Tsi
Total Systems Integration, Inc.
"Tying it all Together"

PO Box 0687, 1263 State Route 598
Galion, Ohio 44833-0687
(419) 468-1855
(419) 462-1606 FAX

May 9, 2012

State of West Virginia
Department of Administration
Purchasing Division
2019 Washington Street
Charleston, WV 25305-0130
Attention: Krista Ferrell
(305) 558-2596

George Dallas
President

(419) 544-1331
georgedallas@total-systems.net

Dear Ms. Ferrell

Total Systems Integration (TSI), Inc. and Hewlett Packard (HP) are pleased to have the opportunity to respond to ISCM0113 RFP. The proposed HP networking products are high quality and high performance, and can help The State of West Virginia Office of Technology achieve enhanced productivity, increased business agility and greater competitive advantage. Key differentiators of the proposed HP networking solution include:

- Lower cost of ownership: The proposed HP networking products, featuring industry-leading warranties with technical support and software upgrades, are engineered for high reliability to industry-standard specifications. No need for expensive Cisco *SmartNet* contracts.
- Product lifecycle management: Our team will work with you to ensure that you are informed of product roadmaps and are able to address end-of-life product issues in a proactive fashion. This will enable the West Virginia Broadband Technology Opportunities Program to maintain a consistent and predictable network infrastructure.
- End-to-end solution: HP can deliver world-class, standards based, networking products, plus deployment and maintenance services
- High Performance, non-blocking Networking equipment, with up to 66% lower TCO with unmatched price/performance, lower power and cooling costs, and reduced complexity and no license costs.

TSI and HP are committed to your project's success and we are confident that our solution addresses your critical business requirements. We look forward to meeting with you to review our capabilities, to discuss the benefits of our proposed solution and to explore the next steps in creating a strong and mutually beneficial business relationship.

You can reach me at (419) 544-1331 or georgedallas@total-systems.net.

Sincerely,

George Dallas

George Dallas
President

RECEIVED
2012 MAY 10 AM 10: 20
WV PURCHASING
DIVISION

Executive Summary

Introduction

Total Systems Integration and Hewlett-Packard Company (HP) appreciate the opportunity to present WVOT with an HP networking solution in response to the West Virginia University Network Upgrade RFP ISCM0113. We understand the technical and business requirements of WVOT and this proposal response will demonstrate the depth of HP's expertise as a world-class network solution provider. The proposal will highlight a comprehensive HP networking product portfolio that includes competitively priced, fully featured, Ethernet Enterprise Switch, that will meet your business needs. In addition, HP offers WVOT the flexibility to choose just the right level of maintenance. Our proposal includes a Five (5) Year, 24 x 7 x 4 ON-SITE warranty with software upgrades, and technical support. In addition we have included a block of hours that can be used to train existing maintenance staff and to assist in the installation and configuration of the system.

Meeting Business Requirements

HP has created a global networking powerhouse and is changing the rules of networking. It is raising the bar through innovation to deliver a differentiated portfolio of edge to core and data center networking solutions, complemented by global service and support capabilities. This expanded portfolio delivers best-in-class solutions that enable clients to harness the power of convergence and accelerate business growth at a lower total cost of ownership. In the updated version of research and advisory firm Gartner Inc.'s 2011 Magic Quadrant for Enterprise LAN (Global)[1], HP is positioned in the Leaders quadrant.

HP is in a unique position to deliver on the promise of the Broadband Initiative Infrastructure, with advanced technology, broad innovation, unparalleled expertise in technology services and enterprise services, and our broad partner ecosystem. Here are some of the benefits that WVOT can realize with an HP networking solution.

- **Open standards-based networking** facilitates incremental migration and leverages the existing expertise of trained network engineers and partners. Allows customer-focused innovation and interoperability instead of vendor lock-in, allowing a choice of best-in-class products and solutions with each purchase.
- **Comprehensive interoperability with tools, best practices and expertise** ensures that you can take advantage of HP networking solutions incrementally with no disruption in existing operations and no rip and replace. This allows customers to evolve networks in a deliberate and safe fashion.
- **Better energy efficiency is achieved** with technologies like variable-speed fans and front-to-back cooling. Our solutions complement HP data center smart grid technologies by driving higher utilization and reducing hardware needs as well as power and cooling requirements.
- **Best- in-class solutions working with industry-leading partners** ("HP AllianceONE") has been pretested and configured to run either within the network fabric infrastructure or by way of dedicated platforms. These include Unified Communications and Collaboration (UC&C) partners like Microsoft, Avaya, and Aastra, application delivery partners, like Riverbed and F5, and a variety of security partners for fast time-to-value.
- **Leading warranties** across our entire Enterprise Networking portfolio contribute to significantly lower Total Cost of Ownership (TCO) and reduce reliance on expensive support contracts.
- **HP FlexNetwork architecture,** the only converged networking architecture that spans from the virtualized data center to the virtual workplace for cloud, multimedia, and mobile services with integrated security solutions. It is the only end-to-end networking architecture that solves legacy network challenges by delivering the scale, security, and manageability needed for cloud-based, video-centric, mobile applications.
- **A fully converged and secure network fabric across voice, video, and data.** Optimized for application delivery and integrated with leaders in application networking, unified communications and other areas, customers can quickly and cost-effectively deploy application services across the extended enterprise.

- **Single-pane-of-glass management deeply integrated with industry-leading IT orchestration software** offers seamless heterogeneous network management and provisioning linked directly to end-user and business demands. HP networking solutions are also integrated with solutions from HP Software to facilitate top-to-bottom management and orchestration across the infrastructure.
- **Secure unified wired and wireless solutions** deliver a seamless experience managed from a single pane of glass across the entire secure campus LAN and branch network.
- **"Intelligent edge capability"** offers centralized command and control at the network edge delivering central policy control to reduce management and security overhead, fewer layers of network hierarchy, and higher throughput and a more efficient network.

HP is transforming networking by delivering a complete portfolio of innovative products, solutions, and services designed to meet the complexities faced by enterprise customers.

The portfolio, with superior technology, delivers a dramatically simpler network infrastructure, flexible application-centric environment, open standards, and proven interoperability to dramatically lower total cost of ownership.

HP customers will better align their application and service delivery needs with user demands across their entire extended enterprise.

By changing the rules of networking and driving toward a converged infrastructure, HP will help free up scarce resources to allow customers to invest in innovation that will drive their IT and business forward.

Customers tell us that the single-vendor paradigm has left their current network infrastructures too complex, too rigid, and too expensive. In addition, emerging compute and delivery models like virtualization and cloud computing are driving even stronger needs for heightened security and IT flexibility.

But because the current status quo left IT with a legacy and proprietary networking environment, many IT organizations lack the ability or resources to address rapid business change.

Moreover, Cisco proposals force organizations to contemplate a complete network and infrastructure refresh—with further **proprietary lock-in** and costly investments and without a coherent vision across network infrastructure, security, and management.

The opportunity has never been better for HP to change the rules of networking by bringing superior technologies and proven deployment experience to a $40B (USD) market.

HP is the only company to offer a full portfolio of standards-based, integrated solutions and services developed specifically to solve the complexities of the extended enterprise. As part of a converged infrastructure solution, HP will help customers dramatically simplify their networks, deliver business services more flexibly, and aggressively contain costs to open up new opportunities for business growth and fulfill the promise of a unified, converged IT infrastructure.

HP's family of data center networking solutions is a dramatically more flexible and scalable alternative to Cisco's Nexus portfolio for a virtualization-enabled, converged infrastructure.

HP's portfolio provides unique features for simpler network designs and reduces the cost of ownership with better energy efficiency and stronger management.

Customers can now build a complete standards-based core-to-edge, non-blocking network with a dramatically streamlined architecture requiring fewer systems and staff and delivering a much lower TCO across both CapEx and OpEx.

Meeting RFQ Requirements

Our fully compliant proposal addresses your technology upgrade needs with HP's best of breed Networking Products.

We provided our 12500 series Layer-3 Core switch with Full Advanced IP Premium Routing software and non-blocking switching fabric. The 12500 series is equipped with nine switch fabrics, six 2000 watt power supplies and can be equipped with a variety of networking interfaces from 10 Gbps Copper and Fiber to 1000 Mbps Copper and Fiber.

The 12500 greatly exceed the capabilities of the RFP's Cisco switch in terms of Capacity, Performance, Power Consumption, Price and Total Cost of Ownership.

The HP 12500 Switch Series is the central building block of the HP FlexFabric Solution. As part of the HP Converged Infrastructure strategy, the HP 12500 series delivers high performance, core switching to the HP FlexNetwork architecture.

The HP 12500 series provides low-latency, high-density 1G/10G ports.

The HP 12500 Switch Series provides an unmatched list of features from L2, L3 IPv4 Unicast/Multicast, L3 IPv6 Unicast/Multicast, MPLS, MPLS VPNs, and VPLS, all features included, **without the cost of a complex software licensing model**. Coupled to the HP Intelligent Resilient Framework (IRF), which allows four chassis to be aggregated into a virtual switching fabric (up to 1152* 10GE line rate from L2 up to VPLS), the HP 12500 Switch Series combines simpler data center designs with the highest level of resiliency in the industry.

The HP 12500 Switch Series is ideal for data center core and broadband network backbones. With its extremely low latency, deep buffering, and 10GbE port density, the HP 12500 Switch Series is capable of handling even the most demanding network traffic loads.

The HP 12500 is 40GbE and 100GbE ready.

HP's IRF is a superior technology versus the Cisco vPC. IRF provides true virtualization of physical devices by unifying the control plane and simplifying management. Cisco's vPC on the other hand maintains two separate control planes and two separate complex configurations. IRF also scales to four devices whereas vPC is limited to two.

HP IRF can be used across up to four 12500 chassis, providing massive scalability with up to 2304 10GbE ports in a single logical device.

Closing

Our fully compliant RFP response takes into consideration all the core values and general goals of our customer base, and solidifies a partnership that:

a. is second to none in the industry, offering lower TCO, better performance and less complexity,

b. that provides advanced functionality at affordable prices, and anticipates WVOT's future needs, and

c. is built upon commitment to providing flexible easy-to-use products, which interoperate easily with other manufacturers' products in the industry.

HP simply provides a partnership unlike any other in the network community with better support and lower Total Cost of Ownership.

State of West Virginia
Department of Administration
Purchasing Division
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

Request for Quotation

RFQ NUMBER	PAGE
ISCM0113	1

ADDRESS CORRESPONDENCE TO ATTENTION OF:
KRISTA FERRELL
304-558-2596

RFQ COPY

VENDOR
TYPE NAME/ADDRESS HERE
Total Systems Integration (TSI), Inc.
1263 SR 598
PO 687
Galion, Ohio 44833-0687
(419) 544-1331
georgedallas@total-systems.net

SHIP TO
DEPARTMENT OF ADMINISTRATION
WVOT NETWORKING SUPERVISOR
1900 KANAWHA BLVD. E.
BUILDING 5, 10TH FLOOR
CHARLESTON, WV
25305 304-558-5472

DATE PRINTED	TERMS OF SALE	SHIP VIA	F.O.B.	FREIGHT TERMS
04/12/2012				

BID OPENING DATE: 05/03/2012 BID OPENING TIME 01:30PM

LINE	QUANTITY	UOP	CAT. NO	ITEM NUMBER	UNIT PRICE	AMOUNT
0001	1	EA		205-43		

PERIPHERAL DEVICES AND ACCESSORIES, COMPUTER SYSTEM

REQUEST FOR QUOTATION
(RFQ)

THE WEST VIRGINIA STATE PURCHASING DIVISION FOR THE AGENCY, THE WEST VIRGINIA OFFICE OF TECHNOLOGY, IS SOLICITING BIDS TO PROVIDE THE AGENCY WITH NETWORK INFRASTRUCTURE EQUIPMENT FOR WEST VIRGINIA UNIVERSITY LOCATED IN MORGANTOWN, WEST VIRIGINIA PER THE ATTACHED SPECIFICATIONS.

THIS SOLICITION IS FOR MATERIALS ONLY. NO INSTALLATION IS REQUIRED.

TECHNICAL QUESTIONS CONCERNING THIS SOLICITATION MUST BE SUBMITTED IN WRITING TO KRISTA FERRELL IN THE WEST VIRGINIA STATE PURCHASING DIVISION VIA FAX AT 304-558-4115 OR VIA EMAIL AT KRISTA.S.FERRELL@WV.GOV.

DEADLINE FOR ALL TECHNICAL QUESTIONS IS 04/24/2012 AT THE CLOSE OF BUSINESS.

ANY TECHNICAL QUESTIONS RECEIVED WILL BE ANSWERED BY FORMAL WRITTEN ADDENDUM TO BE ISSUED BY THE PURCHASING DIVISION AFTER THE DEADLINE HAS LAPSED.

VERBAL COMMUNICATION: ANY VERBAL COMMUNICATION BETWEEN THE VENDOR AND ANY STATE PERSONNEL IS NOT BINDING, INCLUDING THAT MADE AT THE MANDATORY PRE-BID MEETING.

SEE REVERSE SIDE FOR TERMS AND CONDITIONS

SIGNATURE George Dallas
TELEPHONE (419) 544-1331
DATE May 09, 2012

TITLE President
FEIN 34-1757782
ADDRESS CHANGES TO BE NOTED ABOVE

WHEN RESPONDING TO RFQ, INSERT NAME AND ADDRESS IN SPACE ABOVE LABELED 'VENDOR'

GENERAL TERMS & CONDITIONS
REQUEST FOR QUOTATION (RFQ) AND REQUEST FOR PROPOSAL (RFP)

1. Awards will be made in the best interest of the State of West Virginia.

2. The State may accept or reject in part, or in whole, any bid.

3. Prior to any award, the apparent successful vendor must be properly registered with the Purchasing Division and have paid the required $125 fee.

4. All services performed or goods delivered under State Purchase Order/Contracts are to be continued for the term of the Purchase Order/Contracts, contingent upon funds being appropriated by the Legislature or otherwise ~~being made available. In the event funds are not appropriated or otherwise available for these services or goods~~ this Purchase Order/Contract becomes void and of no effect after June 30.

5. Payment may only be made after the delivery and acceptance of goods or services.

6. Interest may be paid for late payment in accordance with the ***West Virginia Code.***

7. Vendor preference will be granted upon written request in accordance with the ***West Virginia Code.***

8. The State of West Virginia is exempt from federal and state taxes and will not pay or reimburse such taxes.

9. The Director of Purchasing may cancel any Purchase Order/Contract upon 30 days written notice to the seller.

10. The laws of the State of West Virginia and the ***Legislative Rules*** of the Purchasing Division shall govern the purchasing process.

11. Any reference to automatic renewal is hereby deleted. The Contract may be renewed only upon mutual written agreement of the parties.

12. BANKRUPTCY: In the event the vendor/contractor files for bankruptcy protection, the State may deem this contract null and void, and terminate such contract without further order.

13. HIPAA BUSINESS ASSOCIATE ADDENDUM: The West Virginia State Government HIPAA Business Associate Addendum (BAA), approved by the Attorney General, is available online at **www.state.wv.us/admin/purchase/vrc/hipaa.html** and is hereby made part of the agreement provided that the Agency meets the definition of a Cover Entity (45 CFR §160.103) and will be disclosing Protected Health Information (45 CFR §160.103) to the vendor.

14. CONFIDENTIALITY: The vendor agrees that he or she will not disclose to anyone, directly or indirectly, any such personally identifiable information or other confidential information gained from the agency, unless the individual who is the subject of the information consents to the disclosure in writing or the disclosure is made pursuant to the agency's policies, procedures, and rules. Vendor further agrees to comply with the Confidentiality Policies and Information Security Accountability Requirements, set forth in http://www.state.wv.us/admin/purchase/privacy/noticeConfidentiality.pdf.

15. LICENSING: Vendors must be licensed and in good standing in accordance with any and all state and local laws and requirements by any state or local agency of West Virginia, including, but not limited to, the West Virginia Secretary of State's Office, the West Virginia Tax Department, and the West Virginia Insurance Commission. The vendor must provide all necessary releases to obtain information to enable the director or spending unit to verify that the vendor is licensed and in good standing with the above entities.

16. ANTITRUST: In submitting a bid to any agency for the State of West Virginia, the bidder offers and agrees that if the bid is accepted the bidder will convey, sell, assign or transfer to the State of West Virginia all rights, title and interest in and to all causes of action it may now or hereafter acquire under the antitrust laws of the United States and the State of West Virginia for price fixing and/or unreasonable restraints of trade relating to the particular commodities or services purchased or acquired by the State of West Virginia. Such assignment shall be made and become effective at the time the purchasing agency tenders the initial payment to the bidder.

I certify that this bid is made without prior understanding, agreement, or connection with any corporation, firm, limited liability company, partnership, or person or entity submitting a bid for the same material, supplies, equipment or services and is in all respects fair and without collusion or Fraud. I further certify that I am authorized to sign the certification on behalf of the bidder or this bid.

INSTRUCTIONS TO BIDDERS

1. Use the quotation forms provided by the Purchasing Division. Complete all sections of the quotation form.

2. Items offered must be in compliance with the specifications. Any deviation from the specifications must be clearly indicated by the bidder. Alternates offered by the bidder as **EQUAL** to the specifications must be clearly defined. A bidder offering an alternate should attach complete specifications and literature to the bid. The Purchasing Division may waive minor deviations to specifications.

3. Unit prices shall prevail in case of discrepancy. All quotations are considered F.O.B. destination unless alternate shipping terms are clearly identified in the quotation.

4. All quotations must be delivered by the bidder to the office listed below prior to the date and time of the bid opening. Failure of the bidder to deliver the quotations on time will result in bid disqualifications: Department of Administration, Purchasing Division, 2019 Washington Street East, P.O. Box 50130, Charleston, WV 25305-0130

5. Communication during the solicitation, bid, evaluation or award periods, except through the Purchasing Division, is strictly prohibited (W.Va. C.S.R. §148-1-6.6).

State of West Virginia
Department of Administration
Purchasing Division
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

Request for Quotation

RFQ NUMBER: ISCM0113

PAGE: 2

ADDRESS CORRESPONDENCE TO ATTENTION OF:
KRISTA FERRELL
304-558-2596

VENDOR:
RFQ COPY
TYPE NAME/ADDRESS HERE
Total Systems Integration (TSI), Inc.
1263 SR 598
PO 687
Galion, Ohio 44833-0687
(419) 544-1331
georgedallas@total-systems.net

SHIP TO:
DEPARTMENT OF ADMINISTRATION
WVOT NETWORKING SUPERVISOR
1900 KANAWHA BLVD. E.
BUILDING 5, 10TH FLOOR
CHARLESTON, WV
25305 304-558-5472

DATE PRINTED	TERMS OF SALE	SHIP VIA	F.O.B.	FREIGHT TERMS
04/12/2012				

BID OPENING DATE: 05/03/2012 BID OPENING TIME 01:30PM

LINE	QUANTITY	UOP	CAT NO	ITEM NUMBER	UNIT PRICE	AMOUNT

ONLY INFORMATION ISSUED IN WRITING AND ADDED TO THE RFQ SPECIFICATIONS BY AN OFFICIAL WRITTEN ADDENDUM BY PURCHASING IS BINDING.

NO CONTACT BETWEEN THE VENDOR AND THE AGENCY IS PERMITTED WITHOUT THE EXPRESS WRITTEN CONSENT OF THE STATE BUYER. VIOLATION MAY RESULT IN REJECTION OF THE BID. THE STATE BUYER LISTED ABOVE IS THE SOLE CONTACT FOR ANY AND ALL INQUIRIES AFTER THIS RFQ HAS BEEN RELEASED.

EXHIBIT 10

REQUISITION NO.: ....

ADDENDUM ACKNOWLEDGEMENT

I HEREBY ACKNOWLEDGE RECEIPT OF THE FOLLOWING CHECKED ADDENDUM(S) AND HAVE MADE THE NECESSARY REVISIONS TO MY PROPOSAL, PLANS AND/OR SPECIFICATION, ETC.

ADDENDUM NO.'S:

NO. 1 ...x....

NO. 2 ...x....

NO. 3 ....

NO. 4 ....

NO. 5 ....

I UNDERSTAND THAT FAILURE TO CONFIRM THE RECEIPT OF THE ADDENDUM(S) MAY BE CAUSE FOR REJECTION OF BIDS.

SEE REVERSE SIDE FOR TERMS AND CONDITIONS

SIGNATURE George Dallas

TELEPHONE (419) 544-1331

DATE May 09, 2012

TITLE President

FEIN 34-1757782

ADDRESS CHANGES TO BE NOTED ABOVE

WHEN RESPONDING TO RFQ, INSERT NAME AND ADDRESS IN SPACE ABOVE LABELED 'VENDOR'

State of West Virginia
Department of Administration
Purchasing Division
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

Request for Quotation

RFQ NUMBER	PAGE
ISCM0113	3

ADDRESS CORRESPONDENCE TO ATTENTION OF:
KRISTA FERRELL
304-558-2596

VENDOR
RFQ COPY
TYPE NAME/ADDRESS HERE
Total Systems Integration (TSI), Inc.
1263 SR 598
PO 687
Galion, Ohio 44833-0687
(419) 544-1331
georgedallas@total-systems.net

SHIP TO
DEPARTMENT OF ADMINISTRATION
WVOT NETWORKING SUPERVISOR
1900 KANAWHA BLVD. E.
BUILDING 5, 10TH FLOOR
CHARLESTON, WV
25305 304-558-5472

DATE PRINTED	TERMS OF SALE	SHIP VIA	F.O.B.	FREIGHT TERMS
04/12/2012				

BID OPENING DATE: 05/03/2012 BID OPENING TIME 01:30PM

LINE	QUANTITY	UOP	CAT. NO	ITEM NUMBER	UNIT PRICE	AMOUNT

VENDOR MUST CLEARLY UNDERSTAND THAT ANY VERBAL REPRESENTATION MADE OR ASSUMED TO BE MADE DURING ANY ORAL DISCUSSION HELD BETWEEN VENDOR'S REPRESENTATIVES AND ANY STATE PERSONNEL IS NOT BINDING. ONLY THE INFORMATION ISSUED IN WRITING AND ADDED TO THE SPECIFICATIONS BY AN OFFICIAL ADDENDUM IS BINDING.

George Dallas
....
SIGNATURE

Total Systems Integration (TSI), Inc.
....
COMPANY

9 May 2012
....
DATE

NOTE: THIS ADDENDUM ACKNOWLEDGEMENT SHOULD BE SUBMITTED WITH THE BID.

REV. 09/21/2009
BANKRUPTCY: IN THE EVENT THE VENDOR/CONTRACTOR FILES FOR BANKRUPTCY PROTECTION, THE STATE MAY DEEM THE CONTRACT NULL AND VOID, AND TERMINATE SUCH CONTRACT WITHOUT FURTHER ORDER.

ANY INDIVIDUAL SIGNING THIS BID IS CERTIFYING THAT: (1) HE OR SHE IS AUTHORIZED BY THE BIDDER TO EXECUTE THE BID OR ANY DOCUMENTS RELATED THERETO ON BEHALF OF THE BIDDER, (2) THAT HE OR SHE IS AUTHORIZED TO BIND THE BIDDER IN A CONTRACTUAL RELATIONSHIP, AND (3) THAT THE BIDDER HAS PROPERLY REGISTERED WITH ANY STATE AGENCIES THAT MAY REQUIRE REGISTRATION.

NOTICE

SEE REVERSE SIDE FOR TERMS AND CONDITIONS

SIGNATURE George Dallas
TELEPHONE (419) 544-1331
DATE May 9, 2012
TITLE President
FEIN 34-1757782
ADDRESS CHANGES TO BE NOTED ABOVE

WHEN RESPONDING TO RFQ, INSERT NAME AND ADDRESS IN SPACE ABOVE LABELED 'VENDOR'

State of West Virginia
Department of Administration
Purchasing Division
2019 Washington Street East
Post Office Box 50130
Charleston, WV 25305-0130

Request for Quotation

RFQ NUMBER	PAGE
ISCM0113	4

ADDRESS CORRESPONDENCE TO ATTENTION OF:
KRISTA FERRELL
304-558-2596

VENDOR
RFQ COPY
TYPE NAME/ADDRESS HERE
Total Systems Integration (TSI), Inc.
1263 SR 598
PO 687
Galion, Ohio 44833-0687
(419) 544-1331
georgedallas@total-systems.net

SHIP TO
DEPARTMENT OF ADMINISTRATION
WVOT NETWORKING SUPERVISOR
1900 KANAWHA BLVD. E.
BUILDING 5, 10TH FLOOR
CHARLESTON, WV
25305 304-558-5472

DATE PRINTED	TERMS OF SALE	SHIP VIA	F.O.B.	FREIGHT TERMS
04/12/2012				

BID OPENING DATE: 05/03/2012 BID OPENING TIME 01:30PM

LINE	QUANTITY	UOP	CAT. NO.	ITEM NUMBER	UNIT PRICE	AMOUNT

A SIGNED BID MUST BE SUBMITTED TO:

DEPARTMENT OF ADMINISTRATION
PURCHASING DIVISION
BUILDING 15
2019 WASHINGTON STREET, EAST
CHARLESTON, WV 25305-0130

THE BID SHOULD CONTAIN THIS INFORMATION ON THE FACE OF THE ENVELOPE OR THE BID MAY NOT BE CONSIDERED:

SEALED BID

BUYER: KRISTA FERRELL-FILE 21

RFQ. NO.: ISCM0113

BID OPENING DATE: 05/03/2012

BID OPENING TIME: 1:30 PM

PLEASE PROVIDE A FAX NUMBER IN CASE IT IS NECESSARY TO CONTACT YOU REGARDING YOUR BID:
GEORGEDALLAS@TOTAL-SYSTEMS.NET

CONTACT PERSON (PLEASE PRINT CLEARLY):
GEORGE DALLAS (419) 544-1331

SEE REVERSE SIDE FOR TERMS AND CONDITIONS

SIGNATURE George Dallas	TELEPHONE (419) 544-1331	DATE May 9, 2012
TITLE President	FEIN 34-1757782	ADDRESS CHANGES TO BE NOTED ABOVE

WHEN RESPONDING TO RFQ, INSERT NAME AND ADDRESS IN SPACE ABOVE LABELED 'VENDOR'

REQUEST FOR QUOTATION
ISCM0113

The Acquisition and Contract Administration Section of the Purchasing Division, hereinafter referred to as "State", is soliciting bids for the Office of Technology, hereinafter referred to as "WVOT", on behalf of the West Virginia University, hereinafter referred to as "WVU" located at 1 Waterfront Place, Morgantown, WV 26506, to acquire network infrastructure equipment.

I. PURPOSE

The WVOT is requesting quotations to provide network infrastructure equipment based on the attached equipment list for an upgrade to the WVU facility. WVU will manage, in collaboration with the Office of Technology, all wide area communications in the State, including data, voice and video. WVU serves the network needs of public libraries, K-12 and technical schools, and state government, both for intrastate applications and Internet access.

II. DEFINITIONS

A. "Vendor": The successful bidder
B. "BTOP": The Broadband Technology Opportunities Program (BTOP) is a grant program associated with the American Recovery and Reinvestment Act (ARRA). The grant program was created to promote the development and adoption of broadband throughout the United States, particularly in unserved and underserved areas.

III. BACKGROUND

The West Virginia Broadband Technology Opportunities Program (BTOP) is deploying broadband services throughout the state of West Virginia. This will allow high quality, affordable, broadband services to schools, libraries, hospitals, public safety agencies, jails and residence of West Virginia.

In order to meet the demands of the robust broadband infrastructure deployment, WVU needs to upgrade their network infrastructure to be reliable for the higher volume of internet access.

IV. GENERAL REQUIREMENTS

1. Warranty: Materials and workmanship hereinafter specified and furnished shall be fully guaranteed by the vendor for five (5) years.

The Vendor's obligation under its manufacturer warranty is limited to the cost of repair of the warranted item or replacement thereof, at the vendor's option. Insurance covering said equipment from damage or loss is to be borne by the Vendor until full acceptance of equipment and services.

2. Equipment Requirements: This RFQ specifies Cisco name brand products, "or equal" and must be able to expand and upgrade with the existing WVU hardware and system architecture. Any alternative products must seamlessly fit into, integrate with and interchange with the existing Cisco infrastructure investment with zero loss of feature, functionality and no infrastructure configuration changes. Vendors who are bidding alternates should so state and include pertinent literature and specifications. Failure to provide information for any alternates may be grounds for rejection of bid.

V. INVOICING AND DELIVERY

1. Invoicing: Invoicing shall be made to the Office of Technology, P.O. Box 50110 Charleston, WV 25305.

2. Delivery Location: Equipment shall be shipped to West Virginia University located at 1 Waterfront Place, Morgantown, WV 26506.

3. Delivery Requirements: Delivery requirements are 30 Days or less ARO (Standard delivery time). Standard delivery shall be F.O.B. destination to the above delivery location. Vendor shall include the cost of standard order delivery charges in its bid pricing and is not permitted to charge the Agency separately for such delivery.

VI. AWARD

The Contract shall be awarded to the Vendor that provides the lowest overall total cost for the items listed on the Equipment Bids Price Sheet.

ISCM0113 Bid Price Sheet

Quantity	Product Number	Description	Unit Price	Extended Price
1	ASR-9010-AC	ASR-9010 AC Chassis		$0.00
2	ASR-9010-FAN	ASR-9010 Fan Tray		$0.00
4	A9K-3KW-AC	3kW AC Power Module		$0.00
4	CAB-AC-C6K-TWLK	Power Cord, 250Vac 16A, twist lock NEMA L6-20 plug, US		$0.00
2	A9K-RSP440-TR	ASR9K Route Switch Processor with 440G/slot Fabric and 6GB		$0.00
1	XR-A9K-PXK9-04.02	Cisco IOS XR IP/MPLS Core Software 3DES for RSP440		$0.00
1	A9K-MOD80-TR	80G Modular Linecard, Packet Transport Optimized		$0.00
1	A9K-MPA-20X1GE	ASR 9000 20-port 1GE Modular Port Adapter		$0.00
1	A9K-MPA-4X10GE	ASR 9000 4-port 10GE Modular Port Adapter		$0.00
4	SFP-GE-L	1000BASE-LX/LH SFP (DOM)		$0.00
6	SFP-GE-T	1000BASE-T SFP (NEBS 3 ESD)		$0.00
6	XFP-10GLR-OC192SR	Multirate XFP module for 10GBASE-LR and OC192 SR-1		$0.00
1	A9K-IVRF-LIC	Infrastructure VRF LC License. Support up to 8 VRFs		$0.00
1	A9K-MOD80-TR	80G Modular Linecard, Packet Transport Optimized		$0.00
2	A9K-MPA-4X10GE	ASR 9000 4-port 10GE Modular Port Adapter		$0.00
2	XFP-10G-MM-SR	10GBASE-SR XFP Module		$0.00
1	A9K-IVRF-LIC	Infrastructure VRF LC License. Support up to 8 VRFs		$0.00
1	ASR-9010-4P-KIT	ASR-9010 4 Post Mounting Kit		$0.00
6	A9K-LC-FILR	A9K Line Card Slot Filler		$0.00
1	CON-OSP-A9KIVRFL	ONSITE 24X7X4 Infrastructure VRF LC License		$0.00
1	CON-OSP-A9KIVRFL	ONSITE 24X7X4 Infrastructure VRF LC License		$0.00
1	CON-OSP-A9KMOD8T	ONSITE 24X7X4 80G Modular Linecard,Pcket Transprt Opt		$0.00
1	CON-OSP-A9KMOD8T	ONSITE 24X7X4 80G Modular Linecard,Pcket Transprt Opt		$0.00
1	CON-OSP-A9KMPA2X	ONSITE 24X7X4 ASR 9000 20-port 1GE Modular Port Adaptr		$0.00
1	CON-OSP-A9KMPA4X	ONSITE 24X7X4 ASR 9000 4-port 10GE Modular Port Adaptr		$0.00
2	CON-OSP-A9KMPA4X	ONSITE 24X7X4 ASR 9000 4-port 10GE Modular Port Adaptr		$0.00
2	CON-OSP-A9KRSP4T	ONSITE 24X7X4 ASR9K Route Switch Processr 440G/slot 6G		$0.00
1	CON-OSP-ASR9010A	ONSITE 24X7X4 ASR-9010 AC Chassis		$0.00
1	CON-OSP-XRA9KPXK	ONSITE 24X7X4 Cisco IOS XR IP/MPLS Core Software 3DES		$0.00
		Subtotal		**$0.00**
		Shipping charges		
		TOTAL		**$0.00**

WE ARE PROPOSING HEWLETT PACKARD EQUIPMENT THAT MEETS OR EXCEED SPECIFICATIONS - SEE ATTACHED BOM AND UNIT PRICES

Rev. 09/08

State of West Virginia

VENDOR PREFERENCE CERTIFICATE

Certification and application* is hereby made for Preference in accordance with ***West Virginia Code***, §5A-3-37. (Does not apply to construction contracts). ***West Virginia Code***, §5A-3-37, provides an opportunity for qualifying vendors to request (at the time of bid) preference for their residency status. Such preference is an evaluation method only and will be applied only to the cost bid in accordance with the ***West Virginia Code***. This certificate for application is to be used to request such preference. The Purchasing Division will make the determination of the Resident Vendor Preference, if applicable.

1. **Application is made for 2.5% resident vendor preference for the reason checked:**

___ Bidder is an individual resident vendor and has resided continuously in West Virginia for four (4) years immediately preceding the date of this certification; **or,**

___ Bidder is a partnership, association or corporation resident vendor and has maintained its headquarters or principal place of business continuously in West Virginia for four (4) years immediately preceding the date of this certification; or 80% of the ownership interest of Bidder is held by another individual, partnership, association or corporation resident vendor who has maintained its headquarters or principal place of business continuously in West Virginia for four (4) years immediately preceding the date of this certification; **or,**

___ Bidder is a nonresident vendor which has an affiliate or subsidiary which employs a minimum of one hundred state residents and which has maintained its headquarters or principal place of business within West Virginia continuously for the four (4) years immediately preceding the date of this certification; **or,**

2. **Application is made for 2.5% resident vendor preference for the reason checked:**

___ Bidder is a resident vendor who certifies that, during the life of the contract, on average at least 75% of the employees working on the project being bid are residents of West Virginia who have resided in the state continuously for the two years immediately preceding submission of this bid; **or,**

3. **Application is made for 2.5% resident vendor preference for the reason checked:**

___ Bidder is a nonresident vendor employing a minimum of one hundred state residents or is a nonresident vendor with an affiliate or subsidiary which maintains its headquarters or principal place of business within West Virginia employing a minimum of one hundred state residents who certifies that, during the life of the contract, on average at least 75% of the employees or Bidder's affiliate's or subsidiary's employees are residents of West Virginia who have resided in the state continuously for the two years immediately preceding submission of this bid; **or,**

4. **Application is made for 5% resident vendor preference for the reason checked:**

___ Bidder meets either the requirement of both subdivisions (1) and (2) or subdivision (1) and (3) as stated above; **or,**

5. **Application is made for 3.5% resident vendor preference who is a veteran for the reason checked:**

___ Bidder is an individual resident vendor who is a veteran of the United States armed forces, the reserves or the National Guard and has resided in West Virginia continuously for the four years immediately preceding the date on which the bid is submitted; **or,**

6. **Application is made for 3.5% resident vendor preference who is a veteran for the reason checked:**

___ Bidder is a resident vendor who is a veteran of the United States armed forces, the reserves or the National Guard, if, for purposes of producing or distributing the commodities or completing the project which is the subject of the vendor's bid and continuously over the entire term of the project, on average at least seventy-five percent of the vendor's employees are residents of West Virginia who have resided in the state continuously for the two immediately preceding years.

Bidder understands if the Secretary of Revenue determines that a Bidder receiving preference has failed to continue to meet the requirements for such preference, the Secretary may order the Director of Purchasing to: (a) reject the bid; or (b) assess a penalty against such Bidder in an amount not to exceed 5% of the bid amount and that such penalty will be paid to the contracting agency or deducted from any unpaid balance on the contract or purchase order.

By submission of this certificate, Bidder agrees to disclose any reasonably requested information to the Purchasing Division and authorizes the Department of Revenue to disclose to the Director of Purchasing appropriate information verifying that Bidder has paid the required business taxes, provided that such information does not contain the amounts of taxes paid nor any other information deemed by the Tax Commissioner to be confidential.

Under penalty of law for false swearing (West Virginia Code, §61-5-3), Bidder hereby certifies that this certificate is true and accurate in all respects; and that if a contract is issued to Bidder and if anything contained within this certificate changes during the term of the contract, Bidder will notify the Purchasing Division in writing immediately.

Bidder: ______________________ **Signed:** ______________________

Date: ______________________ **Title:** ______________________

**Check any combination of preference consideration(s) indicated above, which you are entitled to receive.*

ISCM0113 9

RFQ No. ____________

STATE OF WEST VIRGINIA
Purchasing Division

PURCHASING AFFIDAVIT

West Virginia Code §5A-3-10a states: No contract or renewal of any contract may be awarded by the state or any of its political subdivisions to any vendor or prospective vendor when the vendor or prospective vendor or a related party to the vendor or prospective vendor is a debtor and the debt owed is an amount greater than one thousand dollars in the aggregate.

DEFINITIONS:
"Debt" means any assessment, premium, penalty, fine, tax or other amount of money owed to the state or any of its political subdivisions because of a judgment, fine, permit violation, license assessment, defaulted workers' compensation premium, penalty or other assessment presently delinquent or due and required to be paid to the state or any of its political subdivisions, including any interest or additional penalties accrued thereon.

"Debtor" means any individual, corporation, partnership, association, limited liability company or any other form or business association owing a debt to the state or any of its political subdivisions. "Political subdivision" means any county commission; municipality; county board of education; any instrumentality established by a county or municipality; any separate corporation or instrumentality established by one or more counties or municipalities, as permitted by law; or any public body charged by law with the performance of a government function or whose jurisdiction is coextensive with one or more counties or municipalities. "Related party" means a party, whether an individual, corporation, partnership, association, limited liability company or any other form or business association or other entity whatsoever, related to any vendor by blood, marriage, ownership or contract through which the party has a relationship of ownership or other interest with the vendor so that the party will actually or by effect receive or control a portion of the benefit, profit or other consideration from performance of a vendor contract with the party receiving an amount that meets or exceed five percent of the total contract amount.

EXCEPTION: The prohibition of this section does not apply where a vendor has contested any tax administered pursuant to chapter eleven of this code, workers' compensation premium, permit fee or environmental fee or assessment and the matter has not become final or where the vendor has entered into a payment plan or agreement and the vendor is not in default of any of the provisions of such plan or agreement.

Under penalty of law for false swearing (***West Virginia Code*** §61-5-3), it is hereby certified that the vendor affirms and acknowledges the information in this affidavit and is in compliance with the requirements as stated.

WITNESS THE FOLLOWING SIGNATURE

Vendor's Name: Total Systems Integration (TSI), Inc.

Authorized Signature: Marie C. Dallas Date: 05/09/2012

Secretary/Treasurer

State of Ohio

County of Crawford, to-wit:

Taken, subscribed, and sworn to before me this 9th day of May, 20 12.

My Commission expires ____________, 20___.

AFFIX SEAL HERE **NOTARY PUBLIC** ____________

Original Notarized Form was sent under separate cover

Part Number	Description	QTY	Cost	Extend	
	ISCM0113 Network Equipment				
JF431C	HP 12508 AC Switch Chassis	1	$ 5,455.80	$ 5,455.80	
JC665A	HP X421 Chassis Universal Rck Mntg Kit	1	$ 105.00	$ 105.00	
JC067B	HP 12508 Fabric Module	9	$ 1,675.80	$ 15,082.20	
JC072B	HP 12500 Main Processing Unit	2	$ 3,775.80	$ 7,551.60	
JF429A	HP 12500 2000W AC Power Supply	6	$ 751.80	$ 4,510.80	
JC659A	HP 12500 8-port 10GbE SFP+ LEF Module	2	$ 20,995.80	$ 41,991.60	
JC660A	HP 12500 48-port GbE SFP LEF Module	1	$ 12,595.80	$ 12,595.80	
JD092B	HP X130 10G SFP+ LC SR Transceiver	2	$ 637.98	$ 1,275.96	
JD094B	HP X130 10G SFP+ LC LR Transceiver	6	$ 1,553.58	$ 9,321.48	
JD089B	HP X120 1G SFP RJ45 T Transceiver	6	$ 150.78	$ 904.68	
JD119B	HP X120 1G SFP LC LX Transceiver	4	$ 377.58	$ 1,510.32	

			Hardware	**$ 100,305.24**	
	5 Year Maintenance				
UW992E	HP 5y SupportPlus24 Networks 125xx Svc *7 x 24 x 4 On-Site, Software Upgrades, Tech Support*	1	$ 41,174.25	$ 41,174.25	

			Services	**$ 41,174.25**	

			Total Cost:	**$ 141,479.49**	**<<<<**
				============	
	Optional - Pro Services				
HPOSENG	HP On-Site Engineer, training, configure, support, etc.	1	$ 9,000.00	$ 9,000.00	

HP NETWORKING VS. CISCO

6 Compelling Total Cost of Ownership (TCO) Comparisons

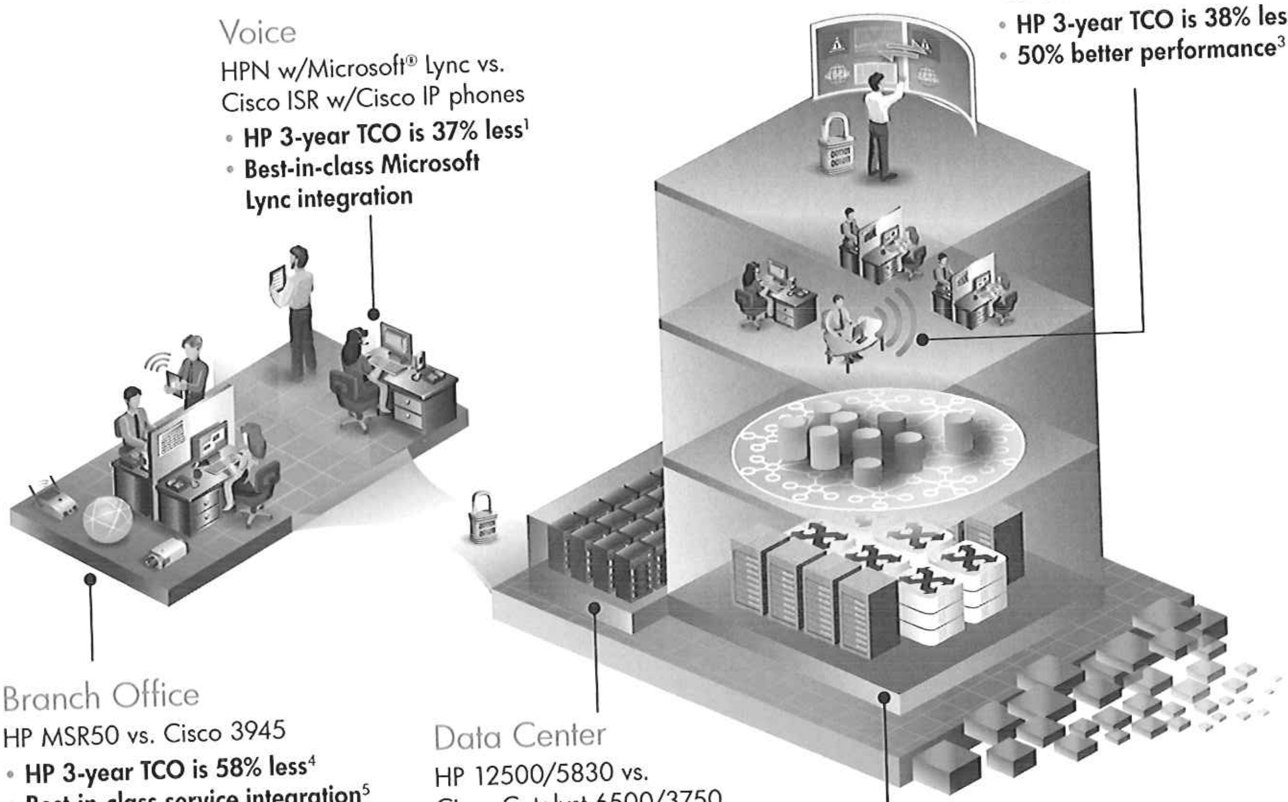

Wireless

HP MSM760 with MSM466 vs. Cisco 5508 with 1142

- **HP 3-year TCO is 38% less**[2]
- **50% better performance**[3]

Voice

HPN w/Microsoft® Lync vs. Cisco ISR w/Cisco IP phones

- **HP 3-year TCO is 37% less**[1]
- **Best-in-class Microsoft Lync integration**

Branch Office

HP MSR50 vs. Cisco 3945

- **HP 3-year TCO is 58% less**[4]
- **Best-in-class service integration**[5]

Data Center

HP 12500/5830 vs. Cisco Catalyst 6500/3750

- **HP 3-year TCO is 50% less**[6]
- **233% better performance**[7]
- **7000% more buffering**[8]

HP 12500/5830 vs. Cisco Nexus 7000/5000/2000

- **HP 3-year TCO is 25% less**[9]
- **100% better routing performance**[10]
- **3900% more buffering**[11]

Campus

HP 10500/3800 vs. Cisco Catalyst 6500/3750

- **HP 3-year TCO is 53% less**[12]
- **376% better performance**[13]

---

[1] Based on 24-user Enterprise Branch solution with HP and Microsoft Lync hardware and software with 3-year maintenance vs. Cisco ISR G2 hardware and software with 3-year maintenance.

[2] Based on list price for HP MSM760 controller and MSM466 access point hardware and software with 3-year maintenance vs. Cisco 5508 and 802.11a/g/n Fixed Unified Access point hardware and software with 3-year maintenance.

[3] Based on MSM466 450 Mbps per radio vs. Cisco 1142 300 Mbps.

[4] Based on list price for HP MSR50-40 hardware and software with 3-year maintenance vs. Cisco 3945 hardware and software with 3-year maintenance.

[5] Service integration provided by AllainceONE partners: Avaya, F5, Microsoft, and Riverbed.

[6] Based on list price for HP 12508 core and HP 5830 TOR hardware and software with 3-year maintenance vs. Cisco Catalyst 6509 core, 6509 distribution, and 4948 TOR hardware and software with 3-year maintenance.

[7] Based on HP 12518 6.66 Tbps throughput capacity vs. Cisco Catalyst 6500 2.0 Tbps throughput capacity.

[8] Based on HP 5830 1.25 Gigabytes buffering vs. Cisco 4948 17.5 Megabytes buffering.

[9] Based on list price for HP 12508 core with HP 5830 TOR hardware and software with 3-year maintenance vs. Cisco Nexus 7010 core, 5548 distribution, and 2248 TOR hardware and software with 3-year maintenance.

[10] Based on HP 12508 960 million pps vs. Cisco Nexus 7010 480 million pps.

[11] Based on HP 5830 1.25 Gigabytes buffering vs. Cisco 2248 32 Megabytes buffering.

[12] Based on list price for HP 10500 core and HP 3800 edge hardware and software with 3-year maintenance vs. Cisco Catalyst 6509 core, 6509 distribution, and 3560 access hardware and software with 3-year maintenance.

[13] Based on HP 10500 1905 million pps and 2.6 Tbps throughput capacity vs. Cisco Catalyst 6509 400 million pps and 2.0 Tbps throughput capacity.

HP 12500 Switch Series

Data sheet

Product overview

The HP 12500 Switch Series comprises a family of powerful, next-generation routing switches with outstanding capacity for the network core or the data center. Besides innovative Intelligent Resilient Framework (IRF) technology that provides unprecedented levels of performance and high availability, 12500 series switches incorporate the Open Application Architecture (OAA), which enables flexible deployment options for new services. These switches also have energy-efficiency features that drive down operational expenses. The 12500 series is ideal for organizations contemplating large-scale data center or campus consolidations, business continuity and disaster recovery sites, metropolitan area network deployments, and other applications requiring a robust, high-performance switching platform.

Key features

- Advanced architecture: midplane, CLOS
- 13.32 Tb switching capacity
- High-density 10 GbE access with 288 1:1, 576 4:1
- 40/100 GbE ready
- Redundant switching fabric, power supply, fan tray

Features and benefits

Quality of Service (QoS)

- **Virtual Output Queue (VOQ):** prevents head-of-line (HOL) blocking per port on peak time and distributes it over a period of time, increasing the switch performance
- **IEEE 802.1p prioritization:** delivers data to devices based on the priority and type of traffic
- **Layer 4 prioritization:** enables prioritization based on TCP/UDP port numbers
- **Broadcast control:** allows limitation of broadcast traffic rate to cut down on unwanted network broadcast traffic
- **Advanced classifier-based QoS:** classifies traffic using multiple match criteria based on Layer 2, 3, and 4 information; applies QoS policies such as setting priority level and rate limit to selected traffic on a per-port or per-VLAN basis
- **Bandwidth shaping:**
  - **Port-based rate limiting:** provides per-port ingress-/egress-enforced maximum bandwidth
  - **Classifier-based rate limiting:** uses ACL to enforce maximum bandwidth for ingress/egress traffic on each port

Data center optimized

- **Very high performance without compromise:** provides 13.32 Tbps, 4320 Mpps (12518 switch), 6.12 Tbps, and 1920 Mpps (12508 switch); leveraging the latest generation of ASICs, the 12500 product family offers outstanding performance and density to build next-generation data centers
- **Very high density 10 GbE connectivity:** the 12518 switch supports up to 576 10-GbE (4:1) or 288 10-GbE (1:1) per physical rack (44RU); the 12508 switch supports up to 256 10-GbE (4:1) or 128 10-GbE (1:1); with two 12508 switches per physical rack (44RU), the density becomes 512 10-GbE (4:1) or 256 10-GbE (1:1)
- **Very high density GbE connectivity:** the 12518 switch supports up to 864 1-GbE (1:1) in a physical (44RU) rack; the 12508 switch supports up to 384 1-GbE (1:1); with two 12508 switches per physical rack (44RU), the density becomes 768 1-GbE (1:1)
- **Four-chassis IRF:** allows the building of large-scale nonblocking, loop-free, metro Layer 2 networks, providing more server access and ultrahigh reliability
- **Scalable system design:** both the 12518 and 12508 switches are built using the latest switching architectures and technologies (CLOS architecture, midplane design), providing the flexibility and scalability for future higher 10 GbE density modules as well as 40 GbE/100 GbE interfaces
- **Ultramodern architecture:** using the latest evolution in switching design, CLOS, the 12500 switch combines performance and ultimate flexibility to provide a smooth evolution path to 25 Tbps; no other switching architecture (Shared Memory/Crossbar) scales to these levels of performance
- **Jumbo Frames:** to accelerate the level of performances, the 12500 switch supports Jumbo Frames (9K) for intra-data-center communication, or for data center to data center traffic (disaster recovery), reducing the amount of time required for data backup and recovery
- **NLB Multicast ARP:** Microsoft® NLB co-works with Multicast ARP to provide servers with load balancing and fault switchover, which lowers costs and investment

Compartmentalization

- **Department protection:** using network virtualization standards (QinQ, VRF, and MPLS), the 12500 switch allows organizations to isolate different business units with different resources (VRFs); using standard-based mechanisms, the network is completely virtualized, reducing cost and operations
- **IEEE 802.1ah Provider Backbone Bridge (Mac in Mac):** Provider Backbone Bridge (PBB) is a Layer 2 VPN technology that allows a complete separation of customer and provider domains by sealing the user MAC in the service provider MAC, which enhances the scalability of an Ethernet network

Management

- **sFlow:** provides scalable, ASIC-based network monitoring and accounting; this allows network operators to gather a variety of sophisticated network statistics and information for capacity planning and real-time network monitoring purposes
- **IEEE 802.1ab LLDP discovery:** advertises and receives management information from adjacent devices on a network
- **USB support:**
  - **File copy:** allows users to copy switch files to and from a USB flash drive

- **Multiple configuration files:** can be stored to the flash image
- **Command-line interface (CLI):** provides a secure, easy-to-use command-line interface for configuring the module via SSH or a switch console; provides direct real-time session visibility
- **Logging:** provides local and remote logging of events via SNMP (v2c and v3) and syslog; provides log throttling and log filtering to reduce the number of log events generated
- **Management interface control:** each of the following interfaces can be enabled or disabled depending on security preferences: console port, telnet port, and SSH port
- **Out-of band-interface:** isolates management traffic from user data plane traffic for complete isolation and total reachability, no matter what happens in the data plane
- **Network management:** Intelligent Management Console (IMC) centrally configures, updates, monitors, and troubleshoots
- **Network management:** SNMP v2c/v3 MIB-II with traps
- **RADIUS accounting:** logs all session details that can be used to generate usage reports or interface to a billing system
- **RMON:** provides advanced monitoring and reporting capabilities for statistics, history, alarms, and events
- **Remote Intelligent Mirroring:** mirrors ingress ACL-selected traffic from a switch port or VLAN to a local or remote switch port anywhere on the network

Connectivity

- **IPv6 native support:**
  - **IPv6 host:** enables switches to be managed and deployed at the IPv6 network's edge
  - **Dual stack (IPv4 and IPv6):** transitions from IPv4 to IPv6, supporting connectivity for both protocols
  - **Multicast Listener Discovery (MLD) snooping:** forwards IPv6 multicast traffic to the appropriate interface
  - **IPv6 ACL/QoS:** supports ACL and QoS for IPv6 network traffic, preventing traffic flooding
  - **IPv6 routing:** supports IPv6 static routes and IPv6 versions of RIP and OSPF routing protocols

Performance

- **13.32 Tbps (12518 switch) and 6.12 Tbps (12508 switch) fully nonblocking CLOS architecture:** includes a high-performance switch design with a nonblocking architecture
- **High-performance bandwidth:** with up to 13.32 Tbps capacity, providing nonblocking throughput for 288 10-GbE ports at Layer 2 and Layer 3 IPv4, Layer 3 IPv6, and MPLS (12518 switch), or 128 10-GbE ports (12508 switch)
- **Hardware-based wire-speed access control lists (ACLs):** feature-rich ACL implementation (TCAM-based) helps ensure high levels of security and ease of administration without impacting network performance
- **High-performance processor system:** the supervisory module uses three different processors to isolate key tasks: control plane (STP, OSPF, BGP, MPLS, etc.), fast recovery protocols (RRPP, BFD, etc.), and chassis management (temperature, power, etc.)

Resiliency and high availability

- **Intelligent Resilient Framework (IRF):** creates virtual resilient switching fabrics, where two or more switches perform as a single Layer 2 switch and Layer 3 router; switches do not have to be co-located and can be part of a disaster-recovery system; servers or switches can be attached using standard LACP for automatic load balancing and high availability; simplifies network operation by eliminating the complexity of Spanning Tree Protocol, Equal-Cost Multipath (ECMP), or VRRP
- **Ultrafast protocol convergence:** enables link connectivity monitoring and reduces network convergence time for RIP, OSPF, BGP, IS-IS, VRRP, MPLS, and IRF
- **Device Link Detection Protocol (DLDP):** monitors link connectivity and shuts down ports at both ends if unidirectional traffic is detected, preventing loops in STP-based networks
- **Complete set of routing protocols (Layer 3 IPv4 and IPv6):** doesn't require customers to think about which protocol is being supported by the 12500 switch; virtually all existing routing protocols (RIP, OSPF, IS-IS, and BGP) are supported for both Layer 3 IPv4 and Layer 3 IPv6; this is also the case for both unicast and multicast, with complete support of PIM-DM, PIM-SM, PIM-SSM, and MSDP
- **Hot patching:** the 12500 switch supports hot patching, allowing in-service patching for some isolated software problems
- **Non Stop Forwarding/Graceful Restart (NSF/GR):** using standardized-based IETF protocols, the 12500 switch provides nonstop forwarding (switching/routing) for Layer 3 routing protocols (Control Plane - OSPF, BGP, and MPLS), providing hitless failover
- **Ultrareliable architecture:** combining hardware redundancy at every layer (power supplies, fans, supervisory modules, etc.) and a multilayered software approach based on the Resilient Virtual Switching Fabric concept (using the IRF technology), the 12500 product family is able to provide the highest level of availability; following design guidelines from HP, customers are now able to build data centers providing an end-to-end availability reaching five 9s
- **Rapid Ring Protection Protocol (RRPP):** provides fast recovery for ring Ethernet-based topology

Layer 2 switching

- **Multiple VLAN Registration Protocol (MVRP):** help to maintain VLAN configuration dynamically based on current network configurations
- **GARP VLAN Registration Protocol:** allows automatic learning and dynamic assignment of VLANs
- **IP multicast snooping and data-driven IGMP:** automatically prevents flooding of IP multicast traffic
- **IEEE 802.1ad QinQ:** increases the scalability of an Ethernet network by providing a hierarchical structure; connects multiple LANs on a high-speed campus or metro network
- **Bridge Protocol Data Unit (BPDU) tunneling:** transmits Spanning Tree Protocol BPDUs transparently, allowing correct tree calculations across service providers, WANs, or MANs
- **VLAN support and tagging:** supports IEEE 802.1Q (4K VLAN IDs)
- **Spanning Tree:** brought by Comware, the 12500 product family supports the entire set of STP protocols (STP, RSTP, and MSTP), facilitating a complete integration with standard networks

Layer 3 routing

- **Layer 3 IPv4 routing:** provides routing of IPv4 at media speed; supports static routes, RIP and RIPv2, OSPF, IS-IS, and BGP
- **RIP and RIPng support:** provides complete support of RIP for both IPv4 and IPv6
- **OSPF and OSPFv3 support:** provides complete support of OSPF for both IPv4 and IPv6
- **IS-IS and IS-ISv6 support:** provides complete support of IS-IS for both IPv4 and IPv6
- **Equal-Cost Multipath (ECMP):** enables multiple equal-cost links in a routing environment to increase link redundancy and scale bandwidth
- **Layer 3 IPv6 routing:** provides routing of IPv6 at media speed; supports static routes, RIPng, OSPFv3, IS-ISv6, and BGP4+
- **IPv6 tunneling:** allows a smooth transition from IPv4 to IPv6 by encapsulating IPv6 traffic over an existing IPv4 infrastructure
- **Complete multicast protocol stack:** PIM-DM, PIM-SM, PIM-SSM, MSDP, and extensions to BGP provide one of the most complete multicast protocol stacks

- **Policy routing:** allows custom filters for increased performance and security; supports ACLs, IP prefix, AS paths, community lists, and aggregate policies
- **MPLS support:** provides extended support of MPLS, including MPLS VPNs and MPLS Traffic Engineering (MPLS TE)
- **VPLS support:** provides extended support of VPLS for data center to data center communication at Layer 2; provides support of hierarchical VPLS for scalability

Security

- **Control Plane Policing (CoPP):** protection against DoS attacks at infrastructure routers and switches; ease of configuration for control plane policies
- **IEEE 802.1X and RADIUS network logins:** control port-based access for authentication and accountability
- **Secure FTP:** allows secure file transfer to and from the switch; protects against unwanted file downloads or unauthorized copying of a switch configuration file
- **Switch management logon security:** can require either RADIUS or TACACS+ authentication for secure switch CLI logon
- **DHCP protection:** blocks DHCP packets from unauthorized DHCP servers, preventing denial-of-service attacks
- **Dynamic ARP protection:** blocks ARP broadcasts from unauthorized hosts, preventing eavesdropping or theft of network data
- **Secure Shell (SSHv2):** encrypts all transmitted data for secure, remote CLI access over IP networks
- **Secure management access:** securely encrypts all access methods (CLI, GUI, or MIB) through SSHv2 and SNMPv3
- **Access control lists (ACLs):** provide IPv4 and IPv6 filtering based on source/destination IP address/subnet and source/destination TCP/UDP port number
- **Media access control (MAC) authentication:** provides simple authentication based on a user's MAC address; supports local or RADIUS-based authentication

Convergence

- **Layer 2, 3, and 4 QoS mechanisms:** support DiffServ priority tagging based on IP address, IP Type of Service (ToS), Layer 3 protocol, TCP/UDP port number, and source port
- **IP multicast snooping and data-driven IGMP:** automatically prevent flooding of IP multicast traffic
- **LLDP-MED:** is a standard extension that automatically configures network devices, including LLDP-capable IP phones
- **Internet Group Management Protocol (IGMP):** is used by IP hosts to establish and maintain multicast groups; supports IGMPv1, v2, and v3; utilizes Any-Source Multicast (ASM) or Source-Specific Multicast (SSM) to manage IPv4 multicast networks
- **Protocol Independent Multicast (PIM):** is used for IPv4 and IPv6 multicast applications; supports PIM Dense Mode (PIM-DM), Sparse Mode (PIM-SM), and Source-Specific Mode (PIM-SSM)
- **Multicast Source Discovery Protocol (MSDP):** is used for inter-domain multicast applications, allowing multiple PIM-SM domains to interoperate
- **Multicast VLAN:** allows multiple VLANs to receive the same IPv4 or IPv6 multicast traffic, reducing network bandwidth demand by eliminating multiple streams to each VLAN

Monitor and diagnostics

- **Port mirroring:** enables traffic on a port to be simultaneously sent to a network analyzer for monitoring
- **CFD (802.1ag):** connectivity fault detection (CFD) provides a Layer 2 link Operations, Administration, and Maintenance (OAM) mechanism used for link connectivity detection and fault locating

Investment protection

- **Modular switch fabric:** provides investment protection by enabling future performance upgrades and increased port density
- **Environmentally friendly:** ROHS support and low power consumption based on the latest technology provide outstanding power efficiency

Warranty and support

- **1-year warranty:** with advance replacement and 10-calendar-day delivery (available in most countries)
- **Electronic and telephone support:** limited electronic and telephone support is available from HP; to reach our support centers, refer to www.hp.com/networking/contact-support; for details on the duration of support provided with your product purchase, refer to www.hp.com/networking/warrantysummary
- **Software releases:** to find software for your product, refer to www.hp.com/networking/support; for details on the software releases available with your product purchase, refer to www.hp.com/networking/warrantysummary

HP 12500 Switch Series

Specifications

	 **HP 12518 AC Switch Chassis (JF430C)** **HP 12518 DC Switch Chassis (JC653A)**	 **HP 12508 AC Switch Chassis (JF431C)** **HP 12508 DC Switch Chassis (JC652A)**
Ports	18 open module slots 2 MPU (for management modules) slots 9 switch fabric slots Supports a maximum of 576 10-GbE ports or 864 Gigabit ports, or a combination	8 open module slots 2 MPU (for management modules) slots 9 switch fabric slots Supports a maximum of 256 10-GbE ports or 384 Gigabit ports, or a combination
Physical characteristics		
Dimensions	29.13(d) x 17.4(w) x 66.38(h) in. (74 x 44.2 x 168.6 cm) (38U height)	29.13(d) x 17.4(w) x 38.39(h) in. (74.0 x 44.2 x 97.5 cm) (22U height)
Weight	352.74 lb. (160 kg)	209.44 lb. (95 kg)
Full configuration weight	639.33 lb. (290 kg)	374.78 lb. (170 kg)
Memory and processor		
Gigabit module	PowerPC @ 667 MHz, 1 GB RAM; packet buffer size: 512 MB RAM (Ingress, shared by 24 1-GbE ports)	PowerPC @ 667 MHz, 1 GB RAM; packet buffer size: 512 MB RAM (Ingress, shared by 24 1-GbE ports)
10G module	PowerPC @ 667 MHz, 1 GB RAM; packet buffer size: 512 MB RAM (Ingress/shared by 2 10-GbE ports)	PowerPC @ 667 MHz, 1 GB RAM; packet buffer size: 512 MB RAM (Ingress/shared by 2 10-GbE ports)
Management module	Primary CPU: PowerPC @ 1000 MHz, 128 MB flash MB, 256 MB compact flash, 4 GB RAM	Primary CPU: PowerPC @ 1000 MHz, 128 MB flash MB, 256 MB compact flash, 4 GB RAM
Fabric	PowerPC @ 400 MHz, 128 MB RAM MB	PowerPC @ 400 MHz, 128 MB RAM MB
Mounting	Mounts in an EIA-standard 19 in. telco rack or equipment cabinet	Mounts in an EIA-standard 19 in. telco rack or equipment cabinet
Performance		
Throughput	4320 million pps	1920 million pps
Routing/Switching capacity	13320 Gbps	6120 Gbps
Environment		
Operating temperature	32°F to 104°F (0°C to 40°C)	32°F to 104°F (0°C to 40°C)
Operating relative humidity	5% to 95%, noncondensing	5% to 95%, noncondensing
Nonoperating/Storage temperature	-40°F to 158°F (-40°C to 70°C)	-40°F to 158°F (-40°C to 70°C)
Nonoperating/Storage relative humidity	5% to 95%, noncondensing	5% to 95%, noncondensing
Electrical characteristics		
Description		Achieved Miercom Certified Green Award 10 GbE modules consume half the power compared to competitive products; redundant, scalable, 90% efficient power supplies deliver high reliability in the data center; new ASIC technology has low power consumption when providing rich features.
Maximum heat dissipation	32859 BTU/hr (34666.24 kJ/hr)	14587 BTU/hr (15389.29 kJ/hr)
Voltage	100-120/200-240 VAC	100-120/200-240 VAC
DC voltage	-48 to -60, rated/-40 to -72, maximum, VDC	-48 to -60, rated/-40 to -72, maximum, VDC
Maximum power rating	10700 W	4750 W
Frequency	50/60 Hz	50/60 Hz
Notes	Maximum power rating and maximum heat dissipation are the worst-case theoretical maximum numbers provided for planning the infrastructure with fully loaded PoE (if equipped), 100% traffic, all ports plugged in, and all modules populated.	Maximum power rating and maximum heat dissipation are the worst-case theoretical maximum numbers provided for planning the infrastructure with fully loaded PoE (if equipped), 100% traffic, all ports plugged in, and all modules populated.
Safety	CE Labeled; cUL Certified; UL Listed; EN 60825-1 Safety of Laser Products-Part 1; EN 60825-2 Safety of Laser Products-Part 2; IEC 60825; IEC 60950-1:2001 (with CB Report); CAN/CSA-C22.2 No. 60950-1-03; Anatel; ULAR; GOST; EN 60950-1/A11; FDA 21 CFR Subchapter J; NOM; UL 60950-1:2003; EN 60950-1:2001; ROHS Compliance	CE Labeled; cUL Certified; UL Listed; EN 60825-1 Safety of Laser Products-Part 1; EN 60825-2 Safety of Laser Products-Part 2; IEC 60825; IEC 60950-1:2001 (with CB Report); CAN/CSA-C22.2 No. 60950-1-03; Anatel; ULAR; GOST; EN 60950-1/A11; FDA 21 CFR Subchapter J; NOM; UL 60950-1:2003; EN 60950-1:2001; ROHS Compliance
Emissions	VCCI Class A; EN 55022 Class A; VCCI V-3/2000.04; ICES-003 Class A; AS/NZS CISPR22 Class A; EMC Directive 2004/108/EC; FCC (CFR 47, Part 15) Class A	VCCI Class A; EN 55022 Class A; VCCI V-3/2000.04; ICES-003 Class A; AS/NZS CISPR22 Class A; EMC Directive 2004/108/EC; FCC (CFR 47, Part 15) Class A
Immunity		
Generic	ETSI EN 300 386 V1.3.3	ETSI EN 300 386 V1.3.3
EN	EN 55024:1998+ A1:2001 + A2:2003	EN 55024:1998+ A1:2001 + A2:2003
ESD	EN 61000-4-2; IEC61000-4-2	EN 61000-4-2; IEC61000-4-2
Radiated	EN 61000-4-3; IEC61000-4-3	EN 61000-4-3; IEC61000-4-3
EFT/Burst	EN 61000-4-4; IEC61000-4-4	EN 61000-4-4; IEC61000-4-4
Surge	EN 61000-4-5; IEC61000-4-5	EN 61000-4-5; IEC61000-4-5
Conducted	EN 61000-4-6; IEC61000-4-6	EN 61000-4-6; IEC61000-4-6
Power frequency magnetic field	IEC 61000-4-8; EN61000-4-8	IEC 61000-4-8; EN61000-4-8
Voltage dips and interruptions	EN 61000-4-11; IEC61000-4-11	EN 61000-4-11; IEC61000-4-11

Specifications (continued)

	HP 12518 AC Switch Chassis (JF430C) HP 12518 DC Switch Chassis (JC653A)	HP 12508 AC Switch Chassis (JF431C) HP 12508 DC Switch Chassis (JC652A)
Harmonics	EN 61000-3-2, IEC 61000-3-2	EN 61000-3-2, IEC 61000-3-2
Flicker	EN 61000-3-3, IEC 61000-3-3	EN 61000-3-3, IEC 61000-3-3
Management	IMC - Intelligent Management Center; command-line interface; out-of-band management (serial RS-232C); SNMP Manager; Telnet; RMON1; FTP; in-line and out-of-band; terminal interface (serial RS-232C); modem interface	IMC - Intelligent Management Center; command-line interface; out-of-band management (serial RS-232C); SNMP Manager; Telnet; RMON1; FTP; in-line and out-of-band; terminal interface (serial RS-232C); modem interface
Services	3-year, 4-hour onsite, 13x5 coverage for hardware (UX046E) 3-year, 4-hour onsite, 24x7 coverage for hardware (UX049E) 3-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 SW phone support and SW updates (UX052E) 3-year, 24x7 SW phone support, software updates (UX055E) Installation with minimum configuration, system-based pricing (UX034E) 4-year, 4-hour onsite, 13x5 coverage for hardware (UX047E) 4-year, 4-hour onsite, 24x7 coverage for hardware (UX050E) 4-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 software phone (UX053E) 4-year, 24x7 SW phone support, software updates (UX056E) 5-year, 4-hour onsite, 13x5 coverage for hardware (UX048E) 5-year, 4-hour onsite, 24x7 coverage for hardware (UX051E) 5-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 software phone (UX054E) 5-year, 24x7 SW phone support, software updates (UX057E) 3 Yr 6 hr Call-to-Repair Onsite (UX058E) 4 Yr 6 hr Call-to-Repair Onsite (UX059E) 5 Yr 6 hr Call-to-Repair Onsite (UX060E) 1-year, 4-hour onsite, 13x5 coverage for hardware (HR489E) 1-year, 4-hour onsite, 24x7 coverage for hardware (HR490E) 1-year, 6 hour Call-To-Repair Onsite for hardware (HR493E) 1-year, 24x7 software phone support, software updates (HR492E) 1-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 software phone support and software updates (HR491E) Refer to the HP website at www.hp.com/networking/services for details on the service-level descriptions and product numbers. For details about services and response times in your area, please contact your local HP sales office.	3-year, 4-hour onsite, 13x5 coverage for hardware (UW984E) 3-year, 4-hour onsite, 24x7 coverage for hardware (UW987E) 3-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 SW phone support and SW updates (UW990E) 3-year, 24x7 SW phone support, software updates (UW993E) Installation with minimum configuration, system-based pricing (UX034E) 4-year, 4-hour onsite, 13x5 coverage for hardware (UW985E) 4-year, 4-hour onsite, 24x7 coverage for hardware (UW988E) 4-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 software phone (UW991E) 4-year, 24x7 SW phone support, software updates (UW994E) 5-year, 4-hour onsite, 13x5 coverage for hardware (UW986E) 5-year, 4-hour onsite, 24x7 coverage for hardware (UW989E) 5-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 software phone (UW992E) 5-year, 24x7 SW phone support, software updates (UW995E) 3 Yr 6 hr Call-to-Repair Onsite (UW996E) 4 Yr 6 hr Call-to-Repair Onsite (UW997E) 5 Yr 6 hr Call-to-Repair Onsite (UW998E) 1-year, 4-hour onsite, 13x5 coverage for hardware (HR494E) 1-year, 4-hour onsite, 24x7 coverage for hardware (HR495E) 1-year, 6 hour Call-To-Repair Onsite for hardware (HR498E) 1-year, 24x7 software phone support, software updates (HR497E) 1-year, 4-hour onsite, 24x7 coverage for hardware, 24x7 software phone support and software updates (HR496E) Refer to the HP website at www.hp.com/networking/services for details on the service-level descriptions and product numbers. For details about services and response times in your area, please contact your local HP sales office.

Specifications (continued)

HP 12518 AC Switch Chassis (JF430C)
HP 12518 DC Switch Chassis (JC653A)

HP 12508 AC Switch Chassis (JF431C)
HP 12508 DC Switch Chassis (JC652A)

Standards and protocols
(applies to all products in series)

BGP
RFC 1657 Definitions of Managed Objects for BGPv4
RFC 1771 BGPv4
RFC 1772 Application of the BGP
RFC 1773 Experience with the BGP-4 Protocol
RFC 1774 BGP-4 Protocol Analysis
RFC 1965 BGP4 confederations
RFC 1997 BGP Communities Attribute
RFC 1998 PPP Gandalf FZA Compression Protocol
RFC 2385 BGP Session Protection via TCP MD5
RFC 2439 BGP Route Flap Damping
RFC 2796 BGP Route Reflection
RFC 2842 Capability Advertisement with BGP-4
RFC 2858 BGP-4 Multi-Protocol Extensions
RFC 2918 Route Refresh Capability

Denial of service protection
RFC 2267 Network Ingress Filtering
Automatic Filtering of well known Denial of Service Packets
CPU DoS Protection
Rate Limiting by ACLs

Device management
RFC 1155 Structure and Mgmt Information (SMIv1)
RFC 1157 SNMPv1/v2c
RFC 1305 NTPv3
RFC 1945 Hypertext Transfer Protocol – HTTP/1.0
RFC 2271 FrameWork
RFC 2452 MIB for TCP6
RFC 2454 MIB for UDP6
RFC 2573 (SNMPv3 Applications)
RFC 2578-2580 SMIv2
RFC 2579 (SMIv2 Text Conventions)
RFC 2580 (SMIv2 Conformance)
RFC 2819 (RMON groups Alarm, Event, History and Statistics only)
RFC 2819 RMON
RFC 3417 (SNMP Transport Mappings)
SNMP v3 and RMON RFC support
SSHv1/SSHv2 Secure Shell
TACACS/TACACS+

General protocols
IEEE 802.1ad Q-in-Q
IEEE 802.1ag Service Layer OAM
IEEE 802.1ah Provider Backbone Bridges
IEEE 802.1D MAC Bridges
IEEE 802.1p Priority
IEEE 802.1Q VLANs
IEEE 802.1s Multiple Spanning Trees
IEEE 802.1v VLAN classification by Protocol and Port
IEEE 802.1w Rapid Reconfiguration of Spanning Tree
IEEE 802.1X PAE
IEEE 802.3ab 1000BASE-T
IEEE 802.3ad Link Aggregation (LAG)
IEEE 802.3ae 10-Gigabit Ethernet
IEEE 802.3ah Ethernet in First Mile over Point to Point Fiber - EFMF
IEEE 802.3i 10BASE-T
IEEE 802.3u 100BASE-X
IEEE 802.3x Flow Control
IEEE 802.3z 1000BASE-X
RFC 768 UDP
RFC 783 TFTP Protocol (revision 2)
RFC 791 IP
RFC 792 ICMP
RFC 793 TCP
RFC 826 ARP
RFC 854 TELNET
RFC 868 Time Protocol
RFC 903 RARP
RFC 951 BOOTP
RFC 959 File Transfer Protocol (FTP)
RFC 1027 Proxy ARP
RFC 1042 IP Datagrams
RFC 1350 TFTP Protocol (revision 2)
RFC 1519 CIDR
RFC 1542 BOOTP Extensions
RFC 1812 IPv4 Routing
RFC 2131 DHCP
RFC 2338 VRRP
RFC 2784 Generic Routing Encapsulation (GRE)
RFC 2865 Remote Authentication Dial In User Service (RADIUS)

IP multicast
RFC 1112 IGMP
RFC 2236 IGMPv2
RFC 2283 Multiprotocol Extensions for BGP-4
RFC 2362 PIM Sparse Mode
RFC 2934 Protocol Independent Multicast MIB for IPv4
RFC 3376 IGMPv3
RFC 3618 Multicast Source Discovery Protocol (MSDP)

IPv6
RFC 1350 TFTP
RFC 1981 IPv6 Path MTU Discovery
RFC 2080 RIPng for IPv6
RFC 2460 IPv6 Specification
RFC 2461 IPv6 Neighbor Discovery
RFC 2462 IPv6 Stateless Address Auto-configuration
RFC 2463 ICMPv6
RFC 2473 Generic Packet Tunneling in IPv6
RFC 2475 IPv6 DiffServ Architecture
RFC 2529 Transmission of IPv6 Packets over IPv4
RFC 2710 Multicast Listener Discovery (MLD) for IPv6
RFC 2740 OSPFv3 for IPv6
RFC 2893 Transition Mechanisms for IPv6 Hosts and Routers
RFC 2925 Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations (Ping only)
RFC 3315 DHCPv6 (client only)
RFC 3484 Default Address Selection for IPv6
RFC 3513 IPv6 Addressing Architecture
RFC 3587 IPv6 Global Unicast Address Format
RFC 3810 Multicast Listener Discovery Version 2 (MLDv2) for IPv6
RFC 4251 SSHv6 Architecture
RFC 4252 SSHv6 Authentication
RFC 4253 SSHv6 Transport Layer
RFC 4254 SSHv6 Connection
RFC 4541 IGMP & MLD Snooping Switch
RFC 4862 IPv6 Stateless Address Auto-configuration

MIBs
IEEE8023-LAG-MIB
RFC 1213 MIB II
RFC 1229 Interface MIB Extensions
RFC 1286 Bridge MIB
RFC 1493 Bridge MIB
RFC 1573 SNMP MIB II
RFC 1643 Ethernet MIB
RFC 1657 BGP-4 MIB
RFC 1724 RIPv2 MIB
RFC 1757 Remote Network Monitoring MIB
RFC 1850 OSPFv2 MIB
RFC 2011 SNMPv2 MIB for IP
RFC 2012 SNMPv2 MIB for TCP
RFC 2013 SNMPv2 MIB for UDP
RFC 2021 RMONv2 MIB
RFC 2096 IP Forwarding Table MIB
RFC 2233 Interfaces MIB
RFC 2273 SNMP-NOTIFICATION-MIB
RFC 2452 IPV6-TCP-MIB
RFC 2454 IPV6-UDP-MIB
RFC 2465 IPv6 MIB
RFC 2466 ICMPv6 MIB
RFC 2571 SNMP Framework MIB
RFC 2572 SNMP-MPD MIB
RFC 2573 SNMP-Target MIB
RFC 2613 SMON MIB
RFC 2618 RADIUS Client MIB
RFC 2620 RADIUS Accounting MIB
RFC 2665 Ethernet-Like-MIB
RFC 2674 802.1p and IEEE 802.1Q Bridge MIB
RFC 2737 Entity MIB (Version 2)
RFC 2787 VRRP MIB
RFC 2819 RMON MIB
RFC 2863 The Interfaces Group MIB
RFC 2925 Ping MIB
RFC 2932IP (Multicast Routing MIB)
RFC 2933 IGMP MIB
RFC 3273 HC-RMON MIB
RFC 3414 SNMP-User based-SM MIB
RFC 3415 SNMP-View based-ACM MIB
RFC 3418 MIB for SNMPv3
RFC 3621 Power Ethernet MIB
RFC 3813 MPLS LSR MIB
RFC 3814 MPLS FTN MIB
RFC 3815 MPLS LDP MIB
RFC 3826 AES for SNMP's USM MIB
RFC 4133 Entity MIB (Version 3)
LLDP-EXT-DOT1-MIB
LLDP-EXT-DOT3-MIB
LLDP-MIB

MPLS
RFC 2205 Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification
RFC 2209 Resource ReSerVation Protocol (RSVP)
RFC 2702 Requirements for Traffic Engineering Over MPLS
RFC 2858 Multiprotocol Extensions for BGP-4
RFC 3031 Multiprotocol Label Switching Architecture
RFC 3032 MPLS Label Stack Encoding
RFC 3036 LDP Specification
RFC 3107 Carrying Label Information in BGP-4
RFC 3209 RSVP-TE: Extensions to RSVP for LSP Tunnels
RFC 3479 Fault Tolerance for the Label Distribution Protocol (LDP)
RFC 3487 Graceful Restart Mechanism for LDP
RFC 4090 Fast Reroute Extensions to RSVP-TE for LSP Tunnels
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs)
RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures
RFC 4447 Pseudowire Setup and Maintenance Using LDP
RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks
RFC 4664 Framework for Layer 2 Virtual Private Networks
RFC 4665 Service Requirements for Layer 2 Provider Provisioned Virtual Private Networks
RFC 4761 Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling
RFC 4762 Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling

Network management
IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
IEEE 802.1D (STP)
RFC 1155 Structure of Management Information
RFC 1157 SNMPv1
RFC 1215 SNMP Generic traps

HP 12500 Switch Series

Specifications (continued)

	HP 12518 AC Switch Chassis (JF430C) HP 12518 DC Switch Chassis (JC653A)	HP 12508 AC Switch Chassis (JF431C) HP 12508 DC Switch Chassis (JC652A)
Standards and protocols (applies to all products in series)	RFC 1905 SNMPv2 Protocol Operations RFC 2211 Controlled-Load Network RFC 2272 SNMPv3 Management Protocol RFC 2273 SNMPv3 Applications RFC 2274 USM for SNMPv3 RFC 2571 SNMP Management Frameworks RFC 2572 SNMPv3 Message Processing RFC 2573 SNMPv3 Applications RFC 2576 Coexistence between SNMP versions RFC 2578 SMIv2 RFC 2819 Four groups of RMON: 1 (statistics), 2 (history), 3 (alarm) and 9 (events) RFC 3164 BSD syslog Protocol RFC 3415 SNMPv3 View-based Access Control Model VACM) ANSI/TIA-1057 LLDP Media Endpoint Discovery (LLDP-MED) SNMPv1/v2c/v3 **OSPF** RFC 1245 OSPF protocol analysis RFC 1246 Experience with OSPF RFC 1587 OSPF NSSA RFC 1765 OSPF Database Overflow	RFC 1850 OSPFv2 Management Information Base (MIB), traps RFC 2328 OSPFv2 RFC 2370 OSPF Opaque LSA Option RFC 3101 OSPF NSSA RFC 3623 Graceful OSPF Restart **QoS/CoS** IEEE 802.1P (CoS) RFC 2212 Guaranteed Quality of Service RFC 2474 DS Field in the IPv4 and IPv6 Headers RFC 2475 DiffServ Architecture RFC 2597 DiffServ Assured Forwarding (AF) RFC 2598 DiffServ Expedited Forwarding (EF) RFC 2697 A Single Rate Three Color Marker RFC 2698 A Two Rate Three Color Marker Bi-directional Rate Shaping **Security** IEEE 802.1X Port Based Network Access Control RFC 1321 The MD5 Message-Digest Algorithm RFC 2082 RIP-2 MD5 Authentication RFC 2104 Keyed-Hashing for Message Authentication RFC 2716 PPP EAP TLS Authentication Protocol RFC 2865 RADIUS Authentication RFC 2866 RADIUS Accounting RFC 2867 RADIUS Accounting Modifications for Tunnel Protocol Support RFC 2868 RADIUS Attributes for Tunnel Protocol Support RFC 2869 RADIUS Extensions RFC 3567 Intermediate System (IS) to IS Cryptographic Authentication Access Control Lists (ACLs) Guest VLAN for 802.1x MAC Authentication SSHv2 Secure Shell Web Authentication **IKEv1** RFC 2865 - Remote Authentication Dial In User Service (RADIUS)

HP 12500 Switch Series accessories

Modules

HP 12500 Main Processing Unit (JC072B)
HP 12500 48-port Gig-T LEB Module (JC074B)
HP 12500 48-port Gig-T LEC Module (JC065B)
HP 12500 48-port GbE SFP LEB Module (JC075B)
HP 12500 48-port GbE SFP LEC Module (JC069B)
HP 12500 48-port GbE SFP LEF Module (JC660A)
HP 12500 8-port 10GbE XFP LEB Module (JC073B)
HP 12500 8-port 10GbE XFP LEC Module (JC068B)
HP 12500 8-port 10GbE SFP+ LEB Module (JC780A)
HP 12500 8-port 10GbE SFP+ LEC Module (JC781A)
HP 12500 8-port 10GbE SFP+ LEF Module (JC659A)
HP 12500 16-port 10GbE SFP+ LEB Module (JC782A)
HP 12500 16-port 10GbE SFP+ LEC Module (JC783A)
HP 12500 32-port 10GbE SFP+ REB Module (JC064B)
HP 12500 32-port 10GbE SFP+ REC Module (JC476B)
HP 12500 Power Monitor Module (JC502A)

Transceivers

HP X120 100M/1G SFP LC LX Transceiver (JF832A)
HP X114 100M SFP LC FX Transceiver (JF833A)
HP X125 1G SFP LC LH40 1310nm Transceiver (JD061A)
HP X120 1G SFP LC LH40 1550nm Transceiver (JD062A)
HP X125 1G SFP LC LH70 Transceiver (JD063B)
HP X120 1G SFP RJ45 T Transceiver (JD089B)
HP X120 1G SFP LC BX 10-U Transceiver (JD098B)
HP X120 1G SFP LC BX 10-D Transceiver (JD099B)
HP X120 1G SFP LC LH100 Transceiver (JD103A)
HP X170 1G SFP LC LH70 1550 Transceiver (JD109A)
HP X170 1G SFP LC LH70 1570 Transceiver (JD110A)
HP X170 1G SFP LC LH70 1590 Transceiver (JD111A)
HP X170 1G SFP LC LH70 1610 Transceiver (JD112A)
HP X170 1G SFP LC LH70 1470 Transceiver (JD113A)
HP X170 1G SFP LC LH70 1490 Transceiver (JD114A)
HP X170 1G SFP LC LH70 1510 Transceiver (JD115A)
HP X170 1G SFP LC LH70 1530 Transceiver (JD116A)
HP X120 1G SFP LC SX Transceiver (JD118B)
HP X120 1G SFP LC LX Transceiver (JD119B)
HP X130 10G XFP LC ZR Transceiver (JD107A)
HP X130 10G XFP LC LR Transceiver (JD108B)
HP X130 10G XFP LC SR Transceiver (JD117B)
HP X135 10G XFP LC ER Transceiver (JD121A)
HP X180 10G XFP LC LH 80km 1538.98nm DWDM Transceiver (JG226A)
HP X180 10G XFP LC LH 80km 1539.77nm DWDM Transceiver (JG227A)
HP X180 10G XFP LC LH 80km 1540.56nm DWDM Transceiver (JG228A)
HP X180 10G XFP LC LH 80km 1542.14nm DWDM Transceiver (JG229A)
HP X180 10G XFP LC LH 80km 1542.94nm DWDM Transceiver (JG230A)
HP X180 10G XFP LC LH 80km 1558.98nm DWDM Transceiver (JG231A)
HP X180 10G XFP LC LH 80km 1559.79nm DWDM Transceiver (JG232A)
HP X180 10G XFP LC LH 80km 1560.61nm DWDM Transceiver (JG233A)
HP X130 10G SFP+ LC SR Transceiver (JD092B)
HP X130 10G SFP+ LC LRM Transceiver (JD093B)
HP X130 10G SFP+ LC LR Transceiver (JD094B)
HP X130 10G SFP+ LC ER 40km Transceiver (JG234A)
HP X240 10G SFP+ to SFP+ 3m Direct Attach Copper Cable (JD097B)
HP X240 10G SFP+ to SFP+ 5m Direct Attach Copper Cable (JG081B)
HP X240 10G SFP+ SFP+ 7m Direct Attach Copper Cable (JC784A)

Cables

HP 12500 Side Cable Management Guide (JC084A)
HP 0.5 m PremierFlex OM3+ LC/LC Optical Cable (BK837A)
HP 1 m PremierFlex OM3+ LC/LC Optical Cable (BK838A)
HP 2 m PremierFlex OM3+ LC/LC Optical Cable (BK839A)
HP 5 m PremierFlex OM3+ LC/LC Optical Cable (BK840A)
HP 15 m PremierFlex OM3+ LC/LC Optical Cable (BK841A)
HP 30 m PremierFlex OM3+ LC/LC Optical Cable (BK842A)
HP 50 m PremierFlex OM3+ LC/LC Optical Cable (BK843A)

Mounting Kit

HP X421 Chassis Universal 4-post Rack Mounting Kit (JC665A)

Appliance

HP 12500 VPN Firewall Module (JC635A)

Memory

HP X600 1G Compact Flash Card (JC684A)
HP 12500 additional 1 GB SDRAM DDR2 (JC071A)

HP 12518 AC Switch Chassis (JF430C)

HP 12518 G2 Fabric Module (JC657A)
HP 12518 Fabric Module (JC066A)
HP 12518 Top and Bottom Cable Guides for AC Powered Switch (JC786A)
HP 12500 2000W AC Power Supply (JF429A)
HP 12500 AC Power Entry Module (JF426A)
HP 12518 Fan Assembly (JC080A)
HP 12518 Optional Air Filter (JC083A)

HP 12508 AC Switch Chassis (JF431C)

HP 12508 Fabric Module (JC067B)
HP 1250x G2 Fabric Module (JC658A)

HP 12500 Switch Series accessories (continued)

HP 12508 Top and Bottom Cable Guides for AC Powered Switch (JC785A)
HP 12500 2000W AC Power Supply (JF429A)
HP 12500 AC Power Entry Module (JF426A)
HP 12508 Fan Assembly (JC081A)
HP 12508 Optional Air Filter (JC082A)

HP 12518 DC Switch Chassis (JC653A)

HP 12518 G2 Fabric Module (JC657A)
HP 12518 Fabric Module (JC066A)
HP 12518 Top and Bottom Cable Guides for DC Powered Switch (JC788A)
HP X210 10-meter JG Connector to Bare 6AWG 37800 Watt 72V DC Power Cable (JG280A)
HP 12500 1800W DC Power Supply (JC651A)
HP 12518 Fan Assembly (JC080A)
HP 12518 Optional Air Filter (JC083A)

HP 12508 DC Switch Chassis (JC652A)

HP 12508 Fabric Module (JC067B)
HP 1250x G2 Fabric Module (JC658A)
HP 12508 Top and Bottom Cable Guides for DC Powered Switch (JC787A)
HP X210 10-meter JG Connector to Bare 6AWG 37800 Watt 72V DC Power Cable (JG280A)
HP 12500 1800W DC Power Supply (JC651A)
HP 12508 Fan Assembly (JC081A)
HP 12508 Optional Air Filter (JC082A)

Products within this series have achieved sufficient scores in each of the rated criteria to achieve the Miercom Certified Green distinction Award. See the Specifications section of this series for more information.

To learn more, visit www.hp.com/networking

4AA3-0666ENW, Created August 2010; Updated March 2012, Rev. 5

Miercom

Lab Testing Summary Report

June 2010

Report 100102B

Product Category:

Power Efficient Ethernet Switches

Vendor Tested:

Products Tested:

HP A12508 Data Center Switch

Key findings and conclusions:

- **10 Gigabit Ethernet modules consume half the power compared to competitive products**
- **New ASIC technology has low power consumption while providing a rich advanced feature set**
- **Intelligent Management Center (IMC) provides complete network management of all devices**
- **Redundant, scalable, 90% efficient power supplies (up to 6) deliver high reliability in the data center**

Hewlett-Packard* engaged Miercom to evaluate the HP A12508 Data Center Switch under the Certified Green Test Program for power consumption and efficiency. We analyzed the overall environmental impact and green features that the A12508 switch offers the data center environment.

The A12508 demonstrated during hands-on testing and by an independent audit that it is a highly flexible, manageable and energy efficient switching solution due to its modular design and use of superior ASIC technology. Using advanced ASICs as well as a unified Comware platform, it also provides leading class remote management and features and helps customers reduce power consumption and related energy costs.

The A12508 is a high density 22 rack unit chassis designed specifically for the data center. It offers exceptional redundancy with six internal power supplies. Only three are required to power a fully-populated chassis and still have 1:1 redundancy. High Availability is provided by two redundant management modules. Two fan trays containing 12 fans each draw cool air from the front, and return warm air trough the rear for hot aisle/cold aisle configurations. The overall noise is reduced due to the efficient design and the variable fan speed. *(continued on page 3)*

Figure 1: Power Consumption Relative to Throughput

Measured with (8) interface modules, each with 8x10GbE ports, (2) management modules, (9) fabric modules, and (3) 2000W power supplies

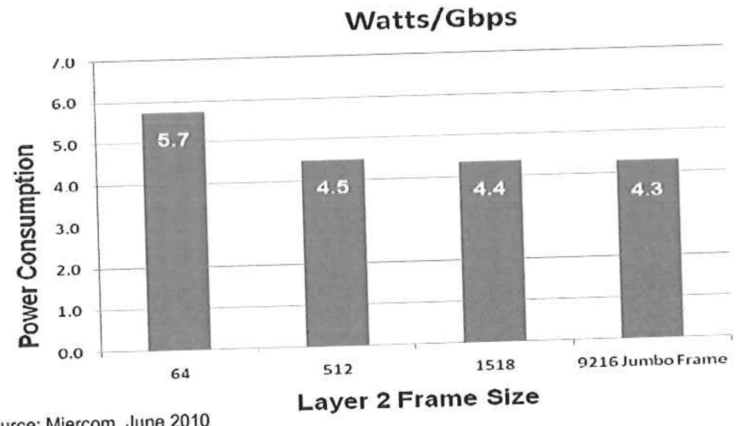

Source: Miercom, June 2010

The HP A12508 Data Center Switch tested consistently low for power consumption relative to throughput for all Gigabit Ethernet (GbE) frame sizes at full line rate.

---

*The HP products referred to in this publication were developed and sold by H3C Technologies Co. Ltd., which was acquired by HP in April 2010. The original report can be found under 3Com at www.miercom.com.

Test Bed Equipment

IP Load Generators

Ixia 1600T and XM2 Multifunction Chassis

HP A12508 Data Center Switch

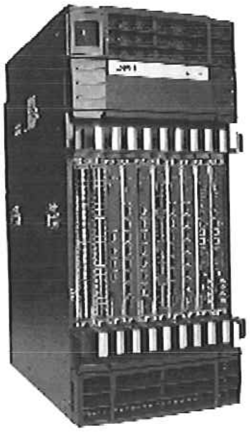

Environmental Testing Tools

Dranetz Encore 61000 Power Analyzer

Miercom Industry Average Green Calculator

Fluke Scope Meter Portable Heat and Power Analyzer

Source: Miercom, June 2010

How We Did It

The HP A12508 Data Center Switch was evaluated for environmental impact by looking at the individual components as well as features and capabilities. Testing focused on the power consumption and efficiency of the product. A full audit was conducted to analyze the overall product-specific environmental impact.

Lab testing was conducted for power consumption under load as well as measurements and audit results verified with site survey assessments.

Measuring Power Consumption: The power consumption of the A12508 was measured at varying frame sizes and link loads that the switch would typically experience in a real world deployment. Power consumption was measured using a Dranetz Encore 61000 Power Analyzer from Dranetz-BMI (www.dranetz-bmi.com). The SUT was loaded with traffic at various rates and packet sizes in accordance with RFC 2544 Benchmarking Methodology for Network Interconnect Development. The SUT was configured with (8) interface modules, each with 8x10GbE ports, (2) management modules, (9) fabric modules, and (3) 2000W power supply modules. The 64 10GbE ports were a combination of XFPs consisting of 3 types: SXP3101SV-02 and SXP3101LX-H2-H3C (both 1310nm), and H8511D3-H3C (850nm). The power consumption was measured as a function of throughput at standard and jumbo frame sizes with 100% link utilization.

Power consumption measurements were taken during system boot-up, idle, and with 70% and 100% load. Power usage was measured, while running Layer 2 and 3 traffic using XM2 and 1600T traffic generators from Ixia (www.ixiacom.com). Traffic was applied to each of the 10 Gbps ports while stressing the product with the features it supports. All power measurements were taken at 220 volts and 50 Hz frequency.

Miercom utilizes Ixia equipment to conduct energy efficiency testing of networking equipment. Ixia's unique approach utilizes coordination of energy measurements with network traffic load – allowing energy consumption to be graphed against network traffic volume. Real-world traffic is generated by Ixia's test platform and test applications, principally IxNetwork for Layer 2-3 routing and switching traffic and IxLoad for layer 4-7 application traffic.

Environmental Analysis: Miercom's environmental review of the HP A12508 Data Center Switch also entailed an examination of the company-wide and product-specific environmental impact reduction efforts. We interviewed HP customers and HP's Green Team regarding the environmental-related features of the equipment and applications.

(continued from page 1) The relative power consumption as a function of throughput at different frame sizes is shown in *Figure 1*. For Layer 2 forwarding, with 100% link utilization at 64 byte frames, power efficiency was measured at 5.7 watts/Gbps. Transmitting using jumbo frames improved efficiency by 24%, to 4.3 watts/Gbps.

Power Efficiency

Figure 2 illustrates the power profile of the HP A12508 switch. Testing was performed and power consumption was measured with the switch in various operational states–first an empty chassis, then with individual modules installed, with fully loaded chassis at idle, connected with active links and with typical and maximum link generated traffic.

The variable speed fan trays are controlled by the management module and allow for most efficient dynamic cooling "as required." Power consumption is 24% lower when the management module is installed since the fans are throttled down from their default maximum speed. The consumption figure for the 8 port 10GbE line module includes the usage for a single management module since this module is a prerequisite for running the line module. Consumption for the 8-port line card alone was 232 watts, *see Figure 2.*

The measurements taken with traffic loads were performed using a chassis configuration consisting of two management modules, eight 10GbE Base-R/W-XFP line cards, each providing 8 ports of 10GbE connectivity, for a total of 64 10GbE interfaces. Nine fabric modules and three 48Volt 2000W power supply modules rounded out the as-tested configuration.

At idle with links up on all ports, the A12508 consumed 2,546 watts. At a typical 70% link utilization load using 64 byte Layer 2 frames, power consumption increased 8% to 2,779 watts. At maximum traffic load, usage was 10% higher than idle, with 2,842 watts recorded. With Layer 3 traffic, using the same rates, consumption was 6% higher –2,794 watts at 70% load; 10% higher –2,878 watts at maximum load.

The HP A12508 utilizes the latest ASIC technology to deliver reduced power consumption. Currently, 65nm silicon technology is used and future incorporation of 45nm silicon is planned to further reduce power consumption, while providing the latest advanced features demanded by customers.

Results of testing showed that the 10GbE modules in our test bed used about half the stated power of comparable modules from other published datasheets. The A12508 used only 5.7 watts/Gbps during testing with 64 byte frames, which decreased to 4.4 watts/Gbps for 1518 byte frames. The A12508 also supports jumbo frames of 9216 bytes. When transmitting using jumbo frames, the switch drew only 4.3 watts/Gbps.

Figure 2: HP A12508 Switch Power Profile

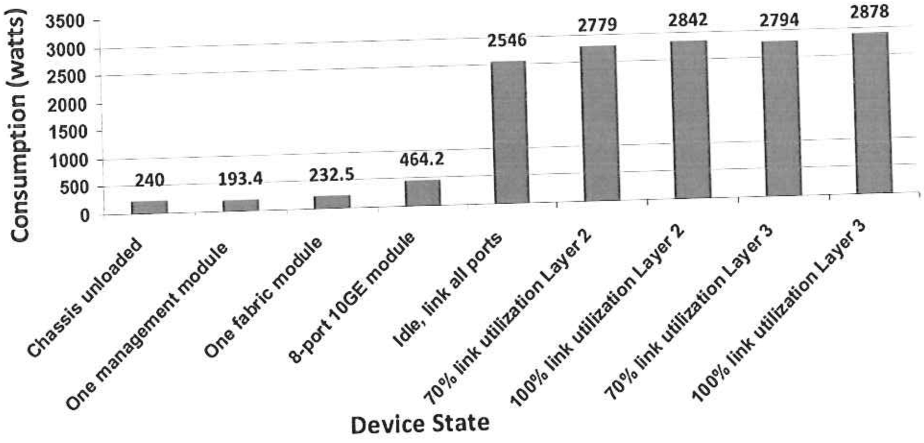

Source: Miercom, June 2010

HP A12508 Switch Power Profile illustrates the power usage with various modules engaged while tested idle thru 100% traffic load.

The power consumption reported for the 8 x 10GbE line modules includes the fabric module usage.

For the link utilization test, a total of 64 10GbE interfaces were used. Traffic consisted of 64 byte frames. Packet inspection of the Layer 3 traffic causes power draw to be slightly higher than for Layer 2 traffic.

Product Efficiency

The HP A12500 is intended for data center placement. It is designed to be flexible, scalable and future-proof. The system architecture delivers 2.2 billion packets-per-second of forwarding performance and aggregates large numbers of Gig Ethernet ports, providing up to 864 line-speed GbE ports, 128 wire-speed 10GbE ports, or 512 non-wire speed 10GbE ports. With a 256 MB buffer for each 10GbE port, the A12500 can support a burst size of data in 200ms to meet the high burst size requirements in large data centers, and support emerging high bandwidth application like video.

The modular architecture enables administrators to reconfigure and upgrade to meet ever-changing data center needs, and new technologies. The switch is compatible with 40-Gigabit Ethernet, 100-Gigabit Ethernet, and Fibre Channel over Ethernet (FCoE).

The HP A12508 comes with a fully distributed architecture separating the forwarding and control planes. This architecture provides 1:1 redundancy for the control plane, and N+1 redundancy for the forwarding plane. The passive backplane design has load sharing fabrics, management modules, and redundant power supplies and fan trays.

All of these core elements are hot-swappable, minimizing the impact of single component failure. Multiple Spanning Tree Protocol, Rapid Spanning Tree Protocol, OSPF Equal Cost Multi Path, and Virtual Router Redundancy Protocol support delivers rapid recovery from device or link failures in the network, minimizing disruptions for business applications. In addition, "always on" high availability is critical.

The HP A12508 shares its unified Comware platform operating system with all other HP switches to provide commonality and consistency of features and operation, which helps to reduce operational costs by providing more efficient network administration. Intelligent Management Center (IMC) is a next generation enterprise management solution. Designed to support the Information Technology Infrastructure Library (ITIL) operational center of excellence best practices model, IMC provides management of not just routers and switches but all networked devices, delivering end-to-end administration through a single interface.

Business Processes

IMC provides network management and administration through a single console interface. This enables end-to-end business administration by combining traditionally separate management tools, policy management, and 3rd party device management for more effective and efficient network administration.

The HP A12500 family of switches is designed to

Figure 3: Intelligent Management Center

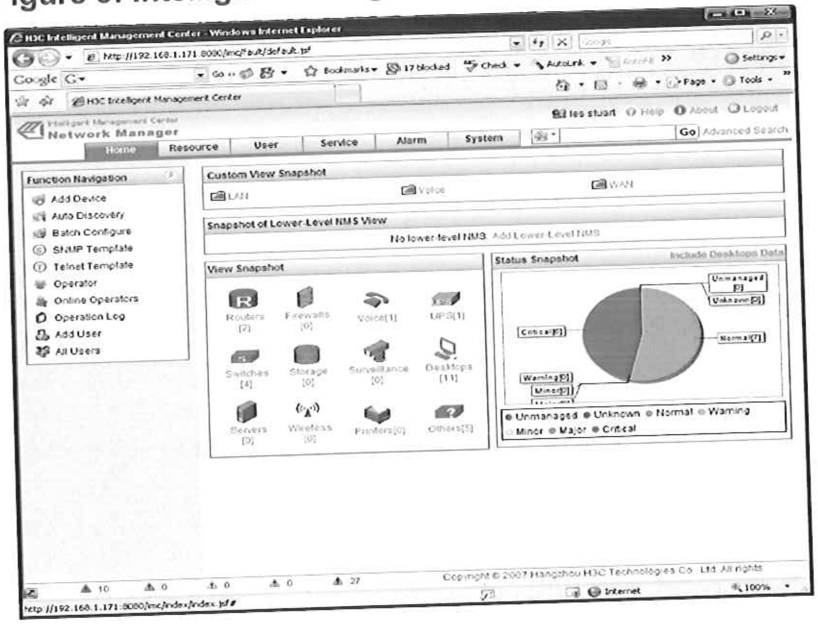

Source: Miercom, June 2010

The HP Intelligent Management Center (IMC) allows for effective remote administration and multi-site management.

IMC provides a single-pane control to control the entire enterprise network topology, not just routers and switches.

reduce power consumption without impacting performance or features. Advanced ASIC design and a common unified Comware platform operating system allow Hewlett-Packard to take advantage of the latest technological innovations known today to reduce power consumption.

Green Innovation

HP's approach of using the most advanced ASIC technology provides the freedom to design switch products using the latest advances in silicon technology to deliver critical data center features and performance, while providing reductions in power usage. For example, distributed temperature management is included on each I/O module to trigger and drive the speed of the cooling fans.

Use of a common unified Comware platform operating system across all models, means that enhancements can be rolled out easily and in a consistent fashion. HP switches use Intelligent Resilient Framework (IRF) to provide distributed high availability and resiliency by extending the control plane of multiple active switches in different geographic locations. IRF eliminates the need for complex redundancy technologies, such as Spanning Tree Protocol (STP) or Virtual Router Redundancy Protocol (VRRP). Enterprises can collapse their network architecture from 3 tiers to 2 tiers, eliminating the additional network layers, the associated power and costs associated with them, and reducing network latency while improving performance.

Affiliations and Standards

HP is a participating member of The Green Grid, and contributes to the leadership of the IEEE 802.3 Ethernet Working Group, the parent body of the IEEE P802.3az an Energy-efficient Ethernet project. This project will provide an interoperability standard which will save power during idle periods on an Ethernet link, and enable some energy savings on attached devices with a Sleep and Wake signal indicating the entry and exit from long idle periods, allowing the devices to enter a low power mode.

Hewlett-Packard is a partner in EnergyStar as well as a stakeholder in reviewing the EnergyStar Requirements for Networking Equipment. They are also an endorser of the EU Data Center Code of Conduct, part of the EU Standby Initiative to improve the energy efficiency of electrical equipment while either Off or in Stand-by. HP has also been certified ISO 14001 compliant.

Business Case

The A12500 employs the latest ASIC technologies and other advanced techniques to deliver industry-leading throughput performance addressing the requirements of the largest data centers in the world. The HP A12500 Data Center Switch is designed with an 80% derating on internal components including the Printed Circuit Board (PCB). This criteria, which is over and above the components manufacturer limits, ensures the longevity of the internal components. Management modules, power supply modules and fan tray modules are all hot-swappable and redundant to provide "always on" high-availability to meet the strict requirements of the data center. The A12508 is compatible with new 40-, 100-Gigabit Ethernet technology, and FCoE. Modular design allows for future upgrades and protects IT investment.

The A12500 using Comware v5, OS provides a comprehensive list of features to help manage fault, configuration, accounting, provisioning locally and remotely. Logs, traces, reports, alarms, traps use protocols and features like SNMPv3, Telnet/SSHv2, FTP/SFTP. It is also supported by Intelligent Management Center (IMC), providing single pane console control of the network infrastructure.

Certified Green

Miercom conducts environmental analysis on products by taking a holistic view of the product life cycle. We consider power efficiency of the product, manufacturing and overall business practices in this analysis. Power consumption and power efficiency are important metrics for comparing products and this data does provide a key component to green analysis, but not the only relevant component.

Miercom believes that a comprehensive environmental analysis such at that conducted for HP is the only credible and relevant approach to evaluating green technologies. This type of study reveals a business case justification for the environmental benefits the products have to offer by virtue of the sustainable benefits in cost savings in power as well as product efficiencies.

Miercom Certified Green

The energy-saving attributes of the HP A12508 Data Center Switch were evaluated by Miercom in accordance with the Certified Green Testing Methodology. The product has been awarded Miercom Certified Green based on the observations and audit analysis.

Based on our hands-on testing and the verified representations made by Hewlett-Packard, Miercom confirms that the HP A12500 family of Data Center Switches is designed to provide enterprise customers superior performance, critical redundancy and availability, and environmentally sound datacenter solutions.

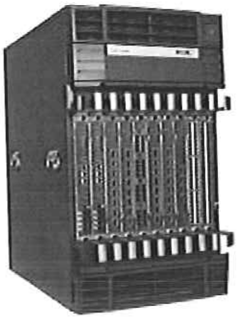

HP A12508
Data Center Switch

3000 Hanover Street
Palo Alto, CA
www.hp.com
1-650-857-1501

About Miercom's Product Testing Services

Hundreds of product-comparison analyses have been published over the years in such leading network trade periodicals as Network World, Business Communications Review - NoJitter, Communications News, xchange, Internet Telephony and other leading publications. Miercom's reputation as the leading, independent product test center is unquestioned.

Miercom's private test services include competitive product analyses, as well as individual product evaluations. Miercom features comprehensive certification and test programs including: Certified Interoperable, Certified Reliable, Certified Secure and Certified Green. Products may also be evaluated under the NetWORKS As Advertised program, the industry's most thorough and trusted assessment for product usability and performance.

Report 100102B reviews@miercom.com www.miercom.com

HP 12500 Data Center Access and Interoperability with Cisco Nexus 7000

Test Plan

Written by Sandra Zimmerman

Edited by Sue Darte

March 2012

Version 9

Table of contents

Introduction

This Test Plan describes the integration of the HP 12500 series switch into the customer's existing architecture. The intended audience for this TCG is HP Networking Solution Architects, HP Networking Technical Consultants, and HP Networking partner technical pre-sales staff.

This document delivers the results of an HP Networking solution for a potential customer. Also, all IP Addresses have been modified to keep their network information confidential. The main goal of the test plan is to identify the architectural, operational and design parameters needed to allow a customer to select network gear based on industry standards. This test plan provides a customer with technical results that prove using an alternate vendor or networking gear solution is possible within their existing environment.

Solution

HPN proposed design to show that HP networking gear can be inserted into an existing architecture and meet all existing requirements. While there are a number of ways to improve reliability and simplify operations using HP technology, HP understands an overall architecture change is generally too disruptive in an existing (non-Greenfield) environment.

This test plan has been designed to address middle-of-row server aggregation to a pair of Cisco Nexus 7000s in the core. The main purpose of this test plan is to familiarize company personnel with the equipment and validate interoperability with Cisco Nexus using the two topologies below.

Network Diagrams

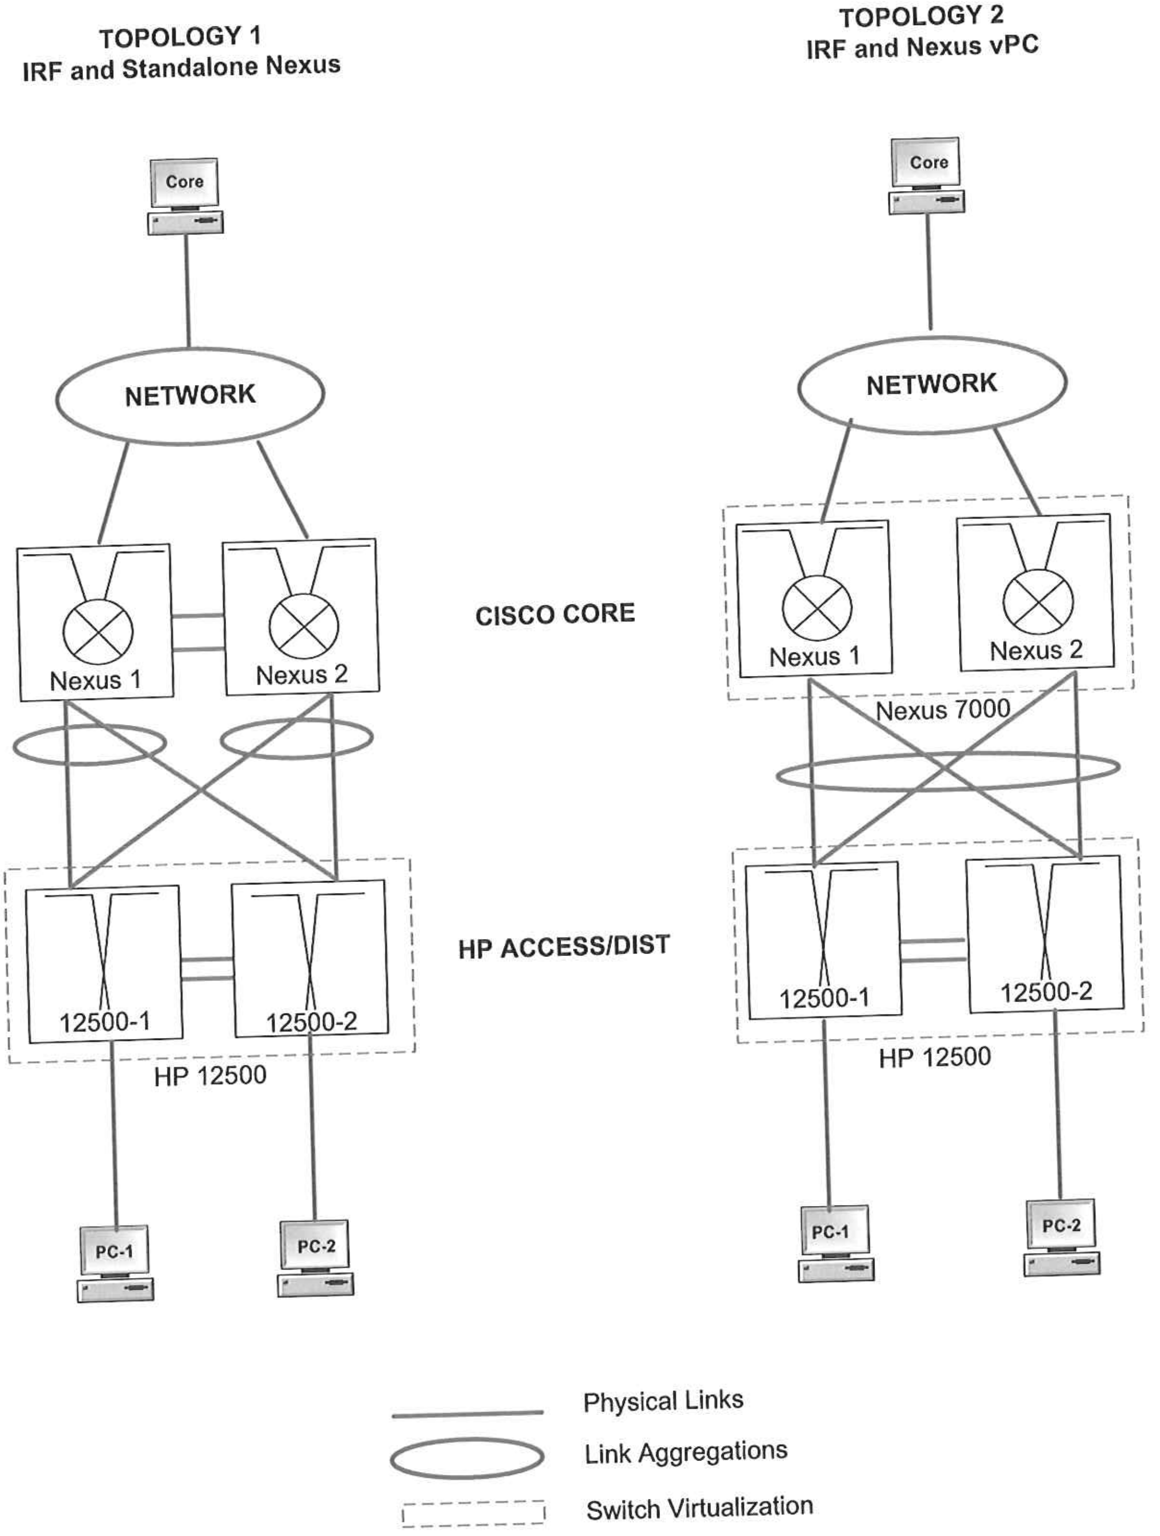

Topology 1 consists of two HP 12500s virtualized into a single switch using HP's Intelligent Resilient Framework (IRF) and two standalone Cisco Nexus 7000s. Topology 2 consists of the HP IRF pair and the two Cisco Nexus 7000s using Cisco's Virtual Port Channel (vPC) technology to simulate switch virtualization to span link-aggregations (port channels) across two switches.

Requirements

In addition to the data center access/server aggregation requirements, the customer wants to include some layer 3 testing for future reference in the event they re-architect their current network or simply want to place HP gear in a distribution/core role where routing is required.

Configuration/Verification (Test Cases)

Demonstrate HP 12500 Configuration Basics and CLI familiarization

CLI Access via Console Port

Objective	The device provides multiple methods of entering the CLI: • Through the console port • Through Telnet • Through SSH with encryption When you use the CLI of a switch for the first time, you must use the Console port – this test case provides those procedures.
Procedure	The console cable that shipped with your device features DB-9 Female to RJ-45 connectors. Procedures: 1. Connect the DB-9 (female) connector of the console cable into the 9-pin serial port of your PC or terminal 2. Connect the RJ-45 connector of the console cable to the console port of the main board of the switch 3. Launch a Terminal Emulation program 4. Configure applicable COM port as 9600, 8, N, 1, None 5. Connect
Expected Results / Objectives	Successful login to CLI

<table>
<tr><td></td><td>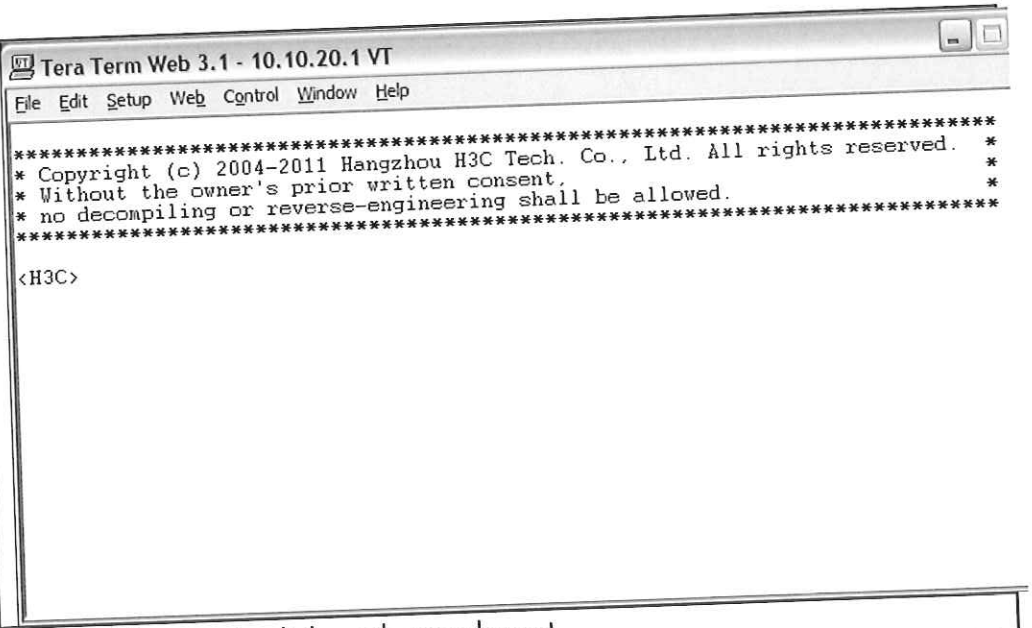

<pre>
****
* Copyright (c) 2004-2011 Hangzhou H3C Tech. Co., Ltd. All rights reserved. *
* Without the owner's prior written consent,   *
* no decompiling or reverse-engineering shall be allowed.   *
****

<H3C>
</pre></td></tr>
<tr><td>Results</td><td>Successful access to cli through console port.</td></tr>
<tr><td>Status (Pass/Fail)</td><td>Pass</td></tr>
<tr><td>Comments</td><td>The flat blue Cisco console cable will also work.

All console ports are active regardless of which Main Board/SUP is active – for first time login best practice is to log into AMB (active main board) which is typically in the lowest slot number.</td></tr>
<tr><td></td><td></td></tr>
</table>

CLI Views and Permission Levels

<table>
<tr><td>Objective</td><td>CLI views are designed to meet various configuration requirements. This case demonstrates The most commonly used views and view operations.</td></tr>
<tr><td>Procedure</td><td>User view
When you first log in to the switch, you are in user view and the prompt is <i><device name></i>. In user view, only a few operations are allowed, for example, display operations, file operations, FTP, and Telnet operations.
Entering system view
To further configure the switch (Cisco config-t equivalent), you need to enter system view. To enter system view, from user view, use the system-view command.
<H3C>
system-view
[H3C]

Note the prompt changed from < > to [ ] – the brackets indicate system-view.</td></tr>
</table>

Interface/Protocols/services views

The switch's CLI views are multi-layered, for example, user view > system view > interface view, VLAN view, etc. When you enter a particular "view" the prompt will reflect which view you are in.

```
[H3C]
int gi 0/1
[H3C-GigabitEthernet0/1]
```

TIP: Typing "display this" or "dis this" will show you the commands that have been applied to that view.

```
[H3C-GigabitEthernet0/1]
dis this
#
interface GigabitEthernet0/1
 port link-mode route
 ip address X.X.20.1 255.255.255.0
#
```

Exiting the current view(s)

```
[H3C-GigabitEthernet0/1]
quit
[H3C]
quit
<H3C>
```

TIP: Ctrl+Z returns you to user-view from any context.

Expected Results / Objectives

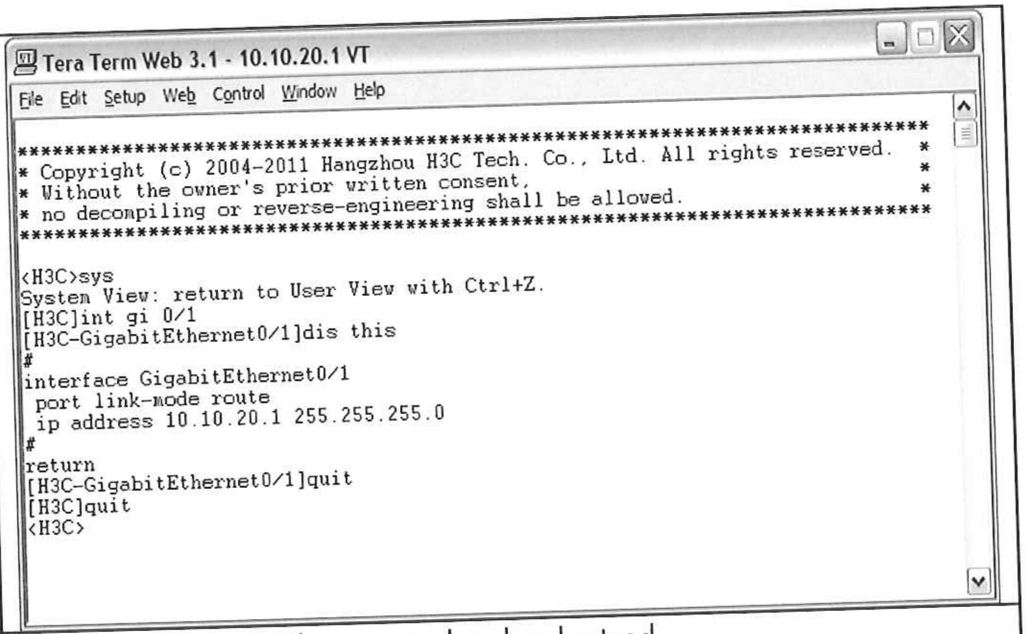

```
******************************************************************************
* Copyright (c) 2004-2011 Hangzhou H3C Tech. Co., Ltd. All rights reserved.  *
* Without the owner's prior written consent,                                 *
* no decompiling or reverse-engineering shall be allowed.                    *
******************************************************************************

<H3C>sys
System View: return to User View with Ctrl+Z.
[H3C]int gi 0/1
[H3C-GigabitEthernet0/1]dis this
#
interface GigabitEthernet0/1
 port link-mode route
 ip address 10.10.20.1 255.255.255.0
#
return
[H3C-GigabitEthernet0/1]quit
[H3C]quit
<H3C>
```

Results	Major CLI differences demonstrated and understood.
Status (Pass/Fail)	Pass

Comments

Fuzzy Match

The switch supports fuzzy matching for efficient input of commands. If, the character string you have typed in the current view can already uniquely identify a keyword, you do not need to type the complete keyword.

Tab and Question Marks

? <enter> in a particular view will provide a list of sub-commands available in that view.

? <enter> after entering a partial command will provide guidance on how the parameters required to complete the command or the sub-commands available.

TAB key in a partial command fills out the rest of the command for you – if there are multiple commands, multiple tabs will scroll through alphabetically.

TAB key during file copy or TFTP will also scroll through file lists which reduces typing when copy files between main boards.

Command Alias and Hot Keys

You can create command alias' and/or Hot Keys for commands with command alias'. The "?" will still provide the parameter options and sub-commands available but the tab key will not complete the command as it does with the actual commands.

Displaying and Executing History Commands

The CLI automatically saves the commands recently used in the history command

	buffer. You can access commands in the history command buffer and execute them again. <u>display history</u> – displays commands saved in the history buffer <u>up arrow or Ctrl+P</u> – displays previous commands one at a time <u>down arrow or Ctrl+N</u> – displays next commands on at a time

Configure VTY and local User-Interfaces

Objective	Demonstrate and configure parameters for vty and local user interfaces

Procedure	In system-view:
	`telnet server enable`
	
	`user-interface vty 0 4`
	`user privilege level 3`
	`protocol inbound all`
	`authentication-mode password`
	`set authentication password cipher h3c`
	
	`local-user test`
	`password simple h3c`
	`authorization-attribute level 3`
	`service-type ssh telnet`
	`service-type ftp`
	
	Configure a VLAN with IP and add an access port
	`int vlan 201`
	`ip address X.X.51.10 24`
	
	Telnet to switch:
	login using local user "test" and pw "h3c"
	
	Display the user-interface and vty interface parameters:
	`user-interface vty 0 4`
	`display this`
	`quit`
	
	Local-user test:
	`display this`
	`quit`

Expected Results / Objectives	Successful login
Results	`[H3C]` `int vlan 201`

```
[H3C-Vlan-interface201]
dis this
#
interface Vlan-interface201
 ip address X.X.51.10 255.255.255.0
#
[H3C-Vlan-interface201]

[H3C]
user-interface vty 0 4

[H3C-ui-vty0-4]
dis this
#
user-interface con 0
user-interface tty 13
user-interface aux 0
user-interface vty 0 4
user privilege level 3
 set authentication password cipher G`M^B<SDBB[Q=^Q`MAF4<1!!
#

[H3C]
local-user test

[H3C-luser-test]
dis this
#
local-user test
 password simple h3c
 authorization-attribute level 3
 service-type ssh telnet
 service-type ftp
#
[H3C-luser-test]
```

Status (Pass/Fail)	Pass
Comments	N/A

Review Device Status

Objective	Display device status for switches, power and software version.
Procedure	display device display power

	display version
Expected Results / Objectives	Modules and power display normal state. Software Revision 1238-P08
Results	missing actual print screen from these display commands Results observations When switch is booting display devices indicate that modules not yet fully loaded are in "fault" state. When they are loaded they move to "normal" state. This is a normal operation/response for this command and "fault" does not indicate a hardware error. (see customer comments below)
Status (Pass/Fail)	Pass contingent on customer verification
Comments	Customer indicates that "fault" state while modules are loading is potentially misleading for support persons not familiar with gear. A different state such "init" or "booting" or "not loaded" – something other than "fault" is preferable.

Initial software upgrade

Objective	Upgrade single chassis
Procedure	Copy new image to each CPU either via TFTP, FTP or USB (see release notes). File system notes: flash storage of active main board/cpu (slot0) = flash:/ flash storage of standby main board/cpu = slot1#flash:/ There are also USB and compact flash available on each main board/cpu – performing the following command will display the choices. User-view: <h3c> copy ? In IRF mode chassis1 or chassis2 is pre-pended – for example chassis1#slot1#flash:/ Set boot loader for new image: boot-loader file flash:/image.bin slot 0 main boot-loader file slot1#flash:/image.bin slot 1 main
Expected Results / Objectives	Switch reboots with new image
Results	[HP12500-Bridge-Aggregation11] dis cur

	version 5.20, Release 1238P08
Status (Pass/Fail)	Pass
Comments	N/A

Startup Config File Settings on Core

Objective	Configure which config file will be used as the startup config file.
Procedure	startup saved-config filename.cfg system-view slave auto-update config
Expected Results / Objectives	Config file saved to new name.
Results	Files saved with slave auto-update as well.
Status (Pass/Fail)	Pass
Comments	N/A

Hostname on Access

Objective	Set System Name
Procedure	Set the hostname for each device based on the diagram. system-view sysname HP12500 save xx.cfg
Expected Results / Objectives	Hostname changed
Results	
Status (Pass/Fail)	Pass
Comments	N/A

Switch Virtualization to Create Single Switch (IRF Pair)

IRF on 12500 Switches Procedure 1

Objective

Configure IRF on each pair of 12500 series devices by converting chassis to IRF mode first and then configuring ports.

This is the recommended procedure to re-configure IRF on switches that have been previously converted.

Key Concepts:

- IRF virtualization physically extends the control plane to a second switch using 10G ports. It is not a virtualization that overlays on top normal data-link or network layers
- The ports dedicated to the inter-switch connection are called IRF-ports and are not useable for any other normal switch operation – they are part of the control plane and allow both switches to assume one switch MAC
- There is one CLI where all chassis/slots/subslots/ports are available
- Link-aggregation/LACP/Portchannels can be spanned across chassis without any modification to the link-aggregation technology
- LACP load balancing does take into consideration local ports first so that the only time data travels across the IRF links is when there is an upstream failure (this is configurable)
- Protocols at all layers operate normally as if the pair is a single switch

Procedure

Assign a unit number to each Chassis.
For unit 1:

```
irf member 1
```

For unit 2:

```
irf member 2
```

Set the Chassis to operate in IRF mode. Reboot the switch when prompted:

```
chassis convert mode irf
```

Set the switches to disable a chassis if both management modules die and to force the slave switch to check the boot file against the master:

```
monitor handshake-timeout disable-port
irf auto-update enable
monitor handshake-timeout disable-port
irf auto-update enable
```

Save the configuration and reboot the switches:

```
quit
save irf.cfg
startup saved-configuration irf.cfg
```

```
reboot
```

Assign IRF priority for each Chassis:

For unit 1:

```
irf member 1 priority 10
```

For unit 2:

```
irf member 1 renumber 2
reboot
```

Wait for reboot

```
irf member 2  priority 30
```

Shutdown the 10 Gbps ports that will form the IRF:

For Unit 1:

```
int TenGigabitEthernet 1/2/0/1
shutdown
quit
int TenGigabitEthernet 1/2/0/2
shutdown
quit
```

For Unit 2:

```
int TenGigabitEthernet 2/2/0/1
shutdown
quit
int TenGigabitEthernet 2/2/0/2
shutdown
quit
```

Assign the 10 Gbps ports to an IRF port group:

On Unit 1:

```
irf-port 1/1
port group interface ten-gigabitethetnet 1/2/0/1 mode enhanced
port group interface ten-gigabitethetnet 1/2/0/2 mode enhanced
quit
```

On Unit 2:

```
irf-port 2/2
port group interface ten-gigabitethetnet 2/2/0/1 mode enhanced
port group interface ten-gigabitethetnet 2/2/0/2 mode enhanced
quit
```

Enable the 10 Gbps ports that will form the IRF:

For Unit 1:

```
int Ten-GigabitEthernet 1/2/0/1
undo shutdown
int Ten-GigabitEthernet 1/2/0/2
```

```
undo shutdown
```

For Unit 2:

```
int Ten-GigabitEthernet 2/2/0/1
undo shutdown
int Ten-GigabitEthernet 2/2/0/2
undo shutdown
```

Save the configuration:

```
quit
save
```

Cable the IRF ports of the two switches. The secondary switch will now request to reboot:

```
Save the configuration
reboot
```

The IRF stack should now be formed. Verify IRF operation:

```
display irf
display irf configuration
display irf topology
```

Expected Results / Objectives

Working IRF/single switch

Results

```
<HP12500>dis irf topo
                                 Topology Info
 -------------------------------------------------------------------------
                  IRF-Port1                        IRF-Port2
 Switch     Link         neighbor         Link         neighbor     Belong To
 2          DIS          --               UP           1            0210-fc02-0000
 1          UP           2                DIS          --           0210-fc02-0000

<HP12500>dis irf conf
 MemberID NewID     IRF-Port1                         IRF-Port2
  1        1        Ten-GigabitEthernet1/4/0/1        disable
                    Ten-GigabitEthernet1/4/0/2
  2        2        disable                           Ten-GigabitEthernet2/4/0/1
                                                      Ten-GigabitEthernet2/4/0/2

<HP12500>dis irf
 Switch    Slot    Role     Priority     CPU-Mac
   1        0      Slave    1            0210-fc01-0000
   1        1      Slave    1            0210-fc01-0001
 *+2        0      Master   30           0210-fc02-0000
   2        1      Slave    30           0210-fc02-0001
--------------------------------------------------

 * indicates the device is the master.
 + indicates the device through which the user logs in.

 The Bridge MAC of the IRF is: 3ce5-a63c-9e00
 Auto upgrade                : yes
 Mac persistent              : always
 Link-delay timer            : 0 ms
 Domain ID                   : 0
 Auto merge                  : no
```

Status (Pass/Fail)

Pass

Comments	N/A

IRF on 12500 Switches Procedure 2

Objective	Configure IRF on each pair of HP 12500 devices by pre-configuring IRF and then converting chassis. This is the recommended procedure to configure IRF on switches that are new/out of the box or have not been previously converted to IRF mode.
Procedure	Ensure ports are shutdown before configuring. Pre-configure IRF ports, member id, priority: irf-port 1 (use 2 for chassis 2) port group interface ten x/x/x port group interface ten x/x/x irf member 1 (use 2 for chassis 2) irf priority x (default is 1, the higher number is higher priority, make one of the chassis a higher number if desired to specify master) save convert to irf mode chassis convert mode irf (reboot) undo shutdown on ten-gig port for master chassis undo shut on ten-gig ports for slave chassis IRF merge will occur and may require a reboot of slave to complete merge The IRF stack should now be formed. Verify IRF operation: display irf display irf configuration display irf topology
Expected Results / Objectives	Working IRF/Single switch

Results

```
<HP12500>dis irf topo
                              Topology Info
 -------------------------------------------------------------------------
                  IRF-Port1                      IRF-Port2
 Switch     Link        neighbor        Link        neighbor     Belong To
 2          DIS         --              UP          1            0210-fc02-0000
 1          UP          2               DIS         --           0210-fc02-0000

<HP12500>dis irf conf
 MemberID NewID     IRF-Port1                        IRF-Port2
  1        1        Ten-GigabitEthernet1/4/0/1       disable
                    Ten-GigabitEthernet1/4/0/2
  2        2        disable                          Ten-GigabitEthernet2/4/0/1
                                                     Ten-GigabitEthernet2/4/0/2

<HP12500>dis irf
 Switch    Slot    Role     Priority     CPU-Mac
```

```
  1        0      Slave    1           0210-fc01-0000
  1        1      Slave    1           0210-fc01-0001
*+2        0      Master   30          0210-fc02-0000
  2        1      Slave    30          0210-fc02-0001
 --------------------------------------------------

 * indicates the device is the master.
 + indicates the device through which the user logs in.

 The Bridge MAC of the IRF is: 3ce5-a63c-9e00
 Auto upgrade                : yes
Mac persistent              : always
 Link-delay timer            : 0 ms
 Domain ID                   : 0
 Auto merge                  : no
<HP12500>dis lacp
                  ^
 % Incomplete command found at '^' position.
<HP12500>dis lacp lo
<HP12500>dis lacp   ?
  system-id  System ID
```

Status (Pass/Fail)	Pass
Comments	N/A

Layer 2 Configuration and Testing

Configure VLANs

Objective

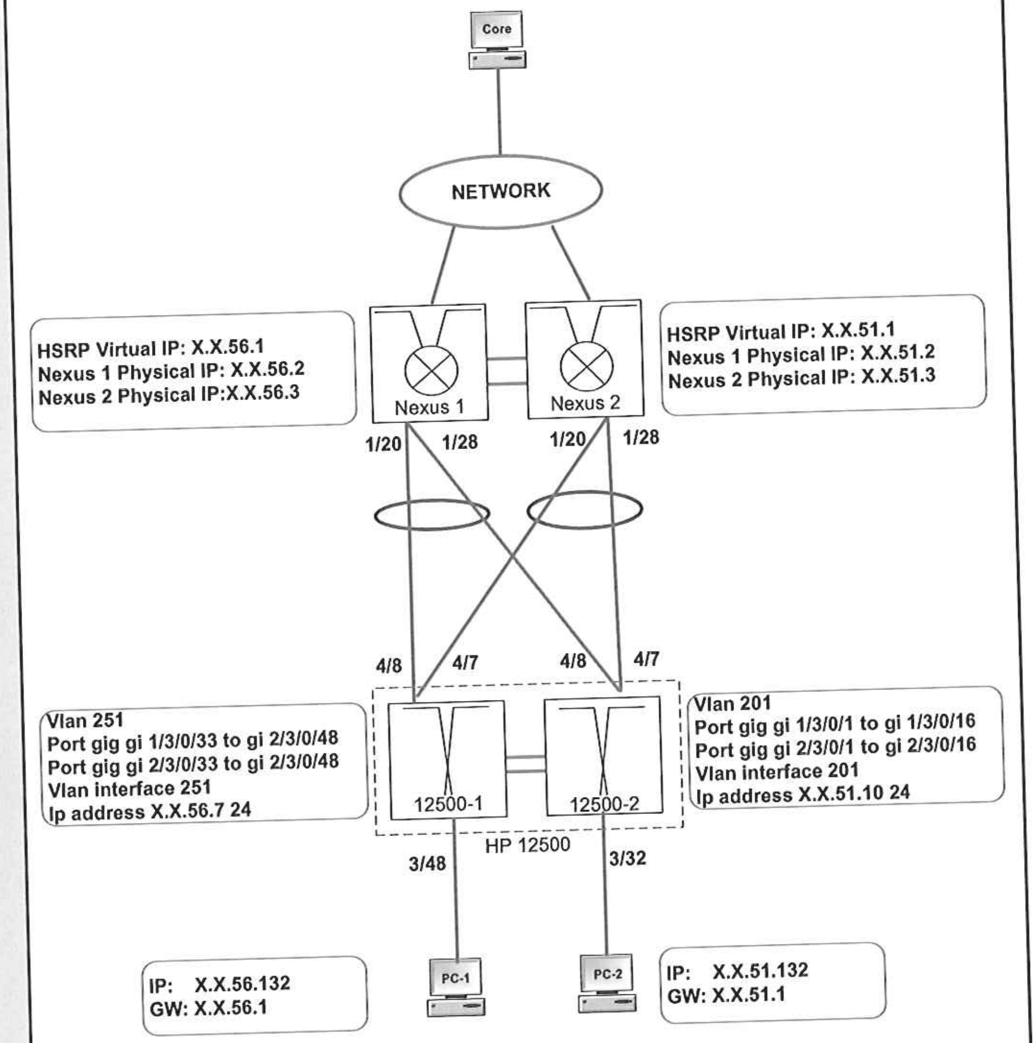

Procedure

Create VLANs, assign access ports:

```
<HP12500>
sys

[HP12500]
vlan 201
[HP12500-vlan-201]
desc Management
port gi 1/3/0/1 to gi 1/3/0/16
port gi 2/3/0/1 to gi 2/3/0/16
quit

[HP12500]
vlan 250
```

```
[HP12500-vlan-250]
desc Prod 1
port gi 1/3/0/17 to gi 1/3/0/32
port gi 2/3/0/17 to gi 2/3/0/32
quit

[HP12500]
vlan 251
[HP12500-vlan-251]
desc Prod 2
port gi 1/3/0/33 to gi 1/3/0/48
port gi 2/3/0/33 to gi 2/3/0/48
quit
```

Create VLAN interfaces, assign IP addresses:

```
[HP12500]
int vlan 201
[HP12500-Interface-Vlan201]
desc Management Vlan
ip address X.X.51.10 24
quit

[HP12500]
int vlan 251
[HP12500-Interface-Vlan251]
desc Management Vlan
address X.X.56.7 24
quit
```

Verify VLAN exists:

```
[HP12500]
dis vlan (Results 1)
```

Verify VLAN interfaces with applicable parameters:

```
[HP12500]
dis ip int br (Results 2)
```

Verify VLAN membership:

```
[HP12500]
dis vlan 201 to 251 (Results 3)
```

Expected Results / Objectives	Results 1 – Will indicate applicable VLANs exist. Results 2 – Will indicate applicable VLAN interfaces and addressing exist. Results 3 – Will indicate VLANs with applicable tagged and untagged port membership.

Results

Results 1

```
[HP12500]
```

```
dis vlan
 Total 4 vlan exist(s).
 The following vlans exist:
  1(default), 201, 250-251

Results 2
[HP12500]
dis ip int br
*down: administratively down
(s): spoofing
Interface                   Physical Protocol IP Address     Description
Vlan-interface201              down     down     X.X.51.10   Vlan-inte...
Vlan-interface250              down     down     unassigned  Vlan-inte...
Vlan-interface251              down     down     X.X.56.7    Vlan-inte...

Results 3
[HP12500]
dis vlan 201 to 251
 VLAN ID: 201
 VLAN Type: static
 Route Interface: configured
 IP Address: X.X.51.10
 Subnet Mask: 255.255.255.0
 Description: VLAN 0201
 Name: VLAN 0201
 Broadcast MAX-ratio: 100%
 Tagged   Ports: none
 Untagged Ports:
    GigabitEthernet0/0/0

 VLAN ID: 250
 VLAN Type: static
 Route Interface: configured
 Description: VLAN 0250
 Name: VLAN 0250
 Broadcast MAX-ratio: 100%
 Tagged   Ports: none
 Untagged Ports: none

 VLAN ID: 251
 VLAN Type: static
 Route Interface: configured
 IP Address: X.X.56.7
 Subnet Mask: 255.255.255.0
 Description: VLAN 0251
 Name: VLAN 0251
 Broadcast MAX-ratio: 100%
 Tagged   Ports: none
 Untagged Ports: none
```

Status (Pass/Fail)	Pass
Comments	N/A

Objective	Configure MST/PVST+ to interoperate as per diagram 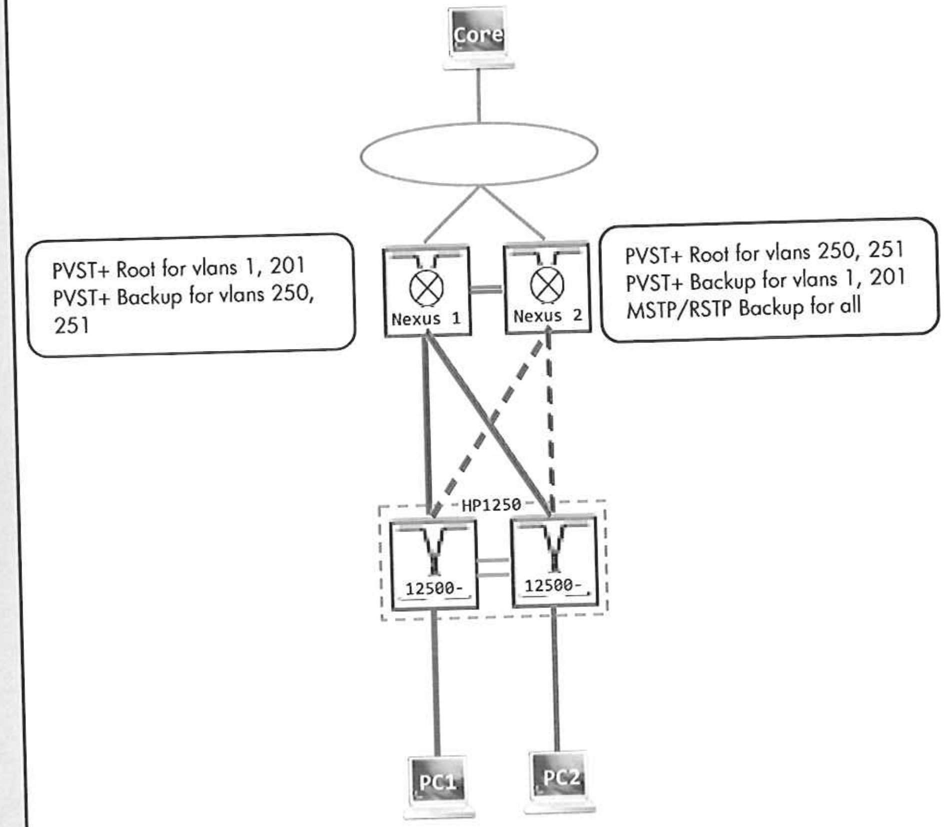
Procedure	Configure pvst+ on cisco switches as per diagram. HP config:

PVST+ Root for vlans 1, 201
PVST+ Backup for vlans 250, 251

PVST+ Root for vlans 250, 251
PVST+ Backup for vlans 1, 201
MSTP/RSTP Backup for all

```
[HP12500]
stp mode mstp
stp pathcost dot1t
stp reg

[HP12500-mst-region]
region-name PVST
revision-level 1
instance 0 vlan 1 to 4094
```

```
active region-configuration
quit

[HP12500]
stp instance 0 priority 61440
stp bpdu-protection
stp enable
```

Vverify region configuration:

```
[HP12500]
stp region-con

[HP12500-mst-region]
dis this (results 1)
```

Verify mst operation (note: LAGs and port channels are required):

```
[HP12500]
dis stp region (results 2)
```

Expected Results / Objectives

Results 1 – Displays region config that matches parameters
Results 2 – Displays region config with digest, instances and VLAN members
Results 3 – (requires LAG and port channel configuration)

Root path for MST 0 should be BRI 10.
Alternate path for MST 0 should be BRI 11.

Results

```
[HP12500]
stp region-con

Results 1
[HP12500-mst-region]
dis this
#
stp region-configuration
 region-name PVST
 revision-level 1
 active region-configuration
#
return
[HP12500-mst-region]

Results 2
[HP12500]
dis stp region-configuration
 Oper configuration
   Format selector      :0
   Region name           :PVST
   Revision level       :1
```

```
 Configuration digest :0xac36177f50283cd4b83821d8ab26de62

 Instance   Vlans Mapped
    0        1 to 4094


Results 3
[HP12500]

dis stp br

MSTID     Port                     Role  STP State     Protection
  0      Bridge-Aggregation10            ROOT  FORWARDING    NONE
  0      Bridge-Aggregation11            ALTE  DISCARDING    NONE
  0      GigabitEthernet1/3/0/1          DESI  FORWARDING    BPDU
  0      GigabitEthernet1/3/0/2          DESI  FORWARDING    BPDU
  0      GigabitEthernet1/3/0/32         DESI  FORWARDING    BPDU
  0      GigabitEthernet2/3/0/1          DESI  FORWARDING    BPDU
  0      GigabitEthernet2/3/0/2          DESI  FORWARDING    BPDU
  0      GigabitEthernet2/3/0/3          DESI  FORWARDING    BPDU
  0      GigabitEthernet2/3/0/4          DESI  FORWARDING    BPDU
  0      GigabitEthernet2/3/0/5          DESI  FORWARDING    BPDU
  0      GigabitEthernet2/3/0/18         DESI  FORWARDING    BPDU
  0      GigabitEthernet2/3/0/45         DESI  FORWARDING    BPDU
  0      Ten-GigabitEthernet1/2/0/5      DESI  FORWARDING    BPDU
  0      Ten-GigabitEthernet1/2/0/6      DESI  FORWARDING    BPDU
  0      Ten-GigabitEthernet2/2/0/5      DESI  FORWARDING    BPDU
  0      Ten-GigabitEthernet2/2/0/6      DESI  FORWARDING    BPDU
```

Status (Pass/Fail)	Pass
Comments	NOTE: Set cost on Nexus HSRP link to be less or equal to port-channels to HPs (to ensure path to root is through HSRP link and not down through access layer and back up). set path costing on HP to dot1t standard set HP root priority to max For additional information, reference document "Migration from Cisco PVST+ to H3C STP.pdf". STP convergence failover resulted in 30 second STP convergence verified by Customer personnel (pvst timers can be adjusted to reduce this to RSTP times but customer was not comfortable with the impact to the entire network by changing those timers).

Configure and Test Spanning Tree – MST

Objective

Configure Cisco and HP for MST

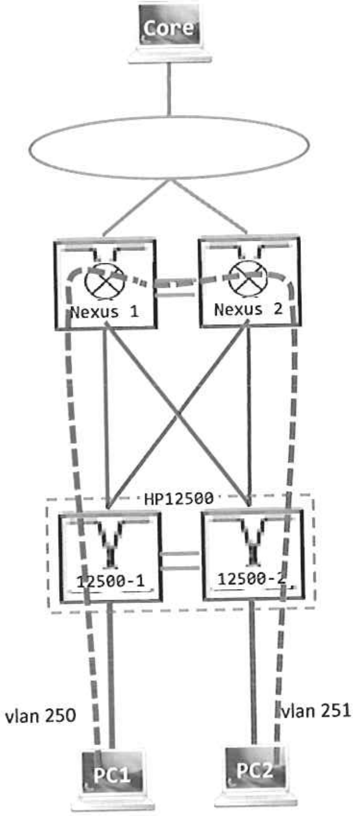

Procedure

Cisco Nexus 1

```
spanning-tree mode mst
spanning-tree mst 0-1 priority 0
spanning-tree mst 2 priority 4096
spanning-tree mst configuration
   name MST
   revision 1
   instance 0-1 vlan 1,201
   instance 2 vlan 250-251
```

Cisco Nexus 2

```
spanning-tree mode mst
spanning-tree mst 0-1 priority 4096
spanning-tree mst 2 priority 0
spanning-tree mst configuration
   name MST
   revision 1
   instance 1 vlan 1,201
   instance 2 vlan 250-251
```

HP12500

```
stp pathcost dot1t
stp mode mstp
stp reg

[HP12500-mst-region]
region-name MST
revision-level 1
instance 1 vlan 1 201
instance 2 vlan 250 to 251
active region-configuration
quit

[HP12500]
stp instance 0 priority 61440
stp bpdu-protection
stp enable
```

Verify region configuration

```
[HP12500]
stp region-con
[HP12500-mst-region]
dis this (results 1)
```

Verify mst digest

```
[HP12500]
dis stp region (results 2)
[Cisco]
show mst config digest
```

Verify mst operation (note: LAGs and port channels are required)

```
[HP12500]
dis stp region (results 3)
```

Expected Results / Objectives

Results 1 – displays region config that matches parameters
Results 2 – hex digest will match on all devices
Results 3 – (requires LAG and port channel configuration)

Root path for MST 0-1 should be BRI 10
Alternate path for MST 0-1 should be BRI 11
Root path for MST 2 should be BRI 11
Alternate path for MST 2 should be BRI 10

Results

Results 1

```
[HP12500-mst-region]
dis this
#
stp region-configuration
 region-name MST
 revision-level 1
 instance 1 vlan 1 201
```

```
instance 2 vlan 250 to 251
active region-configuration
#
 [HP12500-mst-region]

Results 2
[HP12500]
dis stp region-configuration

Oper configuration
  Format selector      :0
  Region name          :MST
  Revision level       :1
  Configuration digest :0x1b1dd7358d9ef13290ff114a060b9d18

  Instance   Vlans Mapped
     0       2 to 200, 202 to 249, 252 to 4094
     1       1, 201
     2       250 to 251
 [HP12500]

Results 3
dis stp br
 MSTID      Port                         Role  STP State     Protection
   0        Bridge-Aggregation10         ROOT  FORWARDING    NONE
   0        Bridge-Aggregation11         ALTE  DISCARDING    NONE
   0        GigabitEthernet1/3/0/1       DESI  FORWARDING    BPDU
   0        GigabitEthernet1/3/0/2       DESI  FORWARDING    BPDU
   0        GigabitEthernet1/3/0/32      DESI  FORWARDING    BPDU
   0        GigabitEthernet2/3/0/1       DESI  FORWARDING    BPDU
   0        GigabitEthernet2/3/0/2       DESI  FORWARDING    BPDU
   0        GigabitEthernet2/3/0/3       DESI  FORWARDING    BPDU
   0        GigabitEthernet2/3/0/4       DESI  FORWARDING    BPDU
   0        GigabitEthernet2/3/0/5       DESI  FORWARDING    BPDU
   0        GigabitEthernet2/3/0/18      DESI  FORWARDING    BPDU
   0        GigabitEthernet2/3/0/45      DESI  FORWARDING    BPDU
   0        Ten-GigabitEthernet1/2/0/5   DESI  FORWARDING    BPDU
   0        Ten-GigabitEthernet1/2/0/6   DESI  FORWARDING    BPDU
   0        Ten-GigabitEthernet2/2/0/5   DESI  FORWARDING    BPDU
   0        Ten-GigabitEthernet2/2/0/6   DESI  FORWARDING    BPDU
   1        Bridge-Aggregation10         ROOT  FORWARDING    NONE
   1        Bridge-Aggregation11         ALTE  DISCARDING    NONE
   1        GigabitEthernet1/3/0/1       DESI  FORWARDING    BPDU
   1        GigabitEthernet1/3/0/2       DESI  FORWARDING    BPDU
   1        GigabitEthernet2/3/0/1       DESI  FORWARDING    BPDU
   1        GigabitEthernet2/3/0/2       DESI  FORWARDING    BPDU
   1        GigabitEthernet2/3/0/3       DESI  FORWARDING    BPDU
   1        GigabitEthernet2/3/0/4       DESI  FORWARDING    BPDU
   1        GigabitEthernet2/3/0/5       DESI  FORWARDING    BPDU
```

```
1      GigabitEthernet2/3/0/45     DESI  FORWARDING    BPDU
2      Bridge-Aggregation10       ALTE  DISCARDING    NONE
2      Bridge-Aggregation11       ROOT  FORWARDING    NONE
2      GigabitEthernet1/3/0/32     DESI  FORWARDING    BPDU
2      GigabitEthernet2/3/0/18     DESI  FORWARDING    BPDU
```

Status (Pass/Fail)	Pass
Comments	STP convergence failover resulted in 3 packet loss – verified by customer.

Configure vPC and Spanning Tree

Objective	Configure vPC on Nexus and remove spanning tree

Procedure

This procedure was completed by customer on the Nexus 7000s

On Nexus
Shut down port channels, configure vPC and create a single port channel (configured like the other ones) to include all four links).

On HP

```
[HP12500]
int bri 10
shut
int bri 11
shut

int bri 12
description Temp LAG to VPC
link-aggregation mode dynamic
quit

int ten 1/2/0/31 to ten1/2/0/32
undo link-ag group
port link-ag group 12
quit

int ten 2/2/0/31 to ten2/2/0/32
undo link-ag group
port link-ag group 12
quit

int bri 12
port link-type trunk
port trunk permit vlan 1 201 205 250 to 251
undo shut
```

```
quit
```

Then you can roll back spanning tree to minimum mstp/rstp config

[HP12500]

```
stp mode mstp
stp pathcost dot1t
stp reg
```

[HP12500-mst-region]

```
region-name MST
revision-level 1
instance 0 vlan 1 to 4094
```

[HP12500-stp-region]

```
active region-configuration
quit
stp instance 0 priority 61440
stp bpdu-protection
stp enable
```

On Cisco Nexus 7000
Roll back Nexus to be MST instance 0 as such - pvst+ will work also,but going foward if you were to move to MST and VPC this is what you would do (unless you had a mixed environment)

```
spanning-tree mode mst
spanning-tree mst 0 priority 0
spanning-tree mst configuration
  name MST
  revision 1
  instance 0 vlan 1-4094
```

Enable the port channels on Nexus if needed and you should be good to go do the following to verify

On HP
[HP12500]

```
dis link-ag ver
```

(you should see bri 12 ports in "S" state and ACDEF on both sides)

```
dis stp br
```

(you should see bri 12 as root path and all else desi)

Expected Results / Objectives	successful vPC link-aggregation across two Nexus chassis remove spanning tree.
Results	vPC was able to span across both Nexus chassis and HP switches did not require spanning-tree – however the Nexus configuration did require spanning tree.

Status (Pass/Fail)	Pass (HP)
Comments	N/A

Configure and Failover 802.1Q LACP Link-Aggregation Groups

Objective

Configure Link Aggregation groups as per diagram

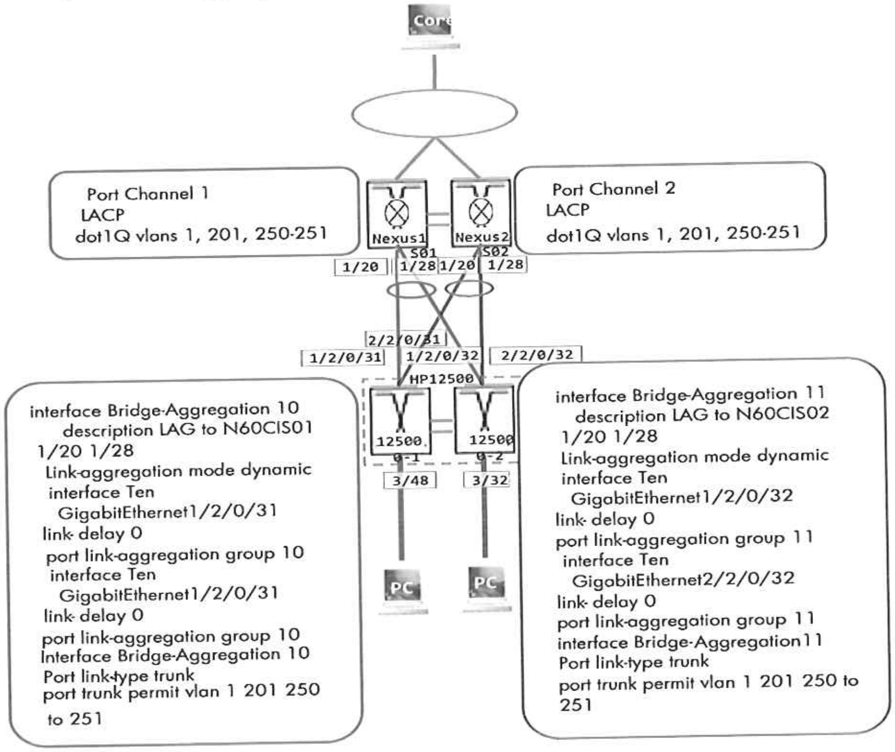

```
interface Bridge-Aggregation 10
    description LAG to N60CIS01
1/20 1/28
 Link-aggregation mode dynamic
 interface Ten
   GigabitEthernet1/2/0/31
link- delay 0
 port link-aggregation group 10
  interface Ten
    GigabitEthernet1/2/0/31
link- delay 0
port link-aggregation group 10
Interface Bridge-Aggregation 10
Port link-type trunk
port trunk permit vlan 1 201 250
 to 251
```

```
interface Bridge-Aggregation 11
   description LAG to N60CIS02
 1/20 1/28
 Link-aggregation mode dynamic
 interface Ten
  GigabitEthernet1/2/0/32
link- delay 0
port link-aggregation group 11
 interface Ten
  GigabitEthernet2/2/0/32
link- delay 0
port link-aggregation group 11
interface Bridge-Aggregation11
Port link-type trunk
port trunk permit vlan 1 201 250 to
251
```

Procedure

Order is important in LAG group configuration:

1. Clear out existing port configs
2. Create LAG group and set it to dynamic if LACP is desired
3. Add the ports to the LAG
4. Do remaining config on the LAG

```
[HP12500]
int bri 10
desc LAG to NEXUS1
link-aggregation mode dynamic
quit

[HP12500]
interface Ten-GigabitEthernet1/2/0/31
link-delay 0
port link-aggregation group 10
quit

[HP12500]
```

```
interface Ten-GigabitEthernet2/2/0/31
link-delay 0
port link-aggregation group 10
quit

[HP12500]
int bri 10
port link-type trunk
port trunk permit vlan 1 201 250 to 251

[HP12500]
int bri 11
desc LAG to NEXUS2
link-aggregation mode dynamic
quit

[HP12500]
interface Ten-GigabitEthernet1/2/0/32
link-delay 0
port link-aggregation group 11
quit

[HP12500]
interface Ten-GigabitEthernet2/2/0/32
link-delay 0
port link-aggregation group 11
quit

[HP12500]
int bri 11
port link-type trunk
port trunk permit vlan 1 201 250 to 251
```

Configure Cisco port channels as per diagram – connect and verify

```
[HP12500]
dis link-ag verbose (results 1)
```

Expected Results / Objectives	Results 1 – link-aggregation group will display port members in "S" for selected state with flags {ACDEF} for both local and remote members

Results

```
[HP12500]
dis link-ag verbose

Aggregation Interface: Bridge-Aggregation10
Aggregation Mode: Dynamic
Loadsharing Type: Shar
System ID: 0x8000, 3ce5-a63c-9e00
Local:
```

```
 Port            Status  Priority Oper-Key  Flag
--------------------------------------------------------------------------------
 XGE1/2/0/31      S       32768    5         {ACDEF}
 XGE2/2/0/31      S       32768    5         {ACDEF}
Remote:
 Actor            Partner Priority Oper-Key  SystemID              Flag
--------------------------------------------------------------------------------
 XGE1/2/0/31      276     32768    0         0x8000, 0024-986d-b3c2 {ACDEF}
 XGE2/2/0/31      284     32768    0         0x8000, 0024-986d-b3c2 {ACDEF}
Aggregation Interface: Bridge-Aggregation11
Aggregation Mode: Dynamic
Loadsharing Type: Shar
System ID: 0x8000, 3ce5-a63c-9e00
Local:
 Port            Status  Priority Oper-Key  Flag
--------------------------------------------------------------------------------
 XGE1/2/0/32      S       32768    6         {ACDEF}
 XGE2/2/0/32      S       32768    6         {ACDEF}
Remote:
 Actor            Partner Priority Oper-Key  SystemID              Flag
--------------------------------------------------------------------------------
 XGE1/2/0/32      276     32768    1         0x8000, 0026-51ce-11c2 {ACDEF}
 XGE2/2/0/32      284     32768    1         0x8000, 0026-51ce-11c2 {ACDEF}
[HP12500]
```

Status (Pass/Fail)	Pass
Comments	LACP failover between VLANs resulted in zero packet loss – verified by customer.

Configure Individual dot1Q Trunk ports

Objective	Configure Individual Trunk Ports
Procedure	Individual trunk ports were not required for this specific scenario – dot1Q trunks were included in the link-aggregation groups – the procedures for individual interfaces are the same as in the link-aggregation group except that the commands are applied to the interface as provided in the example below:

```
<12500>
system-view
[12500]
interface ten 1/2/0/31
description trunk port to EDGE
port link-type trunk
port trunk permit vlan 10 to 11
undo port trunk permit vlan 1
quit

#verify
[12500]
dis this (results 1)
```

Expected Results / Objectives	Results 1 – Displays the commands entered under the interface.
Results	Displays the commands entered under the interface.
Status (Pass/Fail)	N/A
Comments	N/A

Configure Dual Homed Server Ports (link aggregations)

Objective	Configure LAGs for servers
Procedure	Order is important in LAG group configuration: 1. Clear out existing port configs 2. Create LAG group and set it to dynamic if LACP is desired 3. Add the ports to the LAG 4. Do remaining config on the LAG

```
[12500]
interface Bridge-Aggregation3
description LAG to C7000-1 Bay 1 X4 X5
link-aggregation mode dynamic

interface Ten-GigabitEthernet1/2/0/1
port link-aggregation group 3
interface Ten-GigabitEthernet1/2/0/2
port link-aggregation group 3

interface Bridge-Aggregation3
port link-type trunk
port trunk permit vlan 1 201 250 to 251
stp edged-port enable


interface Bridge-Aggregation4
description LAG to C7000-1 Bay 2 X4 X5
link-aggregation mode dynamic

interface Ten-GigabitEthernet2/2/0/1
port link-aggregation group 4
interface Ten-GigabitEthernet2/2/0/2
port link-aggregation group 4

interface Bridge-Aggregation4
port link-type trunk
```

```
port trunk permit vlan 1 201 250 to 251
stp edged-port enable

interface Bridge-Aggregation5
description LAG to C7000-2 Bay 1 X4 X5
link-aggregation mode dynamic

interface Ten-GigabitEthernet1/2/0/3
port link-aggregation group 5
interface Ten-GigabitEthernet1/2/0/4
port link-aggregation group 5

interface Bridge-Aggregation5
port link-type trunk
port trunk permit vlan 1 201 250 to 251
stp edged-port enable

interface Bridge-Aggregation6
description LAG to C7000-2 Bay 2 X4 X5
link-aggregation mode dynamic

interface Ten-GigabitEthernet2/2/0/3
port link-aggregation group 6
interface Ten-GigabitEthernet2/2/0/4
port link-aggregation group 6

port link-type trunk
port trunk permit vlan 1 201 250 to 251
stp edged-port enable
```

Verify after connection

```
display link-aggregation verbose (results 1)
```

Expected Results / Objectives	Results 1 – link-aggregation group will display port members in "S" for selected state with flags {ACDEF} for both local and remote members.
Results	Undo bri shutdowns and display link ag verbose results not captured
Status (Pass/Fail)	Pass
Comments	N/A

Configure Access Ports

Objective	Configure Individual Ports

Procedure	[12500]
	`system-view`
	
	`interface gig 2/3/0/32`
	`description access port`
	`port access vlan 251`
	
	`interface gig 1/3/0/48`
	`description access port`
	`port access vlan 251`
	
	Additionally you can create a manual port group to assign global commands such as port membership and shutdown.
	
	`port-group manual edge-1`
	`group-member Gi 1/0/1 to Gi 1/0/40`
	`quit`

Expected Results / Objectives	PCs attached to ports are in respective VLANs.
Results	PCs attached to ports are in respective VLANs.
Status (Pass/Fail)	Pass
Comments	dis int br and inter (port labels) labels are not consistent same with int bri

Configuration of Management and Access Parameters

NTP on Access

Objective	Configure switches and routers to use two NTP servers

Procedure	[12500]
	`ntp-service unicast-server X.X.17.33 priority`
	`ntp-service unicast-server X.X.56.114`
	
	`display ntp status`
	`display ntp session`
	`display ntp trace`

Expected Results / Objectives	Session established with ntp service.
Results	Configured, tested and documented by customer.
Status (Pass/Fail)	Pass
Comments	N/A

Terminal Settings Access

Objective	Configure synchronous terminal messages to make configuration easier while system messages are scrolling on the screen.

Procedure

```
[12500]
system-view
 info-center synch
```

Expected Results / Objectives	Makes configuration easier while system messages are scrolling on the screen.
Results	configured, tested and documented customer
Status (Pass/Fail)	Pass
Comments	N/A

Remote Access (Telnet, SSH, FTP) Access

Objective	Configure the Switch for Telnet, SSH, and FTP Where X.X. is listed, change to your appropriate address scheme

Procedure

```
[12500]
system-view
acl number 2000
 description SNMP Access List
 rule 1 permit source X.X.177.0 0.0.0.255
 rule 1 comment Allow Subnet Range
 rule 2 permit source X.X.176.0 0.0.1.255
 rule 2 comment Allow Subnet Range
 rule 3 permit source X.X.26.0 0.0.0.127
 rule 3 comment Allow Subnet Range
 rule 4 permit source X.X.63.0 0.0.0.255
 rule 4 comment Allow Subnet Range
 rule 5 deny
```

```
 rule 5 comment deny everything else
acl number 2010 name ACLTEST
 rule 0 deny source X.X.17.180 0
 rule 5 permit source X.X.204.0 0.0.0.255 counting
acl number 2051
 description snmp acl
#
public-key local create rsa
 public-key local create dsa

ssh server enable
undo ssh server compatible-ssh1x
ssh user admin service-type all authentication-type password

ftp server enable
telnet server enable

local-user admin
password simple admin
authorization-attribute level 3
service-type terminal telnet ssh
service-type ftp
quit

user-interface vty 0 4
authentication-mode scheme
protocol inbound all
acl 2051
quit
```

Expected Results / Objectives	access via ssh
Results	access via ssh (configured and tested by customer)
Status (Pass/Fail)	Pass
Comments	N/A

Loop back Address on Access Switch

Objective	Loopback address
Procedure	

```
[12500]
interface LoopBack0
 description loopback0
```

	ip address X.X.X.50 255.255.255.255 ospf cost 10 igmp enable
Expected Results / Objectives	ping loopback
Results	Configured and tested by customer.
Status (Pass/Fail)	Pass
Comments	N/A

HWTacacs on Access Switches

Objective	Configure Radius to be used with IMC.

Procedure

```
[12500]
hwtacacs scheme tac-scheme
 primary authentication X.X.X.53 (use your own IP Addresses)
 secondary authentication X.X.X.54
 primary authorization X.X.X.53
 secondary authorization X.X.X.54
 primary accounting X.X.X.53
 secondary accounting X.X.X.54
 key authentication cc1mst2key3
 key authorization cc1mst2key3
 key accounting cc1mst2key3
 user-name-format without-domain

domain test-domain
 authentication default hwtacacs-scheme tac-scheme local
 authorization default hwtacacs-scheme tac-scheme local
 accounting default hwtacacs-scheme tac-scheme local
 authentication login hwtacacs-scheme tac-scheme local
 authorization login hwtacacs-scheme tac-scheme local
 accounting login hwtacacs-scheme tac-scheme local
 access-limit disable
 state active
 idle-cut disable
 self-service-url disable

domain default enable test-domain
```

Expected Results /	Access via TACACS+ authentication.

Objectives	
Results	Configured and tested by customer.
Status (Pass/Fail)	Pass
Comments	N/A

SNMP Configuration on Access Switches

Objective	Configure the SNMP Agent.
Procedure	

```
<12500>
system-view

[12500]
acl number 2000
 description SNMP Access List
 rule 1 permit source X.X.177.0 0.0.0.255
 rule 1 comment Allow Subnet Range
 rule 2 permit source X.X.176.0 0.0.1.255
 rule 2 comment Allow Subnet Range
 rule 3 permit source X.X.26.0 0.0.0.127
 rule 3 comment Allow Subnet Range
 rule 4 permit source X.X.63.0 0.0.0.255
 rule 4 comment Allow Subnet Range
 rule 5 deny
 rule 5 comment deny everything else

snmp-agent
 snmp-agent local-engineid 800063A2033CE5A63C9E01
 snmp-agent community read BdsW0rld acl 2000
 snmp-agent community write n3verGues5 acl 2000
 snmp-agent sys-info contact CUSTOMER
 snmp-agent sys-info version all
 snmp-agent trap source LoopBack0
#
```

Expected Results / Objectives	snmp access as configured
Results	Configured and tested by customer.
Status (Pass/Fail)	Pass
Comments	N/A

Syslog on Access Switches

Objective	Configure syslog
Procedure	[12500] info-center loghost X.X.116.39 info-center synchronous
Expected Results / Objectives	Verified in the syslog server.
Results	Configured and tested by customer.
Status (Pass/Fail)	Pass
Comments	N/A

QOS/ACL and Port Mirroring Configuration and Testing

An access control list (ACL) is a set of rules (that is, a set of permit or deny statements) for identifying traffic based on matching criteria such as source address, destination address, and port number. The selected traffic will then be permitted or rejected by predefined security policies.

ACLs are widely used in technologies where traffic identification is desired, such as packet filtering and QoS. To restrict packets based on a common rule, the use of ACLs is suggested. These do not replace other security devices such as a Firewall, IPS or UTM.

It is important to remember that the intent of HP 12500 is to forward packets, so the default configuration does not have any ACLs and all packets are forwarded according to the forwarding table.

As with most implementations of an ACL based filter there are four types of ACLs available, Basic, Advanced, Ethernet frame header and User-defined – the table below describes the numbering and the match criteria for each category.

Category	ACL number	Matching criteria
Basic IPv4 ACL	2000 to 2999	Source IP address
Advanced IPv4 ACL	3000 to 3999	Source IP address, destination IP address, protocol carried over IP, and other Layer 3 or Layer 4 protocol header information
Ethernet frame header ACL	4000 to 4999	Layer 2 protocol header fields such as source MAC address, destination MAC address, 802.1p priority, and link layer protocol type
User-defined ACL	5000 to 5999	Customized information of protocol headers such as IP and MPLS headers

QOS/ACL Configuration and Testing

Objective	Test ACL for IPv4 both inbound and outbound directions
Procedure	Create ACL to match or map to classes:

```
[12500]
acl number 3000 name QOS-VOICE
 description This traffic queued as PLATINUM on WAN
 rule 0 permit udp source-port range 16384 32767 dscp ef
 rule 5 permit udp destination-port range 16384 32767 dscp ef

acl number 3005 name QOS-BUSINESS
 description This traffic queued as SILVER on WAN
 rule 0 permit tcp destination-port eq www counting
 rule 5 permit tcp source-port eq www counting
 rule 10 permit tcp destination-port eq 443 counting
 rule 15 permit tcp source-port eq 443
```

Create traffic classifiers – (cisco class-map)

```
traffic classifier QOS-BUSINESS operator and
 if-match acl name QOS-BUSINESS

traffic classifier QOS-VOICE operator and
 if-match acl name QOS-VOICE
```

create behaviors (part of cisco's policy-map)

```
traffic behavior QOS-BUSINESS
 remark dscp af21

traffic behavior QOS-VOICE
 remark dscp ef
```

Create policy – maps classes with behaviors – (other part of Cisco policy-map)

```
qos policy LAN
 classifier QOS-VOICE behavior QOS-VOICE
 classifier QOS-BUSINESS behavior QOS-BUSINESS
```

Testing procedure

mirror port from one PC port to another

```
mirroring-group 2 local
mirroring-group 2 mirroring-port gig2/3/0/48 inbound
mirroring-group 2 monitor-port gig2/3/0/45
```

Send data to be remarked – using Wireshark to detect the marking.

Results indicated no marking – as per R&D, in the internal ASIC procedure, the "Remark" step is after the "Mirroring" step, Queue entering is after the remark step. HP12500 has ingress buffer.

As a result we moved the mirroring-port to the mirror both of the uplinks to the Nexus and observed the packet re-marked correctly via Wireshark.

Expected Results / Objectives	Packet re-marks dscp
Results	Send data to be remarked – using Wireshark to detect the marking. Results indicated no marking – as per R&D, in the internal ASIC procedure, the "Remark" step is after the "Mirroring" step, Queue entering is after the remark step. 12500 has ingress buffer. As a result we moved the mirroring-port to the mirror both of the uplinks to the Nexus and observed the packet re-marked correctly via Wireshark.
Status (Pass/Fail)	Pass
Comments	N/A

L3 OSPF Configuration and Testing

OSPF Configuration and Testing

Objective	Configure OSPF Global Parameters
Procedure	[12500] ospf 1 area 0.0.0.52 network X.X.0.0 0.255.255.255
Expected Results / Objectives	ip routes propagate
Results	Configured and tested by customer.
Status (Pass/Fail)	Pass

Multicast L2 and L3

L2 Multicast (IGMP-Snooping)

Objective	Test L2 Multicast
Procedure	[12500] enable igmp-snooping globally system-view igmp-snooping Enable igmp-snooping in applicable VLANs: vlan 201 igmp-snooping enable vlan 251 igmp-snooping enable In this case the Nexus will route the mcast streams and the HP switches will snoop for mcast streams. Using multicast-hammer tool (default parameters) setup one PC on one VLAN as a server and the other on the other VLAN as the client – note laptops with locked IPS and FW security are prevented from being the "client" and lab laptops must have windows firewall disabled.
Expected Results / Objectives	Client receives mcast streams.

Results	Client received mcast streams.
Status (Pass/Fail)	Pass
Comments	N/A

L3 Multicast (PIM SM)

Objective	Test L2 Multicast
Procedure	Remove igmp-snooping globally: [12500] system-view undo igmp-snooping remove igmp-snooping in applicable vlans vlan 201 undo igmp-snooping enable vlan 251 undo igmp-snooping enable enable multicast routing system-view muliticast routing-enable enable PIM globally system-view pim Enable igmp and PIM on VLAN interfaces: system-view int vlan 201 igmp enable pim sm int vlan 251 igmp enable pim sm In this case the HPs will route the mcast streams. Using multicast-hammer tool (default parameters) setup one PC on one VLAN as a server and the other on the other VLAN as the client – note laptops with

	locked IPS and FW security are prevented from being the "client" and lab laptops must have windows firewall disabled.
Expected Results / Objectives	Client receives mcast streams
Results	Client received mcast streams
Status (Pass/Fail)	Pass

Failover Testing

Link-Aggregation/LACP Failover

Objective	Link-Aggregation/LACP Failover
Procedure	Setup from previous test case

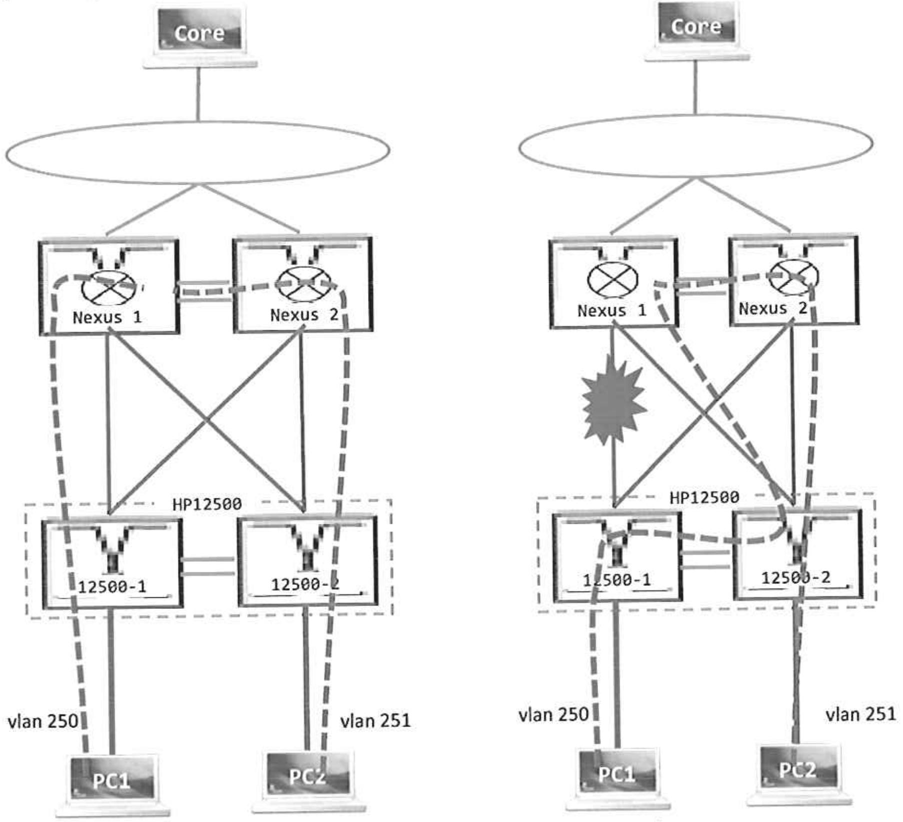

generate UDP traffic from PC1 to PC2 using colasoft packet generator using packet send/spacing rate to 50ms

setup Wireshark on PC2 to filter on all packets coming from IP of PC1 (ip.src == x.x.x.x)

reset packet counters on HP (in user view)

```
<HP12500>
reset counters interface gig
```

start capture on PC2
send test stream of 500 packets
stop capture and verify 500 packets received on PC2

identify which outbound interface

```
<HP12500>
display counters out gig
```

	identify which of the gig interfaces the LACP load balancing chose to use reset packet counters on HP again <HP12500> reset counters interface gig start capture on PC2 – with same filter send the same 500 packets from PC1 about halfway through the send, pull or down the link that the traffic was using after send is complete – verify packets received/displayed on PC2
Expected Results / Objectives	500 packets displayed (zero packet loss at 50ms packet rate).
Results	500 packets displayed/zero packet loss
Status (Pass/Fail)	Pass
Comments	N/A

Spanning-Tree Convergence Failover MST -> PVST

Objective	STP Failover MST/PVST+

Procedure	Setup from previous test case
	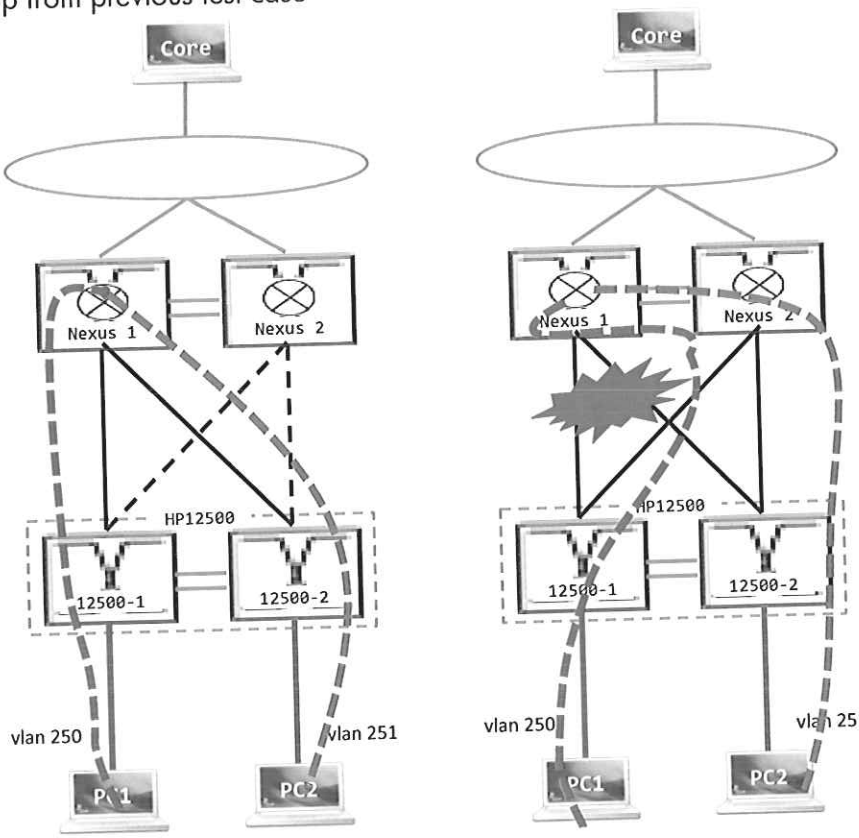
	
	NOTE – HSRP will still dictate Cisco Nexus 1 as router
	Generate pings (continuous) from PC1 to PC2
	Shutdown link-aggregation group
	
	`[12500]`
	`int bri 10`
	`shutdown`
	
	Spanning tree will converge to use alternate path – during this time pings will be lost – after recovery pings will start again – record ping loss and time.

Expected Results / Objectives	Pings will be lost, converge time will depend on PVST timers
Results	6 pings lost ~30 seconds.
Status (Pass/Fail)	Pass
Comments	Lowering the timers on PVST+ will reduce this down to RSTP failover times, however customer was not comfortable with this option as impact on the production network is not predictable at this time.

Spanning-Tree Convergence Failover MST -> MST

Objective	STP Failover MST/MST
Procedure	setup from previous test case 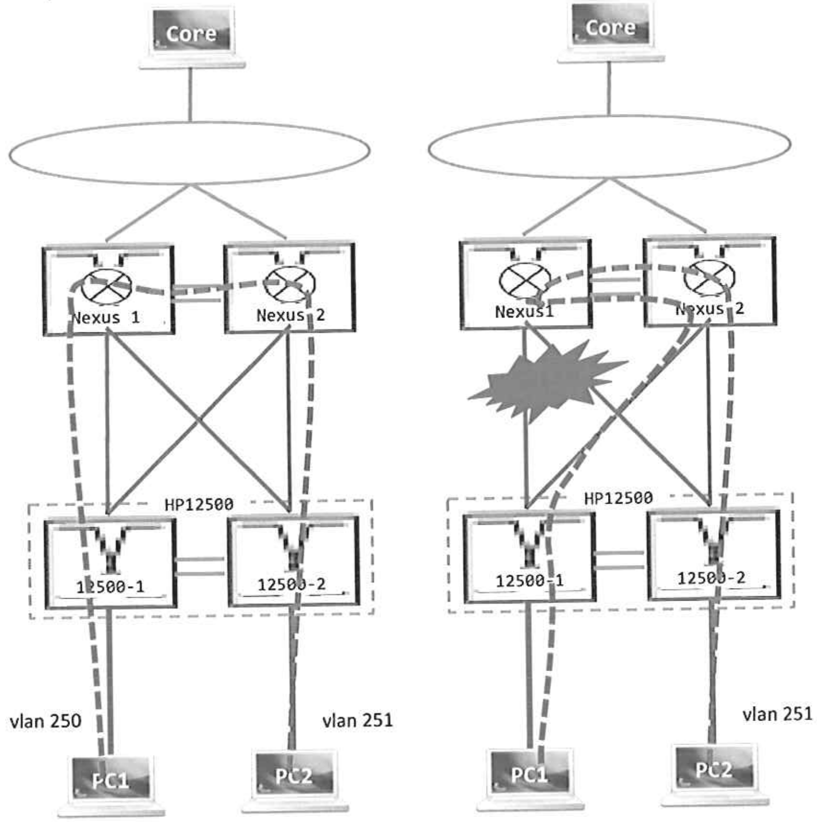 NOTE – HSRP will still dictate Nexus 1 as router generate pings (continuous) from PC1 to PC2 [12500] shutdown link-aggregation group int bri 10 shutdown Spanning tree will converge to use alternate path – during this time pings will be lost – after recovery pings will start again – record ping loss and time.
Expected Results / Objectives	Pings will be lost, converge time will depend on PVST timers.
Results	6 pings lost ~30 seconds.
Status (Pass/Fail)	Pass
Comments	N/A

Active Main Board (Supervisor/CPU) failure

Objective	Fail main board on master chassis
Procedure	Generate and capture traffic as in LACP link-ag failover testing (500 packets). Display device to determine IRF master main board. <12500> display device Start capture and traffic Pull main board out of chassis Record packets received/displayed
Expected Results / Objectives	Zero packet loss
Results	Zero packet loss
Status (Pass/Fail)	Pass
Comments	N/A

Fail Chassis and Observe IRF Merge

Objective	Fail chassis and observe IRF Merge.
Procedure	Power off and on master chassis and observe IRF merge.
Expected Results / Objectives	IRF will be re-established – no second reboot should be required.
Results	IRF will be re-established – no second reboot should be required.
Status (Pass/Fail)	Pass
Comments	N/A

POC Results

Success Criteria Status

Explanations and Operational Descriptions
IRF Operation and Observations with Cisco Nexus vPC

As demonstrated in the POC, IRF is a switch virtualization technology that takes two separate switches, extends the control plane across external 10GB links and creates a single virtual switch. In a broad sense, IRF is a stacking technology – in fact, it is the technology we use to stack the non-chassis based switches.

IRF is much different than Cisco Nexus vPC or even Cisco 6509 VSS. IRF combines two single switches into one virtual switch – before all other network configuration tasks. In fact, nearly all other aspects of the network are oblivious to IRF or the fact that there are two physical switches. One of the few exceptions being the Link-aggregation load balancing algorithm which does consider ports on the local switch first, before crossing data via the IRF links. Additionally, even in a stack environment where there is more than 2 switches/paths – all L2 path costing and L3 routing hops are oblivious to the IRF paths or hops.

Cisco Nexus vPC is not a switch virtualization technology – it's a technology that allows port channels (link-aggregations) to span across two separate chassis. All other protocols behave as if there are two separate switches.

Cisco VSS in the 6509 platform is closer to switch virtualization than vPC but from a control plane perspective they are still two separate switches as described by the way packets are handled below.

Packet Forwarding Example

VSS	**IRF**
The VSS active supervisor engine runs the Layer 2 and Layer 3 protocols and features for the VSS and manages the DFC modules for both chassis. The supervisor in the other chassis is in standby.	Each member device has complete Layer 2 and Layer 3 forwarding capabilities and does not rely on the active main board for updates
Both chassis perform packet forwarding for ingress traffic on their interfaces. If possible, ingress traffic is forwarded to an outgoing interface on the same chassis to minimize data traffic that must traverse the VSL.	When a member device receives a Layer 2/3 packet to be forwarded, it finds the outbound interface (and the next hop) of the packet by searching its Layer 2/3 forwarding table, and then forwards the packet from the outbound interface.
Because the VSS standby chassis is actively forwarding traffic, the VSS active supervisor engine distributes updates to the VSS standby supervisor engine PFC and all VSS standby chassis DFCs.	The outbound interface can be on the local device or on another member device. Forwarding packets from the local device to another member device is unknown to the external, that is, no matter how many

	member devices the Layer 3 packets traverse, the hop count is increased by one only, that is, the packets traverse one network device only.

Spanning Tree Scenarios and Observed Behaviors

This POC configured and tested spanning tree behaviors in three specific scenarios:

- MST → PVST+ Interoperation
- MST → MST operation
- Spanning Tree behaviors between IRF and vPC

PVST+ (per VLAN spanning tree) is a Cisco proprietary protocol that provides similar function to MST (multiple spanning tree protocol) which is the industry standard. In simple terms, Spanning Tree protocol (of any type) has the following characteristics:

- It operates at layer 2 and is only applicable to VLANs and interfaces that are either access (in a single VLAN) or 802.1Q tagged trunk ports that can allow many VLANs.
- It allows networks to be designed with multiple paths (redundancy) without incurring loops.
- It mitigates loops by determining the topology (paths) and blocking the redundant paths.
- If the primary path/device should fail, the alternate/redundant path will become active and communication is restored.

Over the years the standard spanning tree protocol (STP) has been changed to improve failover times, which is called Rapid Spanning Tree Protocol (RSTP). It provides a means to balance L2 traffic across multiple root bridges (typically 2) by pointing a range of VLANs to one root bridge and the other range to the second bridge. Each root would be the backup bridge for the other set of VLANs. This functionality is in PVST+ and MST.

There is not much difference in failover times – both protocols follow the 2-3 seconds of RSTP for converge time – however, that time does increase depending on how many devices/bridges in the spanning tree.

The main differences between the two protocols are as follows:

- PVST+ is easier to configure and is more forgiving when it comes to typos and configuration errors. PVST+ configuration on access switches only have to "know" about its local VLANs. Conversely, MST configurations must describe the same VLANs throughout the entire network or it will not calculate properly and traffic may take non-optimal paths.

- The advantage that MST has over PVST+ is in CPU cycles. PVST+ creates a spanning tree "instance" for every VLAN in the configuration – so for networks with many VLANs, there is a potential to run into processing or VLAN limits, depending on the equipment.

- In an MST configuration the network in scope is considered a "region" and within each a region you create 1 instance per root bridge/path out of a device (typically 2) – you split up the VLANs into these two instances. Instance 1 has a root which also is the backup root for instance 2 (and vice versa).

- There is an instance 0 which is the CIST (similar to the default VLAN) – it's the instance that other MST regions use – it's typically bundled in with the same root as instance 1.

Both Cisco and HP MST configurations are documented in the previous test cases but here are the MST configuration rules of thumb:

- The MST region number (used internally by the protocol) is calculated using the following configurable parameters:
    - Region name
    - Revision number
    - number of instances
    - VLAN membership within those instances

- Every single switch in the MST region must have the exact same parameters for those items above or it will be in a different region and will live in instance 0.

- To avoid future re-configuration of spanning tree – carve up all possible VLANs into instances in large ranges, if possible. PVST+ had an odd/even recommendation but that concept is difficult with MST configuration and leaves much more room for error. For example consider typing all odd characters in the range of 100:
    - vlan 101, 103, 105, 107 . . . 199
- or the range
    - vlan 100 to 199

  In this case you could split up odd even by the leading number – using 100-199 in one instance and 200 – 299 in second.

- Planning out VLAN usage in ranges will lessen the chance of configuration errors and scoping in as many VLANs as possible now will save reconfiguring and convergence events in the future.

NOTE: Ensure that all devices in the network that may participate in spanning tree, including firewalls and/or load balancers support MST and VLAN ranges as described.

PVST+ → MST

In this configuration, you can simply use RSTP or MST with a single instance. MST will interoperate with PVST+ via VLAN 1 only – so VLAN 1 must exist on the links connected to the Cisco root bridges. Also – from the MST or RSTP perspective there will only be one root bridge and the other link will be blocked until failover.

Failover in this scenario was observed at normal spanning tree times of 30 seconds. This can be reduced to RSTP times (3-5 seconds) by changing the PVST+ timers – however more thought would need to be put into the impact to the rest of the devices by changing timers, so caution is advised.

MST → MST

Failover in this scenario was within expectations (4-5 seconds) – Spanning-Tree with vPC.

While spanning tree was successfully disabled with vPC configured – STP was required on the Cisco end or the vPC would not operate.

ISSU as it Works Today

ISSU was first supported on the HP 12500s with code version S12500-CMW520-R1231 released March 18th 2010. Currently, ISSU is only hitless between "compatible" versions of code and completed through a manual process of patching each main board, line card and fabric module individually.

"Compatible" versions of code are not common but can be demonstrated in a lab environment. For that reason – ISSU is not "hitless" (yet) in a production environment.

Hitless ISSU is on the roadmap and is expected in the next major revision of Comware.

Today - IRF ISSU consists of three possible upgrade scenarios:

Compatible: where old and new code versions are completely compatible and each main board and module can be patched without reboots or disruption of service.

Base-Compatible: where old and new versions can co-exist in the same IRF – that means after the Slave chassis is upgraded, IRF comes up and all forwarding is normal. Then when the other chassis reboots to be upgraded there is no spanning tree convergence since at least one link in each lag is still up (on the other chassis) – This type of upgrade was performed in the course of this POC and while there were three reboots, forwarding never stopped. Some packets will get "lost" on the wire but our observations of a switchover of this type is minimal packet loss– 1 ping or so since spanning tree never has to re-converge

Incompatible: where old and new code versions cannot co-exist in same IRF – This means that when the slave switch is upgraded it comes up as a separate switch but with its interfaces down (so they will not forward). When switchover to the new chassis occurs the old active is rebooted and the interfaces on the upgrade chassis become active which triggers a spanning tree convergence.

*Note: Compatibility information can be found in the release notes or by copying the new image to the flash and performing the following command on the swtich:

```
display version comp-matrix file flash:/nameofnewimage.bin
```

The following tables describe the process for each type of ISSU upgrade:

ISSU Compatible Table

COMPATIBLE	ACTION	RESULT
Upgrade Slave CPU on Master Chassis	Upgrades slave CPU on master chassis. Master CPU and Interfaces remain at old version	Slave CPU on active chassis has new version – All traffic and operations function without disruption
Perform switchover from current active CPU to newly upgraded CPU on master chassis	Reboots Active/Master Chassis and newly upgraded chassis becomes new active master Manually apply hotfix's to interface cards	New Master CPU and all interfaces in master chassis have been upgraded, new slave CPU (old master) is still at previous revision IRF is operational, all interfaces forwarding, no spanning tree convergence No traffic interruption
Upgrade all other CPUs one by one and apply hotfix's	Upgrade the remaining CPUs and apply hotfix's to the interface cards in slave chassis	This method ensures non-stop services during the ISSU upgrade This method is applicable only when the result of version compatibility check is compatible

ISSU Base-Compatible Table

BASE-COMPATIBLE	ACTION	RESULT
Upgrade Slave device	Upgrade the current Slave chassis, Slave chassis reboots, all CPUs in chassis and all interfaces in chassis are upgraded at this time	After reboot IRF is operational, all interfaces forwarding, no spanning tree convergence Dual homed traffic using that switch fails over (sub-second) as dictated by LACP specification.
Perform switchover from current Master device to current Slave device	Reboots Active/Master chassis and newly upgraded chassis becomes new active master	After switchover – the old slave chassis is now the new active master, running the new code. The old active master boots up as the slave chassis IRF is operational, all interfaces forwarding, no spanning tree convergence Dual homed traffic using that switch fails over (sub-second) as dictated by LACP specification.
Accept the new master device	Accepts new active/.master upgrade	No reboot – simply allows for commit in next step
Commit/Upgrade the new slave device	Reboots slave chassis (old active/master) upgrades the CPUs and all the interfaces	All chassis, CPUs and interfaces are running the new code and the switch is completely operational Dual homed traffic using that switch fails over (sub-second) as dictated by LACP specification.

ISSU Incompatible Table

INCOMPATIBLE	ACTION	RESULT
Upgrade Slave device	Upgrade the current Slave Chassis, Slave chassis reboots, all CPUs in chassis and all interfaces in chassis are upgraded at this time	After reboot IRF is NOT operational, interfaces are not forwarding (MAD down), no spanning tree convergence Dual homed traffic using that switch fails over (sub-second) as dictated by LACP specification.
Perform switchover from current Master device to current Slave device	Reboots current Active/Master Chassis and the interfaces on newly upgraded chassis become active	When new master chassis links become active, Spanning tree must converges. Ideally within 4-5 seconds depending on spanning tree configuration – legacy STP can take up to 50 seconds to converge. After convergence, all interfaces on that chassis are forwarding – IRF remains down until the old active has completed its reboot/upgrade. After reboot is complete – the old master becomes the new slave and all modules and CPUs have been upgraded.

Define Working Modes of Switch

The system working mode is the perspective of how to allocate the resource by the Active Main Board/CPU System working mode is managed by the CPU and will affect the capabilities and operation of line cards. When there are two or more (IRF system) Main Board/CPUs in the chassis, all the main boards should work in the same working mode, otherwise the slave boards can't work properly or two chassis can't combine to be an IRF system.

If main boards are configured in a different mode from the others (as what happened during our POC) – EACH will have to be changed back manually - the command "system working mode xxxx” is only applicable to the ACTIVE Main Board of the system. The system working mode of all MPUs can be changed in the Bootrom menu (CRTL-B during boot).

The MPU is in standard mode by default.

What does this mean?

The LEB line cards can only work in the standard mode
The LEC line card can work in the Standard/Routee/Bridgee mode
The LEF line card can work in the Standard/Routee/Bridgee/Advanced mode

```
<Sysname> display system working mode
Current system working mode: routee
Working mode after system restart: bridgee
Enhance mode will take effect after system restart
```

So if the Active Main Board is running in the standard mode, all the line cards can work properly. But if the Master MPU is running in the Router/Bridge mode, only the LEC line card can work properly, the LEB line card will not boot.

LEB can only work at "Standard mode",
LEC can be configured at "Standard mode" or "Enhanced mode“,
LEF can be configured at “Advanced mode”

When LEB(standard) is mixed with LEC(Enhanced) and/or LEF(Advanced), the system will work ONLY at "Standard mode".

When LEC(Enhanced) is mixed with LEF(Advanced), the system will work ONLY at either “Standard” or “Enhanced”, Not “Advanced Mode”.

		EB	EC			EF			
System Working Mode		Standard	Standard	Bridgee	Routee	Standard	Bridgee	Routee	Advance
ARP table		12K	12K	16K	64K	12K	16K	64K	
MAC table		64K	128K		64K	256K		128K	256K
IPv4 FIB		256K	256K			256K			1M
IPv6 FIB		64K	64K	128K		64K	128K		512K
MPLS L3VPN		1K	1K	2K		1K	2K		4K
VPLS/PBB		Not support	Not support	2K	4K	Not Support	2K	4K	8K
VLL		Not support	Not support	4K		Not support	4K		30K
ACL	2ASIC	4K@48B	16K@40B			64K@40B			
	4ASIC	8K@48B	32K@40B			128K@40B			
	8ASIC	16K@48B	64K@40B			NA			

Loop Guard Event during Spanning Tree Configuration

During the initial configuration of MST → PVST+ interoperability we had a few moments were loop guard was triggered on the Nexus and it was never exactly clear as to why. The customer and TC did a lot of troubleshooting of multiple events during this configuration. Most of the problems involved what path cost algorithm each vendor was using. It's also possible that something with the vPC on the Nexus wasn't exactly right, customer may be able to provide more clarity from the Nexus perspective, as with all interoperability configurations (especially for the first time), there are generally multiple issues to address.

Spanning Tree, path cost and configuration parameters are described "Spanning Tree Scenarios and Observed Behaviors" section **on page 52.** The bottom line is that when each spanning tree scenario was configured and stabilized, the loop guard problem did not reoccur. Lacking an explanation as to why and when, it was an event that occurred while spanning tree was partially or not correctly configured and now that the procedures have been documented and tested, this should not occur again.

Parts List for 12500

Line#	Part Number	Description	Quantity
1.00	JF431B	HP 12508 Switch Chassis	2
1.01	JF429A	HP 12500 2000W AC Power Supply	12
1.02	JF426A	HP 12518 PEM (Power Electrical Module)	2
1.03	JC081A	HP 12508 Spare Fan Assembly	4
1.04	JC072A	HP 12500 Management Module	4
1.05	JC067B	HP 12508 Fabric Module	18
1.06	JC065A	HP 48-port Gig-T LEC 12500 Module	2
1.07	JD117B	HP X130 10G XFP LC SR Transceiver	16
1.08	JC068A	HP 8-port 10-GbE XFP LEC 12500 Module	2
1.09	JD092B	HP X130 10G SFP+ LC SR Transceiver	32
1.10	JC476A	HP 32-port 10GbE SFP+ LEC Module	2

Current Limitations

See 12500 Documentation for additional information.

Troubleshooting

See 12500 Documentation for additional information.

For more information

To read more about HPN products and solutions, go to http://www.hp.com/networking.

Document Version 9, March 2012

networktest

HP/Cisco Switching and Routing Interoperability Test Results

April 2011

HP/Cisco Enterprise Switching and Routing Interoperability

Executive Summary

HP commissioned Network Test to assess interoperability between its enterprise switch/routers and those of Cisco Systems. Working with an extensive test bed that included core-, distribution-, and access-layer devices, Network Test successfully validated interoperability of 13 key protocols used in enterprise networks. Tests involved IPv4 and IPv6; switching and routing; and unicast and multicast traffic. **In all cases described here, the HP and Cisco switch/routers correctly forwarded traffic.**

The following table summarizes results of interoperability testing.

HP / Cisco Protocol Interoperability

	HP A9505	HP E5406zl	HP A5800
VLAN trunking			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓
L2/L3 jumbo frame handling			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓
Link aggregation			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓
Spanning tree protocol			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓
OSPFv2 for IPv4			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓
OSPFv3 for IPv6			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓
Multicast switching and routing			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓
VRRP			
Cisco Catalyst 3750-E	✓	✓	✓
Cisco Catalyst 4506	✓	✓	✓
Cisco Catalyst 6509	✓	✓	✓

Methodology and Results

Figure 1 below illustrates the test bed used to validate interoperability. The HP and Cisco switch/routers used a three-tier design commonly found in enterprise campus networks, with separate devices at the core, distribution, and access layers. A Spirent TestCenter traffic generator/analyzer emulated clients and servers, and externally verified interoperability of the various protocols.

Except where otherwise noted, tests involved connections between each layer of the network, thus validating interoperability of each protocol using every device on the test bed. Also unless otherwise noted, tests also used multiple redundant connections between switch/routers to exercise link aggregation, spanning tree, and routing protocols.

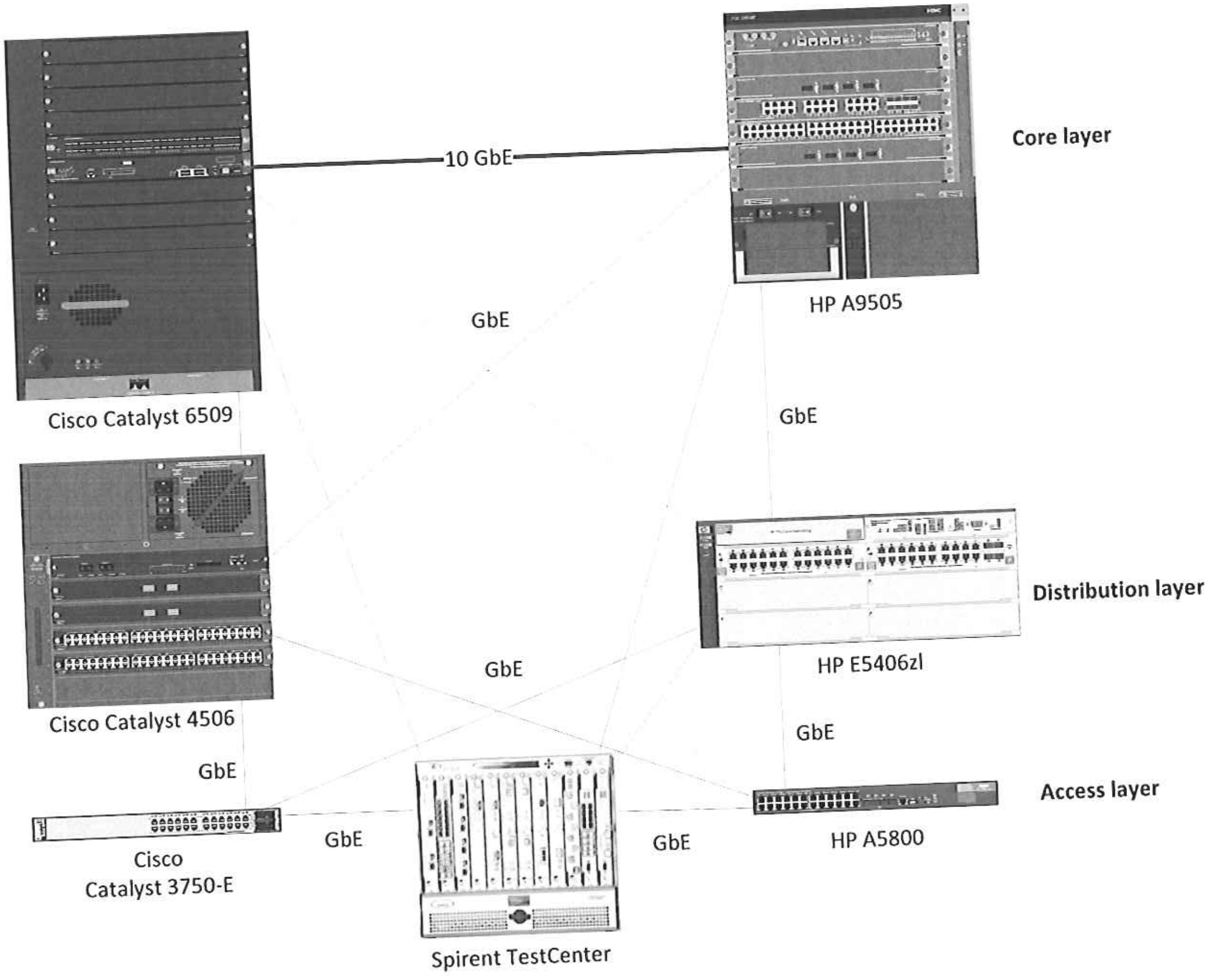

Figure 1: The HP-Cisco Interoperability Test Bed

VLAN Trunking

Network Test evaluated interoperability of IEEE 802.1Q VLAN trunking in three ways: forwarding of allowed tagged traffic; forwarding of allowed untagged (native) traffic; and blocking of disallowed untagged traffic.

Engineers configured five VLANs on each switch, and configured trunk ports between switches to allow traffic from three VLANs as tagged frames and a fourth VLAN as untagged frames. To determine if switches would correctly block disallowed traffic, engineers did not include the fifth VLAN ID in the list of VLANs allowed across trunk ports.

Spirent TestCenter then offered untagged traffic to each HP and Cisco access and distribution switch in a bidirectional pattern. In all cases, traffic counters on the Spirent test instrument verified that **HP and Cisco switches correctly forwarded VLAN traffic that was intended to be forwarded, and did not carry VLAN traffic that was not intended to be forwarded.**

L2/L3 Jumbo Frame Handling

Jumbo Ethernet frames – those larger than the standard maximum length of 1,518 bytes – are commonly used for bulk data-transfer applications such as backups, storage, and disaster recovery. To validate the ability of HP and Cisco switch/routers to exchange jumbo frames, Network Test offered these frames in both switching and routing modes.

In the switching tests, Spirent TestCenter offered 9,216-byte jumbo Ethernet frames using a "partially meshed" topology, meaning all traffic offered to HP devices was destined to Cisco ports and vice-versa. **All HP and Cisco devices correctly switched traffic consisting of jumbo frames.**

In the routing assessment, Network Test enabled OSPF on both HP and Cisco devices. The OSPF protocol will not exchange traffic unless routers agree on the same maximum transmission unit (MTU). In this case, with 9,198-byte IP packets inside 9,216-byte Ethernet frames, **all HP and Cisco devices successfully established OSPF adjacencies as expected, and successfully routed traffic consisting of jumbo frames.**

Link Aggregation

Network Test evaluated the ability of HP and Cisco devices to bundle multiple physical ports into one logical port using the IEEE 802.3ad link aggregation protocol.

Engineers configured the HP and Cisco devices to set up link aggregation groups (LAGs) between the access, distribution, and core layers. Specifically for this test, engineers then disabled any redundant paths through the network, forcing traffic to be forwarded across each LAG. Spirent TestCenter offered bidirectional traffic to eight ports on each access and distribution switch/router, emulating 32 hosts on each port to encourage distribution of flows across the multiple LAG members. **In all cases, the HP and Cisco switches correctly forwarded traffic using link aggregation.**

To validate that link aggregation can supply additional bandwidth by bundling multiple physical ports, engineers configured Spirent TestCenter to offer traffic at an aggregate rate in excess of 1 Gbit/s in each direction, the nominal capacity of a single link. In all cases, the aggregated links carried the additional traffic with zero frame loss.

Spanning Tree Protocol (STP)

The spanning tree protocol serves as a key loop prevention and redundancy mechanism in enterprise networks. Over the years it has been refined with updates such as rapid spanning tree (RSTP) to speed convergence and multiple spanning tree (MSTP) to form a separate spanning tree instance for each VLAN. In addition to these standards-based methods, Cisco switches use proprietary variants called per-VLAN spanning tree plus (PSVT+) and Rapid PVST+.

Network Test verified HP-Cisco interoperability using four variations of spanning tree:

1. RSTP (HP) / PVST+ (Cisco)
2. MSTP (HP and Cisco, using the IEEE 802.1s specification)
3. MSTP (HP) / PVST+ (Cisco)
4. MSTP (HP) / Rapid PVST+ (Cisco)

For each variation, engineers set up redundant connections between all devices, thus forcing spanning tree to select a root bridge and place device ports in either blocked or forwarding states. Engineers then offered traffic to each device using Spirent TestCenter and verified that traffic was received only from an intended port in forwarding state.

While continuing to offer traffic, engineers then tested spanning tree convergence by administratively disabling a port in forwarding state, forcing the spanning tree to bring up ports formerly in blocked state. Engineers verified correct spanning tree operation by observing Spirent TestCenter port counters and by examining the command-line interface (CLI) output for spanning tree on each device. **In all four test cases, spanning tree delivered loop-free operation and seamless failover.**

OSPFv2 for IPv4/OSPFv3 for IPv6

IP routing is a given in enterprise networks, and by far the most commonly used interior gateway protocol is Open Shortest Path First (OSPF).

To validate OSPF interoperability between HP and Cisco devices, engineers enabled OSPF on all switch/routers on the test bed, and then configured Spirent TestCenter to emulate OSPF routers attached to each device. This is a more rigorous and stressful topology than is commonly found in most enterprise networks, where IP routing often is found only on core devices. Here, all switch/routers, including those at the distribution and access layers, brought up OSPF routing sessions and forwarded traffic to and from networks advertised using OSPF.

Engineers conducted these routing tests twice, with IPv4 and IPv6 variations. In IPv4 testing, engineers configured OSPF version 2, while IPv6 testing used the newer OSPFv3 variant of the protocol.[1]

In these tests, Spirent TestCenter emulated OSPF routers attached to each switch/router. After bringing up an OSPF session, these emulated routers used OSPF to advertise networks "behind" them, and then offered traffic to and from these networks.

For this interoperability test to work successfully, HP and Cisco switch/routers would need to share routing information to forward traffic to these emulated networks. That is exactly what happened: **All HP and Cisco devices not only established OSPF sessions over IPv4 and IPv6, but also forwarded all traffic to all networks with zero frame loss observed.**

These results validate OSPF routing interoperability between all HP and Cisco devices, both on IPv4 and IPv6 networks.

Multicast Switching and Routing

Streaming media, conferencing, financial quote services and many other applications are making increasing use of IP multicast. Network Test validated the ability of HP and Cisco equipment to share information about multicast topology both in purely switched and switched/routed environments.

In the switched scenario, engineers configured all HP and Cisco devices in layer-2 mode and enabled IGMP snooping. In the routed scenario, all devices used the Protocol Independent Multicast-Sparse Mode (PIM-SM) routing protocol and OSPF to carry multicast and unicast routing information, respectively.

In both scenarios, a Spirent TestCenter port attached to the Cisco Catalyst 3750-E offered traffic destined to 10 multicast groups while other test ports emulated multicast subscribers to all 10 groups on the HP A5800 and HP E5406zl. Engineers also attached one additional monitor port to each of the HP devices to verify they did not flood multicast frames to non-subscriber ports.

The HP and Cisco devices correctly delivered multicast traffic to subscribers in both switched and routed configurations, and did not flood traffic to non-subscribers.

In addition, Network Test evaluated IGMP snooping support while multicast routing was enabled. When operating in Ethernet switching mode, the HP and Cisco devices use IGMP reports to determine which

---

[1] IETF RFC 2328 describes OSPFv2 and RFC 5340 describes OSPFv3. While the basic mechanics of OSPF are identical in both versions, OSPFv3 introduces new link-state advertisement (LSA) types; removes addressing semantics from OSPF headers; generalizes flooding; and removes OSPF-layer authentication, among other changes.

switch ports have subscribers attached. Working with IGMPv2, engineers verified the HP and Cisco switches correctly populated IGMP snooping tables and forwarded multicast traffic in all cases.

Virtual Router Redundancy Protocol (VRRP)

Network Test verified the ability of HP and Cisco devices to provide router failover using the Virtual Router Redundancy Protocol (VRRP). As defined by the Internet Engineering Task Force (IETF) in RFC 5798, VRRP provides a standard method by which multiple routers select Master and Backup roles, with a Backup router taking over from a Master in the event of a router or link failure.

Testing involved all six HP and Cisco devices as shown in Figure 1 above, with VRRP running on the HP A9505, the HP EE5406zl, and the Cisco Catalyst 6509. The devices running VRRP agreed on a virtual IP (VIP) addresses, verified by examining their respectively CLIs.

Initially, the HP A9505 acted in the Master role and the Cisco Catalyst 6509 acted as Backup. Then engineers configured the Cisco device to take over as Master by changing its priority to force VRRP failover. Again, the two sides agreed on VRRP settings, and traffic counters on Spirent TestCenter showed devices forwarding traffic after the failover.

Engineers then repeated this exercise using the HP EE5406zl and the Cisco Catalyst 6509, and again failover worked as expected.

The results demonstrate that upon failure of an active router or link, HP and Cisco devices work together using VRRP to reroute traffic onto a backup link, with minimal interruption to users and applications.

Conclusion

Interoperability testing was successful in every case where both HP and Cisco devices supported a given protocol. This provides assurance to network professionals considering design or deployment of networks comprised of a mix of HP and Cisco switch/routers.

Appendix A: About Network Test

Network Test is an independent third-party test lab and engineering services consultancy. Our core competencies are performance, security, and conformance assessment of networking equipment and live networks. Our clients include equipment manufacturers, large enterprises, service providers, industry consortia, and trade publications.

Appendix B: Software Releases Tested

This appendix describes the software versions used on the test bed. All tests were conducted in March 2011 at Network Test's facility in Westlake Village, CA, USA.

Component	Version
HP A9505	5.20, Release 1238P08
HP EE5406zl	K.15.03.0007
HP A5800	5.20, Release 1206
Cisco Catalyst 6509	12.2(33)SXI2a
Cisco Catalyst 4506	12.2(20)EWA
Cisco Catalyst 3750-E	12.2(55)SE1
Spirent TestCenter	3.55.5086.0000

Appendix C: Disclaimer

Network Test Inc. has made every attempt to ensure that all test procedures were conducted with the utmost precision and accuracy, but acknowledges that errors do occur. Network Test Inc. shall not be held liable for damages which may result for the use of information contained in this document. All trademarks mentioned in this document are property of their respective owners.

Network Test Inc.
31324 Via Colinas, Suite 113
Westlake Village, CA 91362-6761
USA
+1-818-889-0011
http://networktest.com
info@networktest.com

networktest

HP/Cisco Switching and Routing Interoperability Cookbook

May 2011

TABLE OF CONTENTS

ILLUSTRATIONS

Introduction

Objective

This configuration guide aims to help networking professionals interconnect HP Networking and Cisco Catalyst switches using a variety of protocols commonly found in enterprise campus networks. By following the step-by-step procedures described in this document, it should be possible to verify interoperability and to pass traffic between the two vendors' switches.

Intended audience

This guide is intended for any network architect, administrator, or engineer who needs to interconnect HP and Cisco Ethernet switches.

This guide assumes familiarity with basic Ethernet and TCP/IP networking concepts, as well as at least limited experience with the HP Networking and Cisco IOS command-line interfaces (CLIs). No previous experience is assumed for the protocols discussed in this document.

For basic TCP/IP networking concepts, the standard references are ***Internetworking with TCP/IP, Volume 1*** by Douglas E. Comer and ***TCP/IP Illustrated, Volume 1*** by W. Richard Stevens.

For IP multicast topics, ***Deploying IP Multicast in the Enterprise*** by Thomas A. Maufer is a popular choice.

Devices under test

Using the commands given in this document, Network Test has verified interoperability between the HP A9505, HP E5406zl, and HP A5800 Ethernet switches and Cisco Catalyst 6509, Cisco Catalyst 4506, and Catalyst 3750-E Ethernet switches. Appendix B lists software versions used.

Except where specifically noted, command syntax for HP Networking and Cisco Catalyst switches does not change across product lines. In cases where HP A-series and E-series switches use different command syntax, this is explicitly noted.

Conventions used in this document

The following table lists text and syntax conventions.

<table>
<tr><th>Conventions</th><th>Description</th><th>Examples</th></tr>
<tr><td>Bold Type</td><td>Represents user-inputted text.</td><td>To enter configuration mode, type the <code>system-view</code> command:<pre>
<HP5800> system-view
</pre></td></tr>
<tr><td><code>Fixed-width text like this</code></td><td>Represents output that appears on the terminal screen.</td><td><pre>
<A9505> display stp bridge

 MSTID   Port
Role STP State
Protection

   0   Bridge-
Aggregation20   ROOT
FORWARDING   NONE

   0
GigabitEthernet3/0/11
DESI FORWARDING   NONE

   0
GigabitEthernet3/0/16
DESI FORWARDING   NONE
</pre></td></tr>
<tr><td><i>Italic text like this</i></td><td>Introduces important new terms
Identifies book titles
Identifies RFC and Internet-draft titles</td><td>A policy <i>term</i> is a named structure that defines match conditions and actions.
<i>TCP/IP Illustrated Volume 1</i> by W. Richard Stevens.
RFC 4814, <i>Hash and Stuffing: Overlooked Factors in Network Device Benchmarking</i></td></tr>
</table>

Interoperability testing

For each protocol tested, this document uses a five-section format consisting of objective, technical background, HP configuration, Cisco configuration, and test validation.

Topology

Except where otherwise noted, engineers used the standard test bed shown in Figure 1 below to validate protocol interoperability. The test bed uses the three-tier network design commonly found in campus enterprise networks, with access, distribution, and core layers represented. In this example network, access switches (HP A5800 and Cisco Catalyst C3750-E) connect to distribution switches (HP E5406zl and Cisco Catalyst 4506), which in turn connect to core switches (HP A9505 and Cisco Catalyst 6509). For redundancy, multiple connections exist between switch layers.

Test engineers configured link aggregation between HP A5800 and HP E5406zl switches; between HP E5406zl and HP A9505 switches; between HP A9505 and Cisco Catalyst 6509 switches; between Cisco Catalyst 6509 and Cisco Catalyst 4506 switches; and between Cisco Catalyst 4506 and Cisco Catalyst 3750-E switches. The use of link aggregation is not mandatory, however.

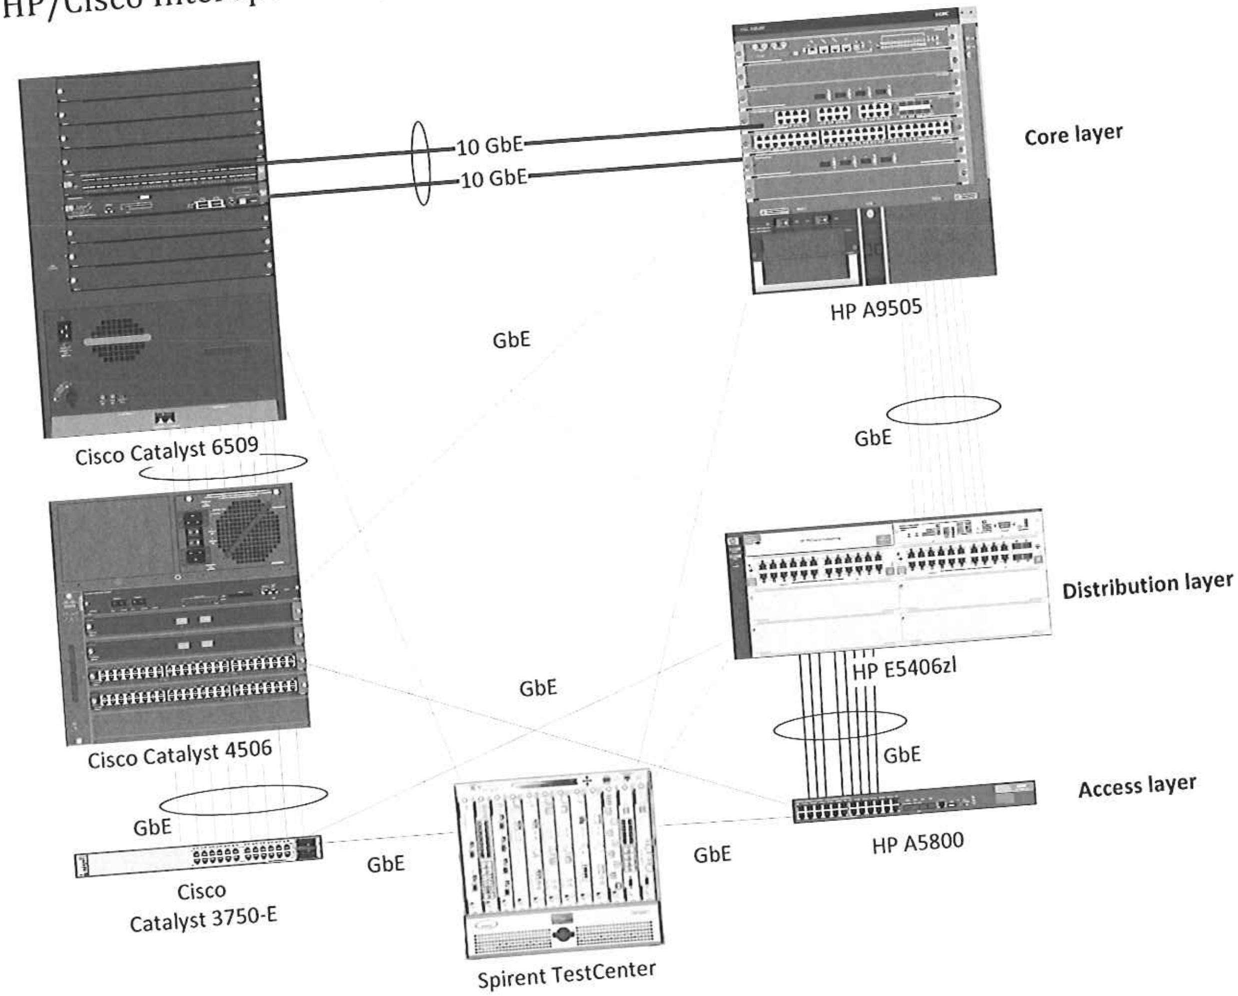

Figure 1: HP-Cisco interoperability test bed

Virtual LAN (VLAN) trunking

Objective

To verify interoperability of IEEE 802.1Q VLAN trunking between HP Networking and Cisco Catalyst switches using tagged traffic.

To verify interoperability of IEEE 802.1Q VLAN trunking between HP Networking and Cisco Catalyst switches using untagged traffic.

Background

The IEEE 802.1Q specification defines a method of defining virtual broadcast domains. A 4-byte VLAN header, usually called a "tag," allows definition of broadcast domains that may differ from physical switch topology. With VLANs, all switch ports are members of the same broadcast domain; with VLAN tagging, a network manager can set up multiple broadcast domains across switches, and restrict broadcasts for different VLANs on different ports.

Topology

This configuration example will validate VLAN trunking interoperability between HP Networking and Cisco Catalyst switches in three ways:

- The switches will forward allowed tagged traffic from multiple VLANs across a trunk port.
- The switches will forward allowed untagged traffic from a native VLAN across a trunk port.
- The switch will not forward disallowed tagged traffic across a trunk port.

The final example above is a negative test to verify that switches with VLAN trunking will forward only traffic explicitly permitted by the switch configurations.

This test used the standard test bed (see Figure 1, above). In this example, all interswitch communication is done via VLAN trunks. The trunk ports on each switch will allow tagged traffic with VLAN IDs from 301 through 303, and untagged traffic from ports with VLAN ID of 300. A fifth VLAN, with an ID of 304, is also defined by the trunk ports are configured not to allow that traffic.

HP A-series commands

First, define VLANs 300 to 304.

```
<HP5800> system-view
[HP5800] vlan 300 to 304
```

Then, define a VLAN trunk port that allows tagged traffic from VLANs 301-303, and native untagged traffic on VLAN 300.

```
[HP5800] interface GigabitEthernet1/0/23
[HP5800-gigabitethernet1/0/23]  port link-mode bridge
[HP5800-gigabitethernet1/0/23]  port link-type trunk
[HP5800-gigabitethernet1/0/23]  undo port trunk permit vlan 1
[HP5800-gigabitethernet1/0/23]  port trunk permit vlan 300 to 303
[HP5800-gigabitethernet1/0/23]  port trunk pvid vlan 300
[HP5800-gigabitethernet1/0/23]  quit
```

Next, define access-mode interfaces allowing untagged traffic for VLANs 300-304.

```
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1]  port link-mode bridge
[HP5800-Gigabitethernet1/0/1]  port access vlan 300
[HP5800-Gigabitethernet1/0/1] interface GigabitEthernet1/0/2
[HP5800-Gigabitethernet1/0/2]  port link-mode bridge
[HP5800-Gigabitethernet1/0/2]  port access vlan 301
[HP5800-Gigabitethernet1/0/2] interface GigabitEthernet1/0/3
[HP5800-Gigabitethernet1/0/3]  port link-mode bridge
[HP5800-Gigabitethernet1/0/3]  port access vlan 302
[HP5800-Gigabitethernet1/0/3] interface GigabitEthernet1/0/4
[HP5800-Gigabitethernet1/0/4]  port link-mode bridge
[HP5800-Gigabitethernet1/0/4]  port access vlan 303
[HP5800-Gigabitethernet1/0/4] interface GigabitEthernet1/0/5
[HP5800-Gigabitethernet1/0/5]  port link-mode bridge
[HP5800-Gigabitethernet1/0/5]  port access vlan 304
[HP5800-Gigabitethernet1/0/5]  quit
[HP5800]  quit
```

HP E-series commands

HP E-series switches combine trunk creation, access ports, and VLAN assignment together into a single VLAN construct. A port that is a member of a single VLAN carrying only untagged traffic is an access port (ports A1-A5 in this example). A port that is a member of multiple VLANs that carries both tagged and untagged traffic is a VLAN trunk port (ports A9-A10 in this example). Here we define VLANs 300-304 and assign ports to them.

```
HP5406ZL# configure
HP5406ZL(config)# vlan 300
HP5406ZL(vlan-300)#    name "VLAN300"
HP5406ZL(vlan-300)#    untagged A1,A9-A10
HP5406ZL(vlan-300)#    ip address 10.1.2.1 255.255.0.0
HP5406ZL(vlan-300)#    exit
```

```
HP5406ZL(config)# vlan 301
HP5406ZL(vlan-301)#     name "VLAN301"
HP5406ZL(vlan-301)#     untagged A2
HP5406ZL(vlan-301)#     ip address 10.2.2.1 255.255.0.0
HP5406ZL(vlan-301)#     tagged A9-A10
HP5406ZL(vlan-301)#     exit
HP5406ZL(config)# vlan 302
HP5406ZL(vlan-302)#     name "VLAN302"
HP5406ZL(vlan-302)#     untagged A3
HP5406ZL(vlan-302)#     ip address 10.3.2.1 255.255.0.0
HP5406ZL(vlan-302)#     tagged A9-A10
HP5406ZL(vlan-302)#     exit
HP5406ZL(config)# vlan 303
HP5406ZL(vlan-303)#     name "VLAN303"
HP5406ZL(vlan-303)#     untagged A4
HP5406ZL(vlan-303)#     ip address 10.4.2.1 255.255.0.0
HP5406ZL(vlan-303)#     tagged A9-A10
HP5406ZL(vlan-303)#     exit
HP5406ZL(config)# vlan 304
HP5406ZL(vlan-304)#     name "VLAN304"
HP5406ZL(vlan-304)#     untagged A5
HP5406ZL(vlan-304)#     ip address 10.5.2.1 255.255.0.0
HP5406ZL(vlan-304)#     exit
HP5406ZL(config)#     exit
```

Cisco commands

The following commands apply to a Cisco Catalyst 6509. The syntax is similar for the Catalyst 3750-E switches and Cisco Catalyst 4506 switches.

First, define VLANs 300 to 304.

```
Cat6509# configure terminal
Cat6509(config)# vlan 300-304
```

Then, define a VLAN trunk port that allows tagged traffic from VLANs 301-303, and native untagged traffic on 300.

```
Cat6509(config)# interface GigabitEthernet4/9
Cat6509(config-if)#   switchport
Cat6509(config-if)#   switchport trunk encapsulation dot1q
Cat6509(config-if)#   switchport trunk native vlan 300
Cat6509(config-if)#   switchport trunk allowed vlan 300-303
Cat6509(config-if)#   switchport mode trunk
Cat6509(config-if)#   exit
```

Next, define access-mode interfaces allowing untagged traffic from VLANs 300-304.

```
Cat6509(config)# interface GigabitEthernet6/0/1
Cat6509(config-if)#   switchport access vlan 300
Cat6509(config-if)#   switchport mode access
Cat6509(config-if)# interface GigabitEthernet6/0/2
```

```
Cat6509(config-if)#  switchport access vlan 301
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)# interface GigabitEthernet6/0/3
Cat6509(config-if)#  switchport access vlan 302
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)# interface GigabitEthernet6/0/4
Cat6509(config-if)#  switchport access vlan 303
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)# interface GigabitEthernet6/0/5
Cat6509(config-if)#  switchport access vlan 304
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  end
```

Validation

The Spirent TestCenter traffic generator/analyzer can be configured to offer traffic between pairs of access-mode interfaces on each switch. In all cases – involving unicast, multicast, or broadcast traffic – traffic will stay local to the VLAN in which it is defined. For example, traffic offered to VLAN 300 on the HP switches will be forwarded only to interfaces in VLAN 300 on the Cisco switches and vice-versa.

If desired, port mirroring can be enabled on either HP or Cisco switches to verify that the trunk ports carry tagged traffic VLAN IDs 301-303 and untagged traffic for VLAN ID 300. As a final verification that VLANs limit broadcast domains, Spirent TestCenter can be configured to offer traffic on access ports with VLAN 304. The trunk ports on all switches will not forward this traffic.

Jumbo frame switching

Objective

To validate the ability of HP Networking and Cisco Catalyst switches to correctly forward bidirectional traffic consisting of jumbo frames.

Background

For many years the IEEE Ethernet specification has defined the maximum length of an Ethernet frame to be 1,518 bytes (or 1,522 bytes with an 802.1Q VLAN tag). The use of jumbo frames – those larger than 1518 bytes – remains nonstandard. However, jumbo frames can improve the performance of applications involving bulk data transfer, such as backup and disaster recovery.

HP and Cisco switches both support 9,216-byte jumbo frames, including Ethernet CRC. This section explains how to configure both vendors' switches to exchange jumbo frames.

Topology

In this example, the Spirent TestCenter traffic generator offers 9,216-byte jumbo Ethernet frames using a "partially meshed" topology, meaning all traffic offered to ports on HP switches are destined to ports on Cisco switches and visa-versa. VLAN trunk ports connect the switches and VLAN access ports at the edge accept untagged jumbo frames. However, the ability to switch jumbo frames does not depend on VLAN tagging. This example would also work with all interfaces passing untagged traffic.

Figure 2 below illustrates the configuration used to validate jumbo frame switching. This test deviates from the standard test bed by the removal of the link aggregation trunks between the Cisco Catalyst 4506 and the Cisco Catalyst 3750-E as well as the link aggregation trunk between the Cisco Catalyst 4506 and the Cisco Catalyst 6509. There is also no connection between the Cisco Catalyst 4506 and the Cisco Catalyst 6509. As noted in the configuration sections below, all interfaces explicitly support switching of jumbo frames. Engineers configured all interswitch trunks to use VLAN trunking, in this case carrying traffic from VLAN 300.

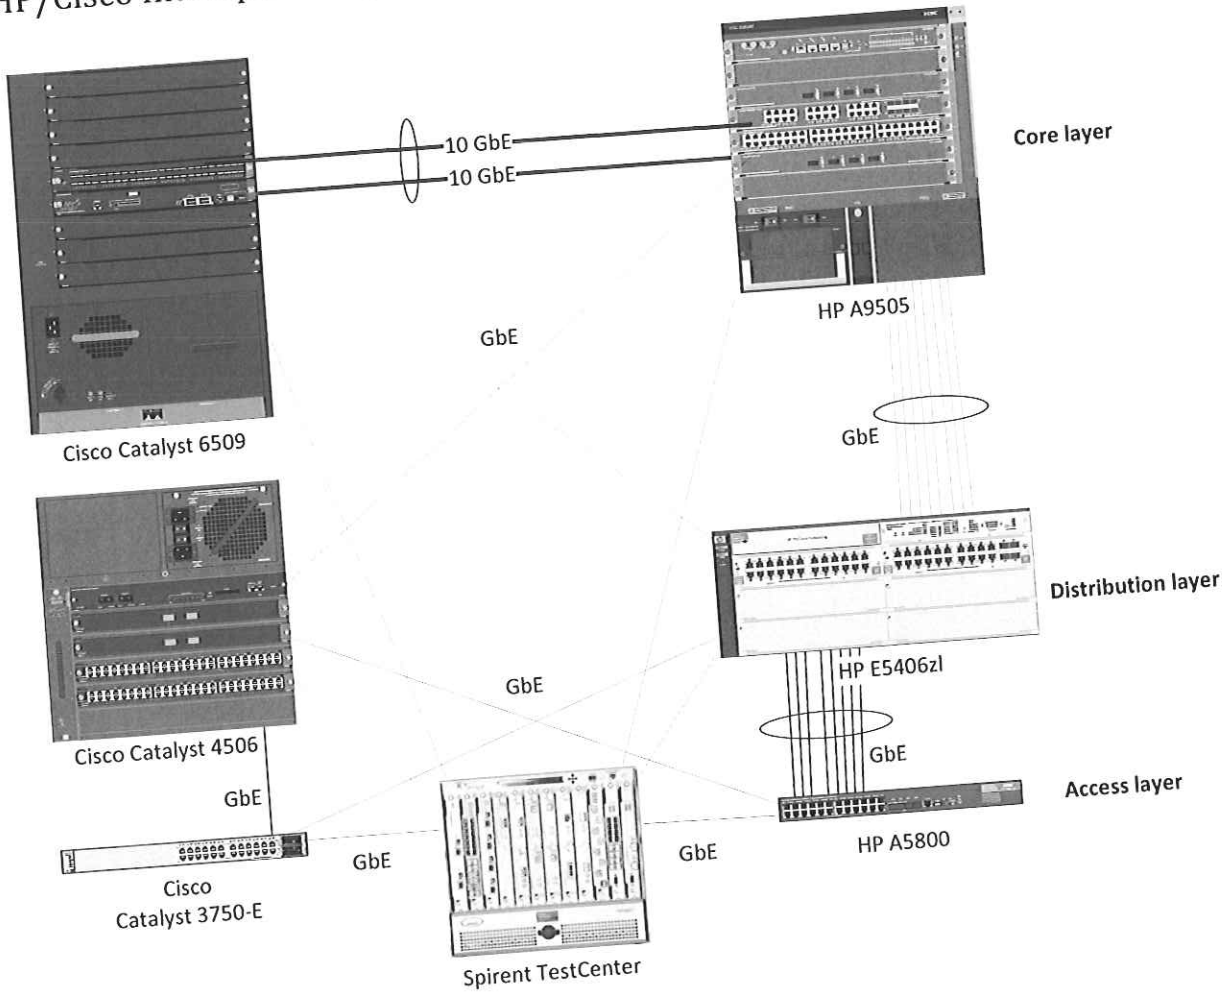

Figure 2: Jumbo frame switching test bed

HP A-series commands

HP A-series switches have jumbo frames enabled by default. The following commands are used to explicitly set the maximum transmission unit (MTU). The MTU is set in the interface configuration context.

```
<HP5800> system-view
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1] port link-mode bridge
[HP5800-Gigabitethernet1/0/1] jumboframe enable 9216
[HP5800-Gigabitethernet1/0/1] quit
[HP5800] quit
```

HP E-series commands

HP E-series switches set the MTU on a per-VLAN basis. When enabled, all ports on that VLAN will forward jumbo frames.

```
HP5406ZL# configure
HP5406ZL(config)# vlan 300
HP5406ZL(vlan-300)#    name "VLAN306"
HP5406ZL(vlan-300)#    untagged A1-A5,A9-A10,Trk1-Trk2
HP5406ZL(vlan-300)#    ip address 10.1.2.1 255.255.0.0
HP5406ZL(vlan-300)#    jumbo
HP5406ZL(vlan-300)#      exit
HP5406ZL(config)#      exit
```

Cisco commands

On Cisco Catalyst 6509 and Cisco Catalyst 4506 switches, jumbo frame support varies by line card. For line cards that support jumbo frames, MTU is set on a per-interface basis.

```
Cat6509# configure terminal
Cat6509(config)# interface GigabitEthernet4/48
Cat6509(config-if)# switchport
Cat6509(config-if)# switchport access vlan 300
Cat6509(config-if)# switchport mode access
Cat6509(config-if)# mtu 9216
Cat6509(config-if)# end
```

On Cisco Catalyst 3750-E switches, MTU is set systemwide:

```
Cat3750E# configure terminal
Cat3750E(config)# system mtu jumbo 9216
Cat3750E(config)# end
```

Validation

Generating jumbo frames between the attached clients and servers will validate the ability of the switches to exchange jumbo traffic. All switches will forward all jumbo frames with zero frame loss.

Jumbo frame routing

Objective

To validate the ability of HP Networking and Cisco Catalyst switches to correctly route bidirectional traffic consisting of jumbo frames.

Background

Some routing protocols, such as open shortest path first (OSPF), require that both routers use the same MTU before exchanging routing information. For Ethernet interfaces, the requirement for matched MTUs applies equally to jumbo frames (those larger than 1518 bytes) as to standard-length frames.

HP Networking and Cisco Catalyst switches both support 9,216-byte jumbo frames, including Ethernet CRC. This section explains how to configure both vendors' devices to set up on an OSPF routing session using jumbo frames.

Topology

In this example, the HP A9505, HP E5406zl, and HP A5800 switches are configured as OSPF routers exchanging jumbo frames with Cisco Catalyst 6509, Cisco Catalyst 4506, and Cisco Catalyst 3750-E switches.

Figure 3 below illustrates the configuration used to validate jumbo frame routing. This test deviates from the standard test bed by the removal of the link aggregation trunks between the Cisco Catalyst 4506 and the Cisco Catalyst 3750-E, and between the Cisco Catalyst 4506 and the Cisco Catalyst 6509. There is also no connection between the Cisco Catalyst 4506 and the Cisco Catalyst 6509. In addition, all devices routed traffic at layer 3 in this test. In this example, OSPF routing sessions are established between all connected devices.

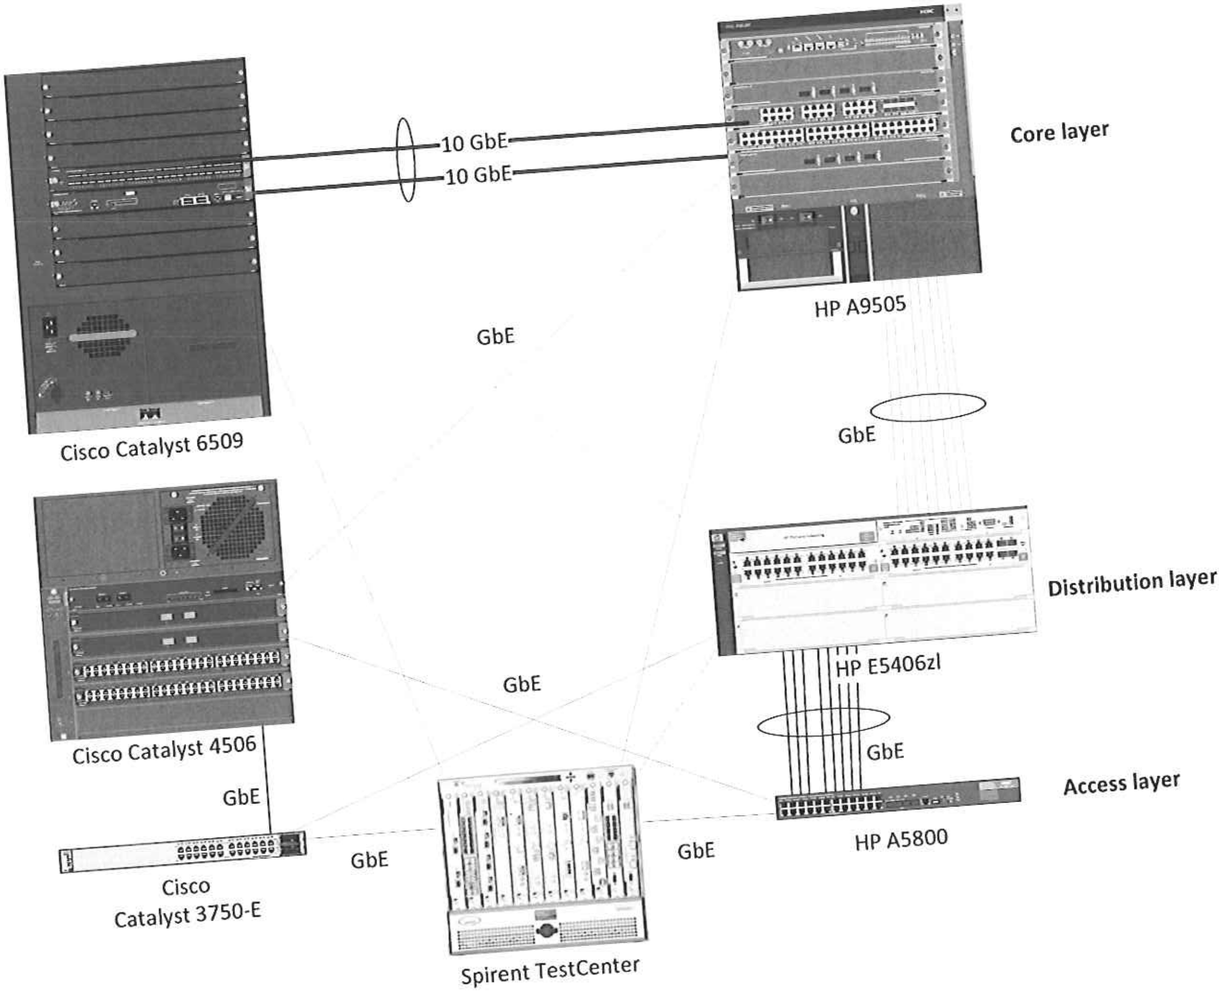

Figure 3: Jumbo frame routing test bed

HP A-series commands

HP A-series switches have jumbo frames enabled by default. The following commands are used to set the jumbo frame MTU. The frame size is set in the interface configuration context.

```
<HP5800> system-view
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1] port link-mode bridge
[HP5800-Gigabitethernet1/0/1] jumboframe enable 9216
[HP5800-Gigabitethernet1/0/1] quit
[HP5800] quit
```

Then OSPF is configured. In this example, the interface is a member of OSPF area 0.

```
[HP5800] ospf 1 router-id 10.0.0.1
[HP5800-OSPF] area 0.0.0.0
[HP5800-OSPF] network 10.0.0.0 0.0.255.255
[HP5800-OSPF] quit
```

HP E-series commands

HP E-series switches set MTU on a per-VLAN basis. When enabled, all ports on that VLAN will forward jumbo frames.

```
HP5406ZL# configure
HP5406ZL(config)# vlan 300
HP5406ZL(vlan-300)#    name "VLAN306"
HP5406ZL(vlan-300)#    untagged A1-A5,A9-A10,Trk1-Trk2
HP5406ZL(vlan-300)#    ip address 10.1.2.1 255.255.0.0
HP5406ZL(vlan-300)#    jumbo
HP5406ZL(vlan-300)#      exit
```

Then set up OSPF routing. In our configuration, the VLAN interfaces were used as the routable interfaces. The `area backbone` command designates OSPF area 0.

```
HP5406(config)# ip routing
HP5406(config)# ip router-id 10.0.32.1
HP5406(config)# router ospf
HP5406(ospf)#    area backbone range 10.0.0.0 255.255.0.0 type summary
HP5406(ospf)#    exit
HP5406(config)# vlan 33
HP5406(vlan-33)#    ip ospf 10.0.33.1 area backbone
HP5406(vlan-33)#    exit
HP5406(config)# vlan 34
HP5406(vlan-34)#    ip ospf 10.0.34.1 area backbone
HP5406(vlan-34)#    exit
HP5406(config)# vlan 35
HP5406(vlan-35)#    ip ospf 10.0.35.1 area backbone
HP5406(vlan-35)#    exit
HP5406(config)# vlan 36
HP5406(vlan-36)#    ip ospf 10.0.36.1 area backbone
HP5406(vlan-36)#    exit
HP5406(config)# vlan 37
HP5406(vlan-37)#    ip ospf 10.0.37.1 area backbone
HP5406(vlan-37)#
HP5406(config)# vlan 38
HP5406(vlan-38)#    ip ospf 10.0.38.1 area backbone
HP5406(vlan-38)#    exit
HP5406(config)# vlan 39
HP5406(vlan-39)#    ip ospf 10.0.39.1 area backbone
HP5406(vlan-39)#    exit
HP5406(config)# vlan 40
HP5406(vlan-40)#    ip ospf 10.0.40.1 area backbone
HP5406(vlan-40)#    exit
HP5406(config)#    exit
```

Cisco commands

On Cisco Catalyst 6509 and Cisco Catalyst 4506 switches, jumbo frame support varies by line card. For those line cards that support jumbo frames, MTU is set on a per-interface basis. Cisco

IOS has separate commands for **`mtu,`** describing the maximum transmission unit for the *Ethernet* frame and for the **`ip mtu,`** describing the MTU for the *IP packet*.

Configure the interface with a jumbo frame size.

```
Cat6509# configure terminal
Cat6509(config)# interface GigabitEthernet4/9
Cat6509(config-if)# ip address 10.0.42.2 255.255.255.0
Cat6509(config-if)# ip mtu 9198
Cat6509(config-if)# exit
```

Then set up OSPF.

```
Cat6509(config)# router ospf 1
Cat6509(config-router)#  log-adjacency-changes
Cat6509(config-router)#  network 10.0.0.0 0.0.255.255 area 0
Cat6509(config-router)#  exit
```

Then set up the VLAN for jumbo frames. This is required to route jumbo frames between VLANs. All interfaces in the VLAN must be set to allow jumbo frames before this command will take effect.

```
Cat6509(config)# interface Vlan193
Cat6509(config-if)# mtu 9216
Cat6509(config-if)#  end
```

On Cisco Catalyst 3750-E switches, MTU is set systemwide:

```
Cat3750E# configure terminal
Cat3750E(config)# system mtu jumbo 9216
Cat3750E(config)# system mtu routing 9198
Cat3750E(config)# router ospf 1
Cat3750E(config-router)#  log-adjacency-changes
Cat3750E(config-router)#  network 10.0.43.2 0.0.0.0 area 0
Cat3750E(config-router)#  network 10.0.75.2 0.0.0.0 area 0
Cat3750E(config-router)#  network 10.0.128.0 0.0.127.255 area 0
Cat3750E(config-router)#  network 192.168.1.0 0.0.0.255 area 0
Cat3750E(config-router)#  network 192.168.2.0 0.0.0.255 area 0
Cat3750E(config-router)#  end
```

Validation

Unless both HP and Cisco interfaces agree on MTU size, OSPF routing adjacencies will remain in ExStart state, and will never transition to OSPF "full" state. To verify that an OSPF adjacency has entered OSPF "full" state on the HP A-series switches, use the **`display ospf peer`** command. To verify that an OSPF adjacency has entered OSPF "full" state on HP E-series and Cisco switches, use the **`show ip ospf neighbor`** command.

The fact that both routers are in Full state indicates they have agreed to exchange IP packets up to 9,198 bytes long (or 9,216 bytes, including Ethernet header and CRC). OSPF routing session establishment will not work unless both sides agree on MTU size.

Link aggregation

Objective

To validate the ability of HP Networking and Cisco Catalyst switches to correctly forward traffic over a logical connection created using IEEE 802.3ad link aggregation.

Background

The IEEE 802.3ad link specification defines a standards-based method for aggregating multiple physical Ethernet links into a single logical link. The logical link, known as a link aggregation group (LAG), is comprised of multiple *members* (pairs of physical interfaces on each switch). LAGs may be defined statically or dynamically, the latter using the link aggregation control protocol (LACP). With LACP enabled, 802.3ad-compliant switches can dynamically add or remove up to eight LAG members.

Link aggregation is useful both for increasing bandwidth beyond the limits of single physical interfaces and, especially when used with LACP, for adding redundancy to network connections.

Topology

In this example, an HP A9505 switch uses two-member LAGs to exchange traffic with a Cisco Catalyst 6509 switch and a Cisco Catalyst 4506 switch. An HP E5406zl switch uses two-member LAGs to exchange traffic with a Cisco Catalyst 6509 switch and a Cisco Catalyst 3750-E switch. An HP A5800 switch uses two-member LAGs to exchange traffic with a Cisco Catalyst 4506 switch.

Figure 4 below shows the topology used to validate link aggregation and LACP functionality. This test deviates from the standard test bed with the additional of several link aggregation groups between the HP E5406zl and the Cisco Catalyst 3750-E, between the HP 5406zl and the Cisco Catalyst 6509, between the HP A9505 and the Cisco Catalyst 4506, and between the HP 5800 and the Cisco C5406. Other connections have been removed between the HP 5406zl and the HP 5800, between the HP 5406zl and HP 9505, between the Cisco Catalyst 3750-E and the Cisco Catalyst 4506, and between the Cisco Catalyst 4506 and the Cisco Catalyst 6509.

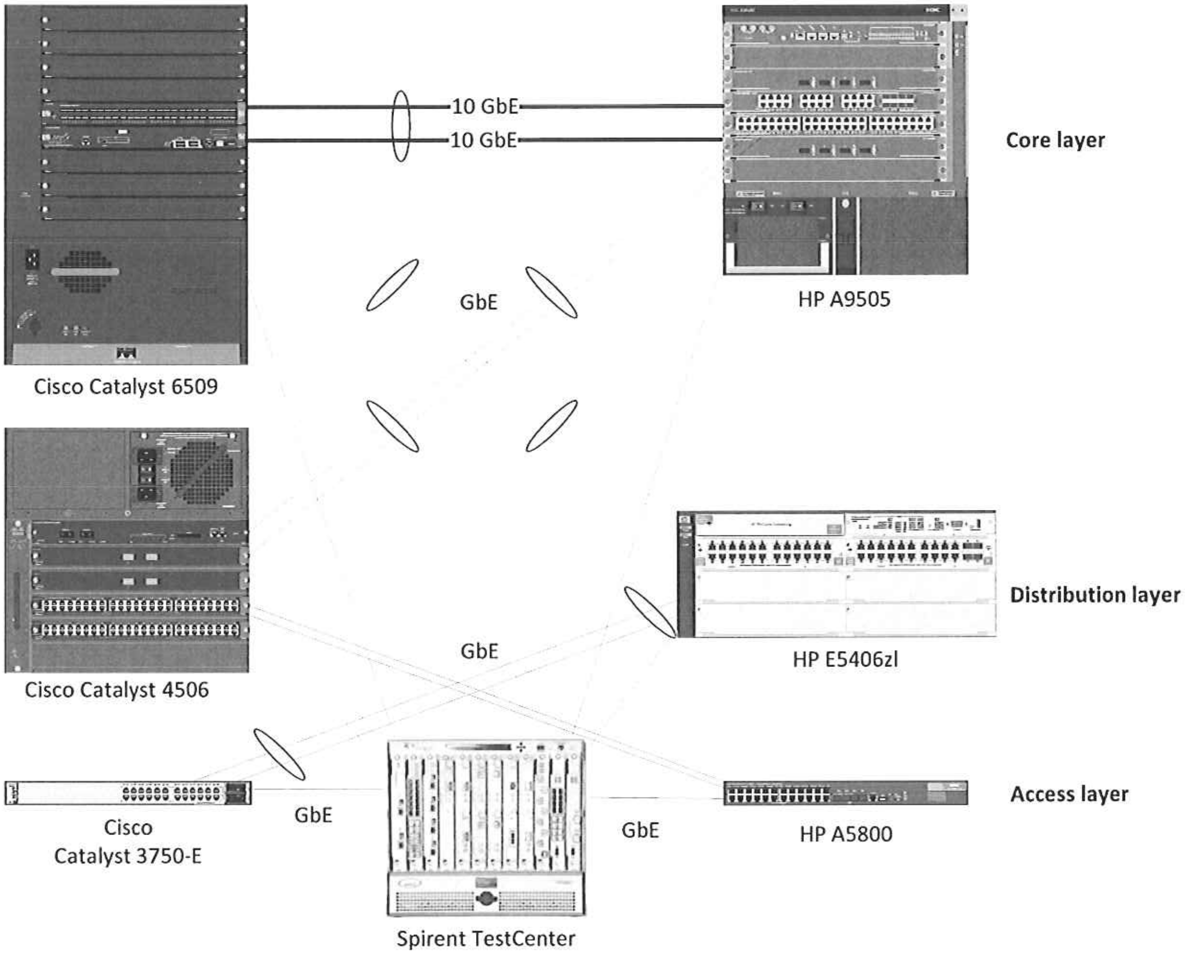

Figure 4: Link aggregation test bed

HP A-series commands

On these HP switches, link aggregation is a two-step process. First a virtual *bridge aggregation* interface is created. Then physical interfaces are associated with the virtual bridge interface. While this example involves a VLAN trunk, a common use of link aggregation, it is not a requirement.

Create the bridge aggregation interface.

```
<HP5800> system-view
[HP5800] interface Bridge-Aggregation31
[HP5800-bridge-aggregation31]  description linkagg_to_4506
[HP5800-bridge-aggregation31]  port link-type trunk
[HP5800-bridge-aggregation31]  undo port trunk permit vlan 1
[HP5800-bridge-aggregation31]  port trunk permit vlan 300 to 303
[HP5800-bridge-aggregation31]  port trunk pvid vlan 300
```

```
[HP5800-bridge-aggregation31]  link-aggregation mode dynamic
[HP5800-bridge-aggregation31]
          link-aggregation load-sharing mode destination-mac source-mac
[HP5800-bridge-aggregation31] quit
```

Then assign physical interfaces to the bridge aggregation virtual interface.

```
[HP5800] interface GigabitEthernet1/0/17
[HP5800-gigabitethernet1/0/17]  port link-mode bridge
[HP5800-gigabitethernet1/0/17]  port link-type trunk
[HP5800-gigabitethernet1/0/17]  undo port trunk permit vlan 1
[HP5800-gigabitethernet1/0/17]  port trunk permit vlan 300 to 303
[HP5800-gigabitethernet1/0/17]  port trunk pvid vlan 300
[HP5800-gigabitethernet1/0/17]  port link-aggregation group 31
[HP5800-gigabitethernet1/0/17] interface GigabitEthernet1/0/18
[HP5800-gigabitethernet1/0/18]  port link-mode bridge
[HP5800-gigabitethernet1/0/18]  port link-type trunk
[HP5800-gigabitethernet1/0/18]  undo port trunk permit vlan 1
[HP5800-gigabitethernet1/0/18]  port trunk permit vlan 300 to 303
[HP5800-gigabitethernet1/0/18]  port trunk pvid vlan 300
[HP5800-gigabitethernet1/0/18]  port link-aggregation group 31
[HP5800-gigabitethernet1/0/18]  quit
[HP5800]  quit
```

HP E-series commands

HP E-series switches create *trunks* to support LACP. A single command creates the trunk and assigns physical members to the trunk.

```
HP5406ZL# configure
HP5406ZL(config)# trunk A9,A12 Trk31 LACP
HP5406ZL(config)# exit
```

Cisco commands

Cisco Catalyst switches, like HP A-series switches, perform a two-step process to create a *Port Channel*. The following commands apply to a Cisco Catalyst 6509. The syntax is similar for the Catalyst 3750-E switches and Cisco Catalyst 4506 switches.

First, create the link aggregation group. Here we also create a VLAN trunk.

```
Cat6509# configure terminal
Cat6509(config)# interface Port-channel1
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport trunk encapsulation dot1q
Cat6509(config-if)#  switchport trunk native vlan 300
Cat6509(config-if)#  switchport trunk allowed vlan 300-303
Cat6509(config-if)#  switchport mode trunk
```

```
Cat6509(config-if)#  exit
```

Next, add interfaces to the link aggregation group. The command **`channel-group 1`** adds an interface to the link aggregation group created in the previous step, while **`mode active`** enables LACP.

```
Cat6509(config-)# interface GigabitEthernet4/1
Cat6509(config-if)# no ip address
Cat6509(config-if)# switchport
Cat6509(config-if)# switchport trunk encapsulation dot1q
Cat6509(config-if)# switchport trunk native vlan 300
Cat6509(config-if)# switchport trunk allowed vlan 300-303
Cat6509(config-if)# switchport mode trunk
Cat6509(config-if)#`channel-group 1 mode active
Cat6509(config)# interface GigabitEthernet4/3
Cat6509(config-if)# no ip address
Cat6509(config-if)# switchport
Cat6509(config-if)# switchport trunk encapsulation dot1q
Cat6509(config-if)# switchport trunk native vlan 300
Cat6509(config-if)# switchport trunk allowed vlan 300-303
Cat6509(config-if)# switchport mode trunk
Cat6509(config-if)# channel-group 1 mode active
Cat6509(config-if)# end
```

Validation

The command **`display link-aggregation summary`** on HP A-series switches will show the status of the bridge aggregation interfaces. On HP E-series switches, the **`show lacp`** command will verify correct operation.

The correct operation of a LAG can with two or more members also can be verified by offering traffic at a rate higher than any single LAG member can carry. If the switch forwards all traffic across the LAG without loss, the LAG is operating properly.

Spanning tree case 1: RSTP/Rapid-PVST+

Objective

To verify interoperability of a rapid spanning tree topology between HP Networking and Cisco Catalyst switches using RSTP and Rapid-PVST+.

To measure convergence time of a rapid spanning tree topology between HP and Cisco after a link failure.

Background

The spanning tree protocol is widely used in campus enterprise networks for loop prevention and redundancy. Rapid spanning tree, defined in IEEE 802.1w, provides much faster convergence time after a link or device failure than the original 802.1D spanning tree specification

Topology

This example uses redundant paths between HP Networking and Cisco Catalyst switches. The default spanning tree mode in Cisco Catalyst switches is that vendor's proprietary per-VLAN spanning tree plus (PVST+) mode, which is interoperable with other vendors' rapid spanning tree implementations.

Figure 1 above shows the RSTP validation test bed. All ports on all switches are access-mode members of the default VLAN. Rapid spanning is enabled on all the HP switches. Cisco's "Rapid-PVST+" is enabled on all the Cisco switches and is interoperable with standard rapid spanning tree. Traffic offered from the Spirent TestCenter generator/analysis verifies the spanning tree topology.

HP A-series commands

Assign all members to be access-mode members of the default VLAN. Here is the command for interface GigabitEthernet1/0/6; the same command would apply to all interfaces participating in the spanning tree topology.

```
<HP5800> system view
[HP5800] interface GigabitEthernet1/0/6
[HP5800-Gigabitethernet1/0/6] port link-mode bridge
[HP5800-Gigabitethernet1/0/6] quit
```

Then enable rapid spanning tree on the HP A-series switches.

```
<HP5800> system-view
[HP5800] stp mode rstp
[HP5800] stp enable
[HP5800] quit
```

HP E-series commands

On HP E-series switches, by default all members are access-mode members of the default VLAN and therefore no per-interface command needs to be done.

Enable rapid spanning tree on the HP E-series switches.

```
HP5406ZL# configure
HP5406ZL(config)# spanning-tree
HP5406ZL(config)# spanning-tree priority 9 force-version rstp-operation
HP5406ZL(config)# exit
```

Cisco commands

First, assign all members to be access-mode members of the default VLAN. The following commands apply to a Cisco Catalyst 6509. The syntax is similar for Cisco Catalyst 3750-E and Cisco Catalyst 4506 switches.

Here is the command for interface GigabitEthernet6/0/3; the same command would apply to all interfaces participating in the spanning tree topology.

```
Cat6509# configure terminal
Cat6509(config)# interface GigabitEthernet6/0/3
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  exit
```

Then enable rapid-pvst mode on the Cisco switches.

```
Cat6509(config)# spanning-tree mode rapid-pvst
Cat6509(config)# end
```

Validation

HP A-series switches can use the command `display stp brief` to verify the state of rapid spanning tree.

```
<A9505>display stp br
 MSTID      Port                        Role  STP State     Protection
   0        Bridge-Aggregation10        ROOT  FORWARDING    NONE
   0        Bridge-Aggregation20        DESI  FORWARDING    NONE
   0        GigabitEthernet3/0/11       ALTE  DISCARDING    NONE
   0        GigabitEthernet3/0/16       DESI  FORWARDING    NONE
```

HP E-series switches uses the command **`show spanning-tree`** to display spanning-tree state.

```
E5406zl# show spanning-tree

 Multiple Spanning Tree (MST) Information

  STP Enabled   : Yes
  Force Version : RSTP-operation
...

                         |           Prio            | Designated    Hello
  Port   Type            | Cost      rity State      | Bridge        Time PtP Edge
  ------ --------- + --------- ---- ---------- + ------------- ---- --- ----
  A1     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A2     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A3     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A4     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A5     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A6     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A7     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A8     100/1000T | 20000     128  Forwarding | 001560-f56200 2    Yes Yes
  A9     100/1000T | 20000     128  Blocking   | 001de6-eb7001 2    Yes No
  A10    100/1000T | 20000     128  Blocking   | 002155-740000 2    Yes No
...

  E2     10GbE-CX4 | 2000      128  Forwarding | 001560-f56200 2    Yes Yes
  F1     10GbE-CX4 | 2000      128  Forwarding | 001560-f56200 2    Yes Yes
  F2     10GbE-CX4 | 2000      128  Forwarding | 001560-f56200 2    Yes Yes
  Trk1             | 20000     64   Blocking   | 002389-11d000 2    Yes No
  Trk2             | 20000     64   Forwarding | 000fe2-f3e292 2    Yes No
```

To verify all switches send traffic only over the spanning tree interfaces in forwarding state, generate a known quantity of frames from Spirent TestCenter or other source and compare switch interface packet counters with those sent and received on each interface. Interfaces in blocking state will receive spanning tree BPDU frames but should transmit no frames.

To determine convergence time, disable one of the spanning tree interfaces in forwarding state while offering a known quantity of frames from Spirent TestCenter or other traffic generator. Convergence time can be derived from frame loss. For example, if Spirent TestCenter generates traffic at a rate of 1,000 frames per second, each dropped frame is equivalent to 1 millisecond of convergence time. If the switches drop 47 frames, then rapid spanning tree convergence time is 47 ms.

Spanning tree case 2: MSTP/PVST+

Objective

To verify interoperability of multiple spanning tree protocol (MSTP) and per-vlan spanning tree protocol plus (PVST+) between HP Networking and Cisco Catalyst switches, respectively.

To measure convergence time of an MSTP-PVST+ topology between HP Networking and Cisco Catalyst switches after a link failure.

Background

As defined in IEEE specification 802.1s, the multiple spanning tree protocol (MSTP) adds loop prevention and redundancy on a per-VLAN basis. With MSTP, individual spanning tree topologies can be configured for each VLAN.

The goal of this exercise is to demonstrate interoperability in a multiple-VLAN environment when the HP Networking and Cisco Catalyst switches use different variations of spanning tree: MSTP on HP and PVST+ on a Cisco Catalyst switch.

Topology

This example uses redundant paths between the HP Networking and Cisco Catalyst switches. VLAN Ids of 300 to 304 have been defined on all switches. MSTP is enabled on all the HP switches, and Rapid PVST+ is enabled on all the Cisco switches.

Figure 1 above illustrates the MSTP-PVST+ validation test bed. The links interconnecting each switch are trunk ports that allow tagged traffic from VLAN IDs 300 to 304. Access ports were configured on the access layer switches, with one port assigned to each VLAN. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies the spanning tree topology in each VLAN.

HP A-series commands

Create VLAN IDs 300 to 304.

```
<HP5800> system-view
[HP5800] vlan 300 to 304
```

Configure access-mode ports for their respective VLANs.

```
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1]  port link-mode bridge
[HP5800-Gigabitethernet1/0/1]  port access vlan 300
[HP5800-Gigabitethernet1/0/1] interface GigabitEthernet1/0/2
[HP5800-Gigabitethernet1/0/2]  port link-mode bridge
[HP5800-Gigabitethernet1/0/2]  port access vlan 301
[HP5800-Gigabitethernet1/0/2] interface GigabitEthernet1/0/3
[HP5800-Gigabitethernet1/0/3]  port link-mode bridge
[HP5800-Gigabitethernet1/0/3]  port access vlan 302
[HP5800-Gigabitethernet1/0/3] interface GigabitEthernet1/0/4
[HP5800-Gigabitethernet1/0/4]  port link-mode bridge
[HP5800-Gigabitethernet1/0/4]  port access vlan 303
[HP5800-Gigabitethernet1/0/4] interface GigabitEthernet1/0/5
[HP5800-Gigabitethernet1/0/5]  port link-mode bridge
[HP5800-Gigabitethernet1/0/5]  port access vlan 304
```

Configure interswitch connections as trunk ports. Here is interface GigabitEthernet1/0/23 as an example.

```
[HP5800] interface GigabitEthernet1/0/23
[HP5800-Gigabitethernet1/0/23]  port link-mode bridge
[HP5800-Gigabitethernet1/0/23]  port link-type trunk
[HP5800-Gigabitethernet1/0/23]  undo port trunk permit vlan 1
[HP5800-Gigabitethernet1/0/23]  port trunk permit vlan 300 to 303
[HP5800-Gigabitethernet1/0/23]  port trunk pvid vlan 300
[HP5800-Gigabitethernet1/0/23]  quit
```

Enable multiple spanning tree. This requires enabling MSTP (the default on HP A-series switches) and configuring one multiple spanning tree instance per VLAN.

```
[HP5800] stp enable
[HP5800] stp region-configuration
[HP5800-mst-region]  instance 2 vlan 300
[HP5800-mst-region]  instance 3 vlan 301
[HP5800-mst-region]  instance 4 vlan 302
[HP5800-mst-region]  instance 5 vlan 303
[HP5800-mst-region]  instance 6 vlan 304
[HP5800-mst-region]  active region-configuration
[HP5800-mst-region]  quit
[HP5800]  quit
```

HP E-series commands

Create the VLANs and assign physical interfaces to them. Interfaces that have both untagged (access ports) and tagged VLAN IDs are VLAN trunks.

```
HP5406ZL# configure
HP5406ZL(config)# vlan 300
HP5406ZL(vlan-300)#    name "VLAN300"
HP5406ZL(vlan-300)#    untagged A1,A9-A10
HP5406ZL(vlan-300)#    ip address 10.1.2.1 255.255.0.0
HP5406ZL(vlan-300)#    exit
```

```
HP5406ZL(config)# vlan 301
HP5406ZL(vlan-300)#    name "VLAN301"
HP5406ZL(vlan-300)#    untagged A2
HP5406ZL(vlan-300)#    ip address 10.2.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(config)#    exit
HP5406ZL(vlan-300)# vlan 302
HP5406ZL(vlan-300)#    name "VLAN302"
HP5406ZL(vlan-300)#    untagged A3
HP5406ZL(vlan-300)#    ip address 10.3.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 303
HP5406ZL(vlan-300)#    name "VLAN303"
HP5406ZL(vlan-300)#    untagged A4
HP5406ZL(vlan-300)#    ip address 10.4.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 304
HP5406ZL(vlan-300)#    name "VLAN304"
HP5406ZL(vlan-300)#    untagged A5
HP5406ZL(vlan-300)#    ip address 10.5.2.1 255.255.0.0
HP5406ZL(vlan-300)#    exit
```

Create the MSTP instances, and assign one VLAN to each instance.

```
HP5406ZL(config)# spanning-tree
HP5406ZL(config)# spanning-tree instance 2 vlan 300
HP5406ZL(config)# spanning-tree instance 3 vlan 301
HP5406ZL(config)# spanning-tree instance 4 vlan 302
HP5406ZL(config)# spanning-tree instance 5 vlan 303
HP5406ZL(config)# spanning-tree instance 6 vlan 304
HP5406ZL(config)# spanning-tree priority 9
HP5406ZL(config)# exit
```

Cisco commands

The following commands apply to a Cisco Catalyst 3750-E switch. The syntax is similar for the Cisco Catalyst 6509 and Cisco Catalyst 4506 switches.

Create the VLANs.

```
Cat3750E# configure terminal
Cat3750E(config)# vlan 300-304
```

Configure access-mode ports for their respective VLANs.

```
Cat3750E(config)# interface GigabitEthernet6/0/1
Cat3750E(config)#  switchport
Cat3750E(config-if)#  switchport access vlan 300
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
```

```
Cat3750E(config-if)# interface GigabitEthernet6/0/2
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 301
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/3
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 302
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/4
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 303
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/5
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 304
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)#  exit
```

Configure interswitch connections as trunk ports. Here is interface GigabitEthernet6/0/10 as an example.

```
Cat3750E(config)# interface GigabitEthernet6/0/10
Cat3750E(config-if)#  switchport trunk encapsulation dot1q
Cat3750E(config-if)#  switchport trunk native vlan 300
Cat3750E(config-if)#  switchport trunk allowed vlan 300-303
Cat3750E(config-if)#  switchport mode trunk
Cat3750E(config-if)#  exit
```

Enable PVST+. On a new switch, PVST+ already is enabled by default.

```
Cat3750E(config)# spanning-tree mode pvst
Cat3750E(config)# end
```

Validation

HP A-series switches can use the command **`display stp brief`** to verify the state of spanning tree on HP A-series switches.

```
<A9505>display stp brief
 MSTID      Port                         Role  STP State     Protection
   0        Bridge-Aggregation10         DESI  FORWARDING    NONE
   0        Bridge-Aggregation20         ROOT  FORWARDING    NONE
   0        GigabitEthernet3/0/11        DESI  FORWARDING    NONE
   0        GigabitEthernet3/0/16        DESI  FORWARDING    NONE
   2        Bridge-Aggregation10         DESI  FORWARDING    NONE
   2        Bridge-Aggregation20         MAST  FORWARDING    NONE
   2        GigabitEthernet3/0/11        DESI  FORWARDING    NONE
   3        Bridge-Aggregation10         DESI  FORWARDING    NONE
```

```
3         Bridge-Aggregation20        MAST  FORWARDING    NONE
3         GigabitEthernet3/0/11       DESI  FORWARDING    NONE
4         Bridge-Aggregation10        DESI  FORWARDING    NONE
4         Bridge-Aggregation20        MAST  FORWARDING    NONE
4         GigabitEthernet3/0/11       DESI  FORWARDING    NONE
5         Bridge-Aggregation10        DESI  FORWARDING    NONE
5         Bridge-Aggregation20        MAST  FORWARDING    NONE
5         GigabitEthernet3/0/11       DESI  FORWARDING    NONE
```

HP E-Series uses the command **`show spanning-tree`** to display the state of spanning tree on HP E-Series switches.

```
E5406zl# show spanning-tree

 Multiple Spanning Tree (MST) Information

  STP Enabled   : Yes
  Force Version : MSTP-operation
...

                     |           Prio             | Designated    Hello
  Port   Type        | Cost      rity State       | Bridge        Time PtP Edge
  ------ --------- + --------- ---- ---------- + ------------- ---- --- ----
  A1     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A2     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A3     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A4     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A5     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A6     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A7     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A8     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A9     100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
  A10    100/1000T | 20000     128  Forwarding | 001560-f56200 2     Yes Yes
...
  E2     10GbE-CX4 | 2000      128  Forwarding | 001560-f56200 2     Yes Yes
  F1     10GbE-CX4 | 2000      128  Forwarding | 001560-f56200 2     Yes Yes
  F2     10GbE-CX4 | 2000      128  Forwarding | 001560-f56200 2     Yes Yes
  Trk1             | 20000     64   Forwarding | 001560-f56200 2     Yes No
  Trk2             | 20000     64   Forwarding | 000fe2-f3e292 2     Yes No
```

To verify all switches send traffic only over the spanning tree interfaces in forwarding state, generate a known quantity of frames from Spirent TestCenter or other source to each VLAN and compare switch interface packet counters with those sent and received on each interface. Interfaces in blocking state will receive spanning tree BPDU frames but should transmit no frames.

To determine convergence time, disable one of the spanning tree interfaces in forwarding state while offering a known quantity of frames from Spirent TestCenter or other traffic generator. Convergence time can be derived from frame loss. For example, if Spirent TestCenter generates traffic at a rate of 1,000 frames per second, each dropped frame is equivalent to 1 millisecond of convergence time. If the switches drop 47 frames, then spanning tree convergence time is 47 ms.

Spanning tree case 3: MSTP/Rapid-PVST+

Objective

To verify interoperability of multiple spanning tree protocol (MSTP) and rapid per-VLAN spanning tree protocol plus (Rapid PVST+) between HP Networking and Cisco Catalyst switches, respectively.

To measure convergence time of an MSTP-Rapid PVST+ topology between HP Networking and Cisco Catalyst switches after a link failure.

Background

As defined in IEEE specification 802.1s, the multiple spanning tree protocol (MSTP) adds loop prevention and redundancy on a per-VLAN basis. With MSTP, individual spanning tree topologies can be configured for each VLAN.

The goal of this exercise is to demonstrate interoperability in a multiple-VLAN environment when the HP Networking and Cisco Catalyst switches use different variations of spanning tree: MSTP on HP switches and Rapid PVST+ on Cisco Catalyst switches.

Topology

This example uses redundant paths between the HP Networking and Cisco Catalyst switches. VLAN Ids of 300 to 304 have been defined on all switches. MSTP is enabled on all HP switches, with Rapid PVST+ defined on all Cisco switches.

Figure 1 above illustrates the MSTP-Rapid PVST+ validation test bed. The links interconnecting each switch are trunk ports that allow tagged traffic from the VLAN IDs 300 to 304. Access ports were configured on the access layer switches, with one port per vlan being configured. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies the spanning tree topology in each VLAN.

HP A-series commands

Create VLAN IDs 300 to 304.

```
<HP5800> system-view
[HP5800] vlan 300 to 304
```

Configure access-mode ports for their respective VLANs.

```
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1]  port link-mode bridge
[HP5800-Gigabitethernet1/0/1]  port access vlan 300
[HP5800-Gigabitethernet1/0/1] interface GigabitEthernet1/0/2
[HP5800-Gigabitethernet1/0/2]  port link-mode bridge
[HP5800-Gigabitethernet1/0/2]  port access vlan 301
[HP5800-Gigabitethernet1/0/2] interface GigabitEthernet1/0/3
[HP5800-Gigabitethernet1/0/3]  port link-mode bridge
[HP5800-Gigabitethernet1/0/3]  port access vlan 302
[HP5800-Gigabitethernet1/0/3] interface GigabitEthernet1/0/4
[HP5800-Gigabitethernet1/0/4]  port link-mode bridge
[HP5800-Gigabitethernet1/0/4]  port access vlan 303
[HP5800-Gigabitethernet1/0/4] interface GigabitEthernet1/0/5
[HP5800-Gigabitethernet1/0/5]  port link-mode bridge
[HP5800-Gigabitethernet1/0/5]  port access vlan 304
[HP5800-Gigabitethernet1/0/5]  quit
```

Configure interswitch connections as trunk ports. Here is interface GigabitEthernet1/0/23 as an example.

```
[HP5800] interface GigabitEthernet1/0/23
[HP5800-Gigabitethernet1/0/23]  port link-mode bridge
[HP5800-Gigabitethernet1/0/23]  port link-type trunk
[HP5800-Gigabitethernet1/0/23]  undo port trunk permit vlan 1
[HP5800-Gigabitethernet1/0/23]  port trunk permit vlan 300 to 303
[HP5800-Gigabitethernet1/0/23]  port trunk pvid vlan 300
[HP5800-Gigabitethernet1/0/23]  quit
```

Enable multiple spanning tree. This requires enable MSTP (the default on HP switches) and configuring one multiple spanning tree instance per VLAN.

```
[HP5800] stp enable
[HP5800] stp region-configuration
[HP5800-mst-region]  instance 2 vlan 300
[HP5800-mst-region]  instance 3 vlan 301
[HP5800-mst-region]  instance 4 vlan 302
[HP5800-mst-region]  instance 5 vlan 303
[HP5800-mst-region]  instance 6 vlan 304
[HP5800-mst-region]  active region-configuration
[HP5800-mst-region]  quit
[HP5800]  quit
```

HP E-series commands

Create the VLANs and assign physical interfaces to them. Interfaces that have both untagged (access ports) and tagged VLAN IDs are VLAN trunks.

```
HP5406ZL# configure
HP5406ZL(config)# vlan 300
HP5406ZL(vlan-300)#    name "VLAN300"
HP5406ZL(vlan-300)#    untagged A1,A9-A10
HP5406ZL(vlan-300)#    ip address 10.1.2.1 255.255.0.0
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 301
HP5406ZL(vlan-300)#    name "VLAN301"
HP5406ZL(vlan-300)#    untagged A2
HP5406ZL(vlan-300)#    ip address 10.2.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(config)#    exit
HP5406ZL(vlan-300)# vlan 302
HP5406ZL(vlan-300)#    name "VLAN302"
HP5406ZL(vlan-300)#    untagged A3
HP5406ZL(vlan-300)#    ip address 10.3.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 303
HP5406ZL(vlan-300)#    name "VLAN303"
HP5406ZL(vlan-300)#    untagged A4
HP5406ZL(vlan-300)#    ip address 10.4.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 304
HP5406ZL(vlan-300)#    name "VLAN304"
HP5406ZL(vlan-300)#    untagged A5
HP5406ZL(vlan-300)#    ip address 10.5.2.1 255.255.0.0
HP5406ZL(vlan-300)#    exit
```

Create the multiple spanning tree protocol instances, and assign one VLAN to each instance.

```
HP5406ZL(config)# spanning-tree
HP5406ZL(config)# spanning-tree instance 2 vlan 300
HP5406ZL(config)# spanning-tree instance 3 vlan 301
HP5406ZL(config)# spanning-tree instance 4 vlan 302
HP5406ZL(config)# spanning-tree instance 5 vlan 303
HP5406ZL(config)# spanning-tree instance 6 vlan 304
HP5406ZL(config)# spanning-tree priority 9tree instance 2 vlan 300
HP5406ZL(config)# spanning-tree instance 3 vlan 301
HP5406ZL(config)# spanning-tree instance 4 vlan 302
HP5406ZL(config)# spanning-tree instance 5 vlan 303
HP5406ZL(config)# spanning-tree instance 6 vlan 304
HP5406ZL(config)# spanning-tree priority 9
```

Cisco commands

The following commands apply to a Cisco Catalyst 3750-E switch. The syntax is similar for the Catalyst 6509 switches and Cisco Catalyst 4506 switches.

Create the VLANs.

```
Cat3750E# configure terminal
Cat3750E(config)# vlan 300-304
```

Configure access-mode ports for the respective VLANs.

```
Cat3750E(config)# interface GigabitEthernet6/0/1
Cat3750E(config)#  switchport
Cat3750E(config-if)#  switchport access vlan 300
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/2
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 301
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/3
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 302
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/4
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 303
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/5
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 304
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)#  exit
```

Configure interswitch connections as trunk ports. Here is interface GigabitEthernet6/0/10.

```
Cat3750E(config)# interface GigabitEthernet6/0/10
Cat3750E(config-if)#  switchport trunk encapsulation dot1q
Cat3750E(config-if)#  switchport trunk native vlan 300
Cat3750E(config-if)#  switchport trunk allowed vlan 300-303
Cat3750E(config-if)#  switchport mode trunk
Cat3750E(config-if)#  exit
```

Enable Rapid PVST+.

```
Cat3750E(config)# spanning-tree mode rapid-pvst
Cat3750E(config)# end
```

Validation

HP A-series switches can use the command **`display stp brief`** to verify the state of spanning tree on HP A-series switches. HP E-series uses the command **`show spanning-tree`** to display the state of spanning tree on HP E-series switches.

To verify switches send traffic only over the spanning tree interfaces in forwarding state, generate a known quantity of frames from Spirent TestCenter to each VLAN and compare switch interface packet counters with those sent and received on each interface. Interfaces in blocking state will receive spanning tree BPDU frames but should transmit no frames.

To determine convergence time, disable one of the spanning tree interfaces in forwarding state while offering a known quantity of frames from Spirent TestCenter or other traffic generator. Convergence time can be derived from frame loss. For example, if Spirent TestCenter generates traffic at a rate of 1,000 frames per second, each dropped frame is equivalent to 1 millisecond of convergence time. If the switches drop 47 frames, then spanning tree convergence time is 47 ms.

Spanning tree case 4: MSTP/MSTP

Objective

To verify interoperability of a multiple spanning tree topology between HP Networking and Cisco Catalyst switches.

To measure convergence time of a multiple spanning tree topology between HP and Cisco switches after a link failure.

Background

As defined in IEEE specification 802.1s, the multiple spanning tree protocol (MSTP) adds loop prevention and redundancy on a per-VLAN basis. With MSTP, a separate spanning tree topology can be configured for each VLAN.

MSTP is the default spanning tree protocol for HP Networking switches. MSTP is enabled by default on HP A-series switches, and disabled by default on HP E-series switches.

Topology

This example uses redundant paths between the HP Networking and Cisco Catalyst switches. VLAN IDs of 300 to 304 have been defined on all switches, and MSTP is enabled on all switches.

Figure 1 above illustrates the MSTP validation test bed. The links interconnecting each switch are trunk ports that allow tagged traffic from the VLAN IDs 300 to 304. Access ports were configured on the access layer switches, with one port per VLAN being configured. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies the spanning tree topology in each VLAN.

HP A-series commands

Create VLAN IDs 300 to 304.

```
<HP5800> system-view
[HP5800] vlan 300 to 304
```

Configure access-mode ports for their respective VLANs.

```
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1]  port link-mode bridge
[HP5800-Gigabitethernet1/0/1]  port access vlan 300
[HP5800-Gigabitethernet1/0/1] interface GigabitEthernet1/0/2
[HP5800-Gigabitethernet1/0/2]  port link-mode bridge
[HP5800-Gigabitethernet1/0/2]  port access vlan 301
[HP5800-Gigabitethernet1/0/2] interface GigabitEthernet1/0/3
[HP5800-Gigabitethernet1/0/3]  port link-mode bridge
[HP5800-Gigabitethernet1/0/3]  port access vlan 302
[HP5800-Gigabitethernet1/0/3] interface GigabitEthernet1/0/4
[HP5800-Gigabitethernet1/0/4]  port link-mode bridge
[HP5800-Gigabitethernet1/0/4]  port access vlan 303
[HP5800-Gigabitethernet1/0/4] interface GigabitEthernet1/0/5
[HP5800-Gigabitethernet1/0/5]  port link-mode bridge
[HP5800-Gigabitethernet1/0/5]  port access vlan 304
[HP5800-Gigabitethernet1/0/5]  quit
```

Configure interswitch connections as trunk ports. Here is interface GigabitEthernet1/0/23 as an example.

```
[HP5800] interface GigabitEthernet1/0/23
[HP5800-Gigabitethernet1/0/23]  port link-mode bridge
[HP5800-Gigabitethernet1/0/23]  port link-type trunk
[HP5800-Gigabitethernet1/0/23]  undo port trunk permit vlan 1
[HP5800-Gigabitethernet1/0/23]  port trunk permit vlan 300 to 303
[HP5800-Gigabitethernet1/0/23]  port trunk pvid vlan 300
[HP5800-Gigabitethernet1/0/23]  quit
```

Enable multiple spanning tree. This requires enabling MSTP (the default on HP A-series switches) and configuring one multiple spanning tree instance per VLAN.

```
[HP5800] stp enable
[HP5800] stp region-configuration
[HP5800-mst-region]  instance 2 vlan 300
[HP5800-mst-region]  instance 3 vlan 301
[HP5800-mst-region]  instance 4 vlan 302
[HP5800-mst-region]  instance 5 vlan 303
[HP5800-mst-region]  instance 6 vlan 304
[HP5800-mst-region]  active region-configuration
[HP5800-mst-region]  quit
[HP5800]  quit
```

HP E-series commands

Create the VLANs and assign physical interfaces to them. Interfaces that have both untagged (access ports) and tagged VLAN IDs are VLAN trunks.

```
HP5406ZL# configure
HP5406ZL(config)# vlan 300
HP5406ZL(vlan-300)#    name "VLAN300"
HP5406ZL(vlan-300)#    untagged A1,A9-A10
HP5406ZL(vlan-300)#    ip address 10.1.2.1 255.255.0.0
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 301
HP5406ZL(vlan-300)#    name "VLAN301"
HP5406ZL(vlan-300)#    untagged A2
HP5406ZL(vlan-300)#    ip address 10.2.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(config)#    exit
HP5406ZL(vlan-300)# vlan 302
HP5406ZL(vlan-300)#    name "VLAN302"
HP5406ZL(vlan-300)#    untagged A3
HP5406ZL(vlan-300)#    ip address 10.3.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 303
HP5406ZL(vlan-300)#    name "VLAN303"
HP5406ZL(vlan-300)#    untagged A4
HP5406ZL(vlan-300)#    ip address 10.4.2.1 255.255.0.0
HP5406ZL(vlan-300)#    tagged A9-A10
HP5406ZL(vlan-300)#    exit
HP5406ZL(config)# vlan 304
HP5406ZL(vlan-300)#    name "VLAN304"
HP5406ZL(vlan-300)#    untagged A5
HP5406ZL(vlan-300)#    ip address 10.5.2.1 255.255.0.0
HP5406ZL(vlan-300)#    exit
```

Create the multiple spanning tree protocol instances, and assign one VLAN to each instance.

```
HP5406ZL(config)# spanning-tree
HP5406ZL(config)# spanning-tree instance 2 vlan 300
HP5406ZL(config)# spanning-tree instance 3 vlan 301
HP5406ZL(config)# spanning-tree instance 4 vlan 302
HP5406ZL(config)# spanning-tree instance 5 vlan 303
HP5406ZL(config)# spanning-tree instance 6 vlan 304
HP5406ZL(config)# spanning-tree priority 9
HP5406ZL(config)# exit
```

Cisco commands

The following commands apply to a Cisco Catalyst 3750-E. The syntax is similar for the Catalyst 6509 switches and Cisco Catalyst 4506 switches.

Create the VLANs.

```
Cat3750E# configure terminal
Cat3750E(config)# vlan 300-304
```

Configure access-mode ports for their respective VLANs.

```
Cat3750E(config)# interface GigabitEthernet6/0/1
Cat3750E(config)#  switchport
Cat3750E(config-if)#  switchport access vlan 300
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/2
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 301
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/3
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 302
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/4
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 303
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)# interface GigabitEthernet6/0/5
Cat3750E(config-if)#  switchport
Cat3750E(config-if)#  switchport access vlan 304
Cat3750E(config-if)#  switchport mode access
Cat3750E(config-if)#  spanning-tree portfast
Cat3750E(config-if)#  exit
```

Configure interswitch connections as trunk ports. Here is interface GigabitEthernet6/0/10 as an example.

```
Cat3750E(config)# interface GigabitEthernet6/0/10
Cat3750E(config-if)#  switchport trunk encapsulation dot1q
Cat3750E(config-if)#  switchport trunk native vlan 300
Cat3750E(config-if)#  switchport trunk allowed vlan 300-303
Cat3750E(config-if)#  switchport mode trunk
Cat3750E(config-if)#  exit
```

Enable multiple spanning tree. This requires enable MSTP and configuring one multiple spanning tree instance per VLAN.

```
Cat3750E(config)# spanning-tree mode mst
Cat3750E(config)# spanning-tree mst configuration
Cat3750E(config-mst)#  instance 1 vlan 300
Cat3750E(config-mst)#  instance 2 vlan 301
Cat3750E(config-mst)#  instance 3 vlan 302
Cat3750E(config-mst)#  instance 4 vlan 303
Cat3750E(config-mst)#  instance 5 vlan 304
Cat3750E(config-mst)#  end
```

Validation

HP A-series switches can use the command **`display stp brief`** to verify the state of spanning tree on HP A-series switches. HP E-series uses the command **`show spanning-tree`** to display the state of spanning tree on HP E-series switches.

To verify all switches send traffic only over the spanning tree interfaces in forwarding state, generate a known quantity of frames from Spirent TestCenter or other source to each VLAN and compare switch interface packet counters with those sent and received on each interface. Interfaces in blocking state will receive spanning tree BPDU frames but should transmit no frames.

To determine convergence time, disable one of the spanning tree interfaces in forwarding state while offering a known quantity of frames from Spirent TestCenter or other traffic generator. Convergence time can be derived from frame loss. For example, if Spirent TestCenter generates traffic at a rate of 1,000 frames per second, each dropped frame is equivalent to 1 millisecond of convergence time. If the switches drop 47 frames, then spanning tree convergence time is 47 ms.

OSPFv2 (OSPF for IPv4)

Objective

To verify that HP Networking and Cisco Catalyst switches are able to establish open shortest path first version 2 (OSPFv2) connections and exchange topology information.

Background

Intended for use on IPv4 networks, OSPFv2 supports IP subnetting and redistribution of routing information learned via other protocols. OSPF also allows session authentication and uses IP multicast for distribution of routing updates.

OSPF uses areas to segment traffic, with area 0 designated as the backbone network. OSPF typically involves coordination among multiple internal routers; area border routers (ABRs) connected to multiple areas; and autonomous system boundary routers (ASBRs).

In addition to standard areas, OSPFv2 also defines two special types of areas: Stubs are areas into which information on external routes is not sent. Instead, the area border router (ABR) generates a default external route into the stub area. A Not-So-Stubby-Area (NSSA) is like a stub area, but it can import external routes into the area for redistribution via OSPF.

Topology

This example uses multiple paths between HP and Cisco devices. Each HP and Cisco switch was configured with multiple networks, which were then advertised by OSPF to its neighbors.

Figure 1 above illustrates the OSPFv2 test bed. Access switches were connected to both distribution switches and the distribution switches were connected to both core switches. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies that traffic is indeed being passed between networks.

HP A-series commands

In this example, switched virtual interfaces (SVIs) are created using VLAN interfaces. Physical interfaces are then mapped to the VLAN interfaces. Routing is done between VLAN interfaces on each switch.

Create the VLANs.

```
<HP5800> system-view
[HP5800] vlan 1
[HP5800] vlan 2 to 9
```

Create the switched virtual interfaces.

```
[HP5800] interface Vlan-interface2
[HP5800-vlan-interface2] ip address 10.0.2.1 255.255.255.0
[HP5800-vlan-interface2] interface Vlan-interface3
[HP5800-vlan-interface3] ip address 10.0.3.1 255.255.255.0
[HP5800-vlan-interface3] interface Vlan-interface4
[HP5800-vlan-interface4] ip address 10.0.4.1 255.255.255.0
[HP5800-vlan-interface4] interface Vlan-interface5
[HP5800-vlan-interface5] ip address 10.0.5.1 255.255.255.0
[HP5800-vlan-interface5] interface Vlan-interface6
[HP5800-vlan-interface6] ip address 10.0.6.1 255.255.255.0
[HP5800-vlan-interface6] interface Vlan-interface7
[HP5800-vlan-interface7] ip address 10.0.7.1 255.255.255.0
[HP5800-vlan-interface7] interface Vlan-interface8
[HP5800-vlan-interface8] ip address 10.0.8.1 255.255.255.0
[HP5800-vlan-interface8] interface Vlan-interface9
[HP5800-vlan-interface9] ip address 10.0.9.1 255.255.255.0
[HP5800-vlan-interface9] quit
```

Associate the physical interfaces with the corresponding SVIs.

```
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1]  port link-mode bridge
[HP5800-Gigabitethernet1/0/1]  port access vlan 2
[HP5800-Gigabitethernet1/0/1] interface GigabitEthernet1/0/2
[HP5800-Gigabitethernet1/0/2]  port link-mode bridge
[HP5800-Gigabitethernet1/0/2]  port access vlan 3
[HP5800-Gigabitethernet1/0/2] interface GigabitEthernet1/0/3
[HP5800-Gigabitethernet1/0/3]  port link-mode bridge
[HP5800-Gigabitethernet1/0/3]  port access vlan 4
[HP5800-Gigabitethernet1/0/3] interface GigabitEthernet1/0/4
[HP5800-Gigabitethernet1/0/4]  port link-mode bridge
[HP5800-Gigabitethernet1/0/4]  port access vlan 5
[HP5800-Gigabitethernet1/0/4] interface GigabitEthernet1/0/5
[HP5800-Gigabitethernet1/0/5]  port link-mode bridge
[HP5800-Gigabitethernet1/0/5]  port access vlan 6
[HP5800-Gigabitethernet1/0/5] interface GigabitEthernet1/0/6
[HP5800-Gigabitethernet1/0/6]  port link-mode bridge
[HP5800-Gigabitethernet1/0/6]  port access vlan 7
[HP5800-Gigabitethernet1/0/6] interface GigabitEthernet1/0/7
[HP5800-Gigabitethernet1/0/7]  port link-mode bridge
[HP5800-Gigabitethernet1/0/7]  port access vlan 8
[HP5800-Gigabitethernet1/0/7] interface GigabitEthernet1/0/8
[HP5800-Gigabitethernet1/0/8]  port link-mode bridge
[HP5800-Gigabitethernet1/0/8]  port access vlan 9
[HP5800-Gigabitethernet1/0/8]  quit
```

Configure OSPF routing.

```
[HP5800] ospf 1 router-id 10.0.0.1
[HP5800-ospf]  area 0.0.0.0
[HP5800-ospf]   network 10.0.0.0 0.0.255.255
[HP5800-ospf]   quit
[HP5800]    quit
```

HP E-series commands

For the HP E-series switches, a single command sets up VLANs and assigns physical interfaces to those VLANs.

Enable IP routing .

```
HP5406ZL> configure
HP5406ZL(config)> ip routing
```

Define and configure VLANs.

```
HP5406ZL(config)> vlan 33
HP5406ZL(vlan-33)>    name "VLAN33"
HP5406ZL(vlan-33)>    untagged A1
HP5406ZL(vlan-33)>    ip address 10.0.33.1 255.255.255.0
HP5406ZL(vlan-33)>    exit
HP5406ZL(vlan-33)> vlan 34
HP5406ZL(vlan-34)>    name "VLAN34"
HP5406ZL(vlan-34)>    untagged A2
HP5406ZL(vlan-34)>    ip address 10.0.34.1 255.255.255.0
HP5406ZL(vlan-34)>    exit
HP5406ZL(vlan-34)> vlan 35
HP5406ZL(vlan-35)>    name "VLAN35"
HP5406ZL(vlan-35)>    untagged A3
HP5406ZL(vlan-35)>    ip address 10.0.35.1 255.255.255.0
HP5406ZL(vlan-35)>    exit
HP5406ZL(vlan-35)> vlan 36
HP5406ZL(vlan-36)>    name "VLAN36"
HP5406ZL(vlan-36)>    untagged A4
HP5406ZL(vlan-36)>    ip address 10.0.36.1 255.255.255.0
HP5406ZL(vlan-36)>    exit
HP5406ZL(vlan-36)> vlan 37
HP5406ZL(vlan-37)>    name "VLAN37"
HP5406ZL(vlan-37)>    untagged A5
HP5406ZL(vlan-37)>    ip address 10.0.37.1 255.255.255.0
HP5406ZL(vlan-37)>    exit
HP5406ZL(vlan-37)> vlan 38
HP5406ZL(vlan-38)>    name "VLAN38"
HP5406ZL(vlan-38)>    untagged A6
HP5406ZL(vlan-38)>    ip address 10.0.38.1 255.255.255.0
HP5406ZL(vlan-38)>    exit
HP5406ZL(vlan-38)> vlan 39
HP5406ZL(vlan-39)>    name "VLAN39"
HP5406ZL(vlan-39)>    untagged A7
HP5406ZL(vlan-39)>    ip address 10.0.39.1 255.255.255.0
```

```
HP5406ZL(vlan-39)>    exit
HP5406ZL(vlan-39)> vlan 40
HP5406ZL(vlan-40)>    name "VLAN40"
HP5406ZL(vlan-40)>    untagged A8
HP5406ZL(vlan-40)>    ip address 10.0.40.1 255.255.255.0
HP5406ZL(vlan-40)>    exit
```

Enable OSPF routing and configure the VLANs for OSPF.

```
HP5406ZL(config)> ip router-id 10.0.32.1
HP5406ZL(config)> router ospf
HP5406ZL(ospf)>    area backbone range 10.0.0.0 255.255.0.0 type summary
HP5406ZL(ospf)>    exit
HP5406ZL(config)> vlan 33
HP5406ZL(vlan-33)>    ip ospf 10.0.33.1 area backbone
HP5406ZL(vlan-33)>    exit
HP5406ZL(vlan-33)> vlan 34
HP5406ZL(vlan-34)>    ip ospf 10.0.34.1 area backbone
HP5406ZL(vlan-34)>    exit
HP5406ZL(vlan-34)> vlan 35
HP5406ZL(vlan-35)>    ip ospf 10.0.35.1 area backbone
HP5406ZL(vlan-35)>    exit
HP5406ZL(vlan-35)> vlan 36
HP5406ZL(vlan-36)>    ip ospf 10.0.36.1 area backbone
HP5406ZL(vlan-36)>    exit
HP5406ZL(vlan-36)> vlan 37
HP5406ZL(vlan-37)>    ip ospf 10.0.37.1 area backbone
HP5406ZL(vlan-37)>    exit
HP5406ZL(vlan-37)> vlan 38
HP5406ZL(vlan-38)>    ip ospf 10.0.38.1 area backbone
HP5406ZL(vlan-38)>    exit
HP5406ZL(vlan-38)> vlan 39
HP5406ZL(vlan-39)>    ip ospf 10.0.39.1 area backbone
HP5406ZL(vlan-39)>    exit
HP5406ZL(vlan-39)> vlan 40
HP5406ZL(vlan-40)>    ip ospf 10.0.40.1 area backbone
HP5406ZL(vlan-40)>    exit
HP5406ZL>    exit
```

Cisco commands

On Cisco Catalyst switches, like HP A-series switches, create the VLANs first and then assign physical interfaces to the VLAN interfaces. The following commands apply to a Cisco Catalyst 6509. The syntax is similar for Cisco Catalyst 3750-E switches and Cisco Catalyst 4506 switches.

First, enable routing.

```
Cat6509# configure terminal
Cat6509(config)# ip routing
```

Then, create the VLAN interfaces.

```
Cat6509(config)# interface Vlan129
Cat6509(config-if)#  ip address 10.0.129.1 255.255.255.0
Cat6509(config-if)# interface Vlan130
Cat6509(config-if)#  ip address 10.0.130.1 255.255.255.0
Cat6509(config-if)# interface Vlan131
Cat6509(config-if)#  ip address 10.0.131.1 255.255.255.0
Cat6509(config-if)# interface Vlan132
Cat6509(config-if)#  ip address 10.0.132.1 255.255.255.0
Cat6509(config-if)# interface Vlan133
Cat6509(config-if)#  ip address 10.0.133.1 255.255.255.0
Cat6509(config-if)# interface Vlan134
Cat6509(config-if)#  ip address 10.0.134.1 255.255.255.0
Cat6509(config-if)# interface Vlan135
Cat6509(config-if)#  ip address 10.0.135.1 255.255.255.0
Cat6509(config-if)# interface Vlan136
Cat6509(config-if)#  ip address 10.0.136.1 255.255.255.0
Cat6509(config-if)#  exit
```

Next, assign physical interfaces to the VLANs.

```
Cat6509(config)# interface GigabitEthernet1/0/1
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 129
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)# interface GigabitEthernet1/0/2
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 130
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)# interface GigabitEthernet1/0/3
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 131
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)# interface GigabitEthernet1/0/4
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 132
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)# interface GigabitEthernet1/0/5
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 133
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)# interface GigabitEthernet1/0/6
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 134
Cat6509(config-if)#  switchport mode access
```

```
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)# interface GigabitEthernet1/0/7
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 135
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)# interface GigabitEthernet1/0/8
Cat6509(config-if)#  no ip address
Cat6509(config-if)#  switchport
Cat6509(config-if)#  switchport access vlan 136
Cat6509(config-if)#  switchport mode access
Cat6509(config-if)#  spanning-tree portfast
Cat6509(config-if)#  exit
```

Finally, enable OSPF routing.

```
Cat6509(config)# router ospf 1
Cat6509(config-ospf)#  log-adjacency-changes
Cat6509(config-ospf)#  network 10.0.0.0 0.0.255.255 area 0
Cat6509(config-ospf)#  end
```

Validation

If the HP and Cisco devices are unable to complete OSPF negotiation, routing adjacencies will remain in the ExStart state. Fully functional adjacencies will be in the Full state. To verify that an OSPF adjacency has entered OSPF Full state on the HP switches, use the **`display ospf peer`** command on A-series switches and the **`show ip ospf neighbor`** command on E-series switches.

OSPFv3 (OSPF for IPv6)

Objective

To verify that HP Networking and Cisco Catalyst switches are able to establish open shortest path first version 3 (OSPFv3) connections and exchange topology information.

Background

OSPFv3 updates the routing protocol for use on IPv6 networks. In a mixed IPv4/IPv6 environment, OSPFv2 must be used in conjunction with OSPFv3.

While the basic mechanics of OSPF are identical in both versions, OSPFv3 introduces new link-state advertisement (LSA) types; removes addressing semantics from OSPF headers; generalizes flooding; and removes OSPF-layer authentication, among other changes. RFC 5340 describes OSPFv3.

Topology

This example uses multiple paths between HP and Cisco devices. Each HP and Cisco switch was configured with multiple networks, which were then advertised by OSPF to its neighbors. The switches have been configured in dual-stack IPv4/IPv6 mode, running both OSPFv2 (to support IPv4 traffic) and OSPFv3 (to support IPv6 traffic).

Figure 1 above illustrates the OSPv3 test bed. Both access switches were connected to both distribution switches, and both distribution switches in turn were connected to both core switches. Traffic offered from the Spirent TestCenter traffic generator/analyzer verifies that traffic is correctly routed between networks.

HP A-series commands

In this example, switched virtual interfaces (SVIs) are created using VLAN interfaces. Physical interfaces are then mapped to the VLAN interfaces. Routing is done between VLAN interfaces on each switch. Unlike OSPFv2, OSPv3 configuration is done on the actual routable interface.

While not required, the configuration is using a dual-stack IPv4/IPv6 setup.

First, enable IPv6.

```
<HP9505> system-view
[HP9505] ipv6
```

Then, configure the VLAN interfaces.

```
[HP9505] vlan 65
[HP9505] vlan 75
[HP9505] interface Vlan-interface65
[HP9505-vlan-interface65]  ipv6 address 2002:9505:0:65::1/64
[HP9505-vlan-interface65]  ospfv3 1 area 0.0.0.0
[HP9505-vlan-interface65]  ip address 10.0.65.1 255.255.255.0
[HP9505-vlan-interface65] interface Vlan-interface75
[HP9505-vlan-interface75]  ipv6 address 2002:9595:4506:75::1/64
[HP9505-vlan-interface75]  ospfv3 1 area 0.0.0.0
[HP9505-vlan-interface75]  ip address 10.0.75.1 255.255.255.0
[HP9505-vlan-interface75]  quit
```

Next, assign the physical interfaces to the VLANs.

```
[HP9505] interface GigabitEthernet3/0/9
[HP9505-Gigabitethernet3/0/9]  port access vlan 65
[HP9505-Gigabitethernet3/0/9] interface GigabitEthernet3/0/11
[HP9505-Gigabitethernet3/0/11]  port access vlan 75
[HP9505-Gigabitethernet3/0/11]  quit
```

Finally, configure the OSPFv3 and OSPFv2 routing processes.

```
[HP9505] ospf 1 router-id 10.0.64.1
[HP9505-ospf]  area 0.0.0.0
[HP9505-ospf]   network 10.0.0.0 0.0.255.255
[HP9505-ospf] ospfv3 1
[HP9505-ospfv3]  router-id 10.0.64.1
[HP9505-ospfv3]  area 0.0.0.0
[HP9505-ospfv3]  quit
[HP9505]  quit
```

HP E-series commands

For HP E-series switches, a single command sets up VLANs and assigns physical interfaces to the VLANs.

First, enable routing.

```
HP5406ZL# configure
HP5406ZL(config)# ip routing
HP5406ZL(config)# ipv6 unicast-routing
```

Next, create the VLAN and assign the associated physical interfaces to it.

```
HP5406ZL(config)# vlan 42
HP5406ZL(vlan-42)#    name "VLAN42"
HP5406ZL(vlan-42)#    untagged A9
HP5406ZL(vlan-42)#    ip address 10.0.42.1 255.255.255.0
HP5406ZL(vlan-42)#    ipv6 address 2002:5406:6509:42::1/64
HP5406ZL(vlan-42)#    exit
```

Then, configure the OSPF processes.

```
HP5406ZL(config)# ip router-id 10.0.32.1
HP5406ZL(config)# router ospf
HP5406ZL(ospf)#    area backbone range 10.0.0.0 255.255.0.0 type summary
HP5406ZL(ospf)#    exit
HP5406ZL(config)# router ospf3
HP5406ZL(ospf3)#    area backbone
HP5406ZL(ospf3)#    enable
HP5406ZL(ospf3)#    exit
```

Finally, enable OSPF on the VLAN.

```
HP5406ZL(config)# vlan 42
HP5406ZL(vlan-42)#    ip ospf 10.0.42.1 area backbone
HP5406ZL(vlan-42)#    ipv6 ospf3 area backbone
HP5406ZL(vlan-42)#    exit
```

Cisco commands

On Cisco Catalyst switches, like HP A-series switches, create the VLANs first and then assign physical interfaces to the VLAN instances. The following commands apply to a Cisco Catalyst 6509. Except where noted, the syntax is similar for the Cisco Catalyst 3750-E and Cisco Catalyst 4506 switches.

First, enable IPv6 routing.

```
Cat6509# configure terminal
Cat6509(config)# ipv6 unicast-routing
```

Configuration syntax is slightly different on the Cisco Catalyst 3750-E. First, configure the system to support IPv6.

```
Cat3750E# configure terminal
Cat3750E(config)# sdm prefer dual-ipv4-and-ipv6 default
Cat3750E(config)# ip routing
Cat3750E(config)# ipv6 unicast-routing
```

Cisco Catalyst 3750-E, Cisco Catalyst 4506, and Cisco Catalyst 6509 switches use the same commands for the remaining steps.

Configure the physical interface. With OSPFv3, the primary configuration is done on the port.

```
Cat6509(config)# interface GigabitEthernet4/9
Cat6509(config-if)#  ip address 10.0.42.2 255.255.255.0
Cat6509(config-if)#  ipv6 address 2002:5406:6509:42::2/64
Cat6509(config-if)#  ipv6 ospf 1 area 0
Cat6509(config-if)#  exit
```

Finally, configure the OSPF router processes. On Cisco Catalyst switches, `ospf` refers to OSPFv2 while `ipv6 ospf` refers to OSPFv3.

```
Cat6509(config)# router ospf 1
Cat6509(config-ospf)#  log-adjacency-changes
Cat6509(config-ospf)#  network 10.0.42.2 0.0.0.0 area 0
Cat6509(config-ospf)#  network 10.0.73.2 0.0.0.0 area 0
Cat6509(config-ospf)#  network 10.0.77.2 0.0.0.0 area 0
Cat6509(config-ospf)#  network 10.0.128.0 0.0.127.255 area 0
Cat6509(config-ospf)# ipv6 router ospf 1
Cat6509(config-ipv6-ospf)#  log-adjacency-changes
Cat6509(config-ipv6-ospf)#  end
```

Validation

If the HP and Cisco devices are unable to complete OSPF negotiation, routing adjacencies will remain in the ExStart state. Fully functional adjacencies will be in the Full state. To verify that an OSPF adjacency has entered OSPF Full state on the HP switches, use the **`display ospf peer`** command on the A-series switches and the **`show ipv6 ospf neighbor`** command on the HP E-series switches.

IP multicast switching

Objective

To verify the ability of HP Networking and Cisco Catalyst switches to correctly forward multicast traffic from a multicast routed network.

Background

In IPv4 networks, Ethernet switches use Internet group management protocol (IGMP) snooping to determine where a switch should forward multicast traffic. With IGMP snooping enabled, a switch listens for IGMP reports from attached multicast subscribers. The switch then maps subscribed multicast group address(es) to the interface on which the subscriber is attached. When the switch receives traffic destined for one or more addresses, it will forward it only to those interfaces from which it has heard membership reports.

IGMP snooping requires the use of either an IGMP querier or an IGMP PIM router. An IGMP querier is useful if no router is available, and in this case acts as the multicast router (mrouter) for VLAN by issuing periodic membership queries. If an IGMP querier is being used, there should only be one querier per VLAN.

Topology

In this example, a Spirent TestCenter port attached to the Cisco Catalyst 3750-E offers traffic destined to 10 multicast groups, while other test ports emulate multicast subscribers on the HP A5800 and HP E5406zl. An IGMP querier is used instead of a multicast router.

Figure 1 above illustrates the topology used to validate IP multicast switching functionality. On HP switches, all subscriber interfaces use the same VLAN for untagged traffic, and IGMP snooping is enabled. On the Cisco Catalyst 6509 and Cisco Catalyst 4506, IGMP snooping is enabled. On the Cisco Catalyst 3750-E, IGMP snooping and IGMP querier are enabled.

HP A-series commands

In this example, all interfaces use the same VLAN for untagged traffic and IGMP snooping is enabled globally:

```
<HP5800> system-view
[HP5800] igmp-snooping
```

IGMP snooping also must be enabled on a per-VLAN basis. Only one VLAN is used in this switching example. In configurations that use additional VLANs, only one querier should be defined per VLAN.

```
[HP5800] vlan 300
[HP5800-Vlan-300] igmp-snooping enable
[HP5800-Vlan-300] igmp-snooping querier
[HP5800-Vlan-300] quit
[HP5800] quit
```

HP E-series commands

On HP E-series switches, IGMP snooping also is enabled on a per-VLAN basis:

```
HP5406ZL# configure
HP5406ZL(config)# vlan 300
HP5406ZL(vlan-300)# ip igmp
HP5406ZL(vlan-300)# exit
```

Cisco commands

The following commands apply to a Cisco Catalyst 6509. Except where noted, the syntax is similar for the Catalyst 3750-E switches and Cisco Catalyst 4506 switches.

First, enable IP multicast routing.

```
Cat6509# configure terminal
Cat6509(config)# ip multicast-routing
```

Cisco Catalyst 3750-E switches use a slightly different syntax.

```
Cat3750# configure terminal
Cat3750(config)# ip routing
Cat3750(config)# ip multicast-routing distributed
```

Cisco Catalyst 3750-E, Cisco Catalyst 4506, and Cisco Catalyst 6509 switches use the same commands for the remaining steps.

Enable IGMP snooping. IGMP snooping is enabled by default on Cisco Catalyst switches for all VLANs. In case it is disabled, it can be enabled with these commands:

```
Cat6509# configure terminal
Cat6509(config)# ip igmp snooping
```

Next, enable an IGMP querier. Only one querier should be defined across all switches that share a common VLAN ID.

```
Cat6509# configure terminal
Cat6509(config)# ip igmp snooping querier
Cat6509(config)# end
```

Validation

Once subscribers attached to the HP A5800 switches and the HP E5406zl switches have joined multicast groups by sending IGMP reports with join messages, multicast traffic for these groups will be forwarded to all subscriber ports.

The HP A-series switch command `display igmp-snooping group` also will verify that the HP A5800 and Cisco devices see one another and can exchange IGMP membership information.

The HP E-series switch command `show ip igmp` also will verify that the HP E5406zl and Cisco devices see one another and can exchange IGMP membership information.

IP multicast routing

Objective

To verify the ability of HP switches to learn multicast routing information from a Cisco device using the protocol independent multicast-sparse mode (PIM-SM) protocol.

To verify the ability of the HP Networking and Cisco Catalyst switches to correctly forward multicast traffic based on routing information learned via PIM-SM.

Background

PIM-SM is a popular choice for multicast routing. Devices running PIM-SM can learn topology information from other PIM-SM routers and make forwarding decisions based on that information.

Topology

This example is similar to that used in the "IP Multicast Switching" section, with two important changes: Routing (including OSPF) is enabled on all switches, and the HP A9505 switch also acts a PIM-SM router.

In this example, a Spirent TestCenter port attached to the Cisco Catalyst 3750-E offers traffic destined to 10 multicast groups on different subnets while other test ports emulated multicast subscribers to all 10 groups on the HP A5800 and HP E5406zl. The Cisco device uses PIM-SM to propagate subnet routing information to other subnets, including the ones on the HP switches, also running PIM-SM, as attached.

The HP switches use PIM-SM and OSPF to propagate routing information. Multicast subscribers attached to VLAN-routed interfaces, each in a different VLAN with each VLAN in a different IP subnet, receive traffic from Sprirent TestCenter. The subscriber interfaces also use IGMP to build a multicast forwarding table.

Figure 1 above illustrates the topology used to validate IP multicast routing functionality. PIM-SM and OSPF routing is enabled on all Cisco and HP devices.

HP A-series commands

First, create the necessary VLANs.

```
<HP5800> system-view
[HP5800] vlan 2 to 10
```

Next, assign IP addresses and enable IGMP on the respective VLANs.

```
[HP5800] interface Vlan-interface2
[HP5800-vlan-interface2] ip address 10.0.2.1 255.255.255.0
[HP5800-vlan-interface2]  igmp enable
[HP5800-vlan-interface2] interface Vlan-interface3
[HP5800-vlan-interface3]  ip address 10.0.3.1 255.255.255.0
[HP5800-vlan-interface3]  igmp enable
[HP5800-vlan-interface3] interface Vlan-interface4
```

```
[HP5800-vlan-interface4]  ip address 10.0.4.1 255.255.255.0
[HP5800-vlan-interface4]  igmp enable
[HP5800-vlan-interface4] interface Vlan-interface5
[HP5800-vlan-interface5]  ip address 10.0.5.1 255.255.255.0
[HP5800-vlan-interface5]  igmp enable
[HP5800-vlan-interface5] interface Vlan-interface6
[HP5800-vlan-interface6]  ip address 10.0.6.1 255.255.255.0
[HP5800-vlan-interface6]  igmp enable
[HP5800-vlan-interface6] interface Vlan-interface7
[HP5800-vlan-interface7]  ip address 10.0.7.1 255.255.255.0
[HP5800-vlan-interface7]  igmp enable
[HP5800-vlan-interface7] interface Vlan-interface8
[HP5800-vlan-interface8]  ip address 10.0.8.1 255.255.255.0
[HP5800-vlan-interface8]  igmp enable
[HP5800-vlan-interface8] interface Vlan-interface9
[HP5800-vlan-interface9]  ip address 10.0.9.1 255.255.255.0
[HP5800-vlan-interface9]  igmp enable
[HP5800-vlan-interface9]  quit
```

Then, assign the interfaces to the respective VLANs.

```
[HP5800] interface GigabitEthernet1/0/1
[HP5800-Gigabitethernet1/0/1] port link-mode bridge
[HP5800-Gigabitethernet1/0/1] port access vlan 2
[HP5800-Gigabitethernet1/0/1] interface GigabitEthernet1/0/2
[HP5800-Gigabitethernet1/0/2] port link-mode bridge
[HP5800-Gigabitethernet1/0/2] port access vlan 3
[HP5800-Gigabitethernet1/0/2] interface GigabitEthernet1/0/3
[HP5800-Gigabitethernet1/0/3] port link-mode bridge
[HP5800-Gigabitethernet1/0/3] port access vlan 4
[HP5800-Gigabitethernet1/0/3] interface GigabitEthernet1/0/4
[HP5800-Gigabitethernet1/0/4] port link-mode bridge
[HP5800-Gigabitethernet1/0/4] port access vlan 5
[HP5800-Gigabitethernet1/0/4] interface GigabitEthernet1/0/5
[HP5800-Gigabitethernet1/0/5] port link-mode bridge
[HP5800-Gigabitethernet1/0/5] port access vlan 6
[HP5800-Gigabitethernet1/0/5] interface GigabitEthernet1/0/6
[HP5800-Gigabitethernet1/0/6] port link-mode bridge
[HP5800-Gigabitethernet1/0/6] port access vlan 7
[HP5800-Gigabitethernet1/0/6] interface GigabitEthernet1/0/7
[HP5800-Gigabitethernet1/0/7] port link-mode bridge
[HP5800-Gigabitethernet1/0/7] port access vlan 8
[HP5800-Gigabitethernet1/0/7] interface GigabitEthernet1/0/8
[HP5800-Gigabitethernet1/0/8] port link-mode bridge
[HP5800-Gigabitethernet1/0/8] port access vlan 9
[HP5800-Gigabitethernet1/0/8] quit
```

Next, enable multicast routing and IGMP.

```
[HP5800] multicast routing-enable
[HP5800] igmp-snooping
```

Then, enable OSPF. Although this step is not strictly necessary for IP multicast routing, some unicast routing protocol or static routing is required.

```
[HP5800] ospf 1 router-id 10.0.0.1
[HP5800-OSPF] area 0.0.0.0
[HP5800-OSPF] network 10.0.0.0 0.0.255.255
[HP5800-OSPF] quit
```

Then, enable PIM-SM and designate a rendezvous point (RP), in this example the VLAN interface on the HP A9505 switch.

```
[HP5800] pim
[HP5800-PIM] static-rp 10.0.73.1
[HP5800-PIM] quit
```

Next, enable PIM-SM on the VLAN used for the trunk line between switches.

```
[HP5800] interface Vlan-interface41
[HP5800-vlan-interface41]  ip address 10.0.41.2 255.255.255.0
[HP5800-vlan-interface41]  igmp enable
[HP5800-vlan-interface41]  pim sm
[HP5800-vlan-interface41]  quit
```

Finally, on the HP A9505, we need to configure the rendezvous point (RP). We will also enable PIM-SM on the same VLAN interface that is being used as the rendezvous point.

```
[HP9505] interface Vlan-interface75
[HP9505-vlan-interface75]  ip address 10.0.75.1 255.255.255.0
[HP9505-vlan-interface75]  igmp enable
[HP9505-vlan-interface75]  pim sm
[HP9505] pim
[HP9505-PIM]  c-rp Vlan-interface73
[HP9505-PIM]  static-rp 10.0.73.1
[HP9505-PIM]  quit
[HP9505]  quit
```

HP E-series commands

First, setup the VLANs that will be used, and assign interfaces to them.

```
HP5406# configure
HP5406(config)# vlan 33
HP5406(vlan-33)#    name "VLAN33"
HP5406(vlan-33)#    untagged A1
HP5406(vlan-33)#    ip address 10.0.33.1 255.255.255.0
HP5406(vlan-33)#    exit
HP5406(config)# vlan 34
HP5406(vlan-34)#    name "VLAN34"
HP5406(vlan-34)#    untagged A2
HP5406(vlan-34)#    ip address 10.0.34.1 255.255.255.0
HP5406(vlan-34)#    exit
HP5406(config)# vlan 35
HP5406(vlan-35)#    name "VLAN35"
HP5406(vlan-35)#    untagged A3
HP5406(vlan-35)#    ip address 10.0.35.1 255.255.255.0
```

```
HP5406(vlan-35)#    exit
HP5406(config)# vlan36
HP5406(vlan-36)#    name "VLAN36"
HP5406(vlan-36)#    untagged A4
HP5406(vlan-36)#    ip address 10.0.36.1 255.255.255.0
HP5406(vlan-36)#    exit
HP5406(config)# vlan 37
HP5406(vlan-37)#    name "VLAN37"
HP5406(vlan-37)#    untagged A5
HP5406(vlan-37)#    ip address 10.0.37.1 255.255.255.0
HP5406(vlan-37)#    exit
HP5406(config)# vlan 38
HP5406(vlan-38)#    name "VLAN38"
HP5406(vlan-38)#    untagged A6
HP5406(vlan-38)#    ip address 10.0.38.1 255.255.255.0
HP5406(vlan-38)#    exit
HP5406(config)# vlan 39
HP5406(vlan-39)#    name "VLAN39"
HP5406(vlan-39)#    untagged A7
HP5406(vlan-39)#    ip address 10.0.39.1 255.255.255.0
HP5406(vlan-39)#    exit
HP5406(config)# vlan 40
HP5406(vlan-40)#    name "VLAN40"
HP5406(vlan-40)#    untagged A8
HP5406(vlan-40)#    ip address 10.0.40.1 255.255.255.0
HP5406(vlan-40)#    exit
```

Next, configure OSPF routing. While OSPF is not strictly necessary, some unicast routing protocol or static routing is required.

```
HP5406(config)# ip routing
HP5406(config)# ip router-id 10.0.32.1
HP5406(config)# router ospf
HP5406(ospf)#    area backbone range 10.0.0.0 255.255.0.0 type summary
HP5406(ospf)#    exit
HP5406(config)# vlan 33
HP5406(vlan-33)#    ip ospf 10.0.33.1 area backbone
HP5406(vlan-33)#    exit
HP5406(config)# vlan 34
HP5406(vlan-34)#    ip ospf 10.0.34.1 area backbone
HP5406(vlan-34)#    exit
HP5406(config)# vlan 35
HP5406(vlan-35)#    ip ospf 10.0.35.1 area backbone
HP5406(vlan-35)#    exit
HP5406(config)# vlan 36
HP5406(vlan-36)#    ip ospf 10.0.36.1 area backbone
HP5406(vlan-36)#    exit
HP5406(config)# vlan 37
HP5406(vlan-37)#    ip ospf 10.0.37.1 area backbone
HP5406(vlan-37)#
HP5406(config)# vlan 38
HP5406(vlan-38)#    ip ospf 10.0.38.1 area backbone
HP5406(vlan-38)#    exit
HP5406(config)# vlan 39
```

```
HP5406(vlan-39)#    ip ospf 10.0.39.1 area backbone
HP5406(vlan-39)#    exit
HP5406(config)# vlan 40
HP5406(vlan-40)#    ip ospf 10.0.40.1 area backbone
HP5406(vlan-40)#    exit
```

Then, enable multicast routing and set the PIM rendezvous point (RP).

```
HP5406(config)# ip multicast-routing
HP5406(config)# router pim
HP5406(pim)#   rp-address 10.0.73.1 224.0.0.0 240.0.0.0
HP5406(pim)#   exit
```

Finally, configure a VLAN to carry traffic between switches, and enable OSPF on that VLAN.

```
HP5406(config)# vlan 43
HP5406(vlan-33)#    name "VLAN43"
HP5406(vlan-33)#    untagged A10
HP5406(vlan-33)#    ip address 10.0.43.1 255.255.255.0
HP5406(vlan-33)#    ip igmp
HP5406(vlan-33)#    exit
HP5406(vlan-33)#    ip ospf 10.0.43.1 area backbone
HP5406(vlan-33)#    ip pim-sparse
HP5406(vlan-33)#       ip-addr any
HP5406(vlan-33)#       exit
HP5406(vlan-33)#    exit
HP5406#    exit
```

Cisco commands

The following commands apply to a Cisco Catalyst 6509. Except where noted, the syntax is similar for Cisco Catalyst 3750-E and Cisco Catalyst 4506 switches.

First, enable IP multicast routing.

```
Cat6509# configure terminal
Cat6509(config)# ip multicast-routing
```

Command syntax is slightly different on the Cisco Catalyst 3750-E.

```
Cat3750E# configure terminal
Cat3750E(config)# ip routing
Cat3750E(config)# ip multicast-routing distributed
```

Cisco Catalyst 3750-E, Cisco Catalyst 4506, and Cisco Catalyst 6509 switches use the same commands for the remaining steps.

Configure interswitch interfaces with an IP address and support for PIM-SM.

```
Cat6509(config)# interface TenGigabitEthernet5/4
Cat6509(config-if)#  ip address 10.0.201.1 255.255.255.0
```

```
Cat6509(config-if)#  ip pim sparse-mode
Cat6509(config-if)#  exit
```

Then, enable OSPF. Although OSPF is not strictly necessary for IP multicast forwarding, some unicast routing protocol or static routing is required.

```
Cat6509(config)# router ospf 1
Cat6509(config-router)#  log-adjacency-changes
Cat6509(config-router)#  network 10.0.0.0 0.0.255.255 area 0
Cat6509(config-router)#  exit
```

4. Configure a PIM rendezvous point (RP). In the case the RP will be on the HP A9505.

```
Cat6509(config)# ip pim rp-address 10.0.73.1
Cat6509(config)# end
```

Validation

Once subscribers attached to the HP switches have joined multicast groups by sending IGMP reports with join messages, any multicast traffic for these groups offered to Interface VLAN73 on the HP9505 will be forward to all subscriber ports on the HP switches.

The HP A-series command `display ip multicast routing-table` will verify that the HP A5800 and Cisco devices see one another and can exchange multicast information. The HP E-series `command show ip mrouter` provides the same verification for HP E5406zl Ethernet switches.

Virtual router redundancy protocol (VRRP) interoperability

Objective

To validate failover functionality of the virtual router redundancy protocol (VRRP) between HP Networking and Cisco Catalyst switches configured as routers.

Background

Two or more routers can make use of VRRP to add redundancy and enhance network availability. With VRRP, all routers share a single virtual IP address. One router acts as the master (active) device, while all others act as backups. If the master router fails (or if a link fails on the interfaces configured with the virtual IP address), one of the backup routers takes over as master.

Topology

In this example, an HP A9505 switch, HP E5406zl, and Cisco Catalyst 6509 switch are all configured to route IP traffic. The interfaces connecting the switches each have unique IP addresses. A shared virtual IP address of 10.0.41.254/24 is used for VRRP, with the HP E5406zl initially acting as the master.

Figure 5 below illustrates the VRRP validation test bed. The HP switches assign an IP address to VLAN 41, and then assign interfaces to that VLAN. However, VRRP also would work if an IP address was assigned directly to the physical interface, as it is with the Cisco Catalyst 6509.

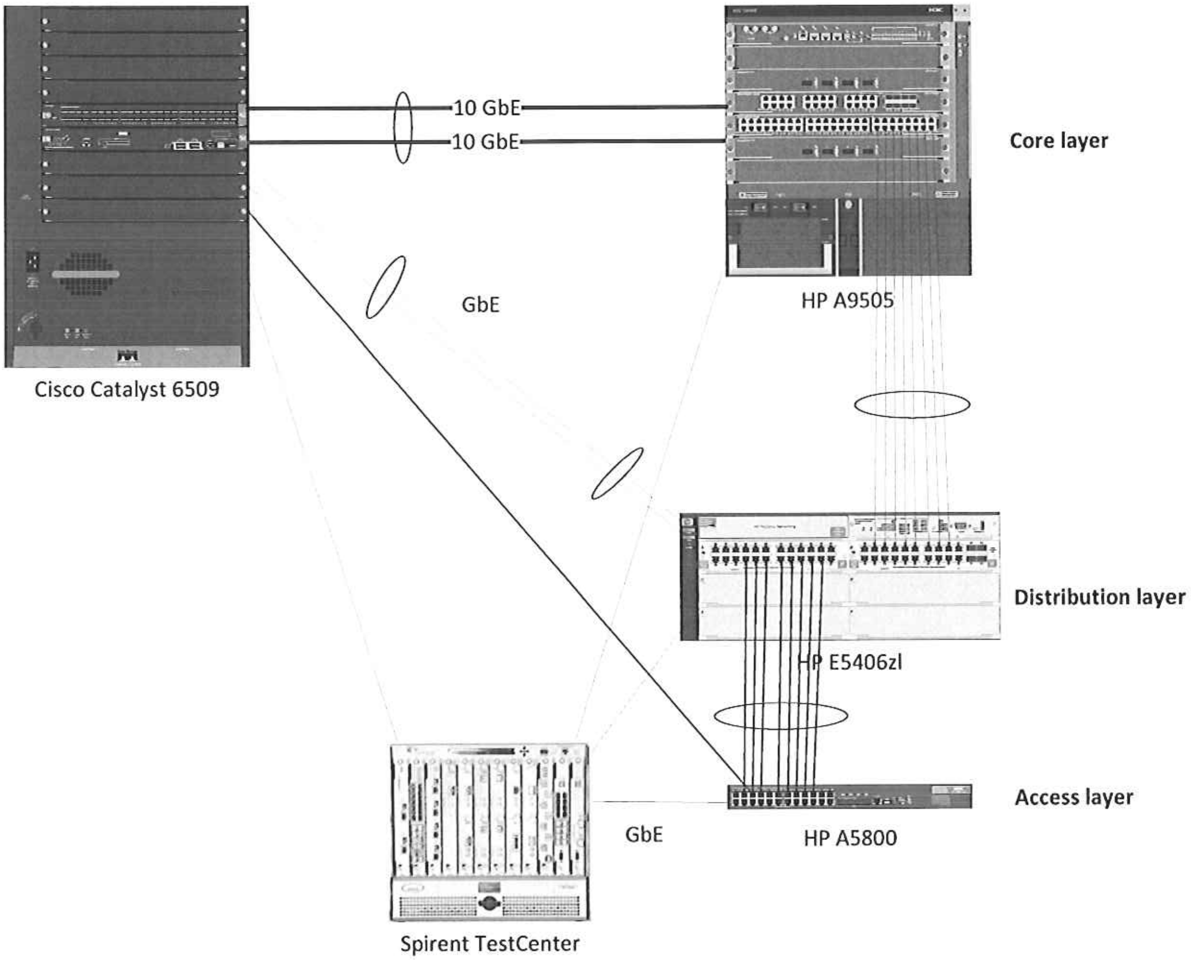

Figure 5: Virtual router redundancy protocol test bed

HP A-series commands

VRRP configuration is done in the interface configuration context. Here, it is done on the VLAN interface.

```
<HP9505> system-view
[HP9505] interface Vlan-interface41
[HP9505-Vlan-interface41]  ip address 10.0.41.1 255.255.255.0
[HP9505-Vlan-interface41]  vrrp vrid 1 virtual-ip 10.0.41.254
[HP9505-Vlan-interface41]  vrrp vrid 1 priority 254
[HP9505-Vlan-interface41]  quit
[HP9505]  quit
```

HP E-series commands

First, create and configure the VLAN interface.

```
HP5406ZL# configure
HP5406ZL(config)# vlan 41
HP5406ZL(vlan-41)#    name "VLAN41"
HP5406ZL(vlan-41)#    untagged A10
HP5406ZL(vlan-41)#    ip address 10.0.41.1 255.255.255.0
HP5406ZL(vlan-41)#    exit
```

Next, configure VRRP.

```
HP5406ZL(config)# router vrrp
HP5406ZL(config)# vlan 41
HP5406ZL(vlan-41)#    vrrp vrid 1
HP5406ZL(vlan-41-vrid-1)#     owner
HP5406ZL(vlan-41-vrid-1)#    virtual-ip-address 10.0.41.254 255.255.255.0
HP5406ZL(vlan-41-vrid-1)#    priority 255
HP5406ZL(vlan-41-vrid-1)#    enable
HP5406ZL(vlan-41-vrid-1)#    exit
HP5406ZL(vlan-41-vrid-1)#  exit
HP5406ZL(vlan-41)#   exit
HP5406ZL#   exit
```

Cisco commands

The following commands apply to a Cisco Catalyst 6509. The syntax is similar for Catalyst 3750-E and Cisco Catalyst 4506 switches.

VRRP configuration is done in the interface configuration context.

```
Cat6509# configure terminal
Cat6509(config)# interface GigabitEthernet4/21
Cat6509(config-if)#  ip address 10.0.41.3 255.255.255.0
```

```
Cat6509(config-if)#  vrrp 1 description VRRP Test
Cat6509(config-if)#  vrrp 1 ip 10.0.41.254
Cat6509(config-if)#  vrrp 1 timers learn
Cat6509(config-if)#  vrrp 1 priority 90
Cat6509(config-if)#  end
```

Validation

Both the HP E5406zl and Cisco Catalyst 6509 support the `show vrrp` command, which will indicate the current VRRP state on each system.

Appendix A: About Network Test

Network Test is an independent third-party test lab and engineering services consultancy. Our core competencies are performance, security, and conformance assessment of networking equipment and live networks. Our clients include equipment manufacturers, large enterprises, service providers, industry consortia, and trade publications.

Appendix B: Sample Configuration Files

This appendix lists URLs for the HP and Cisco switch files used to verify interoperability. These files are freely available for download from a public Network Test server.

A copy of this document, a brief interoperability report, and all HP and Cisco configuration files are available at http://networktest.com/hpiop.

Appendix C: Software Releases Tested

This appendix describes the software versions used on the test bed. All tests were conducted in March 2011 at Network Test's facility in Westlake Village, CA, USA.

Component	Version
HP A9505	5.20, Release 1238P08
HP E5406zl	K.15.03.0007
HP A5800	5.20, Release 1206
Cisco Catalyst 6509	12.2(33)SXI2a
Cisco Catalyst 4506	12.2(20)EWA
Cisco Catalyst 3750-E	12.2(55)SE1
Spirent TestCenter	3.55.5086.0000

Appendix D: Disclaimer

Network Test Inc. has made every attempt to ensure that all test procedures were conducted with the utmost precision and accuracy, but acknowledges that errors do occur. Network Test Inc. shall not be held liable for damages which may result for the use of information contained in this document. All trademarks mentioned in this document are property of their respective owners.

TESTED with SPIRENT

Network Test Inc.
31324 Via Colinas, Suite 113
Westlake Village, CA 91362-6761
USA
+1-818-889-0011
http://networktest.com
info@networktest.com

networktest

FlexNetwork Architecture Delivers Higher Speed, Lower Downtime With HP IRF Technology

August 2011

HP IRF Performance Assessment

Executive Summary

HP commissioned Network Test to assess the performance of Intelligent Resilient Framework (IRF), a method of virtualizing data center switch fabrics for enhanced bandwidth and availability.

On multiple large-scale test beds, IRF clearly outperformed existing redundancy mechanisms such as the spanning tree protocol (STP) and the virtual routing redundancy protocol (VRRP).

Among the key findings of these tests:

- Using VMware's vMotion facility, average virtual machine migration time took around 43 seconds on a network running IRF, compared with around 70 seconds with rapid STP
- IRF virtually doubled network bandwidth compared with STP and VRRP, with much higher throughput rates regardless of frame size
- IRF converged around failed links, line cards, and systems vastly faster than existing redundancy mechanisms such as STP
- In the most extreme failover case, STP took 31 *seconds* to recover after a line card failure; IRF recovered from the same event in 2.5 *milliseconds*
- IRF converges around failed network components far faster than HP's 50-millisecond claim

This document briefly explains IRF concepts and benefits, and then describes procedures and results for IRF tests involving VMware vMotion; network bandwidth; and resiliency.

Introducing IRF

IRF consolidates multiple physical switches so that they appear to the rest of the network as a single logical entity. Up to nine switches can comprise this virtual fabric, which runs on HP's high-end switch/routers[1] and can encompass hundreds or thousands of gigabit and 10-gigabit Ethernet ports.

IRF offers advantages in terms of simpler network design; ease of management; disaster recovery; performance; and resiliency. A virtual fabric essentially "flattens" the data center from three layers into one or two using HP Virtual Connect technology, requiring fewer switches.

Device configuration and management also becomes simpler. Within an IRF domain, configuration of a single primary switch is all that's needed; the primary switch then distributes relevant configuration and protocol information to other switches in the IRF domain. IRF also supports an in-service software upgrade (ISSU) capability that allows individual switches to be taken offline for upgrades without affecting the rest of the virtual fabric.

For disaster recovery, switches within an IRF domain can be deployed across multiple data centers. According to HP, a single IRF domain can link switches up to 70 kilometers (43.5 miles) apart.

---

[1] IRF support is included at no cost on HP 12500, 9500, 7500, 58xx, and 55xx switches.

IRF improves performance and resiliency, as shown in test results described later in this report. A common characteristic of existing data center network designs is their inefficient redundancy mechanisms, such as STP or VRRP.

Both these protocols (along with modern versions of STP such as rapid STP and multiple STP) use an "active/passive" design, where only one pair of interfaces between switches forwards traffic, and all others remain idle until the active link fails. With active/passive mechanisms, half (or more) of all inter-switch links sit idle most of the time. Moreover, both STP and VRRP take a relatively long time to recover from link or component faults, typically on the order of seconds.

IRF uses an "active/active" design that enables switches to forward traffic on all ports, all the time. This frees up bandwidth, boosting performance for all applications. Data centers using virtualization benefit especially well from this design, since the additional bandwidth allows virtual machines to be moved faster between hypervisors. IRF's active/active designs also reduce downtime when link, component, or system failures occur.

IRF Speeds VMware Performance

Over the past few years VMware's vMotion capability has become the "killer app" for large-scale data centers and cloud computing. The ability to migrate virtual machines between physical hosts with zero downtime is a boon to network managers, but also a challenge. As data centers scale up in size and network managers use vMotion to migrate ever-larger numbers of virtual machines, network performance can become a bottleneck. This is an acute concern for disaster recovery and other high-availability applications, where rapid migration of virtual machines is essential.

Network Test and HP engineers constructed a large-scale test bed to compare vMotion performance using IRF and rapid spanning tree protocol (RSTP). With both mechanisms, the goal was to measure the time needed for vMotion migration of 128 virtual machines, each with 8 Gbytes of RAM, running Microsoft SQL Server on Windows Server 2008. Before each migration event, test engineers verified maximum memory usage on each VM, ensuring the most stressful possible load in terms of network utilization. The test migrated virtual machines between VMware ESXi hosts running on a total of 32 HP BL460 G7 blade servers.

In the RSTP case, the network used a typical active/passive design, with some ports forwarding traffic between access and core switches, and others in blocking state (see Figure 1). In this design, RSTP provides excellent loop prevention but also limits available bandwidth; note that half the inter-switch connections shown here are in blocking state.

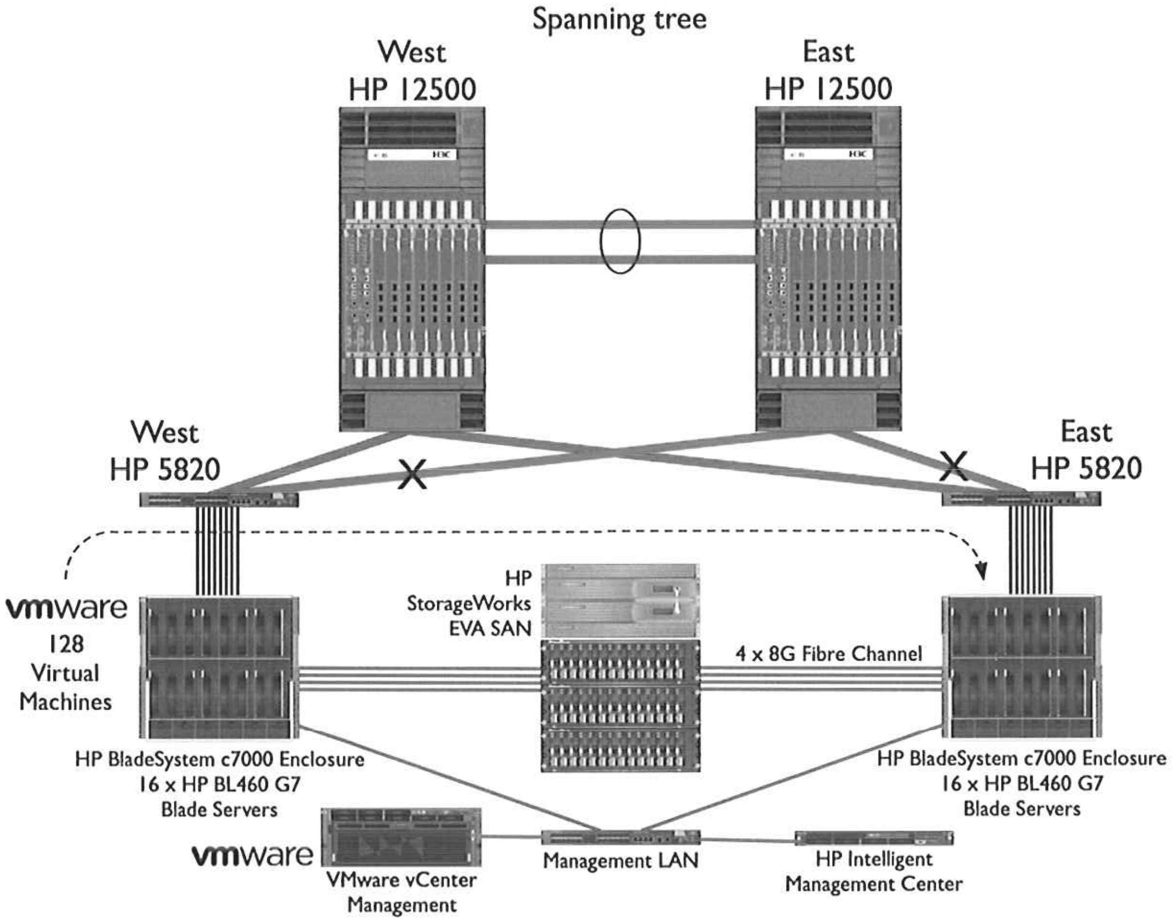

Figure 1: VMware vMotion with RSTP

By virtualizing the core switching infrastructure, IRF increases network capacity (see Figure 2). Here, all inter-switch ports are available, with no loss in redundancy compared with spanning tree. The core switches appear to the rest of the network as a single logical device. That converged device can forward traffic on all links with all attached access switches.

Moreover, IRF can be used together with the link aggregation control protocol (LACP) to add capacity to links between switches (or between switches and servers), again with all ports available all the time. This isn't possible with STP or RSTP since at least half (and possibly more) of all inter-switch links must remain in blocking state at all times.

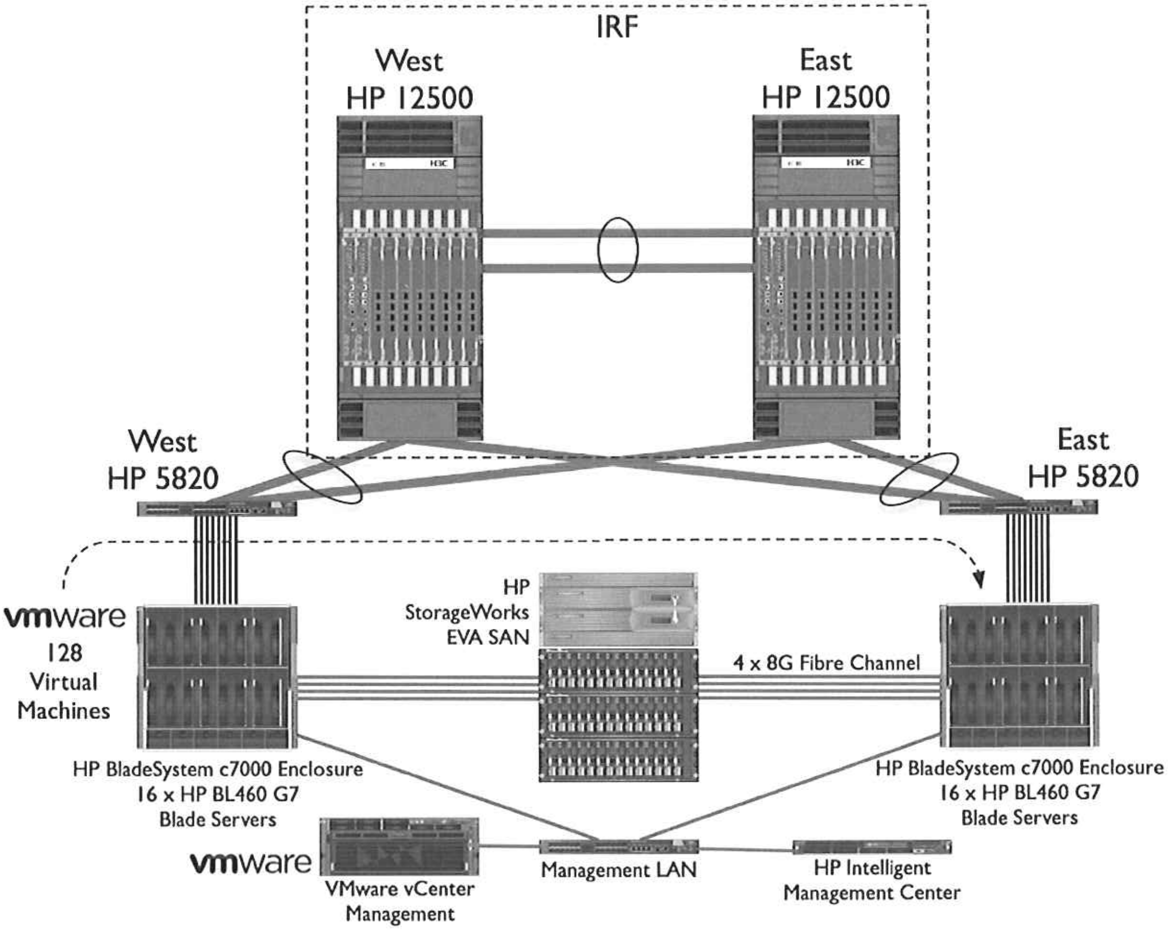

Figure 2: VMware vMotion with IRF

For both RSTP and IRF scenarios, engineers used custom-developed scripts to trigger vMotion migration of 128 virtual machines. In all three trials run, **IRF clearly outperformed RSTP in terms of average vMotion migration time** (see Figure 3). On average, migrations over RSTP took around 70 seconds to complete, while vMotion over IRF took 43 seconds.

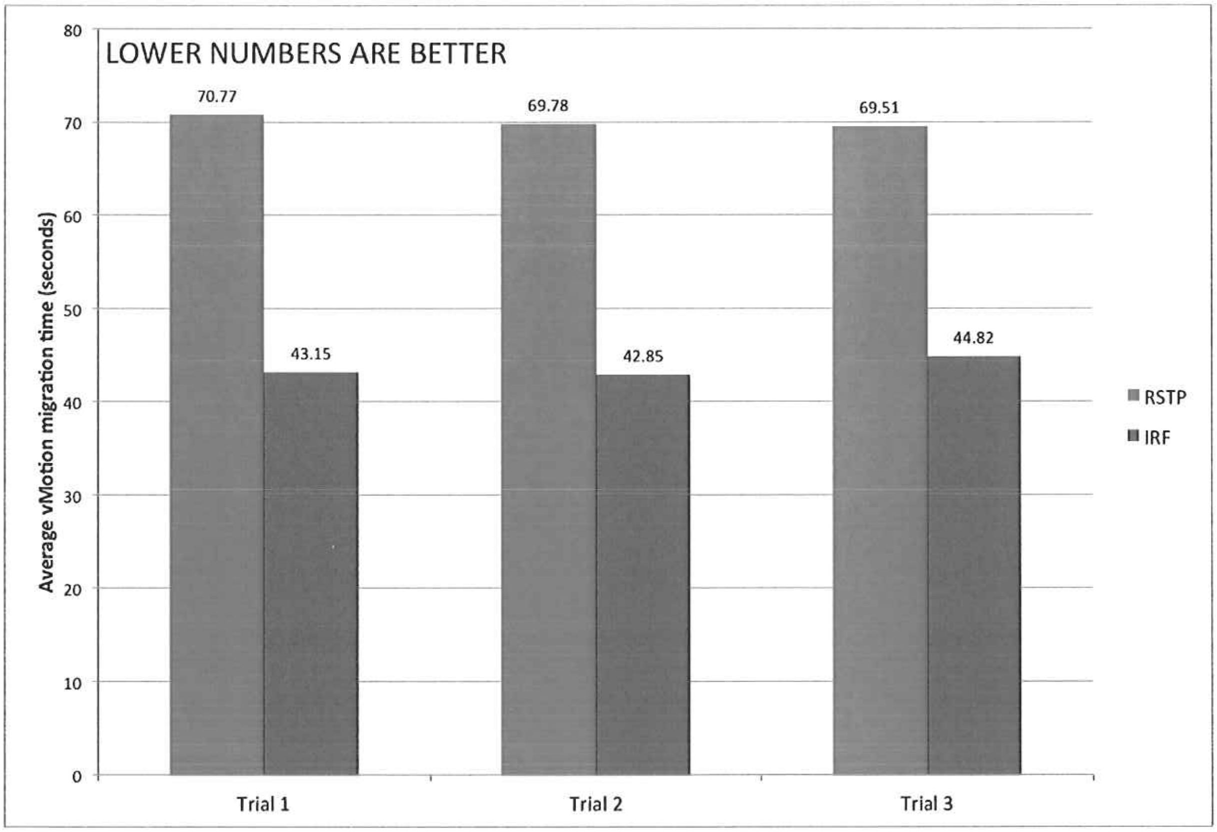

Figure 3: Average vMotion times over RSTP and IRF

Boosting Bandwidth: IRF Speeds Throughput

It's important to note that actual vMotion times depend on many variables, including server and storage resources, virtual machine configuration, and VMware vSphere parameters. As with all application performance tuning, results can vary from site to site. To determine a more basic metric – raw network capacity – Network Test ran additional tests to compare bandwidth available with and without IRF.

Network Test conducted throughput tests comparing IRF with STP at layer 2 and VRRP at layer 3.[2]

Figure 4 below shows the test beds used to compare throughput in the STP and IRF test cases. While both tests involve the same 12500 and 5820 switches, note that half the ports between them are in blocking state in the STP test cases (seen in the left side of Figure 4).

The IRF configuration makes use of all inter-switch links (see in the right of the figure), with the two 12500 switches appearing to the network as a single entity. The test bed for VRRP was similar to that shown here, except that Network Test used 20 gigabit Ethernet connections between each traffic generator and the 5820s to increase emulated host count and ensure more uniform distribution of

---

[2] Network Test did not compare IRF and rapid spanning tree protocol [RSTP] throughput because RSTP results would be identical to those with STP in this particular configuration.

traffic across VRRP connections. Engineers configured the 12500 switches in "bridge extended" mode, which improves performance by allocating additional memory to switching processes.

The IRF configuration uses the link aggregation control protocol (LACP) on connections between the 12500 and 5820 switches. From the perspective of the 5820 access switches, the link-aggregated connection to the core appears to be a single virtual connection. This allows for interesting network designs in disaster-recovery scenarios, for example with IRF and link aggregation connecting switches in different physical locations.

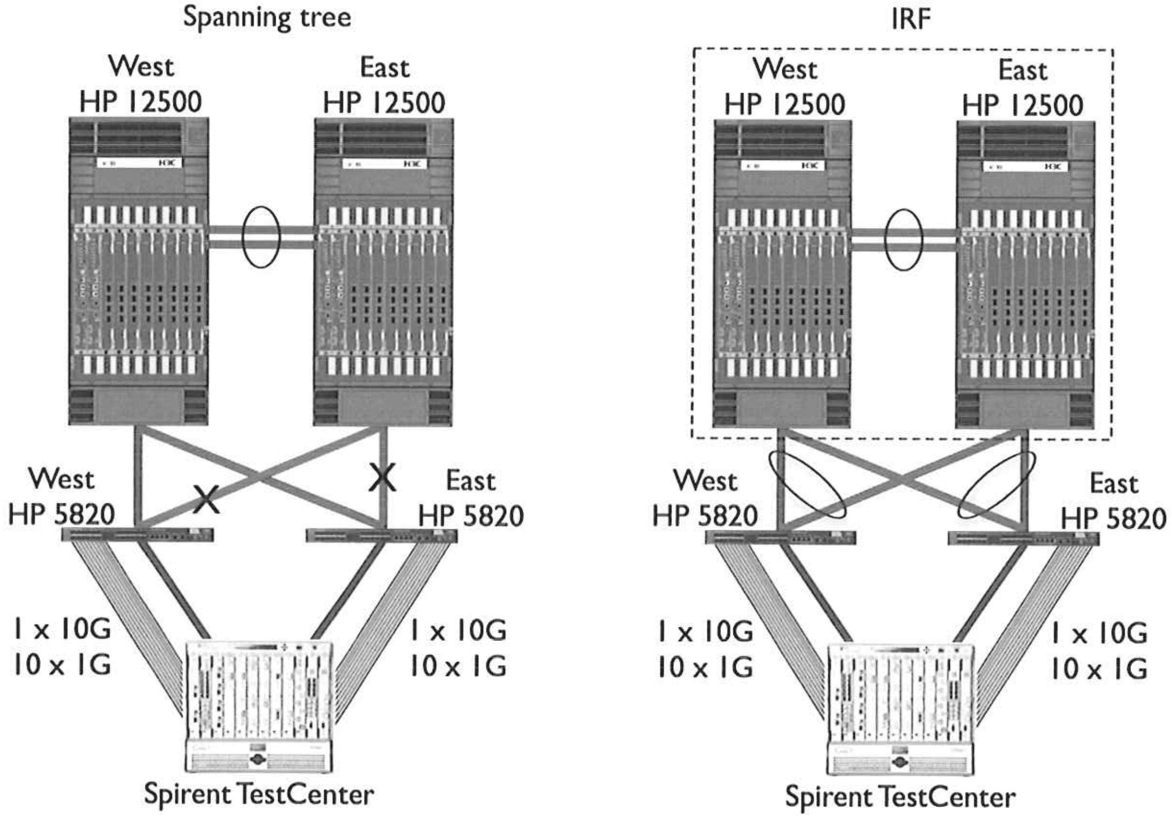

Figure 4: Spanning tree and IRF test beds

Network Test followed the procedures described in RFCs 2544 and 2889 to determine system throughput. Engineers configured the Spirent TestCenter traffic generator/analyzer to offer bidirectional traffic in an "east/west" direction, meaning all frames on the "west" side of the test bed were destined for the "east" side and vice-versa. The aggregate load offered to each side was equivalent to 20 Gbit/s of traffic, equal to the theoretical maximum capacity of the access-core links.

To measure throughput, engineers offer traffic at varying loads, each for a 60-second duration, to determine the highest rate at which the switches would forward all frames with zero frame loss. As defined in RFC 1242, this is the throughput rate.

HP IRF Performance Assessment

Engineers repeated these tests with various frame sizes ranging from 64 bytes (the minimum in Ethernet) through 1,518 bytes (the nominal maximum in Ethernet) through 2,176 bytes (often seen in data centers that use Fibre Channel for storage) through 9,216 (the nonstandard but still widely used jumbo frames common in data centers).

Figure 5 below presents throughput results, expressing the throughput rate as a percentage of the theoretical maximum rate. **For all frame sizes, IRF nearly doubled channel capacity,** delivering near line-rate throughput while the active/passive solutions delivered throughput of only 50 percent of channel capacity.

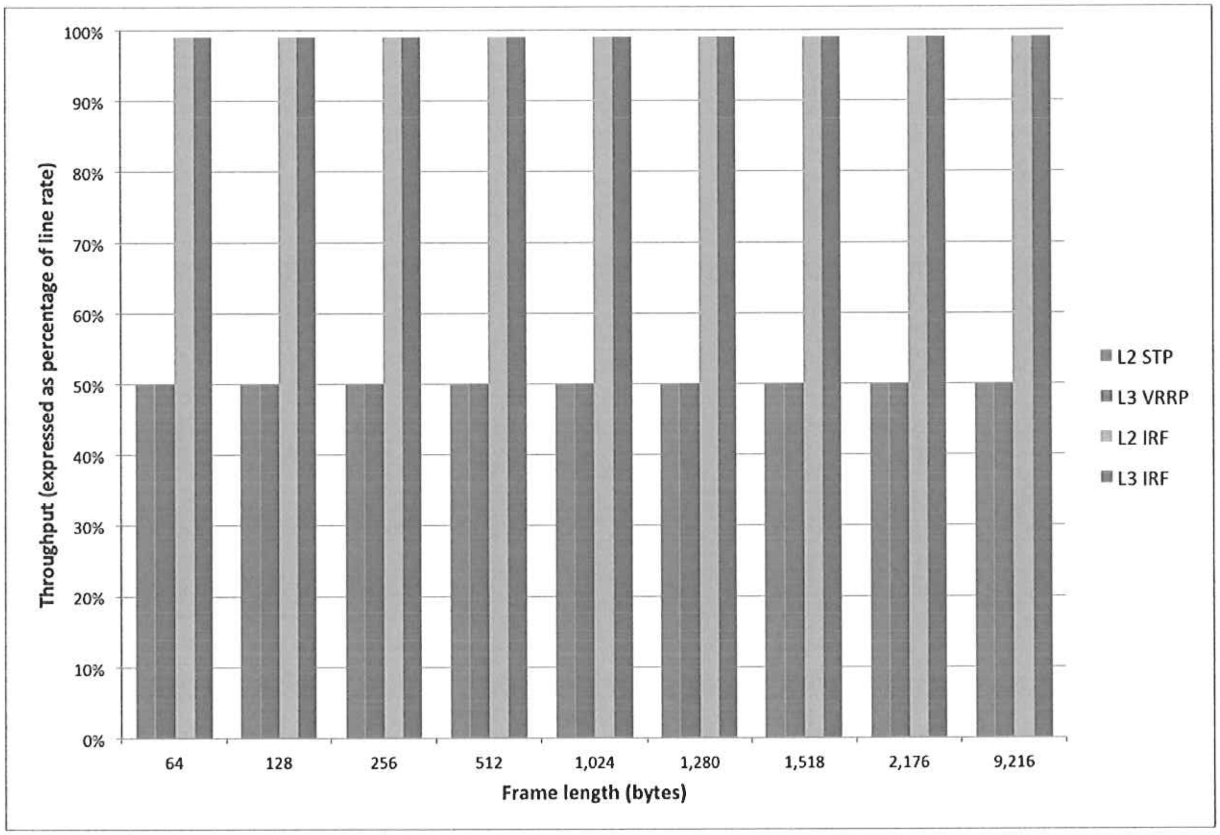

Figure 5: Throughput as a percentage of the theoretical maximum

Note also that **bandwidth utilization nearly doubles with IRF in every test case,** regardless of frame length. **This validates IRF's ability to deliver far more network bandwidth, which in turn speeds performance for all applications, regardless of traffic profile.**

The throughput figures presented in Figure 5 are given as percentages of theoretical line rate. As a unit of measurement, throughput itself is actually a rate and not a percentage. Figure 6 below presents the same data from the throughput tests, this time with throughput expressed in frames per second for each test case. **Regardless of how it's expressed, IRF provides nearly double the network bandwidth as other layer-2 and layer-3 resiliency mechanisms.**

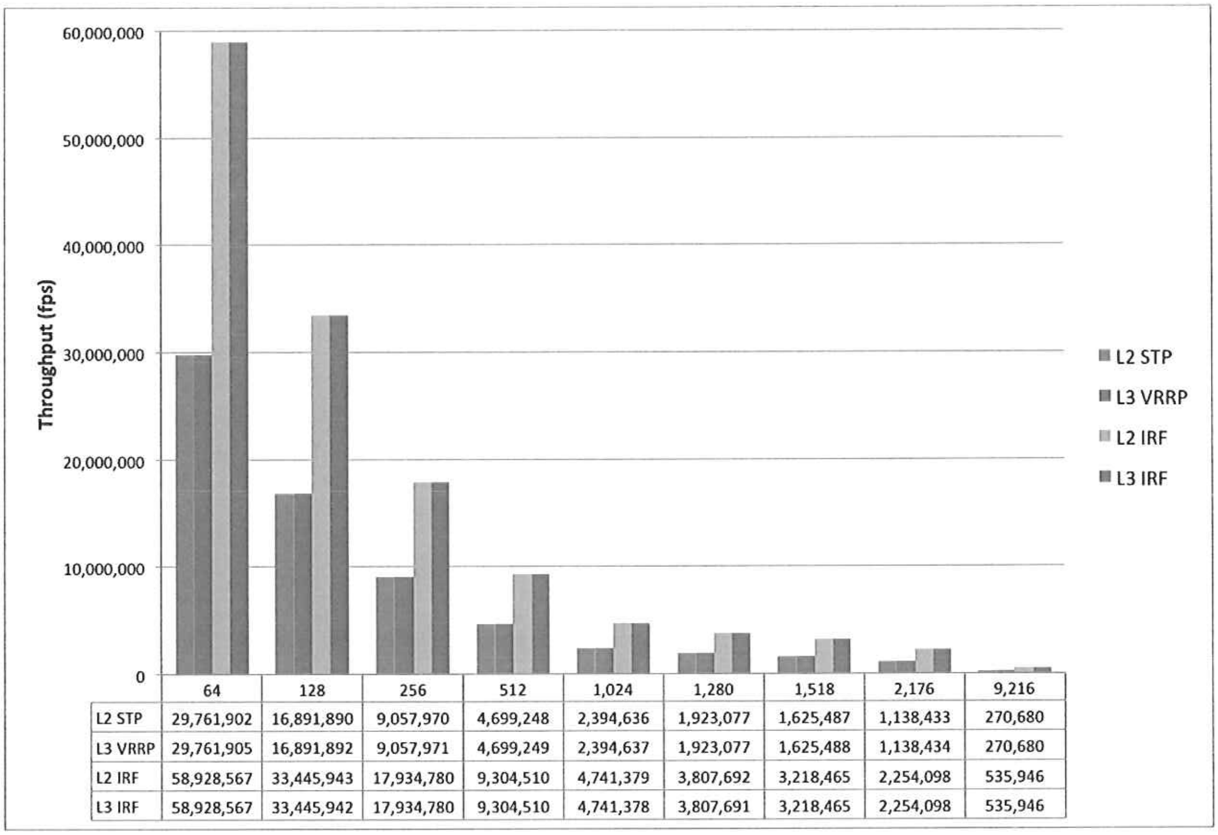

	64	128	256	512	1,024	1,280	1,518	2,176	9,216
L2 STP	29,761,902	16,891,890	9,057,970	4,699,248	2,394,636	1,923,077	1,625,487	1,138,433	270,680
L3 VRRP	29,761,905	16,891,892	9,057,971	4,699,249	2,394,637	1,923,077	1,625,488	1,138,434	270,680
L2 IRF	58,928,567	33,445,943	17,934,780	9,304,510	4,741,379	3,807,692	3,218,465	2,254,098	535,946
L3 IRF	58,928,567	33,445,942	17,934,780	9,304,510	4,741,378	3,807,691	3,218,465	2,254,098	535,946

Figure 6: Comparing STP, VRRP, and IRF throughput

Faster Failovers: IRF Improves Resiliency

While *high performance* is certainly important, *high availability* is an even more critical requirement in enterprise networking. This is especially true in the data center, where even a small amount of downtime can mean significant revenue loss and other disruptions. IRF aims to improve network uptime by recovering from link or component failures far faster than mechanisms such as STP, RSTP, or VRRP.

Networks running STP can take between 30-60 seconds to converge following a single link or component failure. Rapid spanning tree and VRRP are newer and faster, but convergence time can still be significant. HP claims IRF will converge in 50 milliseconds.

To assess that claim, Network Test used the same test bed as in the STP and IRF throughput tests (see Figure 4, again). This time, engineers intentionally caused a failure, and then derived convergence time by examining frame loss as reported by Spirent TestCenter. Engineers tested three failure modes:

- **Link failure:** With traffic active, engineers disconnected the link between the east 12500 and east 5820 switches, forcing traffic to be rerouted onto an alternative path

- **Card failure:** With traffic active, engineers pulled an active line card from the east 12500, forcing traffic to be rerouted
- **System failure:** With traffic active, engineers cut power to the east 12500, forcing traffic to be rerouted through the west 12500

In all cases, Spirent TestCenter offered bidirectional streams of 64-byte frames at exactly 50 percent of line rate throughout the test. This is the most stressful non-overload condition possible; in theory, the system under test is never congested, even during component failure. Thus, any frames dropped during this test were a result of, and only of, path re-computation following a component failure.

Figure 7 below compares convergence times for conventional STP with IRF configured for layer-2 operation. **For all failure modes, IRF converges vastly faster than STP. Further, IRF converges far faster than HP's 50-ms claim in all cases.** In fact, the differences between STP and layer-2 IRF are so large that they cannot be compared on the same scale as with other failover mechanisms.

STP convergence times in this test are, if anything, lower than those typically seen in production networks. In this test, with only two sets of interfaces transitioning between forwarding and blocking states, convergence occurred relatively quickly; in production networks with more ports and switches, STP convergence times typically run on the order of 45 to 60 seconds. IRF convergence times in production may be higher as well, although by a far smaller amount than with STP.

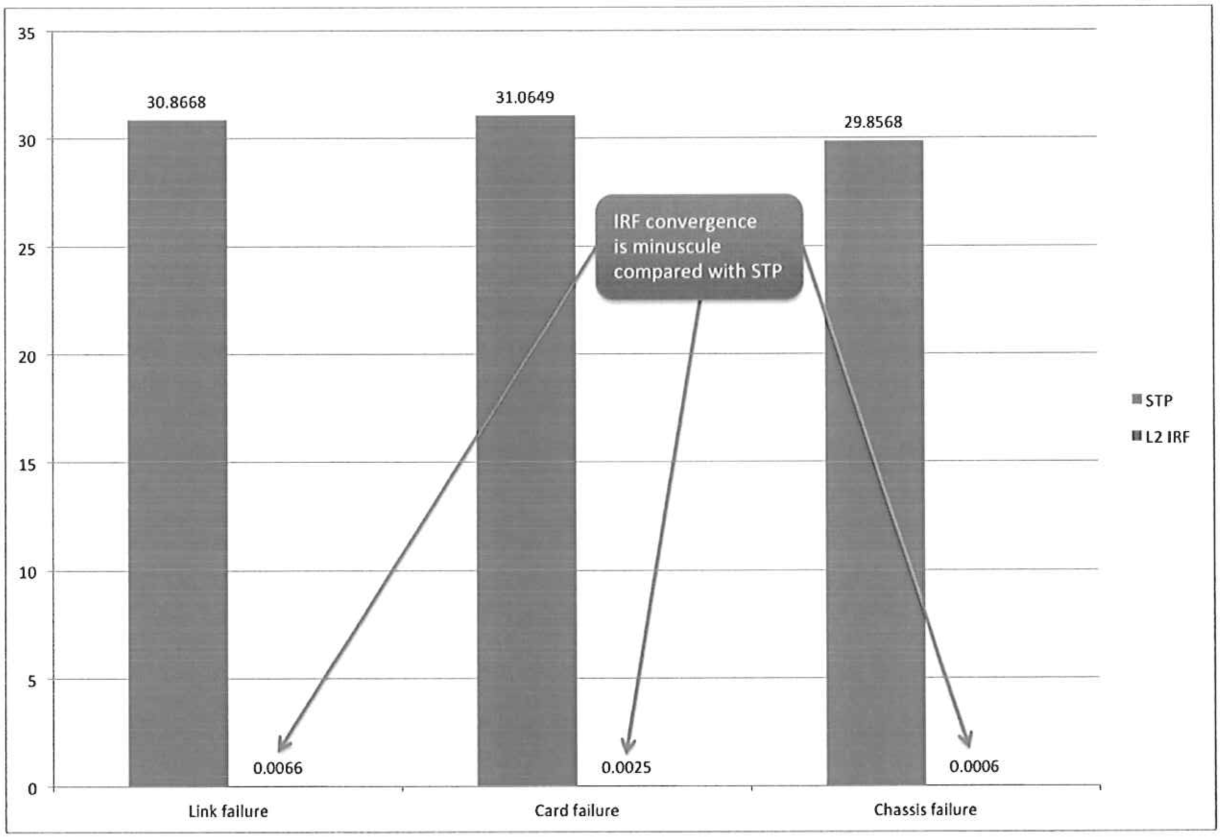

IRF convergence is minuscule compared with STP

Figure 7: STP vs. IRF convergence times

Network Test also compared IRF with rapid spanning tree, the newer and faster mechanism described in IEEE specification 802.1w. **While RSTP converges much faster than STP, it's still no match for IRF in recovering from network outages** (see Figure 8).

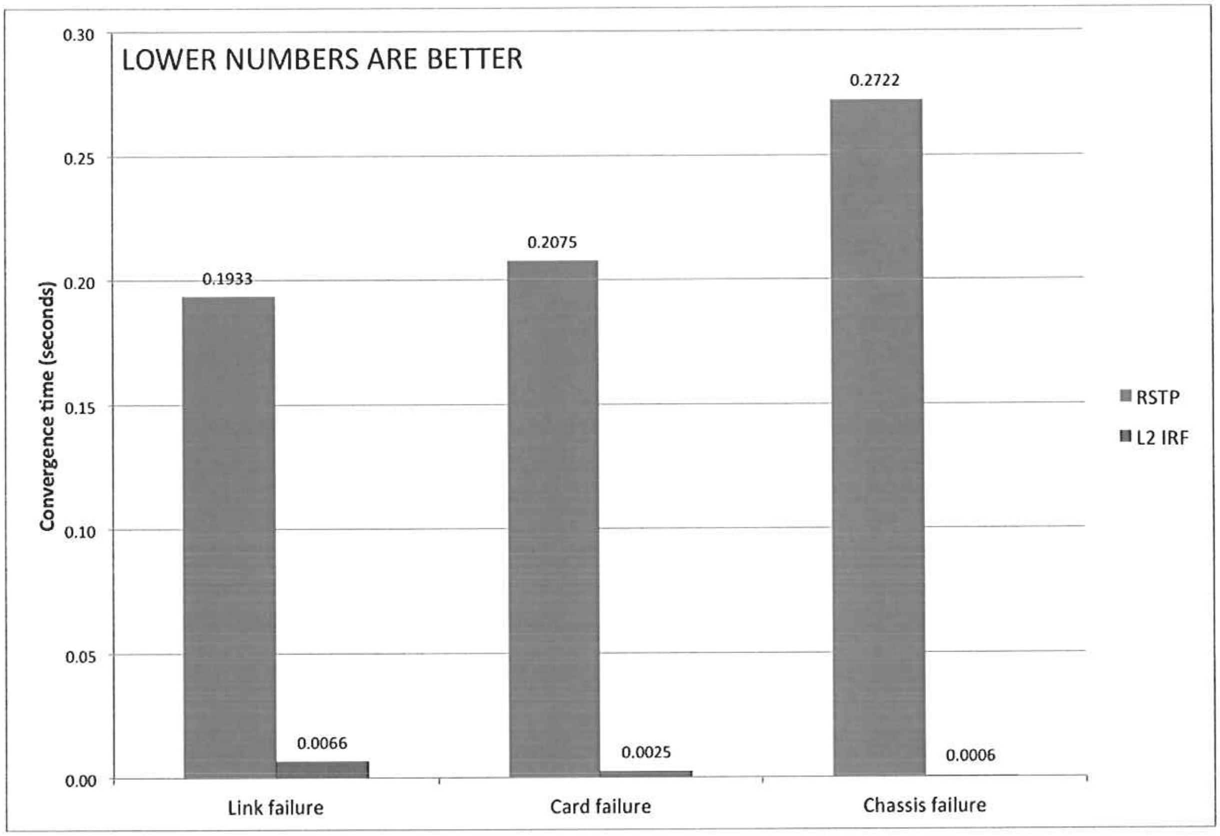

Figure 8: RSTP vs. IRF convergence times

As with STP, the convergence times measured for RSTP were substantially lower than those typically seen on production networks, perhaps due to the small number of links involved. In production settings, RSTP convergence often takes between 1 and 3 seconds following a link or component failure.

The final test compared VRRP and IRF convergence times, with the HP 12500 and HP 5280 both configured in layer-3 mode. In this case, both VRRP and IRF present a single IP address to other devices in the network, and this address migrates to a secondary system when a failure occurs.

Here again, **IRF easily outpaced VRRP when recomputing paths after a network failure** (see Figure 9). VRRP took between 1.9 and 2.2 seconds to converge, compared with times in the single milliseconds or less for IRF.

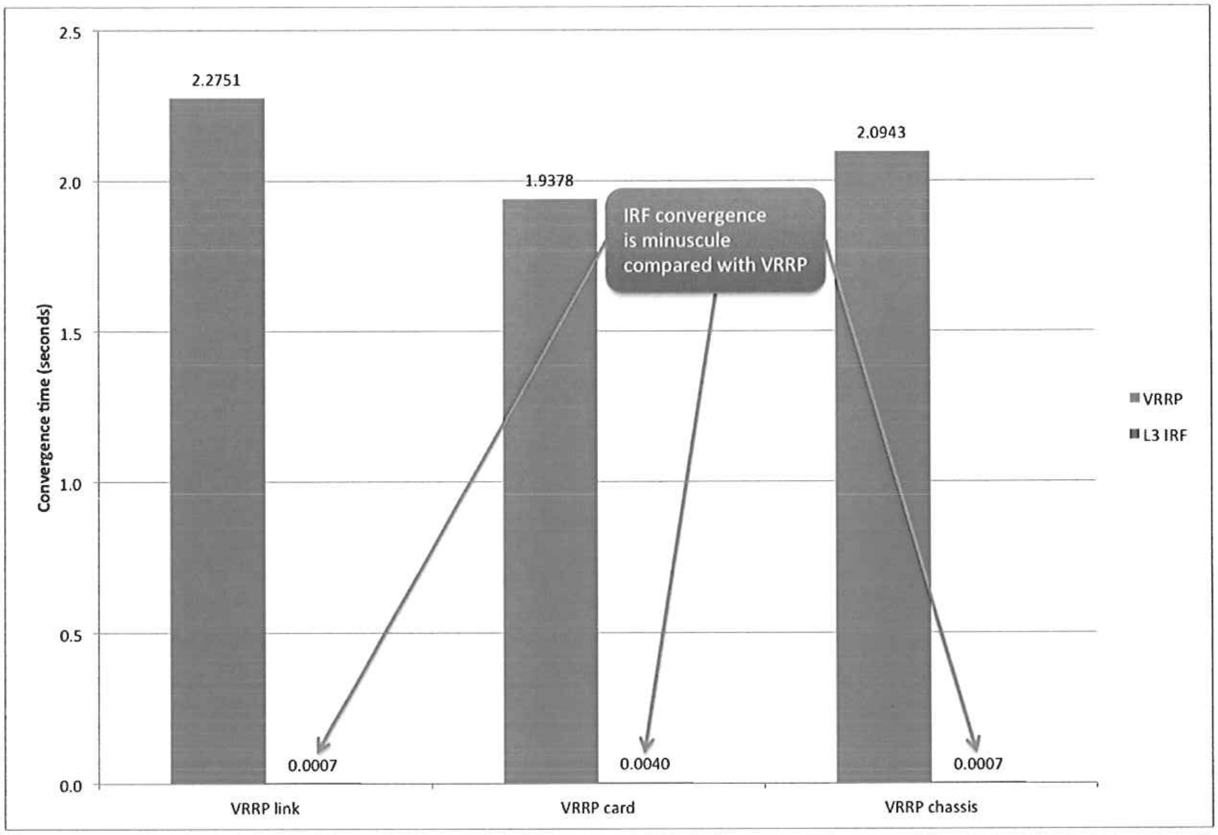

Figure 9: VRRP vs. IRF convergence times

Conclusion

These test results validate IRF's benefits in the areas of network design, performance, and reliability. IRF simplifies network architectures in campus networks and data centers by combining multiple physical switches and presenting them as a single logical fabric to the rest of the network. This approach results in far faster transfer times for virtual machines using VMware vMotion. Performance testing also shows that IRF nearly doubles available bandwidth by virtue of its "active/active" design, compared with "active/passive" designs that tie up switch ports for redundancy. And the results also show huge improvements in convergence times following network failures, both in layer-2 and layer-3 modes, enhancing reliability and improving application performance.

Appendix A: About Network Test

Network Test is an independent third-party test lab and engineering services consultancy. Our core competencies are performance, security, and conformance assessment of networking equipment and live networks. Our clients include equipment manufacturers, large enterprises, service providers, industry consortia, and trade publications.

Appendix B: Software Releases Tested

This appendix describes the software versions used on the test bed. Testing was conducted in July and August 2011 at HP's facilities in Littleton, MA, and Cupertino, CA, USA.

Component	Version
HP 12500	5.20, Release 1335P03
HP 5820	5.20, Release 1211
VMware vSphere	4.1 Update 1
Spirent TestCenter	3.62.0686.0000

Appendix C: Disclaimer

Network Test Inc. has made every attempt to ensure that all test procedures were conducted with the utmost precision and accuracy, but acknowledges that errors do occur. Network Test Inc. shall not be held liable for damages that may result from the use of information contained in this document. All trademarks mentioned in this document are property of their respective owners.

Network Test Inc.
31324 Via Colinas, Suite 113
Westlake Village, CA 91362-6761
USA
+1-818-889-0011
http://networktest.com
info@networktest.com