE trunk to/from DSLAM uplink GigE
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,110,111
switchport mode trunk
```

```
no ip address
```

- b. Prevent unknown unicast traffic from being flooded to the port.

```
switchport block unicast
```



Note Occasionally, unknown unicast or multicast traffic is flooded to a switch port because a MAC address has timed out or has not been learned by the switch. (This condition is especially undesirable for a private VLAN isolated port.) To guarantee that no unicast or multicast traffic is flooded to the port, use the **switchport block unicast** or **switchport block multicast** commands.

- c. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Proceed as in Step 2 of [Establishing 10-GE Interfaces for Transport on DER, page 4-16](#).

```
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
```



Note The cos-map value 1 1 0 is a default setting on 1-GE interfaces.

- f. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- g. Enable PortFast on the interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast
```

- h. Configure the switch to disable any interface that is configured for PortFast and receives a Bridge Protocol Data Unit (BPDU).

```
spanning-tree bpduguard enable
```



Note This guards against a user accidentally connecting a switch to a switchport that is intended for a VoD server or other host. The switchport is disabled and the user must investigate why the port is down. If this command is not used and such an accidental connection were to happen, STP could reconverge and block other connections in the switch.

- i. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

Configuring OSPF Routing for Video and Voice Traffic on AR1

For background and details, refer to [Configuring OSPF Routing for Video and Voice Traffic on DER, page 4-19](#).

The following is configured on AR1.

- Step 1** Define an OSPF routing process to route video traffic. This process associates the transport VLANs for video with the video aggregation VLAN for the for DSLAM1 and other DSLAMs to be served by AR1 (VLAN 110).

```
router ospf 100
  router-id 2.2.2.1
  log-adjacency-changes
  timers throttle spf 10 100 1000
  timers throttle lsa all 1 10 1000
  timers lsa arrival 100
  passive-interface Vlan110
  network 192.168.110.0 0.0.0.255 area 0
  network 192.168.254.2 0.0.0.0 area 0
  network 192.168.254.9 0.0.0.0 area 0
```

- Step 2** Define an OSPF process to route voice traffic. This process associates the transport VLANs for VoIP with the VoIP aggregation VLAN for the for DSLAM1 and other DSLAMs to be served by AR1 (VLAN 111).

```
router ospf 101
  router-id 2.2.2.2
  log-adjacency-changes
  timers throttle spf 10 100 1000
  timers throttle lsa all 1 10 1000
  timers lsa arrival 100
  passive-interface Vlan111
  network 192.168.111.0 0.0.0.255 area 0
  network 192.168.252.2 0.0.0.0 area 0
  network 192.168.252.9 0.0.0.0 area 0
```

Configuring Spanning Tree on AR1



Note See [Configuring Spanning Tree on DER, page 4-20](#).

The following is configured on AR1.

- Step 1** Configure RTSP.

```
spanning-tree mode rapid-pvst
```

- Step 2** Because the transport VLANs in the 10-GE ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 800, 808, 900, 908
```

Configuring AR2

This section addresses the configuration required on the switch labeled AR2 in [Figure 4-1 on page 4-4](#), to route multiple services from AR2 to DER, AR1, and AR3.

See [Configuring DNS Servers, page 4-2](#).

This section addresses the following:

- [Configuring QoS on AR2](#)
- [Establishing and Configuring Interfaces on AR2](#)
- [Configuring OSPF Routing for Video and Voice Traffic on AR2](#)
- [Configuring Spanning Tree on AR2](#)

**Note**

For a complete configuration example, see [Appendix A, “Sample DER and AR Switch Configurations for the 10-GE Symmetric Topology.”](#)

Configuring QoS on AR2

This section presents the following topics:

- [Overview of QoS on a Cisco Catalyst 4500 Series](#)
- [Configuring Marking and Classification on AR2](#)
- [Configuring Mapping on AR2](#)
- [Configuring Queueing on AR2](#)

**Note**

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Overview of QoS on a Cisco Catalyst 4500 Series

This section addresses the configuration of quality of service (QoS) on AR2, through marking, classification, mapping, and queueing, to provide different degrees of quality of service for the different types of services supported by the solution architecture. For example, it is important to ensure the expeditious delivery of video and VoIP traffic, while providing only best-effort delivery for high-speed data (HSD).

By default, the Cisco Catalyst 4500 series switches (including the Cisco Catalyst 4948-10GE) do not trust the incoming QoS markings, and therefore rewrite these bits with zeros. In this solution, packets at the network ingress ports are identified, classified, and marked according to type of traffic. The packets are marked with one of 64 possible Differentiated Services Code Point (DSCP) values at the ingress ports. These in turn are internally mapped to one of eight possible Class of Service (CoS) values. The DSCP values are used to determine the appropriate transmit queue for each packet.

Configuring Marking and Classification on AR2

Do the following to enable marking and classification on AR2.

- Step 1** Enable QoS in global configuration mode.

```
qos
```

- Step 2** Create access lists to identify the different service types in the network.

```
ip access-list extended acl_HSD
  remark Identify HSD traffic
  permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
  remark Identify VoD signaling traffic
  permit ip 192.168.120.0 0.0.0.255 192.168.10.102
  permit ip 192.168.120.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
  remark Identify VoIP traffic
  permit ip 192.168.121.0 0.0.0.255 any
```

- Step 3** Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
```

- Step 4** Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
```

- Step 5** Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

- Step 6** To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
qos trust dscp
```

Configuring Mapping on AR2

Do the following to configure mapping on AR2.

- Step 1** View the Cisco Catalyst 4500 series default DSCP-to-CoS mapping for the different services. Use the **show qos maps dscp-cos** command.

**Note**

At the beginning of this section, we mentioned that there are 64 possible DSCP values and only 8 CoS values. This means that there could be more than one DSCP value for one CoS value. The following command shows the default DSCP-to-CoS mapping on the Cisco Catalyst 4500 series.

**Note**

In the map, d1 corresponds to the *y*-axis value of the table, and d2 to the *x*-axis value.

```
AR2# show qos maps dscp
```

```
DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 04 04 04 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07
```

This table shows the following mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	34	4
VoD high priority	36	4
VoD OOB	24	3
Broadcast video	38	4
VoIP	46	5

Step 2 Change the Cisco Catalyst 4500 series DSCP-to-CoS mapping for the different services to match the specifications of the solution.

The solution specifies the following DSCP-to CoS-mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	38	1
VoD high priority	36	2
VoD OOB	24	3
Broadcast video	34	4
VoIP	46	5

- a. Execute the following command on the Cisco Catalyst 4500 series to modify the DSCP-to-CoS mapping.

```
qos map dscp 38 to cos 1
qos map dscp 36 to cos 2
```

- b. Verify the changes to the DSCP-to-CoS mappings.

```
AR2# show qos maps dscp

DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2 0 1 2 3 4 5 6 7 8 9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 02 04 01 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07
```

Configuring Queueing on AR2

Unlike the Cisco 7600 series and Cisco Catalyst 6500 series, the Cisco Catalyst 4500 series uses the same queueing on all interfaces. Queueing is configured globally.

Do the following to change the DSCP-to-TxQueue mappings on AR2.

- Step 1** View the default DSCP-to-Tx-Queue mapping. The following information was extracted from the **show qos maps dscp** command.

```
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
```

- Step 2** Configure the DSCP-to TxQueue mapping by moving DSCP 34 and 36 to TxQueue2. Additionally, move all DSCPs that are in TxQueue4 to TxQueue2, because TxQueue4 is not used.

```
qos map dscp 34 36 38 48 49 50 51 52 to tx-queue 2
qos map dscp 53 54 55 56 57 58 59 60 to tx-queue 2
qos map dscp 61 62 63 to tx-queue 2
```

- Step 3** Verify the modified DSCP-to-Tx-Queue mapping. The following information was extracted from the `show queueing interface` command.

```
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 02 03 02 03 02 03
4  :   03 03 03 03 03 03 03 03 02 02
5  :   02 02 02 02 02 02 02 02 02 02
6  :   02 02 02 02
```

- Step 4** Configure the TxQueue thresholds.

TxQueue1 uses Weighted Random Early Drop (WRED) for queue-congestion management. Only HSD is queued in this queue, and when the amount of HSD in the queue reaches 75%, random packets are dropped in an attempt to keep the queue from reaching 100% utilization.

```
wrr-queue threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100
```

TxQueue2 uses tail drop for queue congestion management. Low-priority VoD is assigned to the first threshold and is dropped once the queue reaches 45% utilization. High-priority VoD is assigned to the second threshold and is dropped once the queue reaches 85% utilization. VoD signaling, network signaling, and broadcast video are assigned to the third threshold and is dropped once the queue reaches 100% utilization.

```
wrr-queue threshold 2 45 85 100 100 100 100 100
no wrr-queue random-detect 2
```

- Step 5** Configure the bandwidth of the weighted queues.

The weighted queues need to be modified to handle our modified TxQueue mappings. The ratio between TxQueue2 and TxQueue1 is $255/64 = 4$, so TxQueue2 needs four times as much bandwidth as TxQueue1. Therefore, TxQueue1 is allocated 20% of the bandwidth on the interface, and TxQueue2 is allocated 80% of the bandwidth.

```
wrr-queue bandwidth 64 255 0 0 0 0 0
```

- Step 6** Configure the size of the weighted queues.

Each line card has a limited amount of buffer for the transmit queues. For this interface, 40% of the buffer is allocated for TxQueue1, and 50% of the buffer is allocated for TxQueue2.

```
wrr-queue queue-limit 40 50 0 0 0 0 0
```

Establishing and Configuring Interfaces on AR2

Refer to [Figure 4-1 on page 4-4](#).

This section addresses the following:

- [Establishing VLANs for Services on AR2](#)
- [Establishing 10-GE Interfaces for Transport on AR2](#)
- [Establishing 1-GE Interfaces to a DSLAM on AR2](#)

Establishing VLANs for Services on AR2

Before 1-GE interfaces and 10-GE trunks can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-2 on page 4-5](#).)



Note For additional details, see [Establishing VLANs for Services on DER, page 4-11](#), and [Establishing VLANs for Services on AR1, page 4-24](#).

The following is configured on AR2.

Step 1 In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```
vlan 90
name VLAN_90_HSD
```

Step 2 Establish a VLAN for video at the edge.

a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 120
name VLAN_120_Video
```

b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan120
description Video edge VLAN
ip address 192.168.120.1 255.255.255.0
no ip redirects
no ip unreachable
```

c. Enable PIM sparse mode.

```
ip pim sparse-mode
```

d. To ensure consistently fast PIM convergence times, statically join the aggregation VLAN for video at the AR to the multicast groups.

```
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
```

e. Change the load interval from the default of 300.

```
load-interval 30
```

f. Change the ARP timeout from the default.

```
arp timeout 250
```

Step 3 Establish a VLAN for VoIP at the edge.

a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 121
name VLAN_121_VoIP
```

- b.** In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan121
description VoIP edge VLAN
ip address 192.168.121.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30
```

Step 4 Establish VLANs for VoIP transport. The first is to and from AR1.

- a.** In global configuration mode, add the VLAN to the VLAN database.

```
vlan 808
name VLAN_808_VoIP_to/from_AR1
```

- b.** In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan808
description VoIP transport VLAN to/from AR1
ip address 192.168.252.10 255.255.255.252
```

- c.** Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- d.** Change the load interval from the default of 300.

```
load-interval 30
```

- e.** Repeat Step 4a through Step 4d, as appropriate, to establish a VLAN for VoIP transport to and from AR3.

```
vlan 816
name VLAN_816_VoIP_to/from_AR3
```

```
interface Vlan816
description VoIP transport VLAN to/from AR3
ip address 192.168.252.17 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Step 5 Establish VLANs for video transport. The first is to and from AR1.

- a.** In global configuration mode, add the VLAN to the VLAN database.

```
vlan 908
name VLAN_908_Video_to/from_AR1
```

- b.** In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan908
description Video transport VLAN to/from AR1
ip address 192.168.254.10 255.255.255.252
```

- c.** Enable PIM sparse mode.

```
ip pim sparse-mode
```

d. Configure OSPF.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

e. Change the load interval from the default of 300.

```
load-interval 30
```

f. Repeat Step 5a through Step 5e, as appropriate, to establish a VLAN for video transport to and from AR3.

```
vlan 916
name VLAN_916_Video_to/from_AR3

interface Vlan916
description Video transport VLAN to/from AR3
ip address 192.168.254.17 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Establishing 10-GE Interfaces for Transport on AR2

The 10-GE trunk interfaces provide the transport between AR2 and AR1 and AR3.

The following is configured on AR2.


Note

For additional details, see [Establishing 10-GE Interfaces for Transport on DER, page 4-16](#).

Step 1 Establish an interface. The first is to and from AR1.

a. Establish the interface, configure the trunk for 802.1q encapsulation, and assign it to VLANs 90, 808, and 908. (See [Table 4-2 on page 4-5](#).)

```
interface TenGigabitEthernet1/1
description Transport to/from AR1 (TenGig1/3)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,808,908
switchport mode trunk
dampening
load-interval 30
carrier-delay msec 0
```

b. Configure this interface (and all noningress transport interfaces) to trust the incoming DSCP markings. (This maintains the DSCP marking applied at the network ingress interface.)

```
qos trust dscp
```

- c. Set transmit-queue bandwidth thresholds and priority.

```
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
```



Note See [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

- Step 2** Repeat Step 1, as appropriate, to establish an interface to and from AR3 and assign it to VLANs 90, 816, and 916.

```
interface TenGigabitEthernet1/2
  description Transport to/from AR3 (TenGig1/49)
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,816,916
  switchport mode trunk
  dampening
  load-interval 30
  carrier-delay msec 0
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
  spanning-tree cost 10 <---See Note below
```



Note Note that the spanning-tree cost is set to 10 on AR2. This breaks the loop for VLAN 90 (Layer 2) between AR2 and AR3, rather than somewhere else.

Establishing 1-GE Interfaces to a DSLAM on AR2

The only 1-GE interface is a trunk to and from DSLAM2.

The following is configured on AR2.

- Step 1** Establish an interface to DSLAM2.

- a. Establish the interface and assign it to VLANs 90, 120, and 121.

```
interface GigabitEthernet5/1
  description GigE trunk to/from DSLAM uplink GigE
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,120,121
  switchport mode trunk
```

- b. Prevent unknown unicast traffic from being flooded to the port.

```
switchport block unicast
```



Note Occasionally, unknown unicast or multicast traffic is flooded to a switch port because a MAC address has timed out or has not been learned by the switch. (This condition is especially undesirable for a private VLAN isolated port.) To guarantee that no unicast or multicast traffic is flooded to the port, use the **switchport block unicast** or **switchport block multicast** commands.

- c. Change the load interval from the default of 300.

```
load-interval 30
```

- d. Set transmit-queue bandwidth thresholds and priority.

```
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
```

- e. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- f. Enable PortFast on the interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast trunk
```

- g. Configure the switch to disable any interface that is configured for PortFast and receives a Bridge Protocol Data Unit (BPDU).

```
spanning-tree bpduguard enable
```



Note This guards against a user accidentally connecting a switch to a switchport that is intended for a VoD server or other host. The switchport is disabled and the user must investigate why the port is down. If this command is not used and such an accidental connection were to happen, STP could reconverge and block other connections in the switch.

- h. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

Configuring OSPF Routing for Video and Voice Traffic on AR2



Note For background and details, see [Configuring OSPF Routing for Video and Voice Traffic on DER, page 4-19](#).

The following is configured on AR2.

-
- Step 1** Define an OSPF routing process to route video traffic. This process associates the transport VLANs for video with the video aggregation VLAN for the DSLAM (VLAN 120).

```
router ospf 100
router-id 3.3.3.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan120
network 192.168.120.0 0.0.0.255 area 0
network 192.168.254.10 0.0.0.0 area 0
network 192.168.254.17 0.0.0.0 area 0
```

- Step 2** Define an OSPF process to route voice traffic. This process associates the transport VLANs for VoIP with the VoIP aggregation VLAN for the DSLAM (VLAN 121).

```
router ospf 101
router-id 3.3.3.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan121
network 192.168.121.0 0.0.0.255 area 0
network 192.168.252.10 0.0.0.0 area 0
network 192.168.252.17 0.0.0.0 area 0
```

Configuring Spanning Tree on AR2

**Note**

See [Configuring Spanning Tree on DER, page 4-20](#).

The following is configured on AR2.

- Step 1** Configure RTSP.

```
spanning-tree mode rapid-pvst
```

- Step 2** Because the transport VLANs in the 10-GE ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 800, 816, 908, 916
```

Configuring AR3

This section addresses the configuration required on the switch labeled AR3 in [Figure 4-1 on page 4-4](#), to route multiple services from AR3 to AR2 and DER.

See [Configuring DNS Servers, page 4-2](#).

This section addresses the following:

- [Configuring QoS on AR3](#)
- [Establishing and Configuring Interfaces on AR3](#)
- [Configuring OSPF Routing for Video and Voice Traffic on AR3](#)
- [Configuring Spanning Tree on AR3](#)



Note

For a complete configuration example, see [Appendix A, “Sample DER and AR Switch Configurations for the 10-GE Symmetric Topology.”](#)

Configuring QoS on AR3

See [Overview of QoS on a Cisco Catalyst 4500 Series, page 4-31](#).

This section presents the following topics:

- [Configuring Marking and Classification on AR3](#)
- [Configuring Mapping on AR3](#)
- [Configuring Queuing on AR3](#)



Note

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Configuring Marking and Classification on AR3

Do the following to enable marking and classification on AR3.

Step 1 Enable QoS in global configuration mode.

```
qos
```

Step 2 Create access lists to identify the different service types in the network.

```
ip access-list extended aCl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended aCl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip 192.168.130.0 0.0.0.255 192.168.10.102
 permit ip 192.168.130.0 0.0.0.255 192.168.10.103
ip access-list extended aCl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.131.0 0.0.0.255 any
```

Step 3 Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
```

Step 4 Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
```

Step 5 Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note

Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

Step 6 To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
qos trust dscp
```

Configuring Mapping on AR3

Do the following to configure mapping on AR3.

Step 1 View the Cisco Catalyst 4500 series default DSCP-to-CoS mapping for the different services. Use the **show qos maps dscp-cos** command.



Note

At the beginning of this section, we mentioned that there are 64 possible DSCP values and only 8 CoS values. This means that there could be more than one DSCP value for one CoS value. The following command shows the default DSCP-to-CoS mapping on the Cisco Catalyst 4500 series.

**Note**

In the map, d1 corresponds to the y-axis value of the table, and d2 to the x-axis value.

```
AR3# show qos maps dscp
```

```
DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 04 04 04 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07
```

This table shows the following mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	34	4
VoD high priority	36	4
VoD OOB	24	3
Broadcast video	38	4
VoIP	46	5

- Step 2** Change the Cisco Catalyst 4500 series DSCP-to-CoS mapping for the different services to match the specifications of the solution.

The solution specifies the following DSCP-to CoS-mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	38	1
VoD high priority	36	2
VoD OOB	24	3
Broadcast video	34	4
VoIP	46	5

- a. Execute the following command on the Cisco Catalyst 4500 series to modify the DSCP-to-CoS mapping.

```
qos map dscp 38 to cos 1
qos map dscp 36 to cos 2
```

- b. Verify the changes to the DSCP-to-CoS mappings.

```
AR3# show qos maps dscp

DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 02 04 01 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07
```

Configuring Queueing on AR3

Unlike the Cisco 7600 series and Cisco Catalyst 6500 series, the Cisco Catalyst 4500 series uses the same queueing on all interfaces. Queueing is configured globally.

Do the following to change the DSCP-to-TxQueue mappings on AR3.

- Step 1** View the default DSCP-to-Tx-Queue mapping. The following information was extracted from the **show qos maps dscp** command.

```
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
```

- Step 2** Configure the DSCP-to TxQueue mapping by moving DSCP 34, 36, and 38 to TxQueue2. Additionally, move all DSCPs that are in TxQueue4 to TxQueue2, because TxQueue4 is not used.

```
qos map dscp 34 36 38 48 49 50 51 52 to tx-queue 2
qos map dscp 53 54 55 56 57 58 59 60 to tx-queue 2
qos map dscp 61 62 63 to tx-queue 2
```

- Step 3** Verify the modified DSCP-to-TxQueue mapping. The following information was extracted from the **show queueing interface** command.

```
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 02 03 02 03 02 03
4 :    03 03 03 03 03 03 03 03 02 02
5 :    02 02 02 02 02 02 02 02 02 02
6 :    02 02 02 02
```

Step 4 Configure the TxQueue thresholds.

TxQueue1 uses Weighted Random Early Drop (WRED) for queue-congestion management. Only HSD is queued in this queue, and when the amount of HSD in the queue reaches 75%, random packets are dropped in an attempt to keep the queue from reaching 100% utilization.

```
wrr-queue threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
```

TxQueue2 uses tail drop for queue congestion management. Low-priority VoD is assigned to the first threshold and is dropped once the queue reaches 45% utilization. High-priority VoD is assigned to the second threshold and is dropped once the queue reaches 85% utilization. VoD signaling, network signaling, and broadcast video are assigned to the third threshold and is dropped once the queue reaches 100% utilization.

```
wrr-queue threshold 2 45 85 100 100 100 100 100 100
no wrr-queue random-detect 2
```

Step 5 Configure the bandwidth of the weighted queues.

The weighted queues need to be modified to handle our modified TxQueue mappings. The ratio between TxQueue2 and TxQueue1 is $255/64 = 4$, so TxQueue2 needs four times as much bandwidth as TxQueue1. Therefore, TxQueue1 is allocated 20% of the bandwidth on the interface, and TxQueue2 is allocated 80% of the bandwidth.

```
wrr-queue bandwidth 64 255 0 0 0 0 0
```

Step 6 Configure the size of the weighted queues.

Each line card has a limited amount of buffer for the transmit queues. For this interface, 40% of the buffer is allocated for TxQueue1, and 50% of the buffer is allocated for TxQueue2.

```
wrr-queue queue-limit 40 50 0 0 0 0 0
```

Establishing and Configuring Interfaces on AR3

Refer to [Figure 4-1 on page 4-4](#).

This section addresses the following:

- [Establishing VLANs for Services on AR3](#)
- [Establishing 10-GE Interfaces for Transport on AR3](#)
- [Establishing 1-GE Interfaces to a DSLAM on AR3](#)

Establishing VLANs for Services on AR3

Before 1-GE interfaces and 10-GE trunks can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-2 on page 4-5](#).)



Note For additional details, see [Establishing VLANs for Services on DER, page 4-11](#), and [Establishing VLANs for Services on AR1, page 4-24](#).

The following is configured on AR3.

- Step 1** In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```
vlan 90
name VLAN_90_HSD
```

- Step 2** Establish a VLAN for video at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 130
name VLAN_130_Video
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan130
description Video edge VLAN
ip address 192.168.130.1 255.255.255.0
no ip redirects
no ip unreachablees
```

- c. Enable PIM sparse mode

```
ip pim sparse-mode
```

- d. To ensure consistently fast PIM convergence times, statically join the aggregation VLAN for video at the AR to the multicast groups.

```
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

- f. Change the ARP timeout from the default.

```
arp timeout 250
```

- Step 3** Establish a VLAN for VoIP at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 131
name VLAN_131_VoIP
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan131
description VoIP edge VLAN
ip address 192.168.131.1 255.255.255.0
no ip redirects
no ip unreachablees
load-interval 30
```

Step 4 Establish VLANs for VoIP transport. The first is to and from AR2.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 816
name VLAN_816_VoIP_to/from_AR2
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan816
description VoIP transport VLAN to/from AR2
ip address 192.168.252.18 255.255.255.252
```

- c. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Repeat Step 4a through Step 4d to establish a VLAN for VoIP transport to and from DER.

```
vlan 824
name VLAN_824_VoIP_to/from_DER

interface Vlan824
description VoIP transport VLAN to/from DER
ip address 192.168.252.26 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Step 5 Establish VLANs for video transport. The first is to and from AR2.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 916
name VLAN_916_Video_to/from_AR2
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan916
description Video transport VLAN to/from AR2
ip address 192.168.254.18 255.255.255.252
```

- c. Enable PIM sparse mode.

```
ip pim sparse-mode
```

- d. Configure OSPF.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```


- f. Repeat Step 5a through Step 5e to establish a VLAN for video transport to and from DER.

```

vlan 924
name VLAN_924_Video_to/from_DER

interface Vlan924
description Video transport VLAN to/from DER
ip address 192.168.254.26 255.255.255.252 ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30

```

Establishing 10-GE Interfaces for Transport on AR3

The 10-GE trunk interfaces provide the transport between AR3 and AR2 and DER.

The following is configured on AR3.

- Step 1** Establish an interface. The first is to and from AR2.

- a. Establish the interface, configure the trunk for 802.1q encapsulation, and assign it to VLANs 90, 816, and 916. (See [Table 4-2 on page 4-5](#).)

```

interface TenGigabitEthernet1/49
description Transport to/from AR2 (TenGig1/2)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,816,916
switchport mode trunk
dampening
load-interval 30
carrier-delay msec 0

```

- b. Configure this interface (and all noningress transport interfaces) to trust the incoming DSCP markings. (This maintains the DSCP marking applied at the network ingress interface.)

```

qos trust dscp

```

- c. Set transmit-queue bandwidth thresholds and priority.

```

tx-queue 1
bandwidth percent 19
tx-queue 2
bandwidth percent 80
tx-queue 3
priority high
tx-queue 4
bandwidth percent 1

```

- Step 2** Repeat Step 1 to establish an interface to and from DER assign it to VLANs 90, 824, and 924.

```

interface TenGigabitEthernet1/50
description Transport to/from DER (TenGig7/3)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,824,924
switchport mode trunk
dampening
load-interval 30
carrier-delay msec 0

```

Establishing 1-GE Interfaces to a DSLAM on AR3

The only 1-GE interface is a trunk to and from DSLAM3.

The following is configured on AR3.



Note

For additional details, see [Establishing 10-GE Interfaces for Transport on DER, page 4-16](#).

Step 1 Establish an interface to DSLAM1.

- a. Establish the interface and assign it to VLANs 90, 130, and 131.

```
interface GigabitEthernet1/1
  description GigE trunk to/from DSLAM uplink GigE
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,130,131
  switchport mode trunk

  service-policy input setDSCP
  load-interval 30
```

- b. Set transmit-queue bandwidth thresholds and priority.

```
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
```

- c. Prevent unknown unicast traffic from being flooded to the port.

```
switchport block unicast
```



Note

Occasionally, unknown unicast or multicast traffic is flooded to a switch port because a MAC address has timed out or has not been learned by the switch. (This condition is especially undesirable for a private VLAN isolated port.) To guarantee that no unicast or multicast traffic is flooded to the port, use the **switchport block unicast** or **switchport block multicast** commands.

- d. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- e. Enable PortFast on the trunk interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast trunk
```

- f. Configure the switch to disable any interface that is configured for PortFast and receives a Bridge Protocol Data Unit (BPDU).

```
spanning-tree bpduguard enable
```



Note This guards against a user accidentally connecting a switch to a switchport that is intended for a VoD server or other host. The switchport is disabled and the user must investigate why the port is down. If this command is not used and such an accidental connection were to happen, STP could reconverge and block other connections in the switch.

- g. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

Configuring OSPF Routing for Video and Voice Traffic on AR3

For background and details, see [Configuring OSPF Routing for Video and Voice Traffic on DER, page 4-19](#).

The following is configured on AR3.

- Step 1** Define an OSPF routing process to route video traffic. This process associates the transport VLANs for video with the video aggregation VLAN for the DSLAM (VLAN 130).

```
router ospf 100
router-id 4.4.4.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan130
network 192.168.130.0 0.0.0.255 area 0
network 192.168.254.18 0.0.0.0 area 0
network 192.168.254.26 0.0.0.0 area 0
```

- Step 2** Define an OSPF process to route voice traffic. This process associates the transport VLANs for VoIP with the VoIP aggregation VLAN for the DSLAM (VLAN 131).

```
router ospf 101
router-id 4.4.4.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan131
network 192.168.131.0 0.0.0.255 area 0
network 192.168.252.18 0.0.0.0 area 0
network 192.168.252.26 0.0.0.0 area 0
```

Configuring Spanning Tree on AR3



Note See [Configuring Spanning Tree on DER, page 4-20](#).

The following is configured on AR3.

Step 1 Configure RTSP.

```
spanning-tree mode rapid-pvst
```

Step 2 Because the transport VLANs in the 10-GE ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 816, 824, 916, 924
```

Configuring the 1-GE Asymmetric Topology

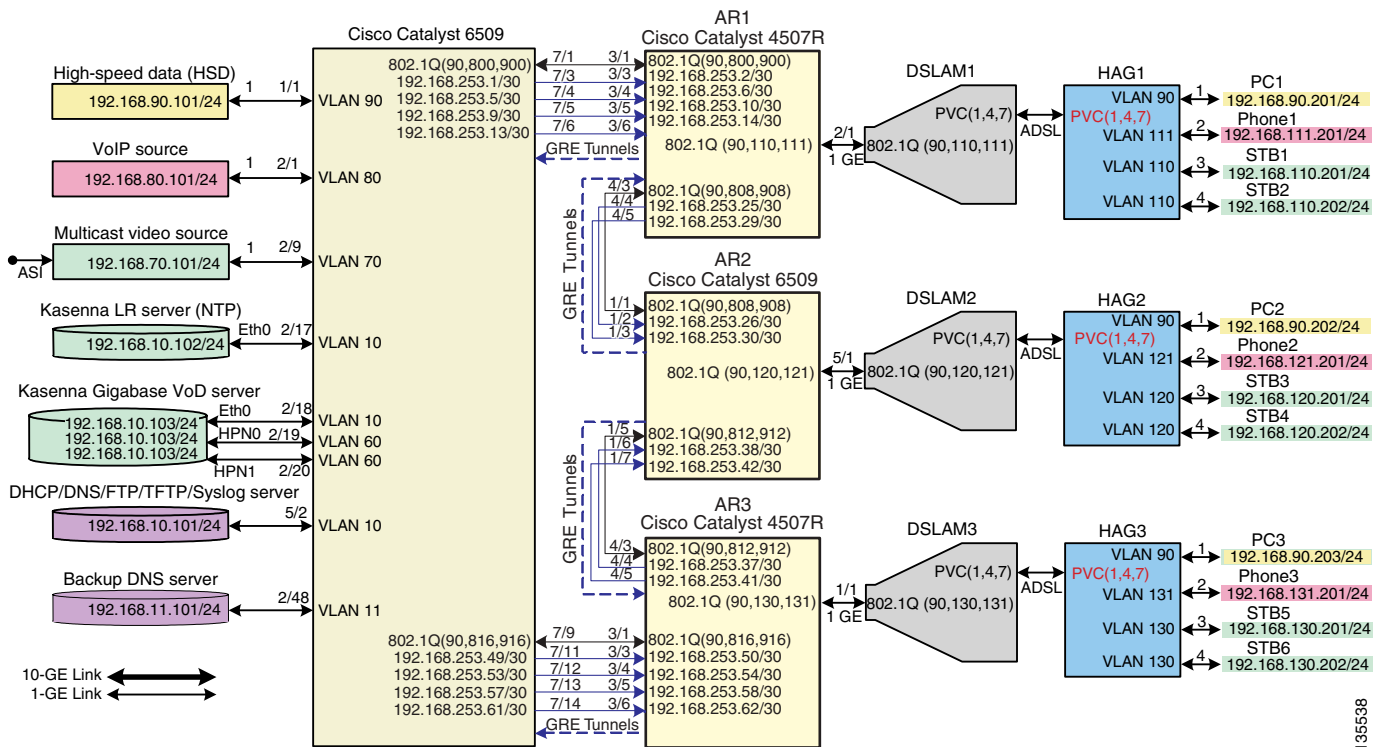
This section presents the following major topics:

- Introduction, page 4-53
- Configuring DER, page 4-57
- Configuring AR1, page 4-74
- Configuring AR2, page 4-89
- Configuring AR3, page 4-100

Introduction

Figure 4-2 illustrates the 1-GE symmetric topology used in the solution. (See [Configuration 2: N x 1-GE Asymmetric Ring](#), page 3-34.) All video traffic sources are on DER. Policy maps are applied to the ingress ports on DER in order to mark the DSCP values of the different service types. Traffic is routed through 1-GE bidirectional links, configured as IEEE 802.1q trunks that carry three VLANs: one for video, one for VoIP, and one for high-speed data (HSD). VoD traffic is also routed through 1-GE unidirectional links that use GRE tunnels for bidirectional connectivity. Multiple OSPF processes are used for the routing protocol. One or two processes advertise routes for the video-related interfaces, and second advertises routes for the VoIP-related interfaces. HSD is carried around the ring on Layer 2.

Figure 4-2 1-GE Asymmetric Topology



The switches in Figure 4-2 use the line cards, hardware versions, and IOS versions listed in Table 4-1 on page 4-5.

Table 4-4 Hardware and IOS Versions for the 1-GE Asymmetric Topology

Switch	Module	Line Card	Hardware Version	IOS Release	Submodule	Hardware Version
DER	1	WS-X6724-SFP	2.0	12.2(18)SXE1	WS-F6700-DFC3BXL	4.0
	2	WS-X6748-GE-TX	1.4			
	5	WS-SUP720-BASE	3.0		WS-F6K-PFC3BXL	1.2
	7	WS-X6816-GBIC	1.7		WS-SUP720 (MFSC)	2.0
					WS-F6K-DFC3BXL	S2.2
AR1	1	WS-X4515	3.1	12.2(25)EWA	—	—
	2	WS-X4306-GB	2.2			
	3					
	5					
AR2	1	WS-X6816-GBIC	1.7	12.2(18)SXE1	WS-F6K-DFC3BXL	2.2
	2	WS-X6724-SFP	1.3		WS-F6700-DFC3BXL	4.0
	5	WS-SUP720-BASE	3.1		WS-F6K-PFC3BXL	1.2
					WS-SUP720 (MSFC)	2.1
AR3	1	WS-X4515	1.2	12.2(25)EWA	—	—
	3	WS-X4306-GB	2.2			
	4					
	5					

Table 4-5 lists VLANs, their descriptions (service types), and IP addresses, for the DER and ARs in Figure 4-1.

Table 4-6 on page 4-55 lists loopback addresses and endpoints for the topology, and describes the associated tunnels.

Table 4-5 VLANs, Descriptions, and IP Addresses for the 1-GE Asymmetric Topology

Node	VLAN	Description	IP Address
DER	10	Management (VoD signaling, DHCP, DNS, FTP, TFTP, Syslog servers)	192.168.10.1/24
	11	Management (backup DNS server)	192.168.11.1/24
	60	VoD server (unicast video aggregation)	192.168.60.1/24
	70	Broadcast video ingress (multicast video aggregation)	192.168.70.1/24
	80	VoIP ingress/egress	192.168.80.1/24
	90	HSD	Layer 2
	800	VoIP transport to/from AR1	192.168.252.1/30
	816	VoIP transport to/from AR3	192.168.252.17/30
	900	Video transport to/from AR1	192.168.254.1/30
916	Video transport to/from AR3	192.168.254.17/30	

Table 4-5 VLANs, Descriptions, and IP Addresses for the 1-GE Asymmetric Topology (continued)

Node	VLAN	Description	IP Address
AR1	90	HSD	Layer 2
	110	Video edge	192.168.110.1/24
	111	VoIP edge	192.168.111.1/24
	800	VoIP transport to/from DER	192.168.252.2/30
	808	VoIP transport to/from AR2	192.168.252.9/30
	900	Video transport to/from DER	192.168.254.2/30
	908	Video transport to/from AR2	192.168.254.9/30
AR2	90	HSD	Layer 2
	120	Video edge	192.168.120.1/24
	121	VoIP edge	192.168.121.1/24
	808	VoIP transport to/from AR1	192.168.252.10/30
	812	VoIP transport to/from AR3	192.168.252.13/30
	908	Video transport to/from AR1	192.168.254.10/30
	912	Video transport to/from AR3	192.168.254.13/30
AR3	90	HSD	Layer 2
	130	Video edge	192.168.130.1/24
	131	VoIP edge	192.168.131.1/24
	812	VoIP transport to/from AR2	192.168.252.14/30
	816	VoIP transport to/from DER	192.168.252.18/30
	912	Video transport to/from AR2	192.168.254.14/30
	916	Video transport to/from DER	192.168.254.18/30

Table 4-6 Loopback and Tunnel Descriptions and IP Addresses for the 1-GE Asymmetric Topology

Node	Loopback	Endpoint for Tunnel No.	Description	IP Address
DER	0	0	Rx side of Tx-only GigabitEthernet7/3	10.10.10.1/32
	4	4	Rx side of Tx-only GigabitEthernet7/4	10.10.10.5/32
	8	8	Rx side of Tx-only GigabitEthernet7/5	10.10.10.9/32
	12	12	Rx side of Tx-only GigabitEthernet7/6	10.10.10.13/32
	48	48	Rx side of Tx-only GigabitEthernet7/11	10.10.10.49/32
	52	52	Rx side of Tx-only GigabitEthernet7/12	10.10.10.53/32
	56	56	Rx side of Tx-only GigabitEthernet7/3	10.10.10.57/32
	60	60	Rx side of Tx-only GigabitEthernet7/14	10.10.10.61/32

Table 4-6 Loopback and Tunnel Descriptions and IP Addresses for the 1-GE Asymmetric Topology (continued)

Node	Loopback	Endpoint for Tunnel No.	Description	IP Address
AR1	0	0	Tx side of Rx-only GigabitEthernet3/3	10.10.10.2/32
	4	4	Tx side of Rx-only GigabitEthernet3/4	10.10.10.6/32
	8	8	Tx side of Rx-only GigabitEthernet3/5	10.10.10.10/32
	12	12	Tx side of Rx-only GigabitEthernet3/6	10.10.10.14/32
	24	24	Tx side of Rx-only GigabitEthernet4/4	10.10.10.25/32
	28	28	Tx side of Rx-only GigabitEthernet4/5	10.10.10.29/32
AR2	24	24	Tx side of Rx-only GigabitEthernet1/2	10.10.10.26/32
	28	28	Tx side of Rx-only GigabitEthernet1/3	10.10.10.30/32
	36	36	Tx side of Rx-only GigabitEthernet1/6	10.10.10.38/32
	40	40	Tx side of Rx-only GigabitEthernet1/7	10.10.10.42/32
AR3	36	36	Tx side of Rx-only GigabitEthernet4/4	10.10.10.37/32
	40	40	Tx side of Rx-only GigabitEthernet4/5	10.10.10.41/32
	48	48	Tx side of Rx-only GigabitEthernet3/3	10.10.10.45/32
	52	52	Tx side of Rx-only GigabitEthernet3/4	10.10.10.54/32
	56	56	Tx side of Rx-only GigabitEthernet3/5	10.10.10.58/32
	60	60	Tx side of Rx-only GigabitEthernet3/6	10.10.10.62/32

Table 4-3 on page 4-6 lists the parameters used to configure the home access gateway (HAG). They are the same as those for the 10-GE symmetric topology.

**Note**

See [HAG Functions, page 3-42](#).

Configuring DER

This section addresses the configuration required on the switch labeled DER in [Figure 4-2 on page 4-53](#), to route multiple services from that switch to the ARs.

See [Configuring DNS Servers, page 4-2](#).

**Note**

A Cisco Catalyst 6509 can also be used, as it uses the same supervisor engine, line cards, and Cisco IOS code as the Cisco 7609 router.

This section addresses the following:

- [Configuring QoS on DER](#)
- [Establishing and Configuring Interfaces on DER](#)
- [Configuring OSPF Routing for Video and Voice Traffic on DER](#)
- [Configuring Spanning Tree on DER](#)

**Note**

For a complete configuration example, see [Appendix B, “Sample DER and AR Switch Configurations for the 1-GE Asymmetric Topology.”](#)

Configuring QoS on DER

This section presents the following topics:

- [Overview of QoS on a Cisco 7600 Series and Cisco Catalyst 6500 Series](#)
- [Configuring Marking and Classification on DER](#)
- [Configuring Mapping on DER](#)

**Note**

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Overview of QoS on a Cisco 7600 Series and Cisco Catalyst 6500 Series

This section addresses the configuration of quality of service (QoS) on the DER, through marking, classification, mapping, and queuing, to provide different degrees of quality of service for the different types of services supported by the solution architecture. For example, it is important to ensure the expeditious delivery of video and VoIP traffic, while providing only best-effort delivery for high-speed data (HSD).

By default, the Cisco 7600 series router and Cisco Catalyst 6500 series switch do not trust the incoming QoS markings, and therefore rewrite these bits with zeros. In this solution, packets at the network ingress ports are identified, classified, and marked according to type of traffic. The packets are marked with one of 64 possible Differentiated Services Code Point (DSCP) values at the ingress ports. These in turn are internally mapped to one of eight possible Class of Service (CoS) values, because CoS is used to determine the appropriate transmit queue for each packet. Queuing is configured on the individual 1-GE interfaces.

**Note**

For more information on class of service, see “White Paper: Cisco IOS Software Features for Differentiated Class of Service for Internetworks,” at the following URL:

http://www.cisco.com/warp/public/cc/pd/iosw/iore/tech/osfea_wp.htm

Configuring Marking and Classification on DER

Do the following to enable marking and classification on DER.

Step 1 Enable QoS in global configuration mode.

```
mls qos
```

Step 2 Create access lists to identify the different service types in the network.

```
ip access-list extended acl_HSD
  permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
  remark Identify VoD signaling traffic
  permit ip host 192.168.10.102 any
  permit ip host 192.168.10.103 any
ip access-list extended acl_VoIP
  remark Identify VoIP traffic
  permit ip 192.168.80.0 0.0.0.255 any
ip access-list extended acl_video_VoD_high
  remark Identify high priority VoD traffic
  permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 5000 9000
  permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 5000 9000
  permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 5000 9000
ip access-list extended acl_video_VoD_low
  remark Identify low priority VoD traffic
  permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 1000 4999
  permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 1000 4999
  permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 1000 4999
ip access-list extended acl_video_broadcast
  remark Identify broadcast video traffic (multicast)
  permit ip 192.168.70.0 0.0.0.255 232.0.0.0 0.255.255.255
```

Step 3 Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_video_VoD_high
  match access-group name acl_video_VoD_high
class-map match-all class_video_VoD_low
  match access-group name acl_video_VoD_low
class-map match-all class_video_broadcast
  match access-group name acl_video_broadcast
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
```

Step 4 Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
  class class_video_broadcast
    set dscp af41
  class class_video_VoD_high
    set dscp af42
  class class_video_VoD_low
    set dscp af43
```

Step 5 Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note

Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

Step 6 To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
mls qos trust dscp
```

Configuring Mapping on DER

Do the following to configure mapping on DER.

Step 1 View the Cisco 7600/Catalyst 6500 default DSCP-to-CoS mapping for the different services. Use the **show mls qos maps dscp-cos** command.



Note

At the beginning of this section, we mentioned that there are 64 possible DSCP values and only 8 CoS values. This means that there could be more than one DSCP value for one CoS value. The following command shows the default DSCP-to-CoS mapping on the Cisco 7600 and Catalyst 6500.



Note

In the map, d1 corresponds to the y-axis value of the table, and d2 to the x-axis value.

```
DER# show mls qos maps dscp-cos
Dscp-cos map: (dscp= d1d2)
d1 : d2 0 1 2 3 4 5 6 7 8 9
-----
0 : 00 00 00 00 00 00 00 00 01 01
1 : 01 01 01 01 01 01 02 02 02 02
2 : 02 02 02 02 03 03 03 03 03 03
3 : 03 03 04 04 04 04 04 04 04 04
4 : 05 05 05 05 05 05 05 05 06 06
5 : 06 06 06 06 06 06 06 06 07 07
6 : 07 07 07 07
```

This table shows the following mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	34	4
VoD high priority	36	4
VoD OOB	24	3
Broadcast video	38	4
VoIP	46	5

- Step 2** Change the Cisco 7600/Catalyst 6500 DSCP-to-CoS mapping for the different services to match the specifications of the solution.

The solution specifies the following DSCP-to CoS-mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	38	1
VoD high priority	36	2
VoD OOB	24	3
Broadcast video	34	4
VoIP	46	5

- a. Execute the following command on the Cisco 7600 and Cisco Catalyst 6500 to modify the DSCP-to-CoS mapping.

```
mls qos map dscp-cos 36 to 2
mls qos map dscp-cos 38 to 1
```

- b. Verify the changes to the DSCP-to-CoS mappings.

```
DER# show mls qos maps dscp-cos
DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2 0 1 2 3 4 5 6 7 8 9
-----
0 : 00 00 00 00 00 00 00 00 01 01
1 : 01 01 01 01 01 01 02 02 02 02
2 : 02 02 02 02 03 03 03 03 03 03
3 : 03 03 04 04 04 04 02 04 01 04
4 : 05 05 05 05 05 05 05 05 06 06
5 : 06 06 06 06 06 06 07 07 07 07
6 : 07 07 07 07
```

Establishing and Configuring Interfaces on DER

Refer to [Figure 4-2 on page 4-53](#).

This section addresses the following:

- [Establishing VLANs for Services on DER](#)
- [Establishing Interfaces for Servers, HSD, and Management on DER](#)
- [Establishing Bidirectional and Unidirectional Interfaces for Transport on DER](#)
- [Establishing Tunnels on DER](#)

Establishing VLANs for Services on DER

Before the 1-GE interfaces can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-5 on page 4-54](#).)

The following is configured on DER.



Tip

For convenience in establish these VLANs and others, you can establish all VLANs in global configuration mode first, then configure all the interfaces in interface configuration mode.

Step 1

Establish a VLAN and VLAN interface for management (including VoD signaling, connectivity with DHCP, DNS, FTP, TFTP, and Syslog servers).

- In global configuration mode, add the VLAN to the VLAN database.

```
vlan 10
name VLAN_10_Management
```

- In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan10
description Management VLAN (VoD signaling, DNS, DHCP, etc)
ip address 192.168.10.1 255.255.255.0
no ip redirects
no ip unreachablees
```

- Change the load interval from the default of 300.

```
load-interval 30
```

- Repeat Step 1a through Step 1c, as appropriate, for the remaining management, unicast video, and and VoIP VLANs. Abbreviated configurations are shown below.

Backup DNS server

```
vlan 11
name VLAN_11_Management
```

```
interface Vlan11
description Management VLAN (Backup DNS)
ip address 192.168.11.1 255.255.255.0
no ip redirects
no ip unreachablees
load-interval 30
```

Unicast video aggregation

```

vlan 60
name VLAN_60_Unicast_Video

interface Vlan60
description VoD server VLAN (Unicast Video)
ip address 192.168.60.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30

```

VoIP

```

vlan 80
name VLAN_80_VoIP

interface Vlan80
description VoIP ingress/egress VLAN
ip address 192.168.80.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30

```

Step 2 Establish a VLAN for multicast video aggregation.

- a.** In global configuration mode, add the VLAN to the VLAN database.

```

vlan 70
name VLAN_70_Multicast_Video

```

- b.** In interface configuration mode, create and configure the VLAN interfaces.

```

interface Vlan70
description Broadcast video source VLAN (Multicast Video)
ip address 192.168.70.1 255.255.255.0
no ip redirects
no ip unreachable

```

- c.** Enable PIM sparse mode. This is the ingress port for broadcast video traffic, which is multicast addressed.

```

ip pim sparse-mode

```

- d.** Change the load interval from the default of 300.

```

load-interval 30

```

Step 3 In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```

vlan 90
name VLAN_90_HSD

```

Step 4 Establish VLANs for VoIP transport. The first is for transport to and from AR1.

- a.** In global configuration mode, add the VLAN to the VLAN database.

```

vlan 800
name VLAN_800_VoIP_to/from_AR1

```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan800
  description VoIP transport to/from AR1
  ip address 192.168.252.1 255.255.255.252
```

- c. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```



Note To avoid the election of the (designated router (DR) and backup designated router (BDR), and prevent the origination of an unnecessary network link state advertisement (LSA), configure the transport VLAN as a point-to-point network. In addition, reduce the interval between OSPF hello messages from 10 seconds to 1 second. This improves reconvergence in the event of failure in the transport or in a neighboring switch.

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Repeat Step 4a through Step 4d for VoIP transport to and from AR3.

```
vlan 816
name VoIP transport to/from AR3

interface Vlan816
  description VoIP transport to/from AR3
  ip address 192.168.252.17 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
```

Step 5 Establish VLANs for video transport. The first is for transport to and from AR1.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 900
name VLAN_900_Video_to/from_AR1
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan900
  description Video transport VLAN to/from AR1
  ip address 192.168.254.1 255.255.255.252
```

- c. Enable PIM sparse mode. This is the ingress port for broadcast video traffic, which is multicast addressed.

```
ip pim sparse-mode
```

- d. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

- f. Repeat Step 5a through Step 5b to establish a VLAN for video transport to and from AR3.

```
vlan 916
name Video_transport_to/from_AR3

interface Vlan916
description Video transport VLAN to/from AR3
ip address 192.168.254.17 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Establishing Interfaces for Servers, HSD, and Management on DER

VoD servers, high-speed data sources, and management resources connect to Layer 2 interfaces on DER, and their traffic is aggregated into the appropriate service VLANs.

The following is configured on DER.

Step 1 Establish an interface.

- a. Establish an interface for high-speed data and assign it to VLAN 90.

```
interface GigabitEthernet1/1
description High speed data ingress/egress port
no ip address
```

- b. Configure the interface as a Layer 2 access port.

```
switchport
switchport access vlan 90
switchport mode access
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```

- d. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- e. Enable PortFast on the interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast
```

- f. Configure the switch to disable any interface that is configured for PortFast and receives a Bridge Protocol Data Unit (BPDU).

```
spanning-tree bpduguard enable
```



Note This guards against a user accidentally connecting a switch to a switchport that is intended for a VoD server or other host. The switchport is disabled and the user must investigate why the port is down. If this command is not used and such an accidental connection were to happen, STP could reconverge and block other connections in the switch.

- g. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

- Step 2** Repeat Step 1a through Step 1g for the remaining server, HSD, and management interfaces and their associated VLANs, changing the value in **switchport access vlan xxx** as appropriate. The abbreviated configurations are shown below.

VoIP traffic

```
interface GigabitEthernet2/1
description VoIP traffic ingress/egress port

switchport access vlan 80
```

Ingress multicast broadcast video

```
interface GigabitEthernet2/9
description Broadcast video source (multicast 232.1.1.1 - 232.1.1.10)

switchport access vlan 70
```

Management for the Kasenna LR server

```
interface GigabitEthernet2/17
description Management port from Kasenna LR Server (Eth0)

switchport access vlan 10
```

Management for the Kasenna VoD pump

```
interface GigabitEthernet2/18
description Kasenna VoD Pump Management

switchport access vlan 10
```

Ingress unicast video from the Kasenna VoD pump (1)

```
interface GigabitEthernet2/19
description Unicast video from Kasenna VoD Pump (HPN0)

switchport access vlan 60
```



Note In Kasenna’s terminology, HPN0 stands for High-Performance Network interface 0.

Ingress unicast video from the Kasenna VoD pump (2)

```
interface GigabitEthernet2/20
description Unicast video from Kasenna VoD Pump (HPN1)

switchport access vlan 60
```

Backup DNS server

```
interface GigabitEthernet2/48
description Backup DNS server

switchport access vlan 11
```

Primary DNS, DHCP, NTP, TFTP, and Syslog servers

```
interface GigabitEthernet5/2
  description Primary DNS/DHCP/NTP/TFTP/Syslog servers

  switchport access vlan 10
```



Note In this case, specify the physical connection on a Gigabit Ethernet interface as RJ-45.

```
media-type rj45
```

Establishing Bidirectional and Unidirectional Interfaces for Transport on DER

The 1-GE interfaces create the ring topology from the DER through the ARs and back to the DER. Both bidirectional and unidirectional trunking interfaces and VoD unidirectional transport are established.

The following is configured on DER.

Step 1 Establish bidirectional transport interfaces.

- a. Establish a bidirectional transport interface to and from AR1.

```
interface GigabitEthernet7/1
  description Transport to/from AR1 (Gig3/1)
  switchport

  switchport mode trunk
  dampening
  no ip address

  carrier-delay msec 0
```

- b. Configure the trunk for 802.1q encapsulation.

```
switchport trunk encapsulation dot1q
```

- c. Assign the trunk to VLANs 90, 800, and 900. (See [Table 4-5 on page 4-54.](#))

```
switchport trunk allowed vlan 90,800,900
```

- d. Change the load interval from the default of 300.

```
load-interval 30
```

Step 2 Configure QoS on the interface.



Note The 1-GE transport links from the DER to the ARs require modifications to the transmit queues. There are eight transmit queues, but this solution uses only three.

- a. View the default CoS to Tx-Queue mapping. The following information was extracted from the **show queueing interface** command.

TxQueue	CoS
1	0, 1
2	2, 3, 4
3	6, 7
4	—
5	—
6	—
7	—
8	5

- b. Configure the CoS-to TxQueue mapping on the transport interfaces. HSD (CoS = 0) remains in TxQueue1 and VoIP (CoS = 5) remains in TxQueue8. The other six CoS values are associated with TxQueue2.

```
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
```



Note TxQueue1 and TxQueue8 use the default mappings. TxQueue2 has three thresholds: Threshold 1 = CoS 1, Threshold 2 = CoS 2, and Threshold 3 = CoS 3, 4, 6, and 7. For details, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

- c. Verify the modified CoS-to-Tx-Queue mapping. The following information was extracted from the **show queueing interface** command.

TxQueue	CoS
1	0
2	1, 2, 3, 4, 6, 7
3	—
4	—
5	—
6	—
7	—
8	5

- d. Configure the TxQueue thresholds.

TxQueue1 uses Weighted Random Early Drop (WRED) for queue-congestion management. Only HSD is queued in this queue, and when the amount of HSD in the queue reaches 75%, random packets are dropped in an attempt to keep the queue from reaching 100% utilization.

```
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect max-threshold 1 100 100
```

TxQueue2 uses tail drop for queue congestion management. Low-priority VoD is assigned to the first threshold and is dropped once the queue reaches 50% utilization. High-priority VoD is assigned to the second threshold and is dropped once the queue reaches 85% utilization. VoD signaling, network signaling, and broadcast video are assigned to the third threshold and is dropped once the queue reaches 100% utilization.

```
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 2 50 100
```

- e. Configure the bandwidth of the weighted queues.

The weighted queues need to be modified to handle our modified TxQueue mappings. The ratio between TxQueue2 and TxQueue1 is $255/64 = 4$, so TxQueue2 needs four times as much bandwidth as TxQueue1. Therefore, TxQueue1 is allocated 20% of the bandwidth on the interface, and TxQueue2 is allocated 80% of the bandwidth.

```
wrr-queue bandwidth 64 255
```

- f. Configure the size of the weighted queues.

Each line card has a limited amount of buffer for the transmit queues. For this interface, 40% of the buffer is allocated for TxQueue1, and 50% of the buffer is allocated for TxQueue2.

```
wrr-queue queue-limit 40 50
```

- g. Configure this interface (and all noningress transport interfaces) to trust the incoming DSCP markings. (This maintains the DSCP marking applied at the network ingress interface.)

```
mls qos trust dscp
```

Step 3 Repeat Step 1 and Step 2, as appropriate, for the bidirectional transport to and from AR3.

```
interface GigabitEthernet7/9
  description Transport to/from AR3 (Gig3/1)
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,816,916
  switchport mode trunk
  dampening
  no ip address
  load-interval 30
  carrier-delay msec 0
  wrr-queue bandwidth 64 255
  wrr-queue queue-limit 40 50
  wrr-queue random-detect min-threshold 1 75 100
  wrr-queue random-detect min-threshold 2 50 100
  wrr-queue random-detect max-threshold 1 100 100
  wrr-queue random-detect max-threshold 2 50 100
  wrr-queue cos-map 1 1 0
  wrr-queue cos-map 2 1 1
  wrr-queue cos-map 2 2 2 3 4 6 7
  mls qos trust dscp
```

Step 4 Establish unidirectional VoD transport interfaces.

- a. Establish the unidirectional VoD transport interface to AR1 and assign the IP address.

```
interface GigabitEthernet7/3
  description VoD transport to AR1 (Gig3/3)
  dampening
  ip address 192.168.253.1 255.255.255.252
```

- b. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```

- d. Disable the link-negotiation protocol on the port.

```
speed nonegotiate
```

- e. Mark the interface as send only.

```
unidirectional send-only
```

- f. Repeat Step 4a through Step 4e, as appropriate, for the remaining unidirectional VoD transport interfaces. Abbreviated interface configurations are shown below.

Second VoD transport interface to AR1

```
interface GigabitEthernet7/4
description VoD transport to AR1 (Gig3/4)

ip address 192.168.253.5 255.255.255.252
```

Third VoD transport interface to AR1

```
interface GigabitEthernet7/5
description VoD transport to AR1 (Gig3/5)

ip address 192.168.253.9 255.255.255.252
```

Fourth VoD interface to AR1

```
interface GigabitEthernet7/6
description VoD transport to AR1 (Gig3/6)

ip address 192.168.253.13 255.255.255.252
```

First VoD interface to AR3

```
interface GigabitEthernet7/11
description VoD transport to AR1 (Gig3/3)

ip address 192.168.253.49 255.255.255.252
```

Second VoD interface to AR3

```
interface GigabitEthernet7/12
description VoD transport to AR1 (Gig3/4)

ip address 192.168.253.53 255.255.255.252
```

Third VoD interface to AR3

```
interface GigabitEthernet7/13
description VoD transport to AR1 (Gig3/5)

ip address 192.168.253.57 255.255.255.252
```

Fourth VoD interface to AR3

```
interface GigabitEthernet7/14
  description VoD transport to AR1 (Gig3/6)

  ip address 192.168.253.61 255.255.255.252
```

Establishing Tunnels on DER

The following is configured on DER.

- Step 1** Create a loopback interface to serve as the tunnel endpoint for the first tunnel and assign an IP address.

```
interface Loopback0
  description Endpoint for Tunnel0
  ip address 10.10.10.1 255.255.255.255
```

- Step 2** Create the interface for the corresponding tunnel. No IP address is required for the tunnel itself.

```
interface Tunnel0
  description Rx-side of Tx-only Gig7/3
  no ip address
```

- Step 3** Configure the source and destination endpoints of the tunnel and assign IP addresses.

```
tunnel source 10.10.10.1
tunnel destination 10.10.10.2
```

- Step 4** Configure UDLR for the tunnel and mark it as receive only.

```
tunnel udlr receive-only GigabitEthernet7/3
```

- Step 5** Refer to [Table 4-7](#) and repeat Steps 1 through Step 4 for the remaining tunnels on DER, making modifications as appropriate.

Table 4-7 Additional Corresponding Loopback and Tunnel Interfaces for DER

Loopback Interface	Tunnel Interface
<pre>interface Loopback4 description Endpoint for Tunnel4 ip address 10.10.10.5 255.255.255.255</pre>	<pre>interface Tunnel4 description Rx-side of Tx-only Gig7/4 no ip address tunnel source 10.10.10.5 tunnel destination 10.10.10.6 tunnel udlr receive-only GigabitEthernet7/4</pre>
<pre>interface Loopback8 description Endpoint for Tunnel8 ip address 10.10.10.9 255.255.255.255</pre>	<pre>interface Tunnel8 description Rx-side of Tx-only Gig7/5 no ip address tunnel source 10.10.10.9 tunnel destination 10.10.10.10 tunnel udlr receive-only GigabitEthernet7/5</pre>
<pre>interface Loopback12 description Endpoint for Tunnel12 ip address 10.10.10.13 255.255.255.255</pre>	<pre>interface Tunnel12 description Rx-side of Tx-only Gig7/6 no ip address tunnel source 10.10.10.13 tunnel destination 10.10.10.14 tunnel udlr receive-only GigabitEthernet7/6</pre>

Table 4-7 Additional Corresponding Loopback and Tunnel Interfaces for DER (continued)

Loopback Interface	Tunnel Interface
<pre>interface Loopback48 description Endpoint for Tunnel48 ip address 10.10.10.49 255.255.255.255</pre>	<pre>interface Tunnel48 description Rx-side of Tx-only Gig7/11 no ip address tunnel source 10.10.10.49 tunnel destination 10.10.10.50 tunnel udlr receive-only GigabitEthernet7/11</pre>
<pre>interface Loopback52 description Endpoint for Tunnel52 ip address 10.10.10.53 255.255.255.255</pre>	<pre>interface Tunnel52 description Rx-side of Tx-only Gig7/12 no ip address tunnel source 10.10.10.53 tunnel destination 10.10.10.54 tunnel udlr receive-only GigabitEthernet7/12</pre>
<pre>interface Loopback56 description Endpoint for Tunnel56 ip address 10.10.10.57 255.255.255.255</pre>	<pre>interface Tunnel56 description Rx-side of Tx-only Gig7/13 no ip address tunnel source 10.10.10.57 tunnel destination 10.10.10.58 tunnel udlr receive-only GigabitEthernet7/13</pre>
<pre>interface Loopback60 description Endpoint for Tunnel60 ip address 10.10.10.61 255.255.255.255</pre>	<pre>interface Tunnel60 description Rx-side of Tx-only Gig7/14 no ip address tunnel source 10.10.10.61 tunnel destination 10.10.10.62 tunnel udlr receive-only GigabitEthernet7/14</pre>

Configuring OSPF Routing for Video and Voice Traffic on DER

There are a number of ways to configure the routing of the multiple services across the 1-GE asymmetric topology. HSD, VoIP, broadcast video, and VoD signaling can be routed across the bidirectional links, while VoD traffic can be routed across both the bidirectional and unidirectional links, or just across the unidirectional links.

In this example, HSD, VoIP, broadcast video, and VoD signaling are routed across the bidirectional links. To accomplish this, the bidirectional link is a trunk that carries three VLANs (90, 8xx, and 9xx). Because VLAN 90 is at Layer 2 around the network, there is no OSPF configuration for HSD. VoIP-related interfaces are advertised across the 8xx VLANs. Broadcast video is multicast, so the path is built from the receiver to the source. Because the DER does not need to know how to route to the destination for broadcast video, we only need to advertise the broadcast video sources across the 9xx VLANs, so that the receivers can build the reverse path back to the broadcast video source. VoD signaling is lower-bitrate, bidirectional traffic, but we still want this traffic to travel across the bidirectional transport links, rather than through the GRE tunnels associated with the unidirectional transport links. To accomplish this, we advertise the VoD server-related interfaces across the 9xx VLANs. In addition, the loopback interfaces that serve as the endpoints of the GRE tunnels for the unidirectional links are also advertised on the 9xx VLANs.

Finally, in this example the VoD traffic is routed across the unidirectional transport links only. VoD traffic is routed from the source to the receivers, so the DER must know how to route to the receivers. To accomplish this, no interfaces need to be advertised from the DER, but the DER needs a routing process associated with the unidirectional transport links to receive routing advertisements for the video aggregation VLANs on the ARs.

Three Open Shortest Path First (OSPF) routing processes must be established:

- OSPF 100—to route the management, broadcast video, and loopbacks over the transport VLANs for video
- OSPF 101—to route VoIP traffic over the transport VLANs for VoIP
- OSPF 102—to route VoD traffic over the unidirectional Layer 3 transport network for VoD

Routing advertisements are enabled on the transport VoD network, but are turned off on the aggregation VLANs by means of the **passive-interface** command.

The following is configured on DER.

Step 1 Define an OSPF routing process to route video traffic.

```
router ospf 100
router-id 1.1.1.1
log-adjacency-changes
```

- a. The OSPF timers are modified to provide fast convergence. The following command enables OSPF SPF throttling: **timers throttle spf** *spf-start spf-hold spf-max-wait*

```
timers throttle spf 10 100 1000
```

- b. The following command sets the rate-limiting values for OSPF link-state advertisement (LSA) generation: **timers throttle lsa all** *start-interval hold-interval max-interval*

```
timers throttle lsa all 1 10 1000
```

- c. The following command controls the minimum interval for accepting the same LSA: **timers lsa arrival** *milliseconds*

```
timers lsa arrival 100
```

If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped.

- d. Apply the **passive-interface** command to the aggregation VLANs.

```
passive-interface Vlan10
passive-interface Vlan11
passive-interface Vlan60
passive-interface Vlan70
```

- e. Advertise the networks in the first OSPF routing process.

```
network 10.10.10.0 0.0.0.255 area 0
network 192.168.10.0 0.0.1.255 area 0
network 192.168.60.0 0.0.0.255 area 0
network 192.168.70.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
```

Step 2 Repeat Step 1, as appropriate, to define a second OSPF process to route VoIP traffic.

```
router ospf 101
router-id 1.1.1.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan80
network 192.168.80.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
```


Step 3 Define a third OSPF process to route VoD transport traffic.

```
router ospf 102
router-id 1.1.1.3
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
```

Configuring Spanning Tree on DER

Because VLAN 90 is at Layer 2 around the 1-GE ring, Spanning Tree Protocol (STP) is needed to guard against loops. To improve convergence time, the four switches are configured for IEEE 802.1w Rapid Spanning Tree Protocol (RTSP), with the root at DER.

Do the following in global configuration mode to configure spanning tree parameters on DER.

Step 1 Configure DER as the root node of the spanning tree for VLAN 90. There are two ways to do this.

a. Use the **root primary** option.

```
spanning-tree vlan 90 root primary
```

or

b. Set the priority to 24576.

```
spanning-tree vlan 90 priority 24576
```

Step 2 Configure RTSP.

```
spanning-tree mode rapid-pvst
```

Step 3 Because the transport VLANs in the ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 800, 816, 900, 916
```

Configuring AR1

This section addresses the configuration required on the switch labeled AR1 in [Figure 4-2 on page 4-53](#), to route multiple services from AR1 to DER and AR2.

See [Configuring DNS Servers, page 4-2](#).

This section addresses the following:

- [Configuring QoS on AR1](#)
- [Establishing and Configuring Interfaces on AR1](#)
- [Configuring OSPF Routing for Video and Voice Traffic on AR1](#)
- [Configuring Spanning Tree on AR1](#)

**Note**

For a complete configuration example, see [Appendix B, “Sample DER and AR Switch Configurations for the 1-GE Asymmetric Topology.”](#)

Configuring QoS on AR1

See [Overview of QoS on a Cisco 7600 Series and Cisco Catalyst 6500 Series, page 4-57](#).

This section presents the following topics:

- [Overview of QoS on a Cisco Catalyst 4500 Series, page 4-74](#)
- [Configuring Marking and Classification on AR1](#)
- [Configuring Mapping on AR1](#)
- [Configuring Queuing on AR1](#)

**Note**

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Overview of QoS on a Cisco Catalyst 4500 Series

This section addresses the configuration of quality of service (QoS) on AR2, through marking, classification, mapping, and queueing, to provide different degrees of quality of service for the different types of services supported by the solution architecture. For example, it is important to ensure the expeditious delivery of video and VoIP traffic, while providing only best-effort delivery for high-speed data (HSD).

By default, the Cisco Catalyst 4500 series switches (including the Cisco Catalyst 4948-10GE) do not trust the incoming QoS markings, and therefore rewrite these bits with zeros. In this solution, packets at the network ingress ports are identified, classified, and marked according to type of traffic. The packets are marked with one of 64 possible Differentiated Services Code Point (DSCP) values at the ingress ports. These in turn are internally mapped to one of eight possible Class of Service (CoS) values. The DSCP values are used to determine the appropriate transmit queue for each packet.

Configuring Marking and Classification on AR1

Do the following to enable marking and classification on AR1.

Step 1 Enable QoS in global configuration mode.

```
qos
```

Step 2 Create access lists to identify the different service types in the network.

```
ip access-list extended acl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip 192.168.110.0 0.0.0.255 192.168.10.102
 permit ip 192.168.110.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.111.0 0.0.0.255 any
```

Step 3 Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
 match access-group name acl_VoIP
class-map match-all class_VoD_signaling
 match access-group name acl_VoD_signaling
class-map match-all class_HSD
 match access-group name acl_HSD
```

Step 4 Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
 description Mark DSCP values for ingress traffic
 class class_VoIP
 set dscp ef
 class class_HSD
 set dscp default
 class class_VoD_signaling
 set dscp cs3
```

Step 5 Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note

Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

Step 6 To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
qos trust dscp
```

Configuring Mapping on AR1

Do the following to configure mapping on AR1.

- Step 1** View the Cisco Catalyst 4500 series default DSCP-to-CoS mapping for the different services. Use the **show qos maps dscp-cos** command.



Note

At the beginning of this section, we mentioned that there are 64 possible DSCP values and only 8 CoS values. This means that there could be more than one DSCP value for one CoS value. The following command shows the default DSCP-to-CoS mapping on the Cisco Catalyst 4500 series.



Note

In the map, d1 corresponds to the y-axis value of the table, and d2 to the x-axis value.

```
AR2# show qos maps dscp

DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 04 04 04 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07
```

This table shows the following mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	34	4
VoD high priority	36	4
VoD OOB	24	3
Broadcast video	38	4
VoIP	46	5

- Step 2** Change the Cisco Catalyst 4500 series DSCP-to-CoS mapping for the different services to match the specifications of the solution.

The solution specifies the following DSCP-to CoS-mappings:

Service Type	DSCP	CoS
HSD	0	0
VoD low priority	38	1
VoD high priority	36	2
VoD OOB	24	3
Broadcast video	34	4
VoIP	46	5

- a. Execute the following command on the Cisco Catalyst 4500 series to modify the DSCP-to-CoS mapping.

```
qos map dscp 38 to cos 1
qos map dscp 36 to cos 2
```

- b. Verify the changes to the DSCP-to-CoS mappings.

```
AR2# show qos maps dscp
```

```
DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 02 04 01 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07
```

Configuring Queueing on AR1

Unlike the Cisco 7600 series and Cisco Catalyst 6500 series, the Cisco Catalyst 4500 series uses the same queueing on all interfaces. Queueing is configured globally.

Do the following to change the DSCP-to-TxQueue mappings on AR1.

- Step 1** View the default DSCP-to-Tx-Queue mapping. The following information was extracted from the **show qos maps dscp** command.

```
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
```

- Step 2** Configure the DSCP-to-TxQueue mapping by moving DSCP 34, 36, and 38 to TxQueue2. Additionally, move all DSCPs that are in TxQueue4 to TxQueue2, because TxQueue4 is not used.

```
qos map dscp 34 36 38 48 49 50 51 52 to tx-queue 2
qos map dscp 53 54 55 56 57 58 59 60 to tx-queue 2
qos map dscp 61 62 63 to tx-queue 2
```

- Step 3** Verify the modified DSCP-to-Tx-Queue mapping. The following information was extracted from the **show queueing interface** command.

```
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 02 03 02 03 02 03
4 :    03 03 03 03 03 03 03 03 02 02
5 :    02 02 02 02 02 02 02 02 02 02
6 :    02 02 02 02
```

- Step 4** Configure the TxQueue thresholds.

TxQueue1 uses Weighted Random Early Drop (WRED) for queue-congestion management. Only HSD is queued in this queue, and when the amount of HSD in the queue reaches 75%, random packets are dropped in an attempt to keep the queue from reaching 100% utilization.

```
wrr-queue threshold 1 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100
```

TxQueue2 uses tail drop for queue congestion management. Low-priority VoD is assigned to the first threshold and is dropped once the queue reaches 45% utilization. High-priority VoD is assigned to the second threshold and is dropped once the queue reaches 85% utilization. VoD signaling, network signaling, and broadcast video are assigned to the third threshold and is dropped once the queue reaches 100% utilization.

```
wrr-queue threshold 2 45 85 100 100 100 100 100
no wrr-queue random-detect 2
```

- Step 5** Configure the bandwidth of the weighted queues.

The weighted queues need to be modified to handle our modified TxQueue mappings. The ratio between TxQueue2 and TxQueue1 is $255/64 = 4$, so TxQueue2 needs four times as much bandwidth as TxQueue1. Therefore, TxQueue1 is allocated 20% of the bandwidth on the interface, and TxQueue2 is allocated 80% of the bandwidth.

```
wrr-queue bandwidth 64 255 0 0 0 0 0
```

- Step 6** Configure the size of the weighted queues.

Each line card has a limited amount of buffer for the transmit queues. For this interface, 40% of the buffer is allocated for TxQueue1, and 50% of the buffer is allocated for TxQueue2.

```
wrr-queue queue-limit 40 50 0 0 0 0 0
```

Establishing and Configuring Interfaces on AR1

Refer to [Figure 4-2 on page 4-53](#).

This section addresses the following:

- [Establishing VLANs for Services on AR1](#)
- [Establishing Bidirectional and Unidirectional Transport Interfaces on AR1](#)
- [Establishing Tunnels on AR1](#)
- [Establishing an Interface to a DSLAM on AR1](#)

Establishing VLANs for Services on AR1

Before bidirectional and unidirectional transport interfaces can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-2 on page 4-5](#).)

The following is configured on AR1.

**Note**

For additional details, see [Establishing VLANs for Services on DER, page 4-61](#).

- Step 1** In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```
vlan 90
name VLAN_90_HSD
```

- Step 2** Establish a VLAN for video at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 110
name VLAN_110_Video
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan110
description Video edge VLAN
ip address 192.168.110.1 255.255.255.0
no ip redirects
no ip unreachablees
```

- c. Enable PIM sparse mode. This is the aggregation VLAN for video traffic to the DSLAMs. Broadcast video is multicast addressed.

```
ip pim sparse-mode
```

- d. To ensure consistently fast PIM convergence times, statically join the aggregation VLAN for video at the AR to the multicast groups.

```
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

- f. Change the ARP timeout from the default.

```
arp timeout 250
```



Note The default timeout for an entry in the ARP cache is 4 hours. The default timeout for an entry in the MAC address table is only 5 minutes. Because video traffic is mostly unidirectional, the MAC address table may not be refreshed within the 5-minute timeout. This causes video traffic to be flooded until the destination MAC address is found. To prevent this, reduce the ARP cache timeout to 250 seconds. This forces the switch to re-ARP for the entries in the ARP cache before the entries in the MAC address table time out, avoiding the disruptive behavior.

Step 3 Establish a VLAN for VoIP at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 111
name VLAN_111_VoIP
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan111
description VoIP edge VLAN
ip address 192.168.111.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30
```

Step 4 Establish a VLAN for VoIP transport to and from DER.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 800
name VLAN_800_VoIP_to/from_DER
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan800
description VoIP transport VLAN to/from DER
ip address 192.168.252.2 255.255.255.252
```

- c. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```


- d. Change the load interval from the default of 300.

```
load-interval 30
```

Step 5 Repeat Step 4, as appropriate, to establish a VLAN for VoIP transport to and from AR2.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 808
name VLAN_808_VoIP_to/from_AR2
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan808
description VoIP transport to/from AR2
ip address 192.168.252.9 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Step 6 Establish a VLAN for video transport to and from DER.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 900
name VLAN_900_Video_to/from_DER
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan900
description Video transport VLAN to/from DER
ip address 192.168.254.2 255.255.255.252
```

- c. Enable PIM sparse mode.

```
ip pim sparse-mode
```

- d. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

Step 7 Repeat Step 6, as appropriate, to establish a VLAN for video transport to and from AR2.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 908
name VLAN_908_Video_to/from_AR2
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan908
description Video transport VLAN to/from AR2
ip address 192.168.254.9 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Establishing Bidirectional and Unidirectional Transport Interfaces on AR1

Bidirectional and unidirectional transport interfaces must be established between AR1 and DER and AR2.

The following is configured on AR1.



Note

For additional details, see [Establishing Bidirectional and Unidirectional Interfaces for Transport on DER, page 4-66](#).

Step 1 Establish bidirectional transport interfaces.

- a. Establish a bidirectional interface to and from DER.

```
interface GigabitEthernet3/1
  description Transport to/from DER (Gig7/1)
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,800,900
  switchport mode trunk
  dampening
  load-interval 30
  carrier-delay msec 0
```

- b. Configure this interface (and all noningress transport interfaces) to trust the incoming DSCP markings. (This maintains the DSCP marking applied at the network ingress interface.)

```
qos trust dscp
```

- c. Set transmit-queue bandwidth thresholds and priority.

```
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
```

- d. Repeat Step 1a through Step 1c, as appropriate, to establish the bidirectional transport interface to AR2.

```
interface GigabitEthernet4/3
  description Transport to/from AR2 (Gig1/1)
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,808,908
  switchport mode trunk
  dampening
  load-interval 30
  carrier-delay msec 0
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
```

Step 2 Establish unidirectional receive-only transport interfaces.

- a. Establish a unidirectional receive-only transport interface to DER. With the exceptions noted, the following is as in Step 1.

```
interface GigabitEthernet3/3
  description Transport from DER (Gig7/3)
  no switchport
  dampening
  ip address 192.168.253.2 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  carrier-delay msec 0

  qos trust dscp
```

- b. Disable the link-negotiation protocol on the port.

```
speed nonegotiate
```

- c. Mark the interface as receive only.

```
unidirectional receive-only
```



Note Transmit-queue bandwidth thresholds and priority do not need to be applied to a receive-only interface.

- d. Repeat Step 2a through Step 2c, as appropriate, for the remaining unidirectional receive-only transport interfaces. Interface configurations are shown below.

Second unidirectional transport from DER

```
interface GigabitEthernet3/4
  description Transport from DER (Gig7/4)
  no switchport
  dampening
  ip address 192.168.253.6 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  carrier-delay msec 0
  speed nonegotiate
  qos trust dscp
  unidirectional receive-only
```

Third unidirectional transport from DER

```
interface GigabitEthernet3/5
  description Transport from DER (Gig7/5)
  no switchport
  dampening
  ip address 192.168.253.10 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  carrier-delay msec 0
  speed nonegotiate
  qos trust dscp
  unidirectional receive-only
```

Fourth unidirectional transport from DER

```

description Transport from DER (Gig7/6)
no switchport
dampening
ip address 192.168.253.14 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
unidirectional receive-only

```

Step 3 Establish unidirectional send-only interfaces.

- a. Establish a unidirectional send-only transport interface to AR2. The following is as in Step 1.

```

interface GigabitEthernet4/4
description Transport to AR2 (Gig1/2)
no switchport
dampening
ip address 192.168.253.25 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1

```

- b. Mark the interface as send only.

```

unidirectional send-only

```

- c. Repeat Step 3a and Step 3b, as appropriate, for the second unidirectional send-only interface to AR2.

```
interface GigabitEthernet4/5
  description Transport to AR2 (Gig1/3)
  no switchport
  dampening
  ip address 192.168.253.29 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  carrier-delay msec 0
  speed nonegotiate
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
  unidirectional send-only
```

Establishing Tunnels on AR1

See [Establishing Tunnels on DER, page 4-70](#).

[Table 4-8 on page 4-85](#) lists the loopback interfaces and corresponding tunnel interfaces configured on AR1.

Table 4-8 Corresponding Loopback and Tunnel Interfaces for AR1

Loopback Interface	Tunnel Interface
<pre>interface Loopback0 description Endpoint for Tunnel0 ip address 10.10.10.2 255.255.255.255</pre>	<pre>interface Tunnel0 description Tx-side of Rx-only Gig3/3 no ip address tunnel source 10.10.10.2 tunnel destination 10.10.10.1 tunnel udlr send-only GigabitEthernet3/3 tunnel udlr address-resolution</pre>
<pre>interface Loopback4 description Endpoint for Tunnel4 ip address 10.10.10.6 255.255.255.255</pre>	<pre>interface Tunnel4 description Tx-side of Rx-only Gig3/4 no ip address tunnel source 10.10.10.6 tunnel destination 10.10.10.5 tunnel udlr send-only GigabitEthernet3/4 tunnel udlr address-resolution</pre>
<pre>interface Loopback8 description Endpoint for Tunnel8 ip address 10.10.10.10 255.255.255.255</pre>	<pre>interface Tunnel8 description Tx-side of Rx-only Gig3/5 no ip address tunnel source 10.10.10.10 tunnel destination 10.10.10.9 tunnel udlr send-only GigabitEthernet3/5 tunnel udlr address-resolution</pre>

Table 4-8 Corresponding Loopback and Tunnel Interfaces for AR1 (continued)

Loopback Interface	Tunnel Interface
<pre>interface Loopback12 description Endpoint for Tunnel12 ip address 10.10.10.14 255.255.255.255</pre>	<pre>interface Tunnel12 description Tx-side of Rx-only Gig3/6 no ip address tunnel source 10.10.10.14 tunnel destination 10.10.10.13 tunnel udldr send-only GigabitEthernet3/6 tunnel udldr address-resolution</pre>
<pre>interface Loopback24 description Endpoint for Tunnel24 ip address 10.10.10.25 255.255.255.255</pre>	<pre>interface Tunnel24 description Rx-side of Tx-only Gig4/4 no ip address tunnel source 10.10.10.25 tunnel destination 10.10.10.26 tunnel udldr receive-only GigabitEthernet4/4</pre>
<pre>interface Loopback28 description Endpoint for Tunnel28 ip address 10.10.10.29 255.255.255.255</pre>	<pre>interface Tunnel28 description Rx-side of Tx-only Gig4/5 no ip address tunnel source 10.10.10.29 tunnel destination 10.10.10.30 tunnel udldr receive-only GigabitEthernet4/5</pre>

Establishing an Interface to a DSLAM on AR1

Do the following to establish an interface to DSLAM1.

The following is configured on AR1.

Step 1 Establish a 1-GE trunk to and from the uplink 1-GE port on the DSLAM.

- a. Configure the trunk for 802.1q encapsulation, and assign the trunk to VLANs 90, 110, and 111.

```
interface GigabitEthernet5/1
description GigE trunk to/from DSLAM uplink GigE
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,110,111

switchport mode trunk
```

- b. Prevent unknown unicast traffic from being flooded to the port.

```
switchport block unicast
```

 **Note**

Occasionally, unknown unicast or multicast traffic is flooded to a switch port because a MAC address has timed out or has not been learned by the switch. (This condition is especially undesirable for a private VLAN isolated port.) To guarantee that no unicast or multicast traffic is flooded to the port, use the **switchport block unicast** or **switchport block multicast** commands.

- c. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Set transmit-queue bandwidth thresholds and priority.

```
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
```

- f. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- g. Enable PortFast on the interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast trunk
```

Step 2 Repeat Step 1 for all additional GE DSLAMs served by the switch.

Configuring OSPF Routing for Video and Voice Traffic on AR1

Refer to [Configuring OSPF Routing for Video and Voice Traffic on DER](#), page 4-71.

The following is configured on AR1.

Step 1 Define an OSPF routing process to route video traffic.

```
router ospf 100
router-id 2.2.2.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 10.10.10.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
maximum-paths 8
```

Step 2 Define a second OSPF process to route VoIP traffic.

```
router ospf 101
router-id 2.2.2.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan111
network 192.168.111.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
```

Step 3 Define a third OSPF process to route VoD transport traffic.

```
router ospf 102
router-id 2.2.2.3
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan110
network 192.168.110.0 0.0.0.255 area 0
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
```

Configuring Spanning Tree on AR1

See [Configuring Spanning Tree on DER, page 4-73](#).

The following is configured on AR1.

Step 1 Configure RTSP.

```
spanning-tree mode rapid-pvst
```

Step 2 Because the transport VLANs in the 10-GE ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 800, 808, 900, 908
```

Configuring AR2

This section addresses the configuration required on the switch labeled AR2 in [Figure 4-2 on page 4-53](#), to route multiple services from AR2 to DER, AR1, and AR3.

See [Configuring DNS Servers, page 4-2](#).

**Note**

A Cisco Catalyst 6509 can also be used, as it uses the same supervisor engine, line cards, and Cisco IOS code as the Cisco 7609 router.

This section addresses the following:

- [Configuring QoS on AR2](#)
- [Establishing and Configuring Interfaces on AR2](#)
- [Configuring OSPF Routing for Video and Voice Traffic on AR2](#)
- [Configuring Spanning Tree on AR2](#)

**Note**

For a complete configuration example, see [Appendix B, “Sample DER and AR Switch Configurations for the 1-GE Asymmetric Topology.”](#)

Configuring QoS on AR2

See [Overview of QoS on a Cisco 7600 Series and Cisco Catalyst 6500 Series, page 4-57](#).

This section presents the following topics:

- [Configuring Marking and Classification on AR2](#)
- [Configuring Mapping on AR2](#)

**Note**

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Configuring Marking and Classification on AR2

Do the following to enable marking and classification on AR2.

Step 1 Enable QoS in global configuration mode.

```
mls qos
```

Step 2 Create access lists to identify the different service types in the network.

```
ip access-list extended acl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip 192.168.120.0 0.0.0.255 192.168.10.102
 permit ip 192.168.120.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.121.0 0.0.0.255 any
```

Step 3 Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
```

Step 4 Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
```

Step 5 Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note

Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

Step 6 To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
mls qos trust dscp
```

Configuring Mapping on AR2

To configure mapping on AR2, refer to [Configuring Mapping on DER, page 4-59](#).

Establishing and Configuring Interfaces on AR2

Refer to [Figure 4-2 on page 4-53](#).

This section addresses the following:

- [Establishing VLANs for Services on AR2](#)
- [Establishing Bidirectional and Unidirectional Transport Interfaces on AR2](#)
- [Establishing Tunnels on AR2](#)
- [Establishing an Interface to a DSLAM on AR2](#)

Establishing VLANs for Services on AR2

Before bidirectional and unidirectional transport interfaces can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-2 on page 4-5](#).)

The following is configured on AR2.

**Note**

For additional details, see [Establishing VLANs for Services on DER, page 4-61](#).

- Step 1** In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```
vlan 90
name VLAN_90_HSD
```

- Step 2** Establish a VLAN for video at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 120
name VLAN_120_Video
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan120
description Video edge VLAN
ip address 192.168.120.1 255.255.255.0
no ip redirects
no ip unreachablees
```

- c. Enable PIM sparse mode. This is the aggregation VLAN for video traffic to the DSLAMs. Broadcast video is multicast addressed.

```
ip pim sparse-mode
```

- d. To ensure consistently fast PIM convergence times, statically join the aggregation VLAN for video at the AR to the multicast groups.

```
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

- f. Change the ARP timeout from the default.

```
arp timeout 250
```

**Note**

The default timeout for an entry in the ARP cache is 4 hours. The default timeout for an entry in the MAC address table is only 5 minutes. Because video traffic is mostly unidirectional, the MAC address table may not be refreshed within the 5-minute timeout. This causes video traffic to be flooded until the destination MAC address is found. To prevent this, reduce the ARP cache timeout to 250 seconds. This forces the switch to re-ARP for the entries in the ARP cache before the entries in the MAC address table time out, avoiding the disruptive behavior.

Step 3 Establish a VLAN for VoIP at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 121
name VLAN_121_VoIP
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan121
description VoIP edge VLAN
ip address 192.168.121.1 255.255.255.0
no ip redirects
no ip unreachable
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```

Step 4 Establish a VLAN for VoIP transport to and from AR1.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 808
name VLAN_808_VoIP_to/from_AR1
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan808
description VoIP transport VLAN to/from AR1
ip address 192.168.252.10 255.255.255.252
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```

- d. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

Step 5 Repeat Step 4, as appropriate, to establish a VLAN for VoIP transport to and from AR3.

```
vlan 812
name VLAN_812_VoIP_to/from_AR3
```

```
interface Vlan812
description VoIP transport VLAN to/from AR3
ip address 192.168.252.14 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Step 6 Establish a VLAN for video transport to and from AR1.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 908
name VLAN_912_Video_to/from_DER
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan908
description Video transport VLAN to/from AR1
ip address 192.168.254.10 255.255.255.252
```

- c. Enable PIM sparse mode.

```
ip pim sparse-mode
```

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Configure OSPF on the transport VLAN interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- Step 7** Repeat Step 6, as appropriate, to establish a VLAN for video transport to and from AR3.

```
vlan 912
name VLAN_908_Video_to/from_AR3

interface Vlan912
description Video transport VLAN to/from AR3
ip address 192.168.254.14 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
```

Establishing Bidirectional and Unidirectional Transport Interfaces on AR2

Bidirectional and unidirectional transport interfaces must be established between AR1 and DER and AR2.

The following is configured on AR2.



Note

For additional details, see to [Establishing Bidirectional and Unidirectional Interfaces for Transport on DER, page 4-66](#).

- Step 1** Establish bidirectional transport interfaces.

- a. Establish a bidirectional interface to and from AR1.

```
interface GigabitEthernet1/1
description Transport to/from AR1 (Gig4/3)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,808,908
switchport mode trunk
no ip address
load-interval 30
carrier-delay msec 0
```

- b. Proceed as in Step 1b through Step 2 of [Establishing Bidirectional and Unidirectional Interfaces for Transport on DER, page 4-66](#).

- c. Establish a bidirectional transport interface to AR3. Note the exception below.

```
interface GigabitEthernet1/5
  description Transport to/from AR3 (Gig4/3)
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,812,912
  switchport mode trunk
  dampening
  no ip address
  load-interval 30
  carrier-delay msec 0

  spanning-tree cost 10 <---See Note below
```



Note Note that the spanning-tree cost is set to 10 on AR2. This breaks the loop for VLAN 90 (Layer 2) between AR2 and AR3, rather than somewhere else.

- d. Proceed as in Step 1b through Step 2 of [Establishing Bidirectional and Unidirectional Interfaces for Transport on DER](#), page 4-66.

Step 2 Establish unidirectional receive-only transport interfaces.

- a. Establish a unidirectional receive-only transport interface to AR1. With the exceptions noted, the following is as in Step 1.

```
interface GigabitEthernet1/2
  description Transport from AR1 (Gig4/4)
  dampening
  ip address 192.168.253.26 255.255.255.252
```

- b. Configure OSPF on the interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```



Note Transmit-queue bandwidth thresholds and priority do not need to be applied to a receive-only interface.

- d. Disable the link-negotiation protocol on the port.

```
speed nonegotiate
```

- e. Mark the interface as receive only.

```
unidirectional receive-only
```

- f. Repeat Step 2a through Step 2e, as appropriate, for the remaining unidirectional receive-only transport interfaces. Interface configurations are shown below.

Second unidirectional transport from AR1

```
interface GigabitEthernet1/3
  description Transport from AR1 (Gig4/5)
  dampening
  ip address 192.168.253.30 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  carrier-delay msec 0
  speed nonegotiate
  wrp-queue bandwidth 64 255
  wrp-queue queue-limit 40 50
  wrp-queue random-detect min-threshold 2 50 100
  wrp-queue random-detect max-threshold 1 100 100
  wrp-queue random-detect max-threshold 2 50 100
  wrp-queue cos-map 1 1 0
  wrp-queue cos-map 2 1 1
  wrp-queue cos-map 2 2 2 3 4 6 7
  mls qos trust dscp
  unidirectional receive-only
```

Third unidirectional transport from AR3

```
interface GigabitEthernet1/6
  description Transport from AR3 (Gig4/4)
  dampening
  ip address 192.168.253.38 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  carrier-delay msec 0
  speed nonegotiate
  wrp-queue bandwidth 64 255
  wrp-queue queue-limit 40 50
  wrp-queue random-detect min-threshold 2 50 100
  wrp-queue random-detect max-threshold 1 100 100
  wrp-queue random-detect max-threshold 2 50 100
  wrp-queue cos-map 1 1 0
  wrp-queue cos-map 2 1 1
  wrp-queue cos-map 2 2 2 3 4 6 7
  mls qos trust dscp
  unidirectional receive-only
```

Fourth unidirectional transport from AR3

```

interface GigabitEthernet1/7
description Transport from AR3 (Gig4/5)
dampening
ip address 192.168.253.42 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional receive-only

```

Establishing Tunnels on AR2

See [Establishing Tunnels on DER, page 4-70](#).

[Table 4-9](#) lists the loopback interfaces and corresponding tunnel interfaces configured on AR2.

Table 4-9 Corresponding Loopback and Tunnel Interfaces for AR2

Loopback Interface	Tunnel Interface
<pre> interface Loopback24 description Endpoint for Tunnel24 ip address 10.10.10.26 255.255.255.255 </pre>	<pre> interface Tunnel24 description Tx-side of Rx-only Gig1/2 no ip address tunnel source 10.10.10.26 tunnel destination 10.10.10.25 tunnel udld send-only GigabitEthernet1/2 tunnel udld address-resolution </pre>
<pre> interface Loopback28 description Endpoint for Tunnel28 ip address 10.10.10.30 255.255.255.255 </pre>	<pre> interface Tunnel28 description Tx-side of Rx-only Gig1/3 no ip address tunnel source 10.10.10.30 tunnel destination 10.10.10.29 tunnel udld send-only GigabitEthernet1/3 tunnel udld address-resolution </pre>

Table 4-9 Corresponding Loopback and Tunnel Interfaces for AR2 (continued)

Loopback Interface	Tunnel Interface
<pre>interface Loopback36 description Endpoint for Tunnel36 ip address 10.10.10.38 255.255.255.255</pre>	<pre>interface Tunnel36 description Tx-side of Rx-only Gig1/6 no ip address tunnel source 10.10.10.38 tunnel destination 10.10.10.37 tunnel udld send-only GigabitEthernet1/6 tunnel udld address-resolution</pre>
<pre>interface Loopback40 description Endpoint for Tunnel40 ip address 10.10.10.42 255.255.255.255</pre>	<pre>interface Tunnel40 description Tx-side of Rx-only Gig1/7 no ip address tunnel source 10.10.10.42 tunnel destination 10.10.10.41 tunnel udld send-only GigabitEthernet1/7 tunnel udld address-resolution</pre>

Establishing an Interface to a DSLAM on AR2

DSLAM2 is an Ericsson GE DSLAM. Note the differences in Step 1f.

The following is configured on AR2.

Step 1 Establish a 1-GE trunk to and from the uplink 1-GE port on the DSLAM.

- a. Configure the trunk for 802.1q encapsulation, and assign the trunk to VLANs 90, 120, and 121.

```
interface GigabitEthernet2/1
description Ericsson DSLAM
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,120,121
switchport mode trunk

no ip address
```

- b. Prevent unknown unicast traffic from being flooded to the port.

```
switchport block unicast
```



Note Occasionally, unknown unicast or multicast traffic is flooded to a switch port because a MAC address has timed out or has not been learned by the switch. (This condition is especially undesirable for a private VLAN isolated port.) To guarantee that no unicast or multicast traffic is flooded to the port, use the **switchport block unicast** or **switchport block multicast** commands.

- c. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- f. Enable PortFast on the interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast trunk
```

- g. Proceed as in Step 1b through Step 2 of [Establishing Bidirectional and Unidirectional Interfaces for Transport on DER](#), page 4-66, but with the following exceptions:

```
wrr-queue bandwidth 64 255 0
```

```
wrr-queue queue-limit 40 50 0
```

```
wrr-queue threshold 1 100 100 100 100 100 100 100 100
```

```
wrr-queue threshold 2 50 100 100 100 100 100 100 100
```

```
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
```

```
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
```

```
no wrr-queue random-detect 2
```

```
wrr-queue cos-map 2 1 1
```

```
wrr-queue cos-map 2 2 2 3 4 6 7
```

- Step 2** Repeat Step 1 for additional Ericsson GE DSLAMs served by the switch.
-

Configuring OSPF Routing for Video and Voice Traffic on AR2

See [Configuring OSPF Routing for Video and Voice Traffic on DER](#), page 4-71.

The following is configured on AR2.

- Step 1** Define an OSPF routing process to route video traffic.

```
router ospf 100
router-id 3.3.3.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 10.10.10.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
maximum-paths 8
```

- Step 2** Define a second OSPF process to route VoIP traffic.

```
router ospf 101
router-id 3.3.3.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan121
network 192.168.121.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
```

Step 3 Define a third OSPF process to route VoD transport traffic.

```
router ospf 102
router-id 3.3.3.3
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan120
network 192.168.120.0 0.0.0.255 area 0
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
```

Configuring Spanning Tree on AR2

See [Configuring Spanning Tree on DER, page 4-20](#).

The following is configured on AR2.

Step 1 Configure RTSP.

```
spanning-tree mode rapid-pvst
```

Step 2 Because the transport VLANs in the 1-GE ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 800, 812, 908, 912
```

Configuring AR3

This section addresses the configuration required on the switch labeled AR3 in [Figure 4-1 on page 4-4](#), to route multiple services from AR3 to AR2 and DER.

See [Configuring DNS Servers, page 4-2](#).

This section addresses the following:

- [Configuring QoS on AR3](#)
- [Establishing and Configuring Interfaces on AR3](#)
- [Configuring OSPF Routing for Video and Voice Traffic on AR3](#)
- [Configuring Spanning Tree on AR3](#)



Note

For a complete configuration example, see [Appendix B, “Sample DER and AR Switch Configurations for the 1-GE Asymmetric Topology.”](#)

Configuring QoS on AR3

See [Overview of QoS on a Cisco Catalyst 4500 Series, page 4-74](#).

This section presents the following topics:

- [Configuring Marking and Classification on AR3](#)
- [Configuring Mapping on AR3](#)
- [Configuring Queuing on AR3](#)



Note

For more information specific to QoS as applied to the solution, see [Appendix C, “Understanding QoS as Implemented in the Solution.”](#)

Configuring Marking and Classification on AR3

Do the following to enable marking and classification on AR3.

Step 1 Enable QoS in global configuration mode.

```
qos
```

Step 2 Create access lists to identify the different service types in the network.

```
ip access-list extended acl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip 192.168.130.0 0.0.0.255 192.168.10.102
 permit ip 192.168.130.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.131.0 0.0.0.255 any
```

Step 3 Create class maps for the access lists created in Step 2.

```
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
```

Step 4 Create a policy map to set the DSCP values of the different classes created in Step 3.

```
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
```

Step 5 Apply the policy map from Step 4 to the ingress interfaces using the following command.

```
service-policy input setDSCP
```



Note

Specific interface examples of this and other interface commands are shown in the interface provisioning sections.

Step 6 To maintain the DSCP marking applied at the network ingress interface, configure all noningress transport interfaces to trust the incoming DSCP markings.

```
qos trust dscp
```

Configuring Mapping on AR3

To configure mapping on AR3, proceed as in [Configuring Mapping on AR1, page 4-76](#).

Configuring Queueing on AR3

To configure queueing on AR3, proceed as in [Configuring Queueing on AR1, page 4-77](#).

Establishing and Configuring Interfaces on AR3

Refer to [Figure 4-2 on page 4-53](#).

This section addresses the following:

- [Establishing VLANs for Services on AR3](#)
- [Establishing Bidirectional and Unidirectional Transport Interfaces on AR3](#)
- [Establishing Tunnels on AR3](#)
- [Establishing an Interface to a DSLAM on AR3](#)

Establishing VLANs for Services on AR3

Before bidirectional and unidirectional transport interfaces can be configured, VLANs for the various services must be created. With the exception of VLAN 90 (high-speed data), these are all Layer 3 VLANs. (Refer to [Table 4-2 on page 4-5](#).)

The following is configured on AR3.



Note

For details, see [Establishing VLANs for Services on DER, page 4-61](#).

- Step 1** In global configuration mode, establish a Layer 2 VLAN for high-speed data (HSD). (No Layer 3 interface is required.)

```
vlan 90
name VLAN_90_HSD
```

- Step 2** Establish a VLAN for video at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 130
name VLAN_130_Video
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan130
description Video edge VLAN
ip address 192.168.130.1 255.255.255.0
no ip redirects
no ip unreachable
```

- c. Enable PIM sparse mode. This is the aggregation VLAN for video traffic to the DSLAMs. Broadcast video is multicast addressed.

```
ip pim sparse-mode
```

- d. To ensure consistently fast PIM convergence times, statically join the aggregation VLAN for video at the AR to the multicast groups.

```
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
```

- e. Change the load interval from the default of 300.

```
load-interval 30
```

- f. Change the ARP timeout from the default.

```
arp timeout 250
```

**Note**

The default timeout for an entry in the ARP cache is 4 hours. The default timeout for an entry in the MAC address table is only 5 minutes. Because video traffic is mostly unidirectional, the MAC address table may not be refreshed within the 5-minute timeout. This causes video traffic to be flooded until the destination MAC address is found. To prevent this, reduce the ARP cache timeout to 250 seconds. This forces the switch to re-ARP for the entries in the ARP cache before the entries in the MAC address table time out, avoiding the disruptive behavior.

Step 3 Establish a VLAN for VoIP at the edge.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 131
name VLAN_131_VoIP
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan131
description VoIP edge VLAN
ip address 192.168.131.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30
```

Step 4 Establish a VLAN for VoIP transport to and from AR2.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 812
name VLAN_812_VoIP_to/from_AR2
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan812
description VoIP transport to/from AR2
ip address 192.168.252.13 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
```

Step 5 Repeat Step 4, as appropriate, to establish a VLAN for VoIP transport to and from DER.

```
vlan 816
name VLAN_808_VoIP_to/from_DER

interface Vlan816
description VoIP transport to/from DER
ip address 192.168.252.18 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
```

Step 6 Establish a VLAN for video transport to and from AR2.

- a. In global configuration mode, add the VLAN to the VLAN database.

```
vlan 912
name VLAN_912_Video_to/from_AR2
```

- b. In interface configuration mode, create and configure the VLAN interface.

```
interface Vlan912
description Video transport to/from AR2
ip address 192.168.254.13 255.255.255.252
```

- c. Enable PIM sparse mode.

```
ip pim sparse-mode
```

- d. Configure OSPF on the interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- Step 7** Repeat Step 6, as appropriate, to establish a VLAN for video transport to and from DER.

```
vlan 916
name VLAN_916_Video_to/from_DER

interface Vlan916
description Video transport to/from DER
ip address 192.168.254.18 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
```

Establishing Bidirectional and Unidirectional Transport Interfaces on AR3

Bidirectional and unidirectional transport interfaces must be established between AR1 and DER and AR2.

The following is configured on AR3.



Note

For additional details, see [Establishing Bidirectional and Unidirectional Interfaces for Transport on DER, page 4-66](#).

- Step 1** Establish bidirectional transport interfaces.

- a. Establish a bidirectional interface to and from DER.

```
interface GigabitEthernet3/1
description Transport to/from DER (Gig7/9)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,816,916
switchport mode trunk
load-interval 30
```

- b. Configure this interface (and all noningress transport interfaces) to trust the incoming DSCP markings. (This maintains the DSCP marking applied at the network ingress interface.)

```
qos trust dscp
```

- c. Set the transmit-queue bandwidth thresholds and priority.

```
tx-queue 1
bandwidth percent 19
tx-queue 2
bandwidth percent 80
tx-queue 3
priority high
tx-queue 4
bandwidth percent 1
```


- d. Repeat Step 1a through Step 1c, as appropriate, to establish the bidirectional transport interface to AR2. Note the exception below.

```
interface GigabitEthernet4/3
  description Transport to/from AR2 (Gig1/5)
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,812,912
  switchport mode trunk
  load-interval 30
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
  spanning-tree cost 10 <---See Note below
```



Note Note that the spanning-tree cost is set to 10 on AR3. This breaks the loop for VLAN 90 (Layer 2) between AR2 and AR3, rather than somewhere else.

Step 2 Establish unidirectional receive-only transport interfaces.

- a. Establish a unidirectional receive-only interface to DER. With the exceptions noted, the following is as in Step 1.

```
interface GigabitEthernet3/3
  description Transport from DER (Gig7/11)
  dampening
  ip address 192.168.253.50 255.255.255.252
```

- b. Configure OSPF on the interface.

```
ip ospf network point-to-point
ip ospf hello-interval 1
```

- c. Change the load interval from the default of 300.

```
load-interval 30
```



Note Transmit-queue bandwidth thresholds and priority do not need to be applied to a receive-only interface.

- d. Disable the link-negotiation protocol on the port.

```
speed nonegotiate
```

- e. Mark the interface as receive only.

```
unidirectional receive-only
```

- f. Repeat Step 2a through Step 2e, as appropriate, for the remaining unidirectional receive-only transport interfaces. Interface configurations are shown below.

Second unidirectional transport from DER

```
interface GigabitEthernet3/4
description Transport from DER (Gig7/12)
no switchport
dampening
ip address 192.168.253.54 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed negotiate
unidirectional receive-only
```

Third unidirectional transport from DER

```
interface GigabitEthernet3/5
description Transport from DER (Gig7/13)
no switchport
dampening
ip address 192.168.253.58 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed negotiate
qos trust dscp
unidirectional receive-only
```

Fourth unidirectional transport from DER

```
interface GigabitEthernet3/6
description Transport from DER (Gig7/14)
no switchport
dampening
ip address 192.168.253.62 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed negotiate
qos trust dscp
unidirectional receive-only
```

Step 3 Establish unidirectional send-only transport interfaces.

- a. Establish a unidirectional send-only transport interface to AR2. With the exception noted, the following is as in Step 1.

```
interface GigabitEthernet4/4
  description Transport to AR2 (Gig1/6)
  no switchport
  ip address 192.168.253.37 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  speed nonegotiate
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
```

- b. Mark the interface as send only.

```
unidirectional send-only
```

- c. Repeat Step 3a and Step 3b, as appropriate, for the second unidirectional send-only transport interface to AR2.

```
interface GigabitEthernet4/5
  description Transport to AR2 (Gig1/7)
  no switchport
  ip address 192.168.253.41 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
  speed nonegotiate
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
  unidirectional send-only
```

Establishing Tunnels on AR3

See [Establishing Tunnels on DER, page 4-70](#).

[Table 4-10 on page 4-108](#) lists the loopback interfaces and corresponding tunnel interfaces configured on AR3.

Table 4-10 Corresponding Loopback and Tunnel Interfaces for AR3

Loopback Interface	Tunnel Interface
<pre>interface Loopback36 description Endpoint for Tunnel36 ip address 10.10.10.37 255.255.255.255</pre>	<pre>interface Tunnel36 description Rx-side of Tx-only Gig4/4 no ip address tunnel source 10.10.10.37 tunnel destination 10.10.10.38 tunnel udld receive-only GigabitEthernet4/4</pre>
<pre>interface Loopback40 description Endpoint for Tunnel40 ip address 10.10.10.41 255.255.255.255</pre>	<pre>interface Tunnel40 description Rx-side of Tx-only Gig4/5 no ip address tunnel source 10.10.10.41 tunnel destination 10.10.10.42 tunnel udld receive-only GigabitEthernet4/5</pre>
<pre>interface Loopback48 description Endpoint for Tunnel48 ip address 10.10.10.50 255.255.255.255</pre>	<pre>interface Tunnel48 description Tx-side of Rx-only Gig3/3 no ip address tunnel source 10.10.10.50 tunnel destination 10.10.10.49 tunnel udld send-only GigabitEthernet3/3 tunnel udld address-resolution</pre>
<pre>interface Loopback52 description Endpoint for Tunnel52 ip address 10.10.10.54 255.255.255.255</pre>	<pre>interface Tunnel52 description Tx-side of Rx-only Gig3/4 no ip address tunnel source 10.10.10.54 tunnel destination 10.10.10.53 tunnel udld send-only GigabitEthernet3/4 tunnel udld address-resolution</pre>
<pre>interface Loopback56 description Endpoint for Tunnel56 ip address 10.10.10.58 255.255.255.255</pre>	<pre>interface Tunnel56 description Tx-side of Rx-only Gig3/5 no ip address tunnel source 10.10.10.58 tunnel destination 10.10.10.57 tunnel udld send-only GigabitEthernet3/5 tunnel udld address-resolution</pre>
<pre>interface Loopback60 description Endpoint for Tunnel60 ip address 10.10.10.62 255.255.255.255</pre>	<pre>interface Tunnel60 description Tx-side of Rx-only Gig3/6 no ip address tunnel source 10.10.10.62 tunnel destination 10.10.10.61 tunnel udld send-only GigabitEthernet3/6 tunnel udld address-resolution</pre>

Establishing an Interface to a DSLAM on AR3

Do the following to establish an interface to DSLAM3.

The following is configured on AR3.

Step 1 Establish a 1-GE trunk to and from the uplink 1-GE port on the DSLAM.

- a. Configure the trunk for 802.1q encapsulation, and assign the trunk to VLANs 90, 110, and 111.

```
interface GigabitEthernet5/1
  description GigE trunk to/from DSLAM uplink GigE
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,130,131

  switchport mode trunk
```

- b. Prevent unknown unicast traffic from being flooded to the port.

```
switchport block unicast
```



Note Occasionally, unknown unicast or multicast traffic is flooded to a switch port because a MAC address has timed out or has not been learned by the switch. (This condition is especially undesirable for a private VLAN isolated port.) To guarantee that no unicast or multicast traffic is flooded to the port, use the **switchport block unicast** or **switchport block multicast** commands.

- c. Apply the “setDSCP” service policy to mark DSCP values in the inbound IP packets.

```
service-policy input setDSCP
```

- d. Change the load interval from the default of 300.

```
load-interval 30
```

- e. Set transmit-queue bandwidth thresholds and priority.

```
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
```

- f. Disable Cisco Discovery Protocol (CDP) on the interface.

```
no cdp enable
```

- g. Enable PortFast on the interface to bypass the listening and learning states in Spanning Tree Protocol (STP). This allows the interface to move immediately from the blocking state to the forwarding state, rather than waiting for STP to converge.

```
spanning-tree portfast trunk
```

Step 2 Repeat Step 1 for all additional GE DSLAMs served by the switch.

Configuring OSPF Routing for Video and Voice Traffic on AR3

Refer to [Configuring OSPF Routing for Video and Voice Traffic on DER](#), page 4-71.

The following is configured on AR3.

Step 1 Define an OSPF routing process to route video traffic.

```
router ospf 100
router-id 4.4.4.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 10.10.10.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
maximum-paths 8
```

Step 2 Define a second OSPF process to route VoIP traffic.

```
router ospf 101
router-id 4.4.4.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan131
network 192.168.131.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
```

Step 3 Define a third OSPF process to route VoD traffic.

```
router ospf 102
router-id 4.4.4.3
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan130
network 192.168.130.0 0.0.0.255 area 0
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
```

Configuring Spanning Tree on AR3

See [Configuring Spanning Tree on DER, page 4-20](#).

The following is configured on AR3.

Step 1 Configure RTSP.

```
spanning-tree mode rapid-pvst
```

Step 2 Because the transport VLANs in the 1-GE ring are point-to-point networks, there is no risk of Layer 2 loops, so STP can be disabled on these VLANs.

```
no spanning-tree vlan 812, 816, 912, 916
```



Monitoring and Troubleshooting

This chapter provides an introduction to monitoring and troubleshooting the Cisco Ethernet switches in the Cisco Gigabit-Ethernet Optimized IPTV/Video over Broadband (GOVoBB) Solution, Release 1.0.

The following major topics are presented:

- [Network Time Protocol \(NTP\), page 5-1](#)
- [Syslog, page 5-2](#)
- [Quality of Service \(QoS\), page 5-4](#)
- [Multicast, page 5-11](#)
- [UDLR and Unidirectional Links, page 5-16](#)
- [References, page 5-18](#)

Network Time Protocol (NTP)

It is important to ensure that all devices in the network are accurately synchronized to the same time source. This allows network events to be correlated (for example, for accounting, event logging, fault analysis, security incident response, and network management). The Network Time Protocol (NTP), RFC 1305, synchronizes timekeeping among a set of distributed time servers and clients.



Note

There are a number of ways to configure NTP, and describing NTP completely is beyond the scope of this document. A number of resources are available on Cisco.com and the Internet regarding NTP configuration.

At a minimum, the Cisco switches should be configured as NTP clients for a reliable time source, by means of the following commands:

```
clock timezone PST -8
clock summer-time PDT recurring 1 Sun Apr 3:00 last Sun Oct 3:00

clock calendar-valid
ntp server <NTP server IP address>
ntp update-calendar
```

Syslog

Cisco IOS Software has the capability to do UNIX system logging (syslog) to a UNIX syslog server. The Cisco UNIX syslog format is compatible with 4.3 Berkeley Standard Distribution (BSD) UNIX. System logging is useful for monitoring interface status, security alerts, environmental conditions, CPU processes, and many other events on the router can be captured and analyzed by means of UNIX syslog. Management platforms such as Cisco Resource Manager Essentials (RME) and Network Analysis Toolkit (NATKit) make powerful use of syslog information to collect inventory and configuration changes.

The following is a summary and description of the recommended IOS configuration for syslog.

Global Syslog Configuration

Configure the following in global configuration mode:

```
no logging console
no logging monitor
logging buffered 16384
logging trap notifications
logging facility local7
logging host-ip
logging source-interface loopback 0
service timestamps debug datetime localtime show-timezone msec
service timestamps log datetime localtime show-timezone msec
```

Interface Syslog Configuration

Configure the following in interface configuration mode on interfaces of interest:

```
logging event link-status
logging event bundle-status
```

Useful Syslog Commands

The following syslog commands are particularly useful:

- [no logging console](#)
- [no logging monitor](#)
- [logging buffered 16384](#)
- [logging trap notifications](#)
- [logging facility local7](#)
- [logging host](#)
- [logging source-interface loopback 0](#)
- [service timestamps debug datetime localtime show-timezone msec](#)
- [logging event](#)

no logging console

By default, all system messages are sent to the system console. Console logging is a high-priority task in Cisco IOS Software. This function was primarily designed to generate error messages to the system operator prior to a system failure. It is recommended that console logging be disabled in all device configurations to avoid a situation where the router/switch might hang while waiting for a response from a terminal. Console messages can, however, be useful during trouble isolation. In these instances, console logging should be enabled by means of the **logging console level** command, to obtain the desired level of message logging. Logging levels range from 0 to 7.

no logging monitor

This command disables logging for terminal lines other than the system console. If monitor logging is required (by means of **logging monitor debugging** or another command option), it should be enabled at the specific logging level required for the activity (see above).

logging buffered 16384

The **logging buffered** command should be added to log system messages in the internal log buffer. The logging buffer is circular. Once the logging buffer is filled, older entries are overwritten by newer entries. The size of the logging buffer is user-configurable and is specified in bytes. The size of the system buffer varies by platform. 16384 is a good default and should provide adequate logging in most cases.

logging trap notifications

This command provides notification (level 5) messaging to the specified syslog server. The default logging level for all devices (console, monitor, buffer, and traps) is debugging (level 7). Leaving the trap logging level at 7 produces many extraneous messages that are of little or no concern to the health of the network. It is recommended that the default logging level for traps be set to 5.

logging facility local7

This command sets the default logging facility/level for UNIX system logging. The syslog server receiving these messages should be configured for the same facility/level.

logging host

This command sets the IP address of the UNIX syslog server.

logging source-interface loopback 0

This command sets the default IP source address for the syslog messages. Hard coding the logging source address makes it easier to identify the host that sent the message.

service timestamps debug datetime localtime show-timezone msec

By default, log messages are not time stamped. Use this command to enable the time stamping of log messages and configure the time stamping of system debug messages. Time stamping enhances real-time debugging by providing the relative timing of logged events. This information is especially useful when

customers send debugging output to technical support personnel for assistance. To enable the time stamping of system debug messages, use the above command in global configuration mode. This only has an affect when debugging is enabled.

logging event

The **logging event link-status** command enables logging related to link status. The **logging event bundle-status** command enables logging related to bundle status.

Quality of Service (QoS)

The following commands are useful in troubleshooting QoS:

- [show class-map](#)
- [show policy-map](#)
- [show qos maps](#)
- [show mls qos maps dscp-cos](#)
- [show qos interface](#)
- [show queueing interface](#)

show class-map

To verify the class map for QoS classification, use the **show class-map** command.

```
DER# show class-map

Class Map match-all class_VoIP (id 1)
  Match access-group name acl_VoIP

Class Map match-any class-default (id 0)
  Match any

Class Map match-all class_video_VoD_high (id 2)
  Match access-group name acl_video_VoD_high

Class Map match-all class_video_VoD_low (id 3)
  Match access-group name acl_video_VoD_low

Class Map match-all class_video_broadcast (id 4)
  Match access-group name acl_video_broadcast

Class Map match-all class_VoD_signaling (id 5)
  Match access-group name acl_VoD_signaling

Class Map match-all class_HSD (id 6)
  Match access-group name acl_HSD
```

show policy-map

To verify the policy map for QoS marking, use the **show policy-map** command.

```
DER# show policy-map
```

```

Policy Map setDSCP
  Description: Mark DSCP values for ingress traffic
  Class class_VoIP
    set dscp ef
  Class class_HSD
    set dscp default
  Class class_VoD_signaling
    set dscp cs3
  Class class_video_broadcast
    set dscp af41
  Class class_video_VoD_high
    set dscp af42
  Class class_video_VoD_low
    set dscp af43

```

show qos maps

On Cisco Catalyst 4500 and Cisco Catalyst 4948-10GE switches, use the **show qos maps** command to verify the DSCP-to-TxQueue and DSCP-to-CoS mappings.

```
AR2# 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

```

<omitted DSCP policing table>

```

DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 02 04 01 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07

```

<omitted CoS to DSCP mapping table>

show mls qos maps dscp-cos

On the Cisco Catalyst 6500 and Cisco 7600 switches, use the **show mls qos maps dscp-cos** command to verify the DSCP-to-CoS mappings.

```
DER# show mls qos maps dscp-cos
```

```

Dscp-cos map:                                     (dscp= d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01

```

```

1 : 01 01 01 01 01 01 02 02 02 02
2 : 02 02 02 02 03 03 03 03 03 03
3 : 03 03 04 04 04 04 02 04 01 04
4 : 05 05 05 05 05 05 05 05 06 06
5 : 06 06 06 06 06 06 07 07 07 07
6 : 07 07 07 07

```

show qos interface

On the Cisco Catalyst 4500 and Cisco Catalyst 4948-10GE switches, use the **show qos interface type slot/module** to verify the QoS state, port trust state, queue bandwidth, priority queue, and queue size.

```
AR2# show qos interface tenGigabitEthernet 1/1
```

```

QoS is enabled globally
Port QoS is enabled
Administrative Port Trust State: 'dscp'
Operational Port Trust State: 'dscp'
Trust device: none
Default DSCP: 0 Default CoS: 0
Appliance trust: none
Tx-Queue   Bandwidth   ShapeRate   Priority   QueueSize
           (bps)         (bps)
1           1900000000   disabled    N/A       2080
2           8000000000   disabled    N/A       2080
3           2500000000   disabled    high      2080
4           1000000000   disabled    N/A       2080

```

show queueing interface

On the Cisco Catalyst 6500 and Cisco 7600 switches, use the **show queueing interface type slot/module** command to verify the queueing strategy, priority queue, WRR bandwidths, queue sizes, thresholds, CoS-to-queue mappings, and queue drops.

```
DER# show queueing interface tenGigabitEthernet 7/1
```

```

Interface TenGigabitEthernet7/1 queueing strategy: Weighted Round-Robin
Port QoS is enabled
Trust state: trust DSCP
Extend trust state: not trusted [COS = 0]
Default COS is 0
Queueing Mode In Tx direction: mode-cos
Transmit queues [type = 1p7q8t]:
Queue Id    Scheduling  Num of thresholds
-----
01          WRR         08
02          WRR         08
03          WRR         08
04          WRR         08
05          WRR         08
06          WRR         08
07          WRR         08
08          Priority    01

WRR bandwidth ratios: 64[queue 1] 255[queue 2] 0[queue 3] 0[queue 4] 0[queue
5] 0[queue
e 6] 0[queue 7]
queue-limit ratios: 40[queue 1] 50[queue 2] 0[queue 3] 0[queue 4] 0[queue
5] 0[queue

```

```
e 6] 0[queue 7]
```

```
queue tail-drop-thresholds
```

```
-----
1   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2   45[1] 85[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
4   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
5   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
6   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
7   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
```

```
queue random-detect-min-thresholds
```

```
-----
1   75[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2   40[1] 70[2] 70[3] 70[4] 70[5] 70[6] 70[7] 70[8]
3   70[1] 70[2] 70[3] 70[4] 70[5] 70[6] 70[7] 70[8]
4   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
5   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
6   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
7   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
```

```
queue random-detect-max-thresholds
```

```
-----
1   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2   70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
4   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
5   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
6   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
7   100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
```

```
WRED disabled queues:      2 4 5 6 7
```

```
queue thresh cos-map
```

```
-----
1   1       0
1   2
1   3
1   4
1   5
1   6
1   7
1   8
2   1       1
2   2       2
2   3       3 4 6 7
2   4
2   5
2   6
2   7
2   8
3   1
3   2
3   3
3   4
3   5
3   6
3   7
3   8
4   1
4   2
4   3
4   4
```

```

4      5
4      6
4      7
4      8
5      1
5      2
5      3
5      4
5      5
5      6
5      7
5      8
6      1
6      2
6      3
6      4
6      5
6      6
6      7
6      8
7      1
7      2
7      3
7      4
7      5
7      6
7      7
7      8
8      1      5

```

```

Queueing Mode In Rx direction: mode-cos
Receive queues [type = 8q8t]:
Queue Id      Scheduling  Num of thresholds
-----

```

```

01      WRR      08
02      WRR      08
03      WRR      08
04      WRR      08
05      WRR      08
06      WRR      08
07      WRR      08
08      WRR      08

```

```

WRR bandwidth ratios: 100[queue 1]  0[queue 2]  0[queue 3]  0[queue 4]  0[queue
5]  0[queue
e 6]  0[queue 7]  0[queue 8]
queue-limit ratios:  100[queue 1]  0[queue 2]  0[queue 3]  0[queue 4]  0[queue
5]  0[queue
e 6]  0[queue 7]  0[queue 8]

```

```

queue tail-drop-thresholds
-----

```

```

1      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
4      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
5      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
6      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
7      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
8      100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]

```

```

queue random-detect-min-thresholds
-----

```

```

1      40[1] 40[2] 50[3] 50[4] 50[5] 50[6] 50[7] 50[8]

```



```

2    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
4    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
5    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
6    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
7    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
8    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]

```

```
queue random-detect-max-thresholds
```

```

-----
1    70[1] 80[2] 90[3] 100[4] 100[5] 100[6] 100[7] 100[8]
2    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
3    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
4    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
5    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
6    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
7    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
8    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]

```

```
WRED disabled queues:      1 2 3 4 5 6 7 8
```

```
queue thresh cos-map
```

```

-----
1    1      0 1 2 3 4 5 6 7
1    2
1    3
1    4
1    5
1    6
1    7
1    8
2    1
2    2
2    3
2    4
2    5
2    6
2    7
2    8
3    1
3    2
3    3
3    4
3    5
3    6
3    7
3    8
4    1
4    2
4    3
4    4
4    5
4    6
4    7
4    8
5    1
5    2
5    3
5    4
5    5
5    6
5    7
5    8
6    1

```

```

6      2
6      3
6      4
6      5
6      6
6      7
6      8
7      1
7      2
7      3
7      4
7      5
7      6
7      7
7      8
8      1
8      2
8      3
8      4
8      5
8      6
8      7
8      8

```

Packets dropped on Transmit:

```

queue      dropped  [cos-map]
-----
1           0  [0 ]
2           0  [1 2 3 4 6 7 ]
3           0  []
4           0  []
5           0  []
6           0  []
7           0  []
8           0  [5 ]

```

Packets dropped on Receive:

```

queue      dropped  [cos-map]
-----
1           0  [0 1 2 3 4 5 6 7 ]
2           0  []
3           0  []
4           0  []
5           0  []
6           0  []
7           0  []
8           0  []

```

Multicast

The following commands are useful in troubleshooting multicast:

- `show ip mroute`
- `show ip mroute ssm`
- `show ip mroute active`
- `show ip pim neighbor`
- `show ip igmp snooping`
- `show ip igmp groups`
- `show ip igmp ssm-mapping`
- `show ip igmp membership`
- `debug ip igmp`
- `debug ip pim`
- `debug domain`

show ip mroute

To see the details of the multicast routing table, use the **show ip mroute** command. The output of this command also shows the legend for the flags.

```
AR3# show ip mroute

IP Multicast Routing Table

Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report, Z - Multicast Tunnel
Y - Joined MDT-data group, y - Sending to MDT-data group

Outgoing interface flags: H - Hardware switched, A - Assert winner

Timers: Uptime/Expires

Interface state: Interface, Next-Hop or VCD, State/Mode

(192.168.70.101, 232.1.5.220), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
Vlan130, Forward/Sparse, 1w1d/00:01:46, H
(192.168.70.101, 232.1.5.221), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
```

```

Outgoing interface list:
Vlan130, Forward/Sparse, 1w1d/00:01:46, H
(192.168.70.101, 232.1.5.222), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
Vlan130, Forward/Sparse, 1w1d/00:01:46, H
(192.168.70.101, 232.1.5.223), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:

```

show ip mroute ssm

To verify the source-specific multicast (SSM) mapping of multicast groups to multicast sources, use the **show ip mroute ssm** command. With this command, you can also verify the path of the multicast ingress and egress interface(s).



Tip

To see the legend for the flags field, you must use the **show ip mroute** command.

```

AR2# show ip mroute ssm

(192.168.70.101, 232.1.5.220), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
Vlan130, Forward/Sparse, 1w1d/00:01:40, H
(192.168.70.101, 232.1.5.221), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
Vlan130, Forward/Sparse, 1w1d/00:01:40, H
(192.168.70.101, 232.1.5.222), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
Vlan130, Forward/Sparse, 1w1d/00:01:40, H
(192.168.70.101, 232.1.5.223), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:

```

```
Vlan130, Forward/Sparse, 1w1d/00:01:40, H
(192.168.70.101, 232.1.5.216), 1w4d/stopped, flags: sTI
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
```

show ip mroute active

To verify the bitrate of a multicast group, use the **show ip mroute active** command.

```
AR2# show ip mroute active

Active IP Multicast Sources - sending >= 4 kbps

Group: 232.255.0.1, (?)
  Source: 192.168.71.105 (1.0.255.232.coronado.net)
    Rate: 334 pps/3517 kbps(1sec), 2829 kbps(last 30 secs), 2703 kbps(life avg)

<rest of the output omitted>
```

show ip pim neighbor

To verify the protocol-independent multicast (PIM) neighbors, use the **show ip pim neighbor** command.

```
AR2# show ip pim neighbor

PIM Neighbor Table
Neighbor          Interface          Uptime/Expires    Ver   DR
Address
192.168.254.9     Vlan908            1d16h/00:01:34    v2    1 / S
192.168.254.18   Vlan916            1d16h/00:01:24    v2    1 / DR S
```

show ip igmp snooping

To verify IGMP snooping on the switch and interfaces, use the **show ip igmp snooping** command.

```
AR2# show ip igmp snooping

Global IGMP Snooping configuration:
-----
IGMP snooping           : Enabled
IGMPv3 snooping         : Enabled
Report suppression     : Enabled
TCN solicit query      : Disabled
TCN flood query count   : 2

Vlan 1:
-----
IGMP snooping           : Enabled
IGMPv2 immediate leave : Disabled
Explicit host tracking   : Enabled
Multicast router learning mode : pim-dvmrp
CGMP interoperability mode : IGMP_ONLY

Vlan 70:
-----
```

```

IGMP snooping                : Enabled
IGMPv2 immediate leave       : Disabled
Explicit host tracking        : Enabled
Multicast router learning mode : pim-dvmrp
CGMP interoperability mode    : IGMP_ONLY

```

<rest of output omitted>

show ip igmp groups

To verify IGMP group membership on a switch, use the **show ip igmp groups** command.

```
AR2# show ip igmp groups
```

```

IGMP Connected Group Membership
Group Address      Interface      Uptime    Expires    Last Reporter
232.255.0.1       Vlan120       1d17h     stopped    0.0.0.0
232.255.0.2       Vlan120       1d17h     stopped    0.0.0.0
232.255.0.3       Vlan120       1d17h     stopped    0.0.0.0
232.255.0.5       Vlan120       1d17h     stopped    0.0.0.0
232.255.0.12      Vlan120       1d17h     stopped    0.0.0.0
224.0.1.40        Vlan120       1d16h     00:02:56  192.168.120.1

```

show ip igmp ssm-mapping

To verify the SSM mapping configuration on the switch, use the **show ip igmp ssm-mapping** command.

```
AR3# show ip igmp ssm-mapping
```

```

SSM Mapping   : Enabled
DNS Lookup    : Enabled
Mcast domain  : coronado.net
Name servers   : 192.168.11.101

```

show ip igmp membership

Another command to verify IGMP group membership, which provides some additional information compared to the previous command, is the **show ip igmp membership** command.

```
AR2# show ip igmp membership
```

```

Flags: A - aggregate, T - tracked
       L - Local, S - static, V - virtual, R - Reported through v3
       I - v3lite, U - Urd, M - SSM (S,G) channel
       1,2,3 - The version of IGMP, the group is in
Channel/Group-Flags:
       / - Filtering entry (Exclude mode (S,G), Include mode (G))
Reporter:
       <mac-or-ip-address> - last reporter if group is not explicitly tracked
       <n>/<m>              - <n> reporter in include mode, <m> reporter in exclude

Channel/Group      Reporter      Uptime    Exp.  Flags  Interface
/*,232.255.0.1     0.0.0.0      1d17h     stop  2MA   V1120
192.168.71.105,232.255.0.1
/*,232.255.0.2     0.0.0.0      1d17h     stop  2MA   V1120
192.168.71.105,232.255.0.2
/*,232.255.0.3     0.0.0.0      1d17h     stop  2MA   V1120
192.168.71.105,232.255.0.3

```

```

/*,232.255.0.5          0.0.0.0      1d17h   stop 2MA  V1120
192.168.71.105,232.255.0.5  1d17h   stop SA  V1120
/*,232.255.0.12       0.0.0.0      1d17h   stop 2MA  V1120
192.168.71.105,232.255.0.12  1d17h   stop SA  V1120
*,224.0.1.40         192.168.120.1 1d16h   02:22 2LA  V1120

```

debug ip igmp

To troubleshoot IGMP issues, use the **debug ip igmp** command. The debug output indicates IGMP membership queries, membership responses, and the conversion of IGMPv2 to IGMPv3 through DNS lookup.

```
AR2# debug ip igmp
```

```

IGMP debugging is on
AR2#
*Aug  8 14:20:53.039: IGMP(0): Received v2 Query on Vlan908 from 192.168.254.9
AR2#
*Aug  8 14:21:16.880: IGMP(0): Send v2 general Query on Vlan120
*Aug  8 14:21:16.880: IGMP(0): Set report delay time to 8.4 seconds for 224.0.1.40 on
Vlan120
*Aug  8 14:21:16.880: IGMP(0): Send v2 general Query on Vlan916
AR2#
*Aug  8 14:21:25.881: IGMP(0): Send v2 Report for 224.0.1.40 on Vlan120
*Aug  8 14:21:25.881: IGMP(0): Received v2 Report on Vlan120 from 192.168.120.1 for
224.0.1.40
*Aug  8 14:21:25.881: IGMP(0): Received Group record for group 224.0.1.40, mode 2 from
192.168.120.1
    for 0 sources
*Aug  8 14:21:25.881: IGMP(0): Updating EXCLUDE group timer for 224.0.1.40
*Aug  8 14:21:25.881: IGMP(0): MRT Add/Update Vlan120 for (*,224.0.1.40) by 0
AR2#
*Aug  8 14:21:39.089: IGMP(0): Convert IGMPv2 static (*, 232.255.0.1) to IGMPv3 with 1
source(s) using DNS
*Aug  8 14:21:39.089: IGMP(0): Convert IGMPv2 static (*, 232.255.0.2) to IGMPv3 with 1
source(s) using DNS
*Aug  8 14:21:39.089: IGMP(0): Convert IGMPv2 static (*, 232.255.0.3) to IGMPv3 with 1
source(s) using DNS
*Aug  8 14:21:39.089: IGMP(0): Convert IGMPv2 static (*, 232.255.0.5) to IGMPv3 with 1
source(s) using DNS
*Aug  8 14:21:39.089: IGMP(0): Convert IGMPv2 static (*, 232.255.0.12) to IGMPv3 with 1
source(s) using DNS

```

debug ip pim

To troubleshoot PIM issues, use the **debug ip pim** command. The output indicates join and prune messages for PIM.

```
AR2# debug ip pim
```

```

PIM debugging is on
AR2#
*Aug  8 14:23:04.149: PIM(0): Building Periodic Join/Prune message for 232.255.0.1
*Aug  8 14:23:04.149: PIM(0): Insert (192.168.71.105,232.255.0.1) join in nbr
192.168.254.18's queue
*Aug  8 14:23:04.149: PIM(0): Building Join/Prune packet for nbr 192.168.254.18
*Aug  8 14:23:04.149: PIM(0): Adding v2 (192.168.71.105/32, 232.255.0.1), S-bit Join
*Aug  8 14:23:04.149: PIM(0): Send v2 join/prune to 192.168.254.18 (Vlan916)

```

debug domain

To troubleshoot domain name server (DNS) lookup issues, use the **debug domain** command.

```
AR2# debug domain
```

```
Aug  8 21:28:34.274: Domain: query for 1.0.255.232.coronado.net type 1 to 192.168.11.101
Aug  8 21:28:34.274: DOM: dom2cache: hostname is 1.0.255.232.coronado.net, RR type=1,
class=1, ttl=43200, n=4
Reply received ok
```

UDLR and Unidirectional Links

The following commands are useful in troubleshooting unidirectional link routing (UDLR) and unidirectional links:

- [show interface tunnel](#)
- [show interface](#)
- [debug tunnel](#)

show interface tunnel

To verify the tunnel for a unidirectional link, use the **show interface tunnel** command. From the output, you can verify the state of the tunnel (up), the encapsulation (GRE), the tunnel source and destination, the associated unidirectional link, and the number of packets through the tunnel.

```
AR1# show interfaces tunnel 0
```

```
Tunnel0 is up, line protocol is up
Hardware is Tunnel
Description: Tx-side of Rx-only Gig3/3
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 10.10.10.2, destination 10.10.10.1, fastswitch TTL 255
Tunnel is send-only UDLR tunnel for receive-only interface GigabitEthernet3/3
Tunnel protocol/transport GRE/IP, key disabled, sequencing disabled
Tunnel TTL 255
Checksumming of packets disabled, fast tunneling enabled
Last input never, output 00:00:00, output hang never
Last clearing of "show interface" counters 1w0d
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
 681125 packets output, 71116160 bytes, 0 underruns
  0 output errors, 0 collisions, 0 interface resets
  0 output buffer failures, 0 output buffers swapped out
```


show interface

To verify a unidirectional link, use the **show interface** *type slot/module* command. From the output, you can verify the state of the interface (up at Layer 1 and Layer 2), the bitrate, output drops, and errors.

```
AR1# show interfaces gigabitEthernet 3/3

GigabitEthernet3/3 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet Port, address is 000c.850e.80bf (bia 000c.850e.80bf)
  Description: Transport to/from L3VPN-PE (Gig7/3)
  Internet address is 192.168.253.2/30
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Carrier delay is 0 msec
  Transmit interface is Tunnel0
  Full-duplex, 1000Mb/s, link type is force-up, media type is 1000BaseSX
  input flow-control is on, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output never, output hang never
  Last clearing of "show interface" counters 1w0d
  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 1000 bits/sec, 1 packets/sec
  30 second output rate 0 bits/sec, 0 packets/sec
  L3 in Switched: ucast: 34536 pkt, 2869265 bytes - mcast: 0 pkt, 0 bytes
  L3 out Switched: ucast: 0 pkt, 0 bytes - mcast: 0 pkt, 0 bytes
  729329 packets input, 75599071 bytes, 0 no buffer
  Received 694750 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
  0 packets output, 0 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
```

debug tunnel

To troubleshoot the GRE tunnel for UDLR, use the **debug tunnel** command. From the output, you can see packets encapsulated in GRE and sent from a source loopback interface to a destination loopback interface.

```
AR1# debug tunnel

Tunnel Interface debugging is on
AR1#
Aug  8 15:09:26.770: Tunnel8: GRE/IP encapsulated 10.10.10.10->10.10.10.9 (linktype=7,
len=104)
Aug  8 15:09:26.770: Tunnel12: GRE/IP encapsulated 10.10.10.14->10.10.10.13 (linktype=7,
len=104)
Aug  8 15:09:26.770: Tunnel16: GRE/IP encapsulated 10.10.10.18->10.10.10.17 (linktype=7,
len=104)
Aug  8 15:09:26.946: Tunnel14: GRE/IP encapsulated 10.10.10.6->10.10.10.5 (linktype=7,
len=104)
Aug  8 15:09:26.946: Tunnel20: GRE/IP encapsulated 10.10.10.22->10.10.10.21 (linktype=7,
len=104)
```

```
Aug  8 15:09:27.022: Tunnel28: GRE/IP to classify 10.10.10.30->10.10.10.29 (len=104  
ttl=254 tos=0xC0)
```

References

The following documents provide practical tips on configuring the switches used in the solution.

- *Best Practices for Catalyst 6500/6000 Series and Catalyst 4500/4000 Series Switches Running Cisco IOS Software*, at the following URL:
http://www.cisco.com/en/US/partner/products/hw/switches/ps700/products_white_paper09186a00801b49a4.shtml#cg24
- *Cisco ISP Essentials: Essential IOS Features Every ISP Should Consider*, by Barry Green and Philip Smith, at the following URL:
<http://wwwin-cons.cisco.com/~philsmi/isp/workshop/afnog2004/inet2000/adv-bgp/iosess29.pdf>



Note

A Cisco Connection Online (CCO) password may be required to access these documents.



Sample DER and AR Switch Configurations for the 10-GE Symmetric Topology

This appendix presents sample distribution edge router (DER) and aggregation router (AR) switch configurations for the symmetric 10-GE topology described in [Configuration 1: 10-GE Layer 3 Symmetric Ring, page 3-33](#). The following configurations are presented:

- [Configuration for DER, page A-1](#)
- [Configuration for AR1, page A-10](#)
- [Configuration for AR2, page A-16](#)
- [Configuration for AR3, page A-21](#)



Note

Refer to [Configuring the 10-GE Symmetric Topology, page 4-4](#).

Configuration for DER

```
upgrade fpd auto
version 12.2
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
service counters max age 10
!
hostname DER
!
boot system disk1:s72033-adventerprisek9_wan-mz.122-18.SXE1.bin
logging snmp-authfail
enable password cisco123
!
no aaa new-model
clock timezone PST -8
clock summer-time PDT recurring
ip subnet-zero
!
!
!
ip cef load-sharing algorithm original
ip multicast-routing
ip igmp ssm-map enable
ip domain multicast conronado.net
no ip domain-lookup
```



```

    name VLAN_800_VoIP_to/from_AR1
    !
vlan 824
    name VLAN_824_VoIP_to/from_AR3
    !
vlan 900
    name VLAN_900_Video_to/from_AR1
    !
vlan 924
    name VLAN_924_Video_to/from_AR3
    !
class-map match-all class_VoIP
    match access-group name acl_VoIP
class-map match-all class_video_VoD_high
    match access-group name acl_video_VoD_high
class-map match-all class_video_VoD_low
    match access-group name acl_video_VoD_low
class-map match-all class_video_broadcast
    match access-group name acl_video_broadcast
class-map match-all class_VoD_signaling
    match access-group name acl_VoD_signaling
class-map match-all class_HSD
    match access-group name acl_HSD
    !
    !
policy-map setDSCP
    description Mark DSCP values for ingress traffic
    class class_VoIP
        set dscp ef
    class class_HSD
        set dscp default
    class class_VoD_signaling
        set dscp cs3
    class class_video_broadcast
        set dscp af41
    class class_video_VoD_high
        set dscp af42
    class class_video_VoD_low
        set dscp af43
    !
    !
    !
interface GigabitEthernet1/1
    description High speed data ingress/egress port
    switchport
    switchport access vlan 90
    switchport mode access
    no ip address
    load-interval 30
    no cdp enable
    spanning-tree portfast
    spanning-tree bpduguard enable
    service-policy input setDSCP
    !
interface GigabitEthernet1/2
    no ip address
    shutdown
    !
interface GigabitEthernet1/3
    no ip address
    shutdown
    !
! <----- interfaces GigabitEthernet1/4 - 14 omitted ----->
!

```

```

interface GigabitEthernet1/15
  no ip address
  shutdown
!
interface GigabitEthernet1/16
  no ip address
  shutdown
!
interface GigabitEthernet2/1
  description VoIP traffic ingress/egress
  switchport
  switchport access vlan 80
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet2/2
  no ip address
  shutdown
!
interface GigabitEthernet2/3
  no ip address
  shutdown
!
! <----- interfaces GigabitEthernet2/4 - 6 omitted ----->
!
interface GigabitEthernet2/7
  no ip address
  shutdown
!
interface GigabitEthernet2/8
  no ip address
  shutdown
!
interface GigabitEthernet2/9
  description Broadcast video source (multicast 232.1.1.1 - 232.1.1.10)
  switchport
  switchport access vlan 70
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet2/10
  no ip address
  shutdown
!
interface GigabitEthernet2/11
  no ip address
  shutdown
!
! <----- interfaces GigabitEthernet2/12 - 14 omitted ----->
!
interface GigabitEthernet2/15
  no ip address
  shutdown
!

```

```
interface GigabitEthernet2/16
  no ip address
  shutdown
!
interface GigabitEthernet2/17
  description Management port from Kasenna LR Server (Eth0)
  switchport
  switchport access vlan 10
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet2/18
  description Kasenna VoD Pump Management
  switchport
  switchport access vlan 10
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet2/19
  description Unicast video from Kasenna VoD Pump (HPN0)
  switchport
  switchport access vlan 60
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet2/20
  description Unicast video from Kasenna VoD Pump (HPN1)
  switchport
  switchport access vlan 60
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet2/21
  no ip address
  shutdown
!
interface GigabitEthernet2/22
  no ip address
  shutdown
!
! <----- interfaces GigabitEthernet2/23 - 45 omitted ----->
!
interface GigabitEthernet2/46
  no ip address
```

```

shutdown
!
interface GigabitEthernet2/47
no ip address
shutdown
!
interface GigabitEthernet2/48
description Backup DNS server
switchport
switchport access vlan 11
switchport mode access
no ip address
load-interval 30
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP
!
interface GigabitEthernet5/1
no ip address
shutdown
!
interface GigabitEthernet5/2
description Primary DNS/DHCP/NTP/TFTP/Syslog servers
switchport
switchport access vlan 10
switchport mode access
no ip address
load-interval 30
media-type rj45
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP
!
interface GigabitEthernet6/1
no ip address
shutdown
!
interface GigabitEthernet6/2
no ip address
shutdown
!
interface TenGigabitEthernet7/1
description Transport to/from AR1 (TenGig1/1)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,800,900
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
wrr-queue bandwidth 64 255 0 0 0 0 0
wrr-queue queue-limit 40 50 0 0 0 0 0
wrr-queue threshold 1 100 100 100 100 100 100 100 100 100
wrr-queue threshold 2 45 85 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
no wrr-queue random-detect 2
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2
wrr-queue cos-map 2 3 3 4 6 7
mls qos trust dscp

```



```

!
interface TenGigabitEthernet7/2
  no ip address
  shutdown
!
interface TenGigabitEthernet7/3
  description Transport to/from AR3 (TenGig1/50)
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,824,924
  switchport mode trunk
  dampening
  no ip address
  load-interval 30
  carrier-delay msec 0
  wrr-queue bandwidth 64 255 0 0 0 0 0
  wrr-queue queue-limit 40 50 0 0 0 0 0
  wrr-queue threshold 1 100 100 100 100 100 100 100 100
  wrr-queue threshold 2 45 85 100 100 100 100 100 100
  wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
  wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
  no wrr-queue random-detect 2
  wrr-queue cos-map 2 1 1
  wrr-queue cos-map 2 2 2
  wrr-queue cos-map 2 3 3 4 6 7
  mls qos trust dscp
!
interface TenGigabitEthernet7/4
  no ip address
  shutdown
!
interface Vlan1
  no ip address
  shutdown
!
interface Vlan10
  description Management VLAN (VoD signaling, Primary DNS, DHCP, etc)
  ip address 192.168.10.1 255.255.255.0
  no ip redirects
  no ip unreachable
  load-interval 30
!
interface Vlan11
  description Management VLAN (Backup DNS)
  ip address 192.168.11.1 255.255.255.0
  no ip redirects
  no ip unreachable
  load-interval 30
!
interface Vlan60
  description VoD server VLAN (Unicast Video)
  ip address 192.168.60.1 255.255.255.0
  no ip redirects
  no ip unreachable
  load-interval 30
!
interface Vlan70
  description Broadcast video source VLAN (Multicast Video)
  ip address 192.168.70.1 255.255.255.0
  no ip redirects
  no ip unreachable
  ip pim sparse-mode
  load-interval 30
!

```

```

interface Vlan80
  description VoIP gateway VLAN
  ip address 192.168.80.1 255.255.255.0
  no ip redirects
  no ip unreachable
  load-interval 30
!
interface Vlan800
  description VoIP transport to/from AR1
  ip address 192.168.252.1 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
!
interface Vlan824
  description VoIP transport to/from AR3
  ip address 192.168.252.25 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
!
interface Vlan900
  description Video transport VLAN to/from AR1
  ip address 192.168.254.1 255.255.255.252
  ip pim sparse-mode
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
!
interface Vlan924
  description Video transport VLAN to/from AR3
  ip address 192.168.254.25 255.255.255.252
  ip pim sparse-mode
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
!
router ospf 100
  router-id 1.1.1.1
  log-adjacency-changes
  timers throttle spf 10 100 1000
  timers throttle lsa all 1 10 1000
  timers lsa arrival 100
  passive-interface Vlan10
  passive-interface Vlan11
  passive-interface Vlan60
  passive-interface Vlan70
  network 192.168.10.0 0.0.1.255 area 0
  network 192.168.60.0 0.0.0.255 area 0
  network 192.168.70.0 0.0.0.255 area 0
  network 192.168.254.1 0.0.0.0 area 0
  network 192.168.254.25 0.0.0.0 area 0
!
router ospf 101
  router-id 1.1.1.2
  log-adjacency-changes
  timers throttle spf 10 100 1000
  timers throttle lsa all 1 10 1000
  timers lsa arrival 100
  passive-interface Vlan80
  network 192.168.80.0 0.0.0.255 area 0
  network 192.168.252.1 0.0.0.0 area 0
  network 192.168.252.25 0.0.0.0 area 0
!

```

```

ip classless
!
no ip http server
ip pim ssm default
!
ip access-list extended acl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip host 192.168.10.102 any
 permit ip host 192.168.10.103 any
ip access-list extended acl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.80.0 0.0.0.255 any
ip access-list extended acl_video_VoD_high
 remark Identify high priority VoD traffic
 permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 5000 9000
 permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 5000 9000
 permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 5000 9000
ip access-list extended acl_video_VoD_low
 remark Identify low priority VoD traffic
 permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 1000 4999
 permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 1000 4999
 permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 1000 4999
ip access-list extended acl_video_broadcast
 remark Identify broadcast video traffic (multicast)
 permit ip 192.168.70.0 0.0.0.255 232.0.0.0 0.255.255.255
!
logging event link-status default
logging trap debugging
logging source-interface Vlan10
logging 192.168.10.101
!
!
!
control-plane
!
!
!
dial-peer cor custom
!
!
!
!
line con 0
 exec-timeout 0 0
 password cisco123
 logging synchronous
 login
line vty 0 4
 exec-timeout 0 0
 password cisco123
 logging synchronous
 login
!
!
ntp clock-period 17179953
ntp update-calendar
ntp server 192.168.10.102 prefer
no cns aaa enable
end

```



```

diagnostic cns publish cisco.cns.device.diag_results
diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy ascending
vlan access-log ratelimit 2000
!
vlan 90
  name VLAN_90_HSD
!
vlan 110
  name VLAN_110_Video
!
vlan 111
  name VLAN_111_VoIP
!
vlan 800
  name VLAN_800_VoIP_to/from_DER
!
vlan 808
  name VLAN_808_VoIP_to/from_AR2
!
vlan 900
  name VLAN_900_Video_to/from_DER
!
vlan 908
  name VLAN_908_Video_to/from_AR2
!
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
!
!
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
!
!
!
interface TenGigabitEthernet1/1
  description Transport to/from DER (TenGig7/1)
  switchport
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,800,900
  switchport mode trunk
  dampening
  no ip address
  load-interval 30
  carrier-delay msec 0
  wrr-queue bandwidth 64 255 0 0 0 0 0
  wrr-queue queue-limit 40 50 0 0 0 0 0
  wrr-queue threshold 1 100 100 100 100 100 100 100 100
  wrr-queue threshold 2 45 85 100 100 100 100 100 100
  wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
  wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
  no wrr-queue random-detect 2
  wrr-queue cos-map 2 1 1

```

```

wrr-queue cos-map 2 2 2
wrr-queue cos-map 2 3 3 4 6 7
mls qos trust dscp
!
interface TenGigabitEthernet1/2
no ip address
shutdown
!
interface TenGigabitEthernet1/3
description Transport to/from AR2 (TenGig1/1)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,808,908
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
wrr-queue bandwidth 64 255 0 0 0 0 0
wrr-queue queue-limit 40 50 0 0 0 0 0
wrr-queue threshold 1 100 100 100 100 100 100 100 100 100
wrr-queue threshold 2 45 85 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
no wrr-queue random-detect 2
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2
wrr-queue cos-map 2 3 3 4 6 7
mls qos trust dscp
!
interface TenGigabitEthernet1/4
no ip address
shutdown
!
interface GigabitEthernet2/1
description GigE trunk to/from DSLAM uplink GigE
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,110,111
switchport mode trunk
switchport block unicast
no ip address
load-interval 30
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
no cdp enable
spanning-tree portfast trunk
service-policy input setDSCP
!
interface GigabitEthernet2/2
no ip address
shutdown
!
interface GigabitEthernet2/3
no ip address
shutdown
!

```

```

! <----- interfaces GigabitEthernet2/4 - 14 omitted ----->
!
interface GigabitEthernet2/15
  no ip address
  shutdown
!
interface GigabitEthernet2/16
  no ip address
  shutdown
!
interface GigabitEthernet5/1
  no ip address
  shutdown
!
interface GigabitEthernet5/2
  no ip address
  shutdown
!
interface Vlan1
  no ip address
  shutdown
!
interface Vlan110
  description Video edge VLAN
  ip address 192.168.110.1 255.255.255.0
  no ip redirects
  no ip unreachableables
  ip pim sparse-mode
  ip igmp static-group 232.1.1.1 source ssm-map
  ip igmp static-group 232.1.1.2 source ssm-map
  ip igmp static-group 232.1.1.3 source ssm-map
  ip igmp static-group 232.1.1.4 source ssm-map
  ip igmp static-group 232.1.1.5 source ssm-map
  ip igmp static-group 232.1.1.6 source ssm-map
  ip igmp static-group 232.1.1.7 source ssm-map
  ip igmp static-group 232.1.1.8 source ssm-map
  ip igmp static-group 232.1.1.9 source ssm-map
  ip igmp static-group 232.1.1.10 source ssm-map
  load-interval 30
  arp timeout 250
!
interface Vlan111
  description VoIP edge VLAN
  ip address 192.168.111.1 255.255.255.0
  no ip redirects
  no ip unreachableables
  load-interval 30
!
interface Vlan800
  description VoIP transport VLAN to/from DER
  ip address 192.168.252.2 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
!
interface Vlan808
  description VoIP transport VLAN to/from AR2
  ip address 192.168.252.9 255.255.255.252
  ip ospf network point-to-point
  ip ospf hello-interval 1
  load-interval 30
!
interface Vlan900
  description Video transport VLAN to/from DER

```

```

ip address 192.168.254.2 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan908
description Video transport VLAN to/from AR2
ip address 192.168.254.9 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
router ospf 100
router-id 2.2.2.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan110
network 192.168.110.0 0.0.0.255 area 0
network 192.168.254.2 0.0.0.0 area 0
network 192.168.254.9 0.0.0.0 area 0
!
router ospf 101
router-id 2.2.2.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan111
network 192.168.111.0 0.0.0.255 area 0
network 192.168.252.2 0.0.0.0 area 0
network 192.168.252.9 0.0.0.0 area 0
!
ip classless
!
no ip http server
ip pim ssm default
!
ip access-list extended acl_HSD
remark Identify HSD traffic
permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
remark Identify VoD signaling traffic
permit ip 192.168.110.0 0.0.0.255 192.168.10.102
permit ip 192.168.110.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
remark Identify VoIP traffic
permit ip 192.168.111.0 0.0.0.255 any
!
logging event link-status default
logging trap debugging
logging source-interface Vlan110
logging 192.168.10.101
!
!
!
control-plane
!
!
!
dial-peer cor custom

```



```
!  
!  
!  
line con 0  
  exec-timeout 0 0  
  password cisco123  
  logging synchronous  
  login  
line vty 0 4  
  exec-timeout 0 0  
  password cisco123  
  logging synchronous  
  login  
!  
!  
ntp clock-period 17180008  
ntp update-calendar  
ntp server 192.168.10.102 prefer  
no cns aaa enable  
end
```

Configuration for AR2

```

version 12.2
no service pad
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
service compress-config
!
hostname AR2
!
boot-start-marker
boot system bootflash:cat4000-i9s-mz.122-25.EWA.bin
boot-end-marker
!
!
redundancy
mode sso
enable password cisco123
!
no aaa new-model
clock timezone PST -8
clock summer-time PDT recurring
hw-module uplink select tengigabitethernet
qos map dscp 34 36 38 48 49 50 51 52 to tx-queue 2
qos map dscp 53 54 55 56 57 58 59 60 to tx-queue 2
qos map dscp 61 62 63 to tx-queue 2
qos map dscp 38 to cos 1
qos map dscp 36 to cos 2
qos
vtp mode transparent
ip subnet-zero
ip cef load-sharing algorithm include-ports destination
ip domain multicast coronado.net
no ip domain-lookup
ip name-server 192.168.10.101
ip name-server 192.168.11.101
!
ip multicast-routing
ip igmp ssm-map enable
!
no file verify auto
!
spanning-tree mode rapid-pvst
spanning-tree portfast bpduguard default
spanning-tree extend system-id
no spanning-tree vlan 808,816,908,916
power redundancy-mode redundant
!
!
!
vlan internal allocation policy ascending
!
vlan 90
name VLAN_90_HSD
!
vlan 120
name VLAN_120_Video
!
vlan 121
name VLAN_121_VoIP
!
vlan 808

```

```
    name VLAN_808_VoIP_to/from_AR1
  !
vlan 816
  name VLAN_816_VoIP_to/from_AR3
  !
vlan 908
  name VLAN_908_Video_to/from_AR1
  !
vlan 916
  name VLAN_916_Video_to/from_AR3
  !
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
  !
  !
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
  !
  !
interface TenGigabitEthernet1/1
  description Transport to/from AR1 (TenGig1/3)
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,808,908
  switchport mode trunk
  dampening
  load-interval 30
  carrier-delay msec 0
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
  !
interface TenGigabitEthernet1/2
  description Transport to/from AR3 (TenGig1/49)
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,816,916
  switchport mode trunk
  dampening
  load-interval 30
  carrier-delay msec 0
  qos trust dscp
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
```

```

spanning-tree cost 10
!
interface GigabitEthernet1/3
!
interface GigabitEthernet1/4
!
interface GigabitEthernet1/5
!
interface GigabitEthernet1/6
!
interface GigabitEthernet5/1
description GigE trunk to/from DSLAM uplink GigE
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,120,121
switchport mode trunk
switchport block unicast
service-policy input setDSCP
load-interval 30
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1
no cdp enable
spanning-tree portfast trunk
!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
! <----- interfaces GigabitEthernet5/4 - 46 omitted ----->
!
interface GigabitEthernet5/47
!
interface GigabitEthernet5/48
!
interface Vlan1
no ip address
!
interface Vlan120
description Video edge VLAN
ip address 192.168.120.1 255.255.255.0
no ip redirects
no ip unreachable
ip pim sparse-mode
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
load-interval 30
arp timeout 250
!
interface Vlan121
description VoIP edge VLAN
ip address 192.168.121.1 255.255.255.0

```

```
no ip redirects
no ip unreachable
load-interval 30
!
interface Vlan808
description VoIP transport VLAN to/from AR1
ip address 192.168.252.10 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan816
description VoIP transport VLAN to/from AR3
ip address 192.168.252.17 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan908
description Video transport VLAN to/from AR1
ip address 192.168.254.10 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan916
description Video transport VLAN to/from AR3
ip address 192.168.254.17 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
router ospf 100
router-id 3.3.3.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan120
network 192.168.120.0 0.0.0.255 area 0
network 192.168.254.10 0.0.0.0 area 0
network 192.168.254.17 0.0.0.0 area 0
!
router ospf 101
router-id 3.3.3.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan121
network 192.168.121.0 0.0.0.255 area 0
network 192.168.252.10 0.0.0.0 area 0
network 192.168.252.17 0.0.0.0 area 0
!
!
no ip http server
!
ip pim ssm default
!
!
ip access-list extended acl_HSD
remark Identify HSD traffic
```

```
permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
remark Identify VoD signaling traffic
permit ip 192.168.120.0 0.0.0.255 192.168.10.102
permit ip 192.168.120.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
remark Identify VoIP traffic
permit ip 192.168.121.0 0.0.0.255 any
!
logging trap debugging
logging source-interface Vlan120
logging 192.168.10.101
!
!
!
line con 0
exec-timeout 0 0
password cisco123
logging synchronous
login
stopbits 1
line vty 0 4
exec-timeout 0 0
password cisco123
logging synchronous
login
!
!
ntp clock-period 17180679
ntp update-calendar
ntp server 192.168.10.102 prefer
!
end
```

Configuration for AR3

```
version 12.2
no service pad
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
service compress-config
!
hostname AR3
!
boot-start-marker
boot system bootflash:cat4000-i5s-mz.122-25.EWA.bin
boot-end-marker
!
enable password cisco123
!
no aaa new-model
clock timezone PST -8
clock summer-time PDT recurring
qos map dscp 34 36 38 48 49 50 51 52 to tx-queue 2
qos map dscp 53 54 55 56 57 58 59 60 to tx-queue 2
qos map dscp 61 62 63 to tx-queue 2
qos map dscp 38 to cos 1
qos map dscp 36 to cos 2
qos
vtp mode transparent
ip subnet-zero
ip cef load-sharing algorithm original
ip domain multicast coronado.net
no ip domain-lookup
ip name-server 192.168.10.101
ip name-server 192.168.11.101
!
ip multicast-routing
ip igmp ssm-map enable
!
no file verify auto
!
spanning-tree mode rapid-pvst
spanning-tree portfast bpduguard default
spanning-tree extend system-id
no spanning-tree vlan 816,824,916,924
power redundancy-mode redundant
!
!
!
vlan internal allocation policy ascending
!
vlan 90
 name VLAN_90_HSD
!
vlan 130
 name VLAN_130_Video
!
vlan 131
 name VLAN_131_VoIP
!
vlan 816
 name VLAN_816_VoIP_to/from_AR2
!
vlan 824
 name VLAN_824_VoIP_to/from_DER
```

```

!
vlan 916
  name VLAN_916_Video_to/from_AR2
!
vlan 924
  name VLAN_924_Video_to/from_DER
!
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
!
!
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
!
!
interface GigabitEthernet1/1
  description GigE trunk to/from DSLAM uplink GigE
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,130,131
  switchport mode trunk
  switchport block unicast
  service-policy input setDSCP
  load-interval 30
  tx-queue 1
    bandwidth percent 19
  tx-queue 2
    bandwidth percent 80
  tx-queue 3
    priority high
  tx-queue 4
    bandwidth percent 1
  no cdp enable
  spanning-tree portfast trunk
!
interface GigabitEthernet1/2
!
interface GigabitEthernet1/3
!
! <----- interfaces GigabitEthernet1/4 - 46 omitted ----->
!
interface GigabitEthernet1/47
!
interface GigabitEthernet1/48
!
interface TenGigabitEthernet1/49
  description Transport to/from AR2 (TenGig1/2)
  switchport trunk encapsulation dot1q
  switchport trunk allowed vlan 90,816,916
  switchport mode trunk
  dampening
  load-interval 30
  carrier-delay msec 0
  qos trust dscp
  tx-queue 1

```



```
        bandwidth percent 19
    tx-queue 2
        bandwidth percent 80
    tx-queue 3
        priority high
    tx-queue 4
        bandwidth percent 1
!
interface TenGigabitEthernet1/50
description Transport to/from DER (TenGig7/3)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,824,924
switchport mode trunk
dampening
load-interval 30
carrier-delay msec 0
qos trust dscp
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1
!
interface Vlan1
no ip address
!
interface Vlan130
description Video edge VLAN
ip address 192.168.130.1 255.255.255.0
no ip redirects
no ip unreachablees
ip pim sparse-mode
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
load-interval 30
arp timeout 250
!
interface Vlan131
description VoIP edge VLAN
ip address 192.168.131.1 255.255.255.0
no ip redirects
no ip unreachablees
load-interval 30
!
interface Vlan816
description VoIP transport VLAN to/from AR2
ip address 192.168.252.18 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan824
description VoIP transport VLAN to/from DER
```

```

ip address 192.168.252.26 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan916
description Video transport VLAN to/from AR2
ip address 192.168.254.18 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan924
description Video transport VLAN to/from DER
ip address 192.168.254.26 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
router ospf 100
router-id 4.4.4.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan130
network 192.168.130.0 0.0.0.255 area 0
network 192.168.254.18 0.0.0.0 area 0
network 192.168.254.26 0.0.0.0 area 0
!
router ospf 101
router-id 4.4.4.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan131
network 192.168.131.0 0.0.0.255 area 0
network 192.168.252.18 0.0.0.0 area 0
network 192.168.252.26 0.0.0.0 area 0
!
no ip http server
!
ip pim ssm default
!
!
ip access-list extended acl_HSD
remark Identify HSD traffic
permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
remark Identify VoD signaling traffic
permit ip 192.168.130.0 0.0.0.255 192.168.10.102
permit ip 192.168.130.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
remark Identify VoIP traffic
permit ip 192.168.131.0 0.0.0.255 any
!
logging trap debugging
logging source-interface Vlan130
logging 192.168.10.101
!
!

```

```
!  
line con 0  
  exec-timeout 0 0  
  password cisco123  
  logging synchronous  
  login  
  stopbits 1  
line vty 0 4  
  exec-timeout 0 0  
  password cisco123  
  logging synchronous  
  login  
!  
!  
ntp clock-period 17180737  
ntp update-calendar  
ntp server 192.168.10.102 prefer  
!  
end
```




Sample DER and AR Switch Configurations for the 1-GE Asymmetric Topology

This appendix presents sample distribution edge router (DER) and aggregation router (AR) switch configurations for the asymmetric 1-GE topology described in [Configuration 2: N x 1-GE Asymmetric Ring, page 3-34](#). The following configurations are presented:

- [Configuration for DER, page B-1](#)
- [Configuration for AR1, page B-15](#)
- [Configuration for AR2, page B-23](#)
- [Configuration for AR3, page B-31](#)



Note

Refer to [Configuring the 1-GE Asymmetric Topology, page 4-53](#).

Configuration for DER

```
upgrade fpd auto
version 12.2
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
service counters max age 10
!
hostname DER
!
boot system disk0:s72033-adventerprisek9_wan-mz.122-18.SXE1.bin
logging snmp-authfail
enable password cisco123
!
no aaa new-model
clock timezone PST -8
clock summer-time PDT recurring
ip subnet-zero
!
!
!
ip multicast-routing
ip igmp ssm-map enable
ip domain multicast coronado.net
no ip domain-lookup
ip domain-name coronado.net
```



```
vlan 800
  name VLAN_800_VoIP_to/from_AR1
  !
vlan 816
  name VLAN_816_VoIP_to/from_AR3
  !
vlan 900
  name VLAN_900_Video_to/from_AR1
  !
vlan 916
  name VLAN_916_Video_to/from_AR3
  !
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_video_VoD_high
  match access-group name acl_video_VoD_high
class-map match-all class_video_VoD_low
  match access-group name acl_video_VoD_low
class-map match-all class_video_broadcast
  match access-group name acl_video_broadcast
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
  !
  !
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
  class class_video_broadcast
    set dscp af41
  class class_video_VoD_high
    set dscp af42
  class class_video_VoD_low
    set dscp af43
  !
  !
  !
interface Loopback0
  description Endpoint for Tunnel0
  ip address 10.10.10.1 255.255.255.255
  !
interface Loopback4
  description Endpoint for Tunnel4
  ip address 10.10.10.5 255.255.255.255
  !
interface Loopback8
  description Endpoint for Tunnel8
  ip address 10.10.10.9 255.255.255.255
  !
interface Loopback12
  description Endpoint for Tunnel12
  ip address 10.10.10.13 255.255.255.255
  !
interface Loopback48
  description Endpoint for Tunnel48
  ip address 10.10.10.49 255.255.255.255
  !
interface Loopback52
```

```

description Endpoint for Tunnel52
ip address 10.10.10.53 255.255.255.255
!
interface Loopback56
description Endpoint for Tunnel56
ip address 10.10.10.57 255.255.255.255
!
interface Loopback60
description Endpoint for Tunnel60
ip address 10.10.10.61 255.255.255.255
!
interface Tunnel0
description Rx-side of Tx-only Gig7/3
no ip address
tunnel source 10.10.10.1
tunnel destination 10.10.10.2
tunnel udlr receive-only GigabitEthernet7/3
!
interface Tunnel4
description Rx-side of Tx-only Gig7/4
no ip address
tunnel source 10.10.10.5
tunnel destination 10.10.10.6
tunnel udlr receive-only GigabitEthernet7/4
!
interface Tunnel8
description Rx-side of Tx-only Gig7/5
no ip address
tunnel source 10.10.10.9
tunnel destination 10.10.10.10
tunnel udlr receive-only GigabitEthernet7/5
!
interface Tunnel12
description Rx-side of Tx-only Gig7/6
no ip address
tunnel source 10.10.10.13
tunnel destination 10.10.10.14
tunnel udlr receive-only GigabitEthernet7/6
!
interface Tunnel48
description Rx-side of Tx-only Gig7/11
no ip address
tunnel source 10.10.10.49
tunnel destination 10.10.10.50
tunnel udlr receive-only GigabitEthernet7/11
!
interface Tunnel52
description Rx-side of Tx-only Gig7/12
no ip address
tunnel source 10.10.10.53
tunnel destination 10.10.10.54
tunnel udlr receive-only GigabitEthernet7/12
!
interface Tunnel56
description Rx-side of Tx-only Gig7/13
no ip address
tunnel source 10.10.10.57
tunnel destination 10.10.10.58
tunnel udlr receive-only GigabitEthernet7/13
!
interface Tunnel60
description Rx-side of Tx-only Gig7/14
no ip address
tunnel source 10.10.10.61

```



```
tunnel destination 10.10.10.62
tunnel udld receive-only GigabitEthernet7/14
!
interface GigabitEthernet1/1
description High speed data ingress/egress port
switchport
switchport access vlan 90
switchport mode access
no ip address
load-interval 30
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP
!
interface GigabitEthernet1/2
no ip address
shutdown
!
interface GigabitEthernet1/3
no ip address
shutdown
!
! <----- interfaces GigabitEthernet1/4 - 22 omitted ----->
!
interface GigabitEthernet1/23
no ip address
shutdown
!
interface GigabitEthernet1/24
no ip address
shutdown
!
interface GigabitEthernet2/1
description VoIP traffic ingress/egress port
switchport
switchport access vlan 80
switchport mode access
no ip address
load-interval 30
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP
!
interface GigabitEthernet2/2
no ip address
shutdown
!
interface GigabitEthernet2/3
no ip address
shutdown
!
! <----- interfaces GigabitEthernet2/4 - 6 omitted ----->
!
interface GigabitEthernet2/7
no ip address
shutdown
!
interface GigabitEthernet2/8
no ip address
shutdown
!
interface GigabitEthernet2/9
```

```

description Broadcast video source (multicast 232.1.1.1 - 232.1.1.10)
switchport
switchport access vlan 70
switchport mode access
no ip address
load-interval 30
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP
!
interface GigabitEthernet2/10
no ip address
shutdown
!
interface GigabitEthernet2/11
no ip address
shutdown
!
! <----- interfaces GigabitEthernet2/12 - 14 omitted ----->
!
interface GigabitEthernet2/15
no ip address
shutdown
!
interface GigabitEthernet2/16
no ip address
shutdown
!
interface GigabitEthernet2/17
description Management port from Kasenna LR Server (Eth0)
switchport
switchport access vlan 10
switchport mode access
no ip address
load-interval 30
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP
!
interface GigabitEthernet2/18
description Kasenna VoD Pump Management
switchport
switchport access vlan 10
switchport mode access
no ip address
load-interval 30
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP
!
interface GigabitEthernet2/19
description Unicast video from Kasenna VoD Pump (HPN0)
switchport
switchport access vlan 60
switchport mode access
no ip address
load-interval 30
no cdp enable
spanning-tree portfast
spanning-tree bpduguard enable
service-policy input setDSCP

```

```
!
interface GigabitEthernet2/20
  description Unicast video from Kasenna VoD Pump (HPN1)
  switchport
  switchport access vlan 60
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet2/21
  no ip address
  shutdown
!
interface GigabitEthernet2/22
  no ip address
  shutdown
!
! <----- interfaces GigabitEthernet2/23 - 45 omitted ----->
!
interface GigabitEthernet2/46
  no ip address
  shutdown
!
interface GigabitEthernet2/47
  no ip address
  shutdown
!
interface GigabitEthernet2/48
  description Backup DNS server
  switchport
  switchport access vlan 11
  switchport mode access
  no ip address
  load-interval 30
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet5/1
  no ip address
  shutdown
!
interface GigabitEthernet5/2
  description Primary DNS/DHCP/NTP/TFTP/Syslog servers
  switchport
  switchport access vlan 10
  switchport mode access
  no ip address
  load-interval 30
  media-type rj45
  no cdp enable
  spanning-tree portfast
  spanning-tree bpduguard enable
  service-policy input setDSCP
!
interface GigabitEthernet7/1
  description Transport to/from AR1 (Gig3/1)
  switchport
  switchport trunk encapsulation dot1q
```

```

switchport trunk allowed vlan 90,800,900
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
!
interface GigabitEthernet7/2
!
interface GigabitEthernet7/3
description VoD transport to AR1 (Gig3/3)
dampening
ip address 192.168.253.1 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional send-only
!
interface GigabitEthernet7/4
description VoD transport to AR1 (Gig3/4)
dampening
ip address 192.168.253.5 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional send-only
!
interface GigabitEthernet7/5
description VoD transport to AR1 (Gig3/5)
dampening

```

```

ip address 192.168.253.9 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional send-only
!
interface GigabitEthernet7/6
description VoD transport to AR1 (Gig3/6)
dampening
ip address 192.168.253.13 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional send-only
!
interface GigabitEthernet7/7
!
interface GigabitEthernet7/8
!
interface GigabitEthernet7/9
description Transport to/from AR3 (Gig3/1)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,816,916
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
!

```

```

interface GigabitEthernet7/10
!
interface GigabitEthernet7/11
description VoD transport to AR3 (Gig3/3)
dampening
ip address 192.168.253.49 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional send-only
!
interface GigabitEthernet7/12
description VoD transport to AR3 (Gig3/4)
dampening
ip address 192.168.253.53 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional send-only
!
interface GigabitEthernet7/13
description VoD transport to AR3 (Gig3/5)
dampening
ip address 192.168.253.57 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional send-only

```

```
!  
interface GigabitEthernet7/14  
  description VoD transport to AR3 (Gig3/6)  
  dampening  
  ip address 192.168.253.61 255.255.255.252  
  ip ospf network point-to-point  
  ip ospf hello-interval 1  
  carrier-delay msec 0  
  speed nonegotiate  
  wrp-queue bandwidth 64 255  
  wrp-queue queue-limit 40 50  
  wrp-queue random-detect min-threshold 1 75 100  
  wrp-queue random-detect min-threshold 2 50 100  
  wrp-queue random-detect max-threshold 1 100 100  
  wrp-queue random-detect max-threshold 2 50 100  
  wrp-queue cos-map 1 1 0  
  wrp-queue cos-map 2 1 1  
  wrp-queue cos-map 2 2 2 3 4 6 7  
  mls qos trust dscp  
  unidirectional send-only  
!  
interface GigabitEthernet7/15  
!  
interface GigabitEthernet7/16  
!  
interface Vlan1  
  no ip address  
  shutdown  
!  
interface Vlan10  
  description Management VLAN (VoD signaling, Primary DNS, DHCP, etc)  
  ip address 192.168.10.1 255.255.255.0  
  no ip redirects  
  no ip unreachablees  
  load-interval 30  
!  
interface Vlan11  
  description Management VLAN (Backup DNS)  
  ip address 192.168.11.1 255.255.255.0  
  no ip redirects  
  no ip unreachablees  
  load-interval 30  
!  
interface Vlan60  
  description VoD server VLAN (Unicast Video)  
  ip address 192.168.60.1 255.255.255.0  
  no ip redirects  
  no ip unreachablees  
  ip pim sparse-mode  
  load-interval 30  
!  
interface Vlan70  
  description Broadcast video ingress VLAN (Multicast Video)  
  ip address 192.168.70.1 255.255.255.0  
  no ip redirects  
  no ip unreachablees  
  ip pim sparse-mode  
  load-interval 30  
!  
interface Vlan80  
  description VoIP ingress/egress VLAN  
  ip address 192.168.80.1 255.255.255.0  
  no ip redirects  
  no ip unreachablees
```

```

load-interval 30
!
interface Vlan800
description VoIP transport to/from AR1
ip address 192.168.252.1 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan816
description VoIP transport to/from AR3
ip address 192.168.252.17 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan900
description Video transport to/from AR1
ip address 192.168.254.1 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan916
description Video transport to/from AR3
ip address 192.168.254.17 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
router ospf 100
router-id 1.1.1.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan10
passive-interface Vlan11
passive-interface Vlan60
passive-interface Vlan70
network 10.10.10.0 0.0.0.255 area 0
network 192.168.10.0 0.0.1.255 area 0
network 192.168.60.0 0.0.0.255 area 0
network 192.168.70.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 101
router-id 1.1.1.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan80
network 192.168.80.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 102
router-id 1.1.1.3
log-adjacency-changes
timers throttle spf 10 100 1000

```



```

timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
!
ip classless
!
no ip http server
ip pim ssm default
!
ip access-list extended acl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip host 192.168.10.102 any
 permit ip host 192.168.10.103 any
ip access-list extended acl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.80.0 0.0.0.255 any
ip access-list extended acl_video_VoD_high
 remark Identify high priority VoD traffic
 permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 5000 9000
 permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 5000 9000
 permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 5000 9000
ip access-list extended acl_video_VoD_low
 remark Identify low priority VoD traffic
 permit udp 192.168.60.0 0.0.0.255 192.168.110.0 0.0.0.255 range 1000 4999
 permit udp 192.168.60.0 0.0.0.255 192.168.120.0 0.0.0.255 range 1000 4999
 permit udp 192.168.60.0 0.0.0.255 192.168.130.0 0.0.0.255 range 1000 4999
ip access-list extended acl_video_broadcast
 remark Identify broadcast video traffic (multicast)
 permit ip 192.168.70.0 0.0.0.255 232.0.0.0 0.255.255.255
!
logging event link-status default
logging trap debugging
logging source-interface Vlan10
logging 192.168.10.101
!
!
!
control-plane
!
!
!
dial-peer cor custom
!
!
!
!
line con 0
 exec-timeout 0 0
 password cisco123
 logging synchronous
 login
line vty 0 4
 exec-timeout 0 0
 password cisco123
 logging synchronous
 login
!
!
ntp clock-period 17179996
ntp update-calendar

```

■ Configuration for DER

```
ntp server 192.168.10.102 prefer
no cns aaa enable
end
```

Configuration for AR1

```
version 12.2
no service pad
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
service internal
service compress-config
!
hostname AR1
!
boot-start-marker
boot system bootflash:cat4000-i5s-mz.122-25.EWA.bin
boot-end-marker
!
!
redundancy
notification-timer 60000
mode sso
main-cpu
auto-sync standard
enable password cisco123
!
no aaa new-model
clock timezone PST -8
clock summer-time PDT recurring
qos map dscp 34 36 38 48 49 50 51 52 to tx-queue 2
qos map dscp 53 54 55 56 57 58 59 60 to tx-queue 2
qos map dscp 61 62 63 to tx-queue 2
qos map dscp 38 to cos 1
qos map dscp 36 to cos 2
qos
vtp mode transparent
ip subnet-zero
ip cef load-sharing algorithm include-ports destination
ip domain multicast coronado.net
no ip domain-lookup
ip name-server 192.168.10.101
ip name-server 192.168.11.101
!
ip multicast-routing
ip igmp ssm-map enable
!
no file verify auto
!
spanning-tree mode rapid-pvst
spanning-tree portfast bpduguard default
spanning-tree extend system-id
no spanning-tree vlan 800,808,900,908
power redundancy-mode redundant
!
!
!
process-max-time 20
vlan internal allocation policy ascending
!
vlan 90
name VLAN_90_HSD
!
vlan 110
name VLAN_110_Video
!
```

```

vlan 111
  name VLAN_111_VoIP
!
vlan 800
  name VLAN_800_VoIP_to/from_DER
!
vlan 808
  name VLAN_808_VoIP_to/from_AR2
!
vlan 900
  name VLAN_900_Video_to/from_DER
!
vlan 908
  name VLAN_908_Video_to/from_AR2
!
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
!
!
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
!
!
interface Tunnel0
  description Tx-side of Rx-only Gig3/3
  no ip address
  tunnel source 10.10.10.2
  tunnel destination 10.10.10.1
  tunnel udlr send-only GigabitEthernet3/3
  tunnel udlr address-resolution
!
interface Tunnel4
  description Tx-side of Rx-only Gig3/4
  no ip address
  tunnel source 10.10.10.6
  tunnel destination 10.10.10.5
  tunnel udlr send-only GigabitEthernet3/4
  tunnel udlr address-resolution
!
interface Tunnel8
  description Tx-side of Rx-only Gig3/5
  no ip address
  tunnel source 10.10.10.10
  tunnel destination 10.10.10.9
  tunnel udlr send-only GigabitEthernet3/5
  tunnel udlr address-resolution
!
interface Tunnel12
  description Tx-side of Rx-only Gig3/6
  no ip address
  tunnel source 10.10.10.14
  tunnel destination 10.10.10.13
  tunnel udlr send-only GigabitEthernet3/6
  tunnel udlr address-resolution

```

```
!  
interface Tunnel24  
  description Rx-side of Tx-only Gig4/4  
  no ip address  
  tunnel source 10.10.10.25  
  tunnel destination 10.10.10.26  
  tunnel udld receive-only GigabitEthernet4/4  
!  
interface Tunnel28  
  description Rx-side of Tx-only Gig4/5  
  no ip address  
  tunnel source 10.10.10.29  
  tunnel destination 10.10.10.30  
  tunnel udld receive-only GigabitEthernet4/5  
!  
interface Loopback0  
  description Endpoint for Tunnel0  
  ip address 10.10.10.2 255.255.255.255  
!  
interface Loopback4  
  description Endpoint for Tunnel4  
  ip address 10.10.10.6 255.255.255.255  
!  
interface Loopback8  
  description Endpoint for Tunnel8  
  ip address 10.10.10.10 255.255.255.255  
!  
interface Loopback12  
  description Endpoint for Tunnel12  
  ip address 10.10.10.14 255.255.255.255  
!  
interface Loopback24  
  description Endpoint for Tunnel24  
  ip address 10.10.10.25 255.255.255.255  
!  
interface Loopback28  
  description Endpoint for Tunnel28  
  ip address 10.10.10.29 255.255.255.255  
!  
interface FastEthernet1  
  no ip address  
  speed auto  
  duplex auto  
!  
interface GigabitEthernet1/1  
!  
interface GigabitEthernet1/2  
!  
interface GigabitEthernet3/1  
  description Transport to/from DER (Gig7/1)  
  switchport trunk encapsulation dot1q  
  switchport trunk allowed vlan 90,800,900  
  switchport mode trunk  
  dampening  
  load-interval 30  
  carrier-delay msec 0  
  qos trust dscp  
  tx-queue 1  
    bandwidth percent 19  
  tx-queue 2  
    bandwidth percent 80  
  tx-queue 3  
    priority high  
  tx-queue 4
```

```

        bandwidth percent 1
    !
interface GigabitEthernet3/2
!
interface GigabitEthernet3/3
description Transport from DER (Gig7/3)
no switchport
dampening
ip address 192.168.253.2 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet3/4
description Transport from DER (Gig7/4)
no switchport
dampening
ip address 192.168.253.6 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet3/5
description Transport from DER (Gig7/5)
no switchport
dampening
ip address 192.168.253.10 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet3/6
description Transport from DER (Gig7/6)
no switchport
dampening
ip address 192.168.253.14 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet4/1
!
interface GigabitEthernet4/2
!
interface GigabitEthernet4/3
description Transport to/from AR2 (Gig1/1)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,808,908

```

```
switchport mode trunk
dampening
load-interval 30
carrier-delay msec 0
qos trust dscp
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1
!
interface GigabitEthernet4/4
description Transport to AR2 (Gig1/2)
no switchport
dampening
ip address 192.168.253.25 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed negotiate
qos trust dscp
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1
unidirectional send-only
!
interface GigabitEthernet4/5
description Transport to AR2 (Gig1/3)
no switchport
dampening
ip address 192.168.253.29 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed negotiate
qos trust dscp
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1
unidirectional send-only
!
interface GigabitEthernet4/6
!
interface GigabitEthernet5/1
description GigE trunk to/from DSLAM uplink GigE
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,110,111
switchport mode trunk
switchport block unicast
```

```

service-policy input setDSCP
load-interval 30
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1
no cdp enable
spanning-tree portfast trunk
!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
interface GigabitEthernet5/5
!
interface GigabitEthernet5/6
!
interface Vlan1
no ip address
!
interface Vlan110
description Video edge VLAN
ip address 192.168.110.1 255.255.255.0
no ip redirects
no ip unreachable
ip pim sparse-mode
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
load-interval 30
arp timeout 250
!
interface Vlan111
description VoIP edge VLAN
ip address 192.168.111.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30
!
interface Vlan800
description VoIP transport to/from DER
ip address 192.168.252.2 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan808
description VoIP transport to/from AR2
ip address 192.168.252.9 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1

```



```
load-interval 30
!
interface Vlan900
description Video transport to/from DER
ip address 192.168.254.2 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan908
description Video transport to/from AR2
ip address 192.168.254.9 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
router ospf 100
router-id 2.2.2.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 10.10.10.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 101
router-id 2.2.2.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan111
network 192.168.111.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 102
router-id 2.2.2.3
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan110
network 192.168.110.0 0.0.0.255 area 0
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
!
no ip http server
!
ip pim ssm default
!
ip access-list extended acl_HSD
remark Identify HSD traffic
permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
remark Identify VoD signaling traffic
permit ip 192.168.110.0 0.0.0.255 192.168.10.102
permit ip 192.168.110.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
remark Identify VoIP traffic
permit ip 192.168.111.0 0.0.0.255 any
```

```
!  
logging trap debugging  
logging source-interface Vlan110  
logging 192.168.10.101  
!  
!  
!  
line con 0  
  exec-timeout 0 0  
  password cisco123  
  logging synchronous  
  login  
line vty 0 4  
  exec-timeout 0 0  
  password cisco123  
  logging synchronous  
  login  
!  
!  
scheduler runtime netinput 100  
no scheduler max-sched-time  
ntp clock-period 17179583  
ntp update-calendar  
ntp server 192.168.10.102 prefer  
!  
end
```



```

diagnostic cns subscribe cisco.cns.device.diag_commands
!
vlan internal allocation policy ascending
vlan access-log ratelimit 2000
!
vlan 90
  name VLAN_90_HSD
!
vlan 120
  name VLAN_120_Video
!
vlan 121
  name VLAN_121_VoIP
!
vlan 808
  name VLAN_808_VoIP_to/from_AR1
!
vlan 812
  name VLAN_812_VoIP_to/from_AR3
!
vlan 908
  name VLAN_908_Video_to/from_AR1
!
vlan 912
  name VLAN_912_Video_to/from_AR3
!
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
!
!
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
!
!
!
interface Loopback24
  description Endpoint for Tunnel24
  ip address 10.10.10.26 255.255.255.255
!
interface Loopback28
  description Endpoint for Tunnel28
  ip address 10.10.10.30 255.255.255.255
!
interface Loopback36
  description Endpoint for Tunnel36
  ip address 10.10.10.38 255.255.255.255
!
interface Loopback40
  description Endpoint for Tunnel40
  ip address 10.10.10.42 255.255.255.255
!
interface Tunnel24
  description Tx-side of Rx-only Gig1/2
  no ip address

```

```
tunnel source 10.10.10.26
tunnel destination 10.10.10.25
tunnel udld send-only GigabitEthernet1/2
tunnel udld address-resolution
!
interface Tunnel28
description Tx-side of Rx-only Gig1/3
no ip address
tunnel source 10.10.10.30
tunnel destination 10.10.10.29
tunnel udld send-only GigabitEthernet1/3
tunnel udld address-resolution
!
interface Tunnel36
description Tx-side of Rx-only Gig1/6
no ip address
tunnel source 10.10.10.38
tunnel destination 10.10.10.37
tunnel udld send-only GigabitEthernet1/6
tunnel udld address-resolution
!
interface Tunnel40
description Tx-side of Rx-only Gig1/7
no ip address
tunnel source 10.10.10.42
tunnel destination 10.10.10.41
tunnel udld send-only GigabitEthernet1/7
tunnel udld address-resolution
!
interface GigabitEthernet1/1
description Transport to/from AR1 (Gig4/3)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,808,908
switchport mode trunk
no ip address
load-interval 30
carrier-delay msec 0
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
!
interface GigabitEthernet1/2
description Transport from AR1 (Gig4/4)
dampening
ip address 192.168.253.26 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
```

```

wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet1/3
description Transport from AR1 (Gig4/5)
dampening
ip address 192.168.253.30 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet1/4
!
interface GigabitEthernet1/5
description Transport to/from AR3 (Gig4/3)
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,812,912
switchport mode trunk
dampening
no ip address
load-interval 30
carrier-delay msec 0
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
spanning-tree cost 10
!
interface GigabitEthernet1/6
description Transport from AR3 (Gig4/4)
dampening
ip address 192.168.253.38 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100

```

```

wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet1/7
description Transport from AR3 (Gig4/5)
dampening
ip address 192.168.253.42 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
wrr-queue bandwidth 64 255
wrr-queue queue-limit 40 50
wrr-queue random-detect min-threshold 1 75 100
wrr-queue random-detect min-threshold 2 50 100
wrr-queue random-detect max-threshold 1 100 100
wrr-queue random-detect max-threshold 2 50 100
wrr-queue cos-map 1 1 0
wrr-queue cos-map 2 1 1
wrr-queue cos-map 2 2 2 3 4 6 7
mls qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet1/8
no ip address
shutdown
!
interface GigabitEthernet1/9
no ip address
shutdown
!
! <----- interfaces GigabitEthernet1/10 - 14 omitted ----->!
!
interface GigabitEthernet1/15
no ip address
shutdown
!
interface GigabitEthernet1/16
no ip address
shutdown
!
interface GigabitEthernet2/1
description Ericsson DSLAM
switchport
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,120,121
switchport mode trunk
switchport block unicast
no ip address
load-interval 30
wrr-queue bandwidth 64 255 0
wrr-queue queue-limit 40 50 0
wrr-queue threshold 1 100 100 100 100 100 100 100 100
wrr-queue threshold 2 50 100 100 100 100 100 100 100
wrr-queue random-detect min-threshold 1 75 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
no wrr-queue random-detect 2
wrr-queue cos-map 2 1 1

```

```

wrr-queue cos-map 2 2 2 3 4 6 7
no cdp enable
spanning-tree portfast trunk
service-policy input setDSCP
!
interface GigabitEthernet2/2
no ip address
shutdown
!
interface GigabitEthernet2/3
no ip address
shutdown
!
! <----- interfaces GigabitEthernet1/10 - 14 omitted ----->!
!
interface GigabitEthernet2/23
no ip address
shutdown
!
interface GigabitEthernet2/24
no ip address
shutdown
!
interface GigabitEthernet5/1
no ip address
shutdown
!
interface GigabitEthernet5/2
no ip address
shutdown
!
interface Vlan1
no ip address
shutdown
!
interface Vlan120
description Video edge VLAN
ip address 192.168.120.1 255.255.255.0
no ip redirects
no ip unreachable
ip pim sparse-mode
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
ip igmp static-group 232.1.1.11 source ssm-map
load-interval 30
arp timeout 250
!
interface Vlan121
description VoIP edge VLAN
ip address 192.168.121.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30
!

```



```
interface Vlan808
description VoIP transport VLAN to/from AR1
ip address 192.168.252.10 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan812
description VoIP transport VLAN to/from AR3
ip address 192.168.252.14 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan908
description Video transport VLAN to/from AR1
ip address 192.168.254.10 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
interface Vlan912
description Video transport VLAN to/from AR3
ip address 192.168.254.14 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
!
router ospf 100
router-id 3.3.3.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 10.10.10.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 101
router-id 3.3.3.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan121
network 192.168.121.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 102
router-id 3.3.3.3
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan120
network 192.168.120.0 0.0.0.255 area 0
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
!
ip classless
!
```

```

no ip http server
ip pim ssm default
!
ip access-list extended acl_HSD
 remark Identify HSD traffic
 permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
 remark Identify VoD signaling traffic
 permit ip 192.168.120.0 0.0.0.255 192.168.10.102
 permit ip 192.168.120.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
 remark Identify VoIP traffic
 permit ip 192.168.121.0 0.0.0.255 any
!
logging trap debugging
logging source-interface Vlan120
logging 192.168.10.101
!
!
!
control-plane
!
!
!
dial-peer cor custom
!
!
!
line con 0
 exec-timeout 0 0
 password cisco123
 logging synchronous
 login
line vty 0 4
 exec-timeout 0 0
 password cisco123
 logging synchronous
 login
!
!
ntp clock-period 17179835
ntp update-calendar
ntp server 192.168.10.102 prefer
no cns aaa enable
end

```

Configuration for AR3

```
version 12.2
no service pad
service timestamps debug datetime msec localtime
service timestamps log datetime msec localtime
no service password-encryption
service compress-config
!
hostname AR3
!
boot-start-marker
boot system bootflash:cat4000-i5s-mz.122-25.EWA.bin
boot-end-marker
!
!
redundancy
 mode sso
 main-cpu
  auto-sync standard
enable password cisco123
!
no aaa new-model
clock timezone PST -8
clock summer-time PDT recurring
qos map dscp 38 to cos 1
qos map dscp 36 to cos 2
qos
vtp mode transparent
ip subnet-zero
ip cef load-sharing algorithm include-ports destination
ip domain multicast coronado.net
no ip domain-lookup
ip name-server 192.168.10.101
ip name-server 192.168.11.101
!
ip multicast-routing
ip igmp ssm-map enable
!
no file verify auto
!
spanning-tree mode rapid-pvst
spanning-tree portfast bpduguard default
spanning-tree extend system-id
no spanning-tree vlan 812,816,912,916
power redundancy-mode redundant
!
!
!
vlan internal allocation policy ascending
!
!
vlan 90
 name VLAN_90_HSD
!
vlan 130
 name VLAN_130_Video
!
vlan 131
 name VLAN_131_VoIP
!
vlan 812
 name VLAN_812_VoIP_to/from_AR2
```

```

!
vlan 816
  name VLAN_816_VoIP_to/from_DER
!
vlan 912
  name VLAN_912_Video_to/from_AR2
!
vlan 916
  name VLAN_916_Video_to/from_DER
!
class-map match-all class_VoIP
  match access-group name acl_VoIP
class-map match-all class_VoD_signaling
  match access-group name acl_VoD_signaling
class-map match-all class_HSD
  match access-group name acl_HSD
!
!
policy-map setDSCP
  description Mark DSCP values for ingress traffic
  class class_VoIP
    set dscp ef
  class class_HSD
    set dscp default
  class class_VoD_signaling
    set dscp cs3
!
!
interface Tunnel36
  description Rx-side of Tx-only Gig4/4
  no ip address
  tunnel source 10.10.10.37
  tunnel destination 10.10.10.38
  tunnel udlr receive-only GigabitEthernet4/4
!
interface Tunnel40
  description Rx-side of Tx-only Gig4/5
  no ip address
  tunnel source 10.10.10.41
  tunnel destination 10.10.10.42
  tunnel udlr receive-only GigabitEthernet4/5
!
!
interface Tunnel48
  description Tx-side of Rx-only Gig3/3
  no ip address
  tunnel source 10.10.10.50
  tunnel destination 10.10.10.49
  tunnel udlr send-only GigabitEthernet3/3
  tunnel udlr address-resolution
!
interface Tunnel52
  description Tx-side of Rx-only Gig3/4
  no ip address
  tunnel source 10.10.10.54
  tunnel destination 10.10.10.53
  tunnel udlr send-only GigabitEthernet3/4
  tunnel udlr address-resolution
!
interface Tunnel56
  description Tx-side of Rx-only Gig3/5
  no ip address
  tunnel source 10.10.10.58
  tunnel destination 10.10.10.57

```

```
tunnel udld send-only GigabitEthernet3/5
tunnel udld address-resolution
!
interface Tunnel60
description Tx-side of Rx-only Gig3/6
no ip address
tunnel source 10.10.10.62
tunnel destination 10.10.10.61
tunnel udld send-only GigabitEthernet3/6
tunnel udld address-resolution
!
interface Loopback36
description Endpoint for Tunnel36
ip address 10.10.10.37 255.255.255.255
!
interface Loopback40
description Endpoint for Tunnel40
ip address 10.10.10.41 255.255.255.255
!
interface Loopback48
description Endpoint for Tunnel48
ip address 10.10.10.50 255.255.255.255
!
interface Loopback52
description Endpoint for Tunnel52
ip address 10.10.10.54 255.255.255.255
!
interface Loopback56
description Endpoint for Tunnel56
ip address 10.10.10.58 255.255.255.255
!
interface Loopback60
description Endpoint for Tunnel60
ip address 10.10.10.62 255.255.255.255
!
interface FastEthernet1
no ip address
shutdown
speed auto
duplex auto
!
interface GigabitEthernet1/1
!
interface GigabitEthernet1/2
!
interface GigabitEthernet3/1
description Transport to/from DER (Gig7/9)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,816,916
switchport mode trunk
load-interval 30
qos trust dscp
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80
tx-queue 3
    priority high
tx-queue 4
    bandwidth percent 1
!
interface GigabitEthernet3/2
!
interface GigabitEthernet3/3
```

```

description Transport from DER (Gig7/11)
no switchport
dampening
ip address 192.168.253.50 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
unidirectional receive-only
!
interface GigabitEthernet3/4
description Transport from DER (Gig7/12)
no switchport
dampening
ip address 192.168.253.54 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
unidirectional receive-only
!
interface GigabitEthernet3/5
description Transport from DER (Gig7/13)
no switchport
dampening
ip address 192.168.253.58 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet3/6
description Transport from DER (Gig7/14)
no switchport
dampening
ip address 192.168.253.62 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
carrier-delay msec 0
speed nonegotiate
qos trust dscp
unidirectional receive-only
!
interface GigabitEthernet4/1
!
interface GigabitEthernet4/2
!
interface GigabitEthernet4/3
description Transport to/from AR2 (Gig1/5)
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,812,912
switchport mode trunk
load-interval 30
qos trust dscp
tx-queue 1
    bandwidth percent 19
tx-queue 2
    bandwidth percent 80

```

```
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
spanning-tree cost 10
!
interface GigabitEthernet4/4
description Transport to AR2 (Gig1/6)
no switchport
ip address 192.168.253.37 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
speed nonegotiate
qos trust dscp
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
unidirectional send-only
!
interface GigabitEthernet4/5
description Transport to AR2 (Gig1/7)
no switchport
ip address 192.168.253.41 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
load-interval 30
speed nonegotiate
qos trust dscp
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
unidirectional send-only
!
interface GigabitEthernet4/6
!
interface GigabitEthernet5/1
description GigE trunk to/from DSLAM uplink GigE
switchport trunk encapsulation dot1q
switchport trunk allowed vlan 90,130,131
switchport mode trunk
switchport block unicast
service-policy input setDSCP
load-interval 30
tx-queue 1
  bandwidth percent 19
tx-queue 2
  bandwidth percent 80
tx-queue 3
  priority high
tx-queue 4
  bandwidth percent 1
no cdp enable
spanning-tree portfast trunk
```

```

!
interface GigabitEthernet5/2
!
interface GigabitEthernet5/3
!
interface GigabitEthernet5/4
!
interface GigabitEthernet5/5
!
interface GigabitEthernet5/6
!
interface Vlan1
no ip address
!
interface Vlan130
description Video edge VLAN
ip address 192.168.130.1 255.255.255.0
no ip redirects
no ip unreachable
ip pim sparse-mode
ip igmp static-group 232.1.1.1 source ssm-map
ip igmp static-group 232.1.1.2 source ssm-map
ip igmp static-group 232.1.1.3 source ssm-map
ip igmp static-group 232.1.1.4 source ssm-map
ip igmp static-group 232.1.1.5 source ssm-map
ip igmp static-group 232.1.1.6 source ssm-map
ip igmp static-group 232.1.1.7 source ssm-map
ip igmp static-group 232.1.1.8 source ssm-map
ip igmp static-group 232.1.1.9 source ssm-map
ip igmp static-group 232.1.1.10 source ssm-map
load-interval 30
arp timeout 250
!
interface Vlan131
description VoIP edge VLAN
ip address 192.168.131.1 255.255.255.0
no ip redirects
no ip unreachable
load-interval 30
!
interface Vlan812
description VoIP transport to/from AR2
ip address 192.168.252.13 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
!
interface Vlan816
description VoIP transport to/from DER
ip address 192.168.252.18 255.255.255.252
ip ospf network point-to-point
ip ospf hello-interval 1
!
interface Vlan912
description Video transport to/from AR2
ip address 192.168.254.13 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point
ip ospf hello-interval 1
!
interface Vlan916
description Video transport to/from DER
ip address 192.168.254.18 255.255.255.252
ip pim sparse-mode
ip ospf network point-to-point

```



```
ip ospf hello-interval 1
!
router ospf 100
router-id 4.4.4.1
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
network 10.10.10.0 0.0.0.255 area 0
network 192.168.254.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 101
router-id 4.4.4.2
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan131
network 192.168.131.0 0.0.0.255 area 0
network 192.168.252.0 0.0.0.255 area 0
maximum-paths 8
!
router ospf 102
router-id 4.4.4.3
log-adjacency-changes
timers throttle spf 10 100 1000
timers throttle lsa all 1 10 1000
timers lsa arrival 100
passive-interface Vlan130
network 192.168.130.0 0.0.0.255 area 0
network 192.168.253.0 0.0.0.255 area 0
maximum-paths 8
!
no ip http server
!
ip pim ssm default
!
ip access-list extended acl_HSD
remark Identify HSD traffic
permit ip 192.168.90.0 0.0.0.255 any
ip access-list extended acl_VoD_signaling
remark Identify VoD signaling traffic
permit ip 192.168.130.0 0.0.0.255 192.168.10.102
permit ip 192.168.130.0 0.0.0.255 192.168.10.103
ip access-list extended acl_VoIP
remark Identify VoIP traffic
permit ip 192.168.131.0 0.0.0.255 any
!
logging trap debugging
logging source-interface Vlan130
logging 192.168.10.101
!
line con 0
exec-timeout 0 0
password cisco123
logging synchronous
login
line vty 0 4
exec-timeout 0 0
password cisco123
logging synchronous
login
!
```

■ Configuration for AR3

```
!  
ntp clock-period 17179522  
ntp update-calendar  
ntp server 192.168.10.102 prefer  
!  
end
```



Understanding QoS as Implemented in the Solution

This chapter presents a more detailed understanding of Quality of Service (QoS) as it is implemented in the solution, and presents the following topics:

- [Introduction, page C-1](#)
- [DiffServ Classification, page C-3](#)
- [Class Selector Values, page C-3F](#)
- [Assured Forwarding, page C-4](#)
- [Express Forwarding, page C-4](#)
- [Default Class, page C-5](#)
- [DSCP-to-CoS Mapping, page C-5](#)
- [Queueing and Thresholds, page C-6](#)
- [1-GE Asymmetric and 10-GE Symmetric Topologies: Known Threshold Parameters, page C-9](#)



Note

For the architectural design aspects of QoS, refer to [QoS Architecture, page 3-46](#).


Introduction

Each packet that enters the network at the DER or AR is classified according to the service type. Classification is done through extended ACLs and a policy map that specifies the Differentiated Services Code Point (DSCP) value (0–63) to assign to the incoming packet. The noningress ports of the network are configured to trust the incoming Differentiated Services Code Point (DSCP) values.


On interfaces configured as Layer-2 IEEE 802.1q trunks, all traffic is encapsulated in 802.1q frames except for traffic that is in the native VLAN. The IEEE 802.1q frame headers have a 2-byte Tag Control Information field that carries the Class of Service (CoS) value (0–7) in the three most-significant bits, which are called the User Priority bits. [Figure C-1 on page C-2](#) illustrates the Layer 2 and Layer 3 QoS fields.

Figure C-1 Layer 2 and Layer 3 QoS Fields


Encapsulated Packet




Layer 2 ISL Frame



Layer 2 802.1Q/P Frame



Layer 3 IPv4 Packet



Different switch models handle DSCP differently. In the Cisco 7600 series and Cisco Catalyst 6500 series switches, the DSCP values are mapped to CoS values. The CoS values are in turn mapped to different transmit queues in the egress interface, and thresholds are used for buffer management. In the Cisco Catalyst 4000, DSCP values are mapped to transmit queues in the egress interface. Weighted Round Robin (WRR) and Strict Priority (SP) is used to service the transmit queues.

Table C-1 summarizes the service types and their corresponding DSCP, Ethernet CoS, and ATM Class values as used in the solution.

Table C-1 Solution Service Types and Values

Service	DSCP	Ethernet CoS	ATM Class
Multicast video	AF41	4	VBR-rt (variable bit rate, real time)
Unicast video (50%)	AF42	2	
	AF43	1	
VoIP	EF	5	CBR (constant bit rate)
Signaling	CS3	3	VBR-rt
HSD	Default	0	UBR (unspecified bit rate)

DiffServ Classification

The six most-significant bits of the DiffServ field are the DSCP values. [The last two bits in the DiffServ field were not defined within the DiffServ field architecture, but were later specified as Explicit Congestion Notification (ECN) bits.] As illustrated in [Table C-2](#), the DiffServ standard uses the three most-significant bits (DS5, DS4, DS3) for precedence bits, and offers finer granularity through the use of the next three bits.

Table C-2 DiffServ Bits and Values

Bit	0	1	2	3	4	5	6	7
Value	DS5	DS4	DS3	DS2	DS1	DS0	ECN	ECN

Four types of forwarding behavior are used: Class Selector (CS), Assured Forwarding (AF), Express Forwarding (EF), and Default.

Class Selector Values

Class Selector values have the format $xx000$, where x is 0 or 1. These values enable backward compatibility with the older IP-Precedence scheme. For this solution, we recommend that VoIP or video signaling be marked with a Class of Service (CS) of 3. This class guarantees a minimum rate and has a DSCP value of 24 (0x011000). [Table C-3](#) shows the class selector values for the DiffServ precedence levels.

Table C-3 Class Selector Values for DiffServ Precedence Levels

Class Selector Value	DiffServ Precedence Level and Description
0x111000 CS7 DSCP 56	Link layer and routing protocol keepalive
0x110000 CS6 DSCP 48	IP routing protocols
0x101000 CS5 DSCP 40	Express forwarding
0x100000 CS4 DSCP 32	Class 4
0x011000 CS3 DSCP 24	Class 3
0x010000 CS2 DSCP 16	Class 2

Table C-3 Class Selector Values for DiffServ Precedence Levels (continued)

Class Selector Value	DiffServ Precedence Level and Description
0x001000 CS1 DSCP 8	Class 1
0x000000 CS0 DSCP 0	Best effort

Assured Forwarding

Assured Forwarding (AF) provides for the delivery of IP packets in four independently forwarded AF classes. Within each AF class, an IP packet can be assigned one of three different levels of drop precedence. This class is used when a service (application) requires a high probability of packets being forwarded, so long as the aggregate traffic from each site does not exceed the subscribed information rate (profile). Each of the four AF classes allocates a certain amount of forwarding resources, such as buffer space and bandwidth in each network node. When congestion occurs, the drop precedence of a packet determines the relative importance of the packet within the AF class.

For this solution, we recommend that all video be marked as with Precedence Level 4 (known as the Assured Forwarding class), with three levels of granularity. Therefore, multicast video has the lowest allowed drops (AF41), and unicast video has the second lowest (AF42). [Table C-4](#) illustrates the drop thresholds and values for the Assured Forwarding classes.

Table C-4 Drop Thresholds and Values for the Assured Forwarding Classes

Drop Threshold	Class 1	Class 2	Class 3	Class 4
Low	001010 AF11 DSCP 10	010010 AF21 DSCP 18	011010 AF31 DSCP 26	100010 AF41 DSCP 34
Medium	001100 AF12 DSCP 12	010100 AF22 DSCP 20	011100 AF32 DSCP 28	100100 AF42 DSCP 36
High	001110 AF13 DSCP 14	010110 AF23 DSCP 22	011110 AF33 DSCP 30	100110 AF43 DSCP 38

Express Forwarding

Express forwarding is a low-loss, low-latency, low-jitter, assured-bandwidth, end-to-end service. Expedited Forwarding (EF) is defined as within the range for express forwarding. The DSCP value for EF is 46 (0x101110). For this solution, we recommend that VoIP traffic be marked as EF.

Default Class

The Default class is a best-effort class with no guaranteed rates. The DSCP value for Default is 0 (0x000000). For this solution, we recommend that high-speed data (HSD) be marked as Default. [Table C-5](#) summarizes the recommended DSCP markings for the various services.

Table C-5 Recommended DSCP Markings

Service	DSCP
Multicast video	AF41
Unicast video 1 (50%)	AF42
Unicast video 2 (50%)	AF43
VoIP	EF
Signaling	CS3
HSD	Default

DSCP-to-CoS Mapping

The DSCP values in the IP packets are mapped to a CoS field in the Ethernet frame. These tags form part of the IEEE 802.1q header. On the DSL links, the ATM classes and CoS values are associated with the different services. The different services should have the CoS values shown in [Table C-6](#).

Table C-6 Solution DSCP-to-CoS Mappings

Service	DSCP	Ethernet CoS	ATM Class	ATM Priority
Multicast video	AF41 (34)	4	VBR-rt (variable bit rate, real time)	5
Unicast video (50%)	AF42 (36)	2		5
	AF43 (38)	1		5
VoIP	EF (46)	5	CBR (constant bit rate)	6
Signaling	CS3 (24)	3	VBR-rt	5
HSD	Default (0)	0	UBR (unspecified bit rate)	0

For Ethernet interfaces, the default DSCP-to-CoS mapping is used, as shown in [Table C-7 on page C-6](#).

Table C-7 Default Ethernet DSCP-to-CoS Mappings

DSCP	Ethernet CoS
0–7	0
8–15	1
16–23	2
24–31	3
32–39	4
40–47	5
48–55	6
56–63	7

Queueing and Thresholds

Each port in the switch has a series of input and output queues that are used as temporary storage areas for data. The queues are implemented in hardware ASICs for each port by means of two types of queues: a special strict-priority (SP) queue, and weighted queues. The SP queues are used for latency-sensitive traffic such as VoIP and are serviced until they are empty. The weighted queues are serviced only when the SP queue is empty. If a weighted queue is being serviced and a frame arrives in the SP queue, the scheduling of frames from the lower queues stops in order to process the frame in the SP queue. Only when the SP queue is empty does the scheduling of packets from the lower queue(s) recommence. Frames are placed in the different queues according to the CoS field in the frame.

Weighted Round Robin (WRR) is the scheduling algorithm used to empty the transmit queues. For each queue a weighting is used to dictate how much data is emptied from the queue before moving onto the next queue.

The Cisco 7600 series and Cisco Catalyst 6000 series also have a means of ensuring that buffers do not overflow. Thresholds are defined in the queues at which the congestion management algorithm can start dropping data. Different CoS values in the same queue can be mapped to different thresholds. Once the buffers have been fully utilized, tail-drop buffer management is used to drop newly arriving frames.

Queueing Structures on Cisco 7600 Series and Cisco Catalyst 6000 Series Line Cards

Each line card on the Cisco 7600 series and the Cisco Catalyst 6000 series has a unique queuing structure, as discussed below.

- 1-GE Line Card (WS-X6816)

This line card uses the queue structure 1P1Q4T on the receive side and 1P2Q2T on the transmit side. Each GE ASIC on the line card supports a total of 2 Mbytes of buffer space, which is equally divided among each of the four GE ports under the control of that ASIC. This means that each GE port supports 512 kbytes of buffering for both receive and transmit. The receive-side buffer holds 73 kbytes, and the transmit-side buffer holds 439 kbytes.

- 1-GE Line Cards (WS-X6724, WS-X6748)

These line cards use the queue structure 1Q8T on the receive side and 1P3Q8T on the transmit side. When a DFC3a is used, two receive queues are available, the new one being a SP queue. Each GE port supports 1.3 Mbytes of buffer space. The receive-side buffer holds 166 kbytes, and the transmit-side buffer holds 1.2 Mbytes.

- 10-GE Line Card (WS-X6704-10GE)

This line card uses the queue structure 8Q8T on the receive side and 1P7Q8T on the transmit side. An SP queue is supported on the transmit side only. Each 10-GE port supports 16 Mbytes. The receive-side buffer holds 2 Mbytes, and the transmit-side buffer holds 14 Mbytes.

Queueing Structures on Cisco Catalyst 4000 Series Line Cards and Ports

Unlike the Cisco 7600 series and the Cisco Catalyst 6500 series, the Cisco Catalyst 4000 series line cards share the same queueing structure. The only difference is in the queue sizes for the 1-GE and 10-GE ports.

- 1-GE Line Card (WS-X4306-GB)

This line card has one receive queue and four transmit queues. One transmit queue can be configured as an SP queue. Each transmit queues holds 1920 packets.

- 10-GE Ports (WS-X4315-GB)

This SupV with 10-GE ports has one receive queue and four transmit queues. One transmit queue can be configured as an SP queue. Each transmit queue holds 2080 packets.

We recommend using three transmit queues for the six types of traffic on the network. In the lowest-priority queue, HSD should be carried. In the middle priority queue, VoD signaling, VoD high priority, VoD low priority, broadcast video, and network signaling should be carried. In the strict priority queue, VoIP should be carried. For more information about the DSCP and CoS values used in the solution, see [DSCP-to-CoS Mapping, page C-5](#).

The amount of buffering the network queueing provides for average video traffic is calculated below.

10-GE Symmetric Topology: Known and Unknown Queue Parameters

For the symmetric topology, 10-GE links are used in the transport. The amount of video traffic in the 10-GE transport approaches 7 Gbps. On the Cisco 7600 series and the Cisco Catalyst 6500 series, the transmit queue (TxQueue) for video is configured for 7 Mbytes. On the Cisco Catalyst 4500 series, the transmit queue for video is configured for approximately 3.2 Mbytes. Known and unknown queue parameters for the 10-GE symmetric topology are listed below.

Known Parameters	Value
<i>link_speed</i>	10,000,000,000 bits/sec
<i>video_queue_size_Cat6k</i>	7,000,000 bytes (56,000,000 bits)
<i>video_queue_size_Cat4k</i>	3,200,000 bytes (25,600,000 bits)
<i>video_link_utilization</i>	70%

Unknown Parameters	Value
<i>avg_buffering</i>	? sec
<i>video_queue_size_Cat6k</i>	$link_speed * avg_buffering * video_link_utilization$
<i>avg_buffering_Cat6k</i>	$(link_speed * video_link_utilization) / video_queue_size_Cat6k$
	$(10,000,000,000 \text{ bits/sec} * 0.70) / 56,000,000 \text{ bits}$
	= 1/125 sec
	= 8 msec
<i>avg_buffering_Cat4k</i>	$(link_speed * video_link_utilization) / video_queue_size_Cat4k$
	$(10,000,000,000 \text{ bits/sec} * 0.70) / 25,600,000 \text{ bits}$
	= 1/274 sec
	= ~4 msec (3.65)

1-GE Asymmetric Topology: Known and Unknown Queue Parameters

For the asymmetric topology, 1-GE links are used in the transport. The maximum amount of video traffic on any of the bidirectional links is approximately 700 Mbps. The amount of video traffic on the unidirectional links approaches 700 Mbps. Known and unknown queue parameters for the 1-GE asymmetric topology are listed below.

Known Parameters	Value
<i>link_speed</i>	1,000,000,000 bits/sec
<i>video_queue_size_Cat6k</i>	256,000 bytes
	(2,048,000 bits)
<i>video_queue_size_Cat4k</i>	2,972,160 bytes
	(~24,000,000 bits)
<i>video_link_utilization</i>	70%

Unknown Parameters	Value
<i>avg_buffering</i>	? sec
<i>video_queue_size_1GE</i>	$link_speed * avg_buffering * video_link_utilization$
<i>avg_buffering_Cat6k</i>	$(link_speed * video_link_utilization) / video_queue_size_Cat6k$
	$(1,000,000,000 \text{ bits/sec} * 0.70) / 2,048,000 \text{ bits}$
	= 1/342 sec
	= ~3 msec

Unknown Parameters	Value
<i>avg_buffering_Cat4k</i>	$(link_speed * video_link_utilization) / video_queue_size_Cat4k$
	$(1,000,000,000 \text{ bits/sec} * 0.70) / 24,000,000 \text{ bits}$
	= 1/3428 sec
	= ~30 msec

1-GE Asymmetric and 10-GE Symmetric Topologies: Known Threshold Parameters

For the 10-GE links, the following thresholds should be configured on the transmit queue that holds broadcast video, VoD, VoD signaling, and network signaling. These thresholds should be configured for tail drop.

For the Cisco 7600 series and Cisco Catalyst 6500 series that have thresholds in the queues, we recommend that two thresholds be used in the middle-priority queue to drop VoD low-priority and VoD high-priority traffic before dropping broadcast video traffic. For the 10-GE links, the following two thresholds should be configured on the transmit queue that holds broadcast video, VoD low, VoD high, VoD signaling, and network signaling. These thresholds should be configured for tail drop.

Known Parameters (10 GE)

Parameter	Value
<i>link_speed</i>	10,000,000,000 bits/sec
<i>avg_buffering</i>	8 msec
<i>VoD_low_link_utilization</i>	45%
<i>VoD_high_link_utilization</i>	85%
<i>VoD_low_threshold</i>	$link_speed * avg_buffer * VoD_low_link_utilization$
	= 10,000,000,000 bits/sec * 0.008 sec * 45%
	= 36,000,000 bits
	= 4,500,000 bytes
<i>VoD_high_threshold</i>	$link_speed * avg_buffer * (VoD_low_link_utilization + VoD_high_link_utilization)$
	= 10,000,000,000 bits/sec * 0.008 sec * 85%
	= 68,000,000 bits
	= 8,500,000 bytes

For the 1-GE links, broadcast video is separated from the VoD traffic, so two thresholds should be configured on the transmit queue of the unidirectional links that hold VoD low, VoD high, VoD signaling, and network signaling.

$$VoD_low_threshold = 50\%$$

$VoD_high_threshold = 100\%$

These thresholds should be configured for tail drop.


Two types of queue management are used for HSD. On the Cisco 7600 series and the Cisco Catalyst 6500 series, Weighted Random Early Drop (WRED) should be configured between 75% and 100% queue utilization. On the Catalyst 4500 series, tail drop should be configured at 100%.

For VoIP, tail drop at 100% utilization should be configured on all three types of switches.

The following figures show the queue size, CoS-to-TxQueue mapping, queue thresholds, and WRR for the 10-GE symmetric and 1-GE asymmetric topologies with the following line cards:

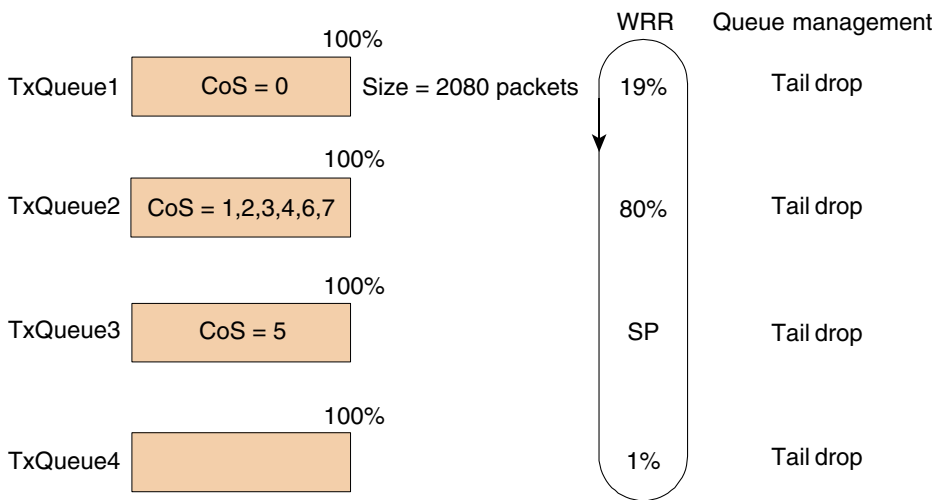
- WS-X6704-10GE (Figure C-2 on page C-10)
- WS-X4513-10GE (Figure C-3 on page C-10)
- WS-X6816 (Figure C-4 on page C-11)
- WS-X4306-GB (Figure C-5 on page C-11)

Figure C-2 WS-X6704-10GE Queueing Structure (10-GE Symmetric)




135485

Figure C-3 WS-X4513-10GE Queueing Structure (10-GE Symmetric)



135486

Figure C-4 WS-X6816 Queueing Structure (1-GE Asymmetric)



Note

Because VoIP, HSD, broadcast video, and VoD signaling are carried on bidirectional links only, no VoD traffic populates TxQueue 2 on the bidirectional links. Because VoD traffic is carried on the unidirectional links only, no VoIP or HSD traffic populates TxQueue 2 on the unidirectional links.

Figure C-5 WS-X4306-GB Queueing Structure (1-GE Asymmetric)

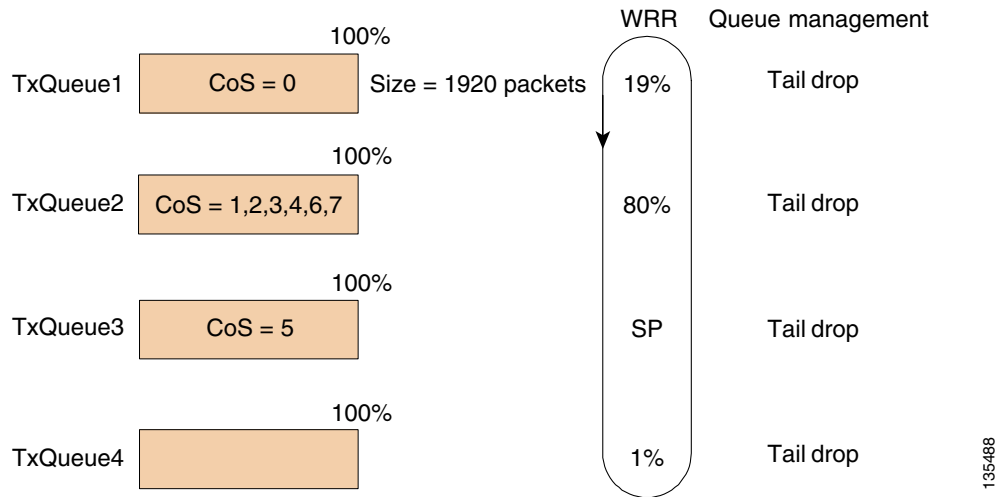


Table C-8 on page C-11 shows the queueing and thresholds used in the solution for a Cisco Catalyst 6000 series. The thresholds for the Cisco Catalyst 4000 series are the same as the queue length.

Table C-8 Solution Queueing and Thresholds for a Cisco Catalyst 6000 Series

Service	Transmit Queue/Percent	Threshold for 1 GE, kB	Threshold for 10 GE, MB
HSD	TxQueue 1/20	~132	4.5

Table C-8 *Solution Queueing and Thresholds for a Cisco Catalyst 6000 Series*

Service	Transmit Queue/Percent	Threshold for 1 GE, kB	Threshold for 10 GE, MB
Multicast video	TxQueue 2/80	220	7
VoD high		220	~6
VoD low		110	3.15
Signaling		220	7
VoIP	TxQueue 3/SP ¹	44	1

1. Strict priority



Configuring DSL Equipment

This chapter presents key details of configuring the DSL equipment as used in the solution, and presents the following topics:

- [Network Diagram, page D-1](#)
- [Hardware and Software Versions, page D-3](#)
- [Configuring Ericsson Components, page D-4](#)
- [Special Issues, page D-15](#)



Note

Ericsson DSL equipment was tested in this solution. This appendix does not provide detailed information about Ericsson products. Refer to Ericsson user documentation for further information.

Network Diagram

[Figure D-1 on page D-2](#) illustrates an example network of Ericsson DSL equipment. A Public Ethernet Manager (PEM) terminal communicates with an Ethernet Controller Node (here an ECN320), which in turn aggregates traffic from one or more Ethernet DSLAM Nodes (here an EDN312xp DSLAM). The DSLAM, in turn, communicates with an HM340d home access gateway (HAG).

[Table D-1 on page D-2](#) lists the VLANs, their descriptions, and addresses for the ECN320 and EDN312xp DSLAM. [Table D-2 on page D-3](#) lists the configuration parameters for the HM340d.

Figure D-1 Example Ericsson Network

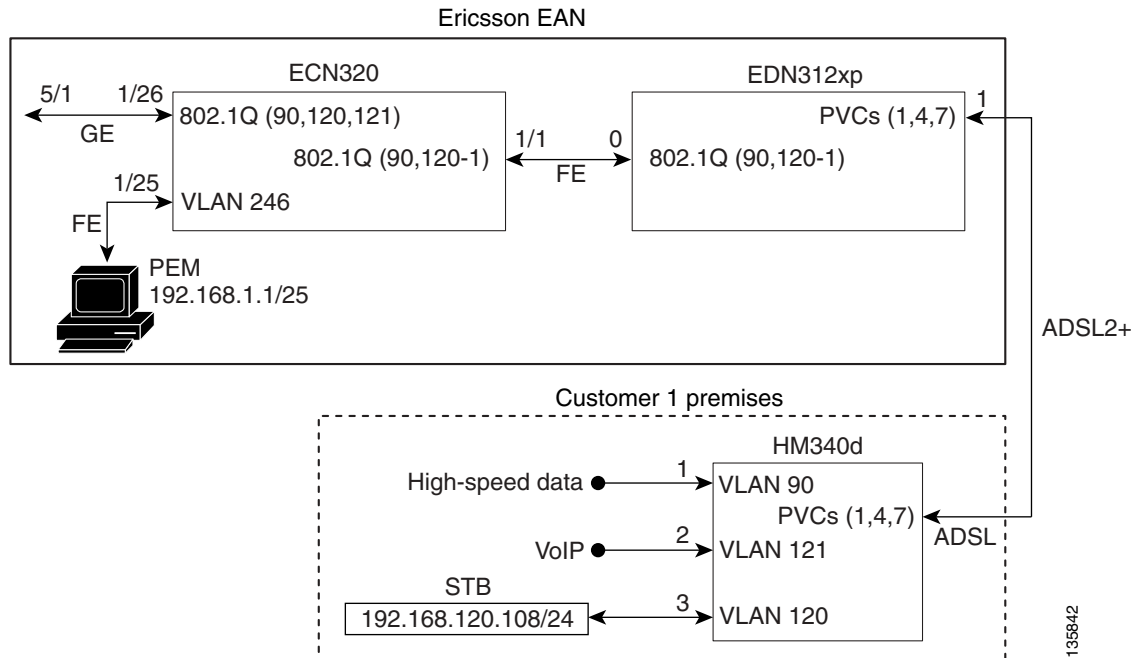


Table D-1 ECN320 and EDN312xp DSLAM VLANs, Descriptions, and IP Addresses

Node	VLAN	Description	IP Address
ECN320	90	High-speed data	Layer 2
	120	Video	Layer 2
	121	VoIP	Layer 2
	246	External interface	192.168.1.100/25
	247	Internal interface	10.0.100.1/16
	248	Untagged	10.1.100.1/24
EDN312xp DSLAM	90	High-speed data	Layer 2
	120	Video	Layer 2
	121	VoIP	Layer 2
	247	Internal interface	10.0.100.1/16
	248	Untagged	10.1.100.1/24

Table D-2 HM340d Configuration Parameters

Traffic	VLAN	HAG Ports	PVC ¹	VPI ²	VCI ³	Encapsulation	Service Class	PCR ⁴	SCR ⁵	MBS ⁶
HSD	90	0	1	8	35	LLC	UBR	—	—	—
VoIP	121	1	4	0	51		CBR	—	300	—
Video	120	2	7	8	59		VBR-RT	1200	600	10

1. Permanent virtual connection
2. Virtual path identifier
3. Virtual connection identifier
4. Peak cell rate
5. Sustained cell rate
6. Maximum burst size

Hardware and Software Versions

Table D-3 on page D-3 lists the hardware and software versions for the Ericsson equipment.

Table D-3 Hardware and Software Versions for Ericsson Equipment

Equipment	Hardware Version	Software Version
Switch	ECN320, R01	CXC 132 7380 R3C06
DSLAM	EDN312xp, R1	CXC 132 8112 R2C05
HAG	HM340dp Home Access Gateway, ZAT 759 89/1, R1A	CXC 132 7758 R2A

Configuring Ericsson Components

The following tasks are presented in the general order in which they should occur:

- [Configuring the Switch, page D-4](#)
- [Configuring the DSLAM, page D-4](#)
- [Configuring the HAG, page D-6](#)
- [Creating Line Configurations, page D-8](#)
- [Creating Services and Profiles, page D-9](#)
- [Creating User Profiles and Adding Services, page D-12](#)

Configuring the Switch

To configure the Ericsson ECN320 switch, use Hyperterminal or a similar application to set parameters as follows:

Interface	Area	Parameter and Setting
Management interface toward PEM	External interface	vlan = 246
		IP = 192.168.1.100
		Netmask = 255.255.255.0
Management interface toward internal nodes	Internal interface	vlan = 247
		IP = 10.0.100.1
		Netmask = 255.255.0.0
		Untagged vlan = 248
		IP = 10.0.100.1
	Netmask = 255.255.255.0	

To save the completed configuration, use the following command:

config save-configuration

Configuring the DSLAM

To configure the access network on the EDN312xp DSLAM, use Ericsson's Network Configuration Manager Application and set parameters as follows:

Choose ...	Area	Parameter and Setting
<i>Network > Line Terminations and Regions</i>	Region	Region Name = Root

Choose ...	Area	Parameter and Setting
Network Elements >	DHCP Server	Name = DHCPServer
		Region = (root)
		Lease Time = 11520
	Domain File Server	IP addr = 192.168.1.1
		FTP User = ftpuser
		FTP PW = ericsson
		Remote Storage Login = eda-mp
		Remote Storage PW = ericsson
	Region = (root)	
	NTP Server	IP addr = 192.168.1.1
	PEM ¹ Domain Service	IP addr = 192.168.1.1
		Region = (root)
Networks >	IP Network	ID = 192.168.1.0
		Mask = 255.255.255.0
		GW = 192.168.1.20
		Max Ethernet Frame Size = 1526
	Domain Subnets	Name = Subnet1
		DHCP Server = DHCPServer
		Domain File Server = 192.168.1.1
		PEM Domain Service = DomainService
		NTP Server = 192.168.1.1
	IP Ranges	Network ID = 192.168.1.0
		Network Mask = 255.255.255.0
		Lower Limit = 192.168.1.50
		Upper Limit = 192.168.1.100

1. Public Ethernet Manager, Ericsson's DSLAM configuration application.



Note

For more information on these configurations, see the document *Ericsson EDA Network and System Administration*.

Configuring the HAG

Two files are used to configure the Ericsson HM340d HAG:

- **atm.conf** describes the ATM permanent virtual circuits (PVCs) that are configured in the HAG, allowing PVC Ethernet frames to be bridged in accordance with RFC 2684. This file is used for both the user profiles that are created. (See [Creating Services and Profiles, page D-9](#).)
- **bridge.conf** maps the ports on the HAG to a specific VLAN/PVC number. This file is copied and edited as appropriate for both the user profiles that are created.



Note

For HAG configuration parameters, see [Table D-2 on page D-3](#). For more information, refer to the document *Ericsson Service Gateway HM340d Operator's Guide*.

Because the HAG configuration files used in the solution are not the Ericsson defaults, you must edit the default files to configure the HAG to forward the three VLANs and services. This information is part of the DSLAM service configurations, and must also be included in the HAG configuration.



Note

Data is on port 1, voice is on port 2, and video is on ports 3 and 4. The DSLAM ports are physically labeled 1 through 4 on the outside of the HAG, although in the file `bridge.conf` these numbers correspond to 0 through 3.

Edit the default files to conform to the following.

atm.conf

The following shows the `atm.conf` file used for both user profiles.

```
# atm.conf -- ATM PVC configuration

# Each line in this file will result in a ATM PVC being configured, and on this PVC
# ethernet frames will be bridged (RFC 2684).

# ATM PVC Interface number 0 (zero) is the management PVC.
```

PVC	VPI	VCI	Encap	Service Class	Parms
0	12	35	llc	nrtvbr	300 150 10
1	8	35	llc	ubr	
2	0	35	llc	ubr	
3	0	43	llc	ubr_pcr	600
4	0	51	llc	cbr	300
5	8	51	llc	nrtvbr	600 300 10
6	8	43	llc	rtvbr	600 300 10
7	8	59	llc	rtvbr	1200 600 10

bridge.conf

The following shows the bridge.conf file used for Profile1.

```
# bridge.conf -- virtual/software ethernet bridge configuration

# The information in this file determines which logical ethernet bridges should be
present.

# Each line is a bridge with the members as a space-separated list, where each member is
either a PVC or a tagged or untagged ethernet port. A PVC member is listed as "pvcN" where
N is the ATM PVC identifier from the /etc/atm.conf configuration file. An untagged port
member is listed as "portN", and a tagged port as "tagged-portN", where N is the port
number (0-3, inclusive).

Each logical port (PVC, tagged or untagged ethernet port) may only be a member of one
bridge. If one untagged port (for example "port2") is used, the corresponding tagged port
("tagged-port2") may not be used, and vice versa.

# VLAN id          Members
90      port0      pvc1
120     port2      pvc7
121     port1      pvc4
```

The following shows the bridge.conf file used for Profile2. Note the addition of port 3.

```
# bridge.conf -- virtual/software ethernet bridge configuration

# The information in this file determines which logical ethernet bridges should be
present.

# Each line is a bridge with the members as a space-separated list, where each member is
either a PVC or a tagged or untagged ethernet port. A PVC member is listed as "pvcN" where
N is the ATM PVC identifier from the /etc/atm.conf configuration file. An untagged port
member is listed as "portN", and a tagged port as "tagged-portN", where N is the port
number (0-3, inclusive).

Each logical port (PVC, tagged or untagged ethernet port) may only be a member of one
bridge. If one untagged port (for example "port2") is used, the corresponding tagged port
("tagged-port2") may not be used, and vice versa.

# VLAN id          Members
90      port0      pvc1
120     port2 port3 pvc7
121     port1      pvc4
```

Creating Line Configurations

Line configurations are required to establish communication between the DSLAM and the HAG. A separate line configuration is used by each profile.

Using the Ericsson PEM configuration application, choose *Service Configuration > DSL Line*, and set (or confirm) parameters as follows:

Profile	Area	Parameter	Setting
1	Channel 0	Name	OneVideoVoiceDataLow
		Transmission mode	Autodetect
		Min. bit rate downstream	7008
		Min. bit rate upstream	512
		Max. bit rate downstream	24000
		Max. bit rate upstream	1408
		Interleave delay downstream	0
		Interleave delay upstream	0
	Line	Transmit PSD	Priority to rate
		Target SNR margin downstream	6.0
		Target SNR margin upstream	6.0
		Max. SNR margin downstream	6.0
		Max. SNR margin upstream	6.0
		Rate adaptation mode	Disabled
2	Channel 0	Name	TwoVideoVoiceDataLow
		Transmission mode	Autodetect
		Min. bit rate downstream	7008
		Min. bit rate upstream	512
		Max. bit rate downstream	24000
		Max bit rate upstream	1408
		Interleave delay downstream	0
		Interleave delay upstream	0
	Line	Transmit PSD	Priority to rate
		Target SNR margin downstream	6.0
		Target SNR margin upstream	6.0
		Max. SNR margin downstream	6.0
		Max. SNR margin upstream	6.0
		Rate adaptation mode	Disabled

Creating Services and Profiles

Using the PEM configuration application, create services and user profiles for video, voice, and data. These services create the bridge between the Ethernet VLAN services for video, voice and data and the ATM PVC (VPI/VCI pairs).

Creating Services and Profiles for Video

Creating a Video Service

To create a video service using the Ericsson PEM configuration application, choose *Service Configuration > Action > Create New*, and set (or confirm) parameters as follows:

Parameter	Setting
Service Name	Video
Customer Service type	Video
CPE access method	Static IP
Relay agent configuration	Not used
IP settings	Enable IGMP snooping (checked)
Broadcast Allowed	Not checked
Default Gateway	192.168.120.1
Enable Mac forced forwarding	Checked
Enable virtual Mac address	Checked
Connections allowed	2
ATM Service Class	VBR-rt
VPI	8
VCI	59
Enable upstream policing	Checked
VLAN Usage	Service VLAN preconfigured to all switches
Ethernet Priority	5
VLAN ID	120

Creating Video Bandwidth Profiles

Create two different bandwidth configurations (profiles) for video. These can be applied to the video service configuration depending on the profile the user is using.

To create video bandwidth profiles using the Ericsson PEM configuration application, choose *Service Configuration > Video > Bandwidth > Create*, and set parameters as follows:

Profile	Parameter	Setting
1	Name	VideoLowBW
	PCR Down/Up	6016/512
	SCR Down/Up	5014/128
	MBS Down/Up	30/30
2	Name	VideoBW
	PCR Down/Up	10016/512
	SCR Down/Up	10016/128
	MBS Down/Up	30/30

Creating Services and Profiles for Voice

Creating a Voice Service

To create a voice service using the Ericsson PEM configuration application, choose *Service Configuration > Action > Create New*, and set (or confirm) parameters as follows:

Parameter	Setting
Service Name	Voice
Customer Service type	Voice
CPE access method	Static IP
Relay agent configuration	Not used
IP settings	Enable IGMP snooping (not checked)
Broadcast Allowed	Not checked
Default Gateway	192.168.121.1
Enable Mac forced forwarding	Checked
Enable virtual Mac address	Checked
Connections allowed	1
ATM Service Class	CBR
VPI	0
VCI	51
Enable upstream policing	Checked
VLAN Usage	Service VLAN preconfigured to all switches

Parameter	Setting
Ethernet Priority	6
VLAN ID	121

Creating a Voice Bandwidth Profile

Create a single bandwidth configuration (profile) for voice. This can be applied to the voice service configuration for both Profile 1 and Profile 2.

To create video bandwidth profiles using the Ericsson PEM configuration application, choose *Service Configuration > Voice > Bandwidth > Create*, and set parameters as follows:

Profile	Parameter	Setting
1, 2	Name	VoiceBW
	Down/Up	320/320
	IP address	192.168.121.107/24

Creating Services and Profiles for Data

Creating a Data Service

To create a data service using the Ericsson PEM configuration application, choose *Service Configuration > Action > Create New*, and set (or confirm) parameters as follows:

Parameter	Setting
Service Name	Data
Customer Service type	Data
CPE access method	Transparent LAN
Relay agent configuration	Not used
IP settings	Enable IGMP snooping (not checked)
Broadcast Allowed	N/A
Enable Mac forced forwarding	N/A
Enable virtual Mac address	N/A
ATM Service Class	UBR
VPI	8
VCI	35
Enable upstream policing	Checked
VLAN Usage	Service VLAN preconfigured to all switches
Ethernet Priority	0
VLAN ID	90

**Note**

No IP address is required because a transparent VLAN for data service is used. A filter is not applicable.

Creating Data Bandwidth Profiles

Create two different bandwidth configurations (profiles) for data. These can be applied to the data service configuration depending on the profile the user is using.

To create video bandwidth profiles using the Ericsson PEM configuration application, choose *Service Configuration > Data > Bandwidth > Create*, and set parameters as follows:

Profile	Parameter	Setting
1	Name	DataLowBW
	PCR Down/Up	1152/512
	SCR Down/Up	N/A
	MBS Down/Up	N/A
2	Name	DataBW
	PCR Down/Up	1728/512
	SCR Down/Up	N/A
	MBS Down/Up	N/A

Creating User Profiles and Adding Services

Line and service configurations must be completed before you can user profiles.

The following tasks use the Ericsson PEM configuration application to create two user profiles and add video, voice, and data services.

Creating Profile 1

Do the following to create Profile 1 and add services.

-
- Step 1** Create the profile.
- Choose *Service Configuration > End User > New EDA End-User*.
 - Under Customer number, enter **User101**.
 - Choose *End User > Line Setup*.
 - Under Line Configuration, select OneVideoVoiceDataLow.
- Step 2** Add video service.
- In the Add Customized Services window, click Add.
 - From the drop-down menu, choose Video.
 - Under Bandwidth, choose VideoLowBW.
 - For Static IP Address, enter **192.168.120.109**.
 - For a filter, choose FilterAll.



Note This filter is created in [Creating an IP Filter, page D-15](#).

- Step 3** Add voice service.
- In the Add Customized Services window, click Add.
 - From the drop-down menu, choose Voice.
 - Under Bandwidth, choose VoiceBW.
 - For the IP Address, enter **192.168.121.107**.
 - For a filter, choose FilterAll.

- Step 4** Add data service.
- In the Add Customized Services window, click Add.
 - From the drop-down menu, choose Data.
 - Under Bandwidth, choose DataLowBW.



Note No IP address is required because a transparent VLAN for data service is used. A filter is not applicable.

- Set the EDN Name and EDF position used by the PEM to identify the line configuration for this user:
EAN Name: **ECN320-192-168-1-100**
MDF Position: **1.0.1**
- Select Line Activate and Apply to activate User101 with the line and service configuration.



Note The connection status LED on the PEM should be green.

Creating Profile 2

Do the following to create Profile 2 and add services.

-
- Step 1** Create the profile.
- Choose *Service Configuration > End User > New EDA End-User*.
 - Under Customer number, enter **User102**.
 - Choose *End User > Line Setup*.
 - Under Line Configuration, select **TwoVideoVoiceDataLow**.

- Step 2** Add video service.
- In the Add Customized Services window, click Add.
 - From the drop-down menu, choose Video.
 - Under Bandwidth, choose VideoBW.
 - For Static IP Address, enter **192.168.120.108, 192.168.120.110**
 - For a filter, choose FilterAll.

- Step 3** Add voice service.
- In the Add Customized Services window, click Add.
 - From the drop-down menu, choose Voice.
 - Under Bandwidth, choose VoiceBW.
 - For the IP Address, enter **192.168.121.107**.
 - For a filter, choose FilterAll.

- Step 4** Add data service.
- In the Add Customized Services window, click Add.
 - From the drop-down menu, choose Data.
 - Under Bandwidth, choose DataBW.



Note No IP address is required because a transparent VLAN for data service is used. A filter is not applicable.

- Set the EDN Name and EDF position used by the PEM to identify the line configuration for this user:
EDN Name: **ECN320-192-168-1-100**
MDF Position: **1.0.2**
- Select Line Activate and Apply to activate User102 with the line and service configuration.



Note The connection status LED on the PEM should be green.

Creating an IP Filter

If a static IP address is used as part of a video, voice, or data service configuration, an IP filter must be applied for the static IP address to work. Ericsson does not provide a default filter that allows all addresses in the downstream direction to be passed through to the HAG. At least one IP address must be entered into the filter, with that IP address to be marked as “allow” or “deny.” Because the **range** command is not supported in the filter configuration in the downstream direction, each address through which traffic is allowed to pass must be entered individually into the filter.

A workaround is to create a filter that denies only one IP address in the downstream direction. The IP address to deny can be any IP address that will not be used to send to, or receive from, the HAG attached to the DSLAM line port for this service configuration. This solution is easier than attempting to add all the IP addresses of all devices that will be sending to, or receiving from, the device attached to the port of the HAG.

Do the following to create an IP filter.

-
- Step 1** Using the Ericsson PEM configuration application, choose *Service Configuration > New EDA Filter*.
 - Step 2** Under Configuration name, enter **FilterAll**.
 - Step 3** Uncheck the box labeled “ICMP security enabled.”
 - Step 4** Create an upstream filter to allow a range of IP addresses.
 - a. Select the Up Stream tab.
 - b. Enter **192.168.0.0 – 255.255.255.255**
 - c. Click Allow.
 - d. Click OK.
 - Step 5** Create a downstream filter to deny one IP address and allow all other addresses.
 - a. Select the Down Stream tab.
 - b. To create a filter that allows any IP addresses except the following (any IP address not used in the system), enter **172.2.2.2**.
 - c. Click Deny.
 - d. Click OK.
 - Step 6** Assign the filter to the desired service and line configuration.
-

Special Issues

Note the following special issues:

1. If multicast (broadcast) video is to be delivered to the STB through the DSLAM, the Service Configuration for Video must have IGMP snooping enabled.
2. At the time of this printing, Ericsson DSL equipment does not support IGMP version 3. If IGMPv3 commands are sent to the Ericsson equipment, messages are discarded and the broadcast is not played through the STB. Consequently, Cisco switches connected to the Ericsson ECN320 switch must send IGMPv2 commands to the Ericsson equipment.

