


## Increasing Capacity and Improving Availability

Ralf Krause

### Link aggregation



#### Link aggregation

- Increasing capacity between switches
- Load sharing
- Static vs. dynamic

*Layer 2 redundancy: STP and RSTP*

*Multiple Spanning Tree*

*HP Switch Meshing*

Rev. 4.11b 2

## Requirements for link aggregation



**Link aggregation is also known as “port trunking” in HP ProCurve documentation**


### Requirements for port trunking:

- HP ProCurve 2500, 2600, 2800, and 4100gl series, and 6108 switches allow up to four links to be aggregated
- HP ProCurve 5300xl series switches allow up to eight links to be aggregated
- The links in a port trunk must:
  - Be coterminal, i.e., they must begin together and end together
  - Support the same mode and flow control options

Rev. 4.11b

3

## Layer 2 conversations



- For load-sharing purposes, a conversation is unidirectional, consisting of a source/destination address (SA/DA) pair
- Traffic between any two hosts consists of two conversations:
  - Transmissions that originate with Server A and are destined for the Backup Server have the SA/DA pair: 3B,56
  - Transmissions that originate with the Backup Server and are destined for Server B have the SA/DA pair: 56,3B

Rev. 4.11b

4

## Conversation-based load sharing



### For outbound traffic destined for opposite side of trunk, the switch:

- Compares each frame's source and destination addresses (S,D) with table entries
- Frames that are part of an existing conversation are forwarded over the link that carried previous frames in that conversation
- For frames that have no matching (S,D) table entry:
  - The switch chooses one of the links in the trunk
  - It creates a new (S,D) table entry indicating the chosen link
  - The frame and all subsequent frames in the conversation are transmitted over the link indicated in the table entry


Rev. 4.11b

5


## Bi-directional load sharing



### Two separate conversations (3B,56 and 2B,56) travel over different links

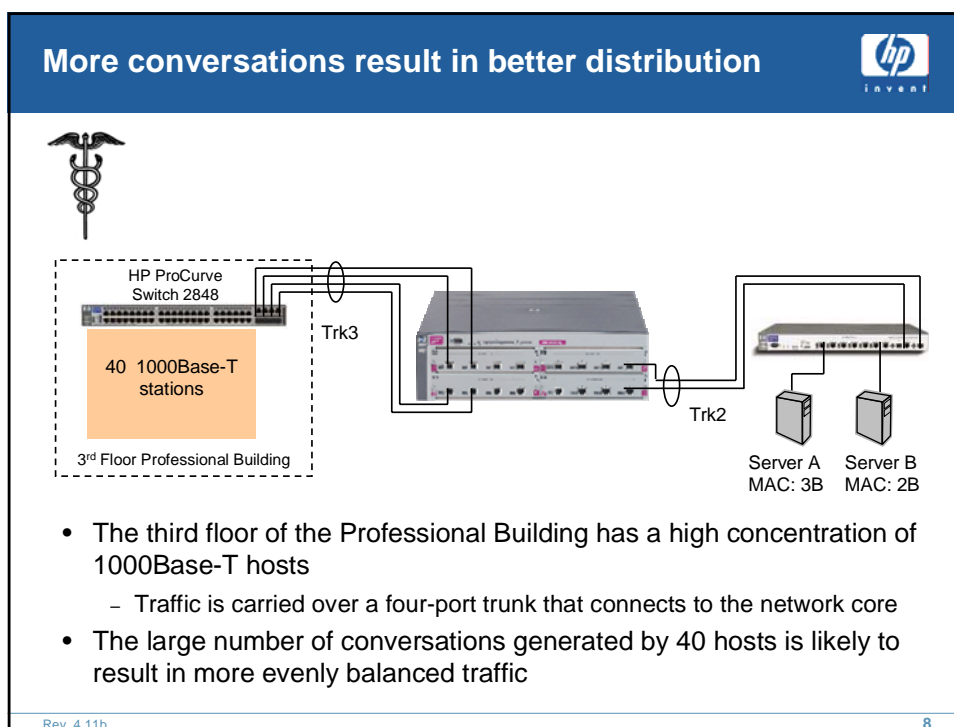
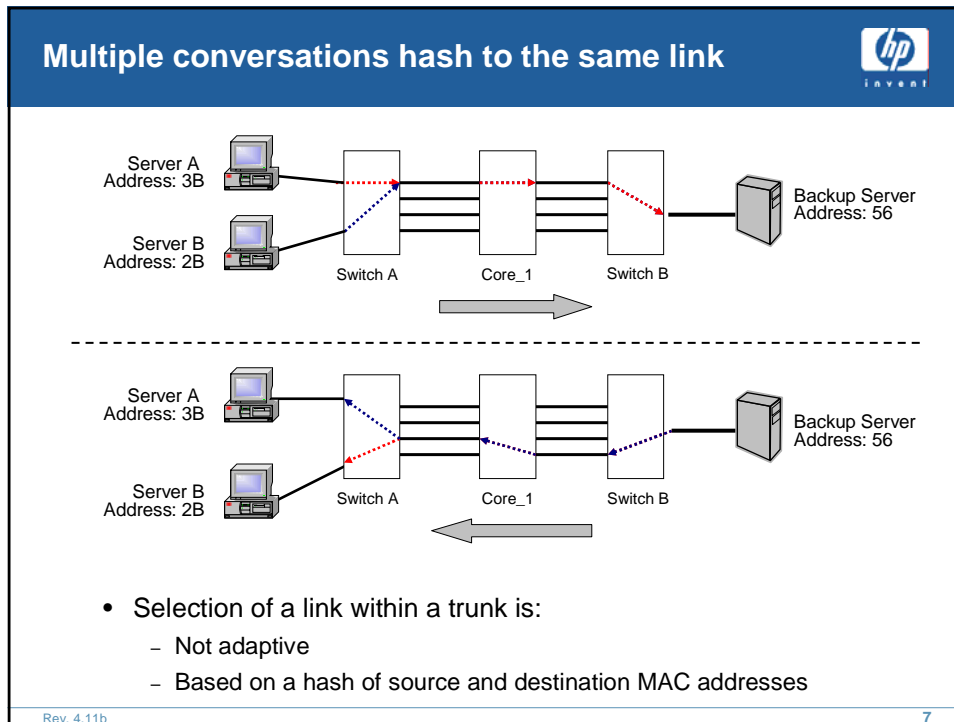


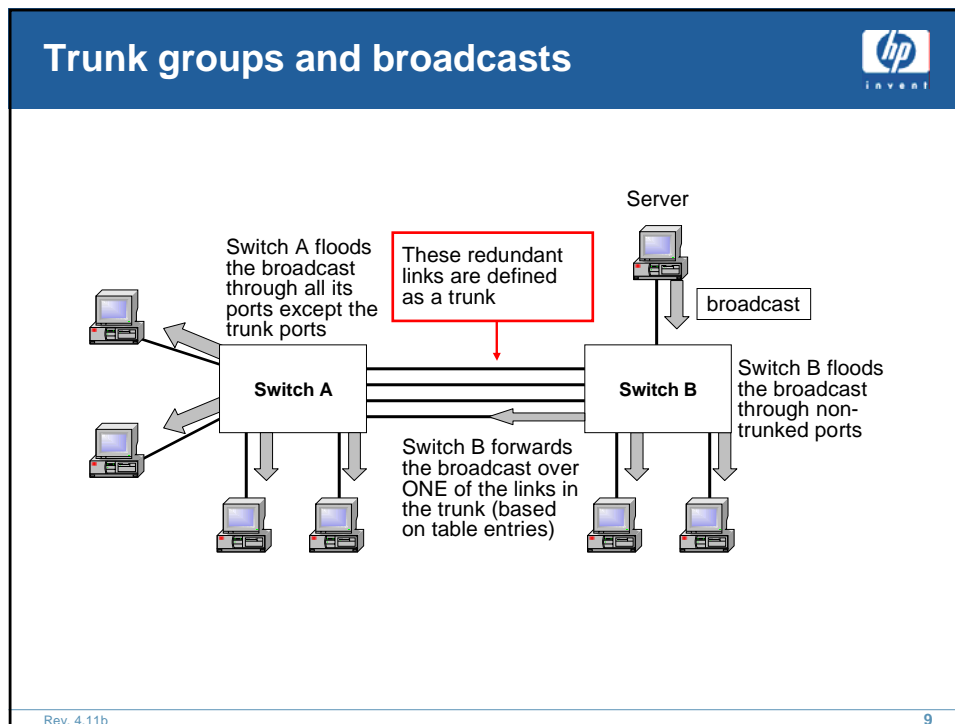
### Two separate conversations (56,3B and 56,2B) travel over different links



Rev. 4.11b

6





## Port trunking methods

### HP Port Trunking

- Does not use a protocol to set up the trunk
- Port trunking is compatible with other trunking methods because it is statically defined

### Fast EtherChannel (FEC)

- FEC is a Cisco standard with widespread compatibility with other switches and multiple-adaptor servers

### Link Aggregation Control Protocol (LACP)

- LACP is defined by IEEE standard 802.3ad
- Both sides may be statically defined; however, LACP also supports a dynamic method for recognizing aggregated links

**All three methods use both source and destination addresses for load sharing**

Rev. 4.11b

10

## HP ProCurve product support for link aggregation



- HP ProCurve 2500 series switches support one trunk group with four ports
- HP ProCurve 2600, 2800, and 4100gl series and 6108 switches can have six trunk groups
  - Each trunk may include up to four ports
- HP ProCurve 5300xl series switches can have 36 trunk groups
  - Each trunk may have up to eight ports
- Port trunking on HP ProCurve 2800 series switches requires that all of the ports in a single trunk must be in the same 12-port group of ports
  - Ports 1-4, for example, may be in the same trunk
  - Ports 1,3,15,16 may not be in the same trunk

Rev. 4.11b

11

## Configuring port trunking




```
Edge_1(config)# trunk ?
[ethernet] PORT-LIST Specify the ports that are to be added/removed from a
trunk
Edge_1(config)# trunk 25,26 ?
trk1 Trunk group 1
trk2 Trunk group 2
...
Edge_1(config)# trunk 25,26 trk1
trunk Do not use any protocol to create or maintain the trunk
fec Use Port Aggregation protocol (PAgP)
lacp Use IEEE 802.1ad Link Aggregation Protocol
<cr>
Edge_1(config)# trunk 25,26 trk1 trunk
Edge_1(config)#
```

```
Core_1(config)# trunk c1,c2 trk1 trunk
```

Rev. 4.11b

12

## Impact of port trunking on VLAN status



**VLAN membership before trunk definition:**

```
Core_1# show run
...
vlan 1
  name "DEFAULT_VLAN"
  untagged 13-26
  exit
vlan 50
  name "VLAN50"
  untagged 1-12
  tagged 25
  ...
```

Before creating a trunk, port 25 is a tagged member of VLAN 50

**Define trunk:**

```
Core_1(config)# trunk 25,26 Trk1 trunk
```

**VLAN membership after trunk definition:**


```
Core_1# show run
...
vlan 1
  name "DEFAULT_VLAN"
  untagged 13-24, Trk1
  exit
vlan 50
  name "VLAN50"
  untagged 1-12
  ...
Core_1(config)# vlan 50 tagged Trk1
...
```

On trunk creation, Trk1 becomes an untagged member of the default VLAN  
Port 25 is no longer associated with VLAN 50

Trk1 must be assigned as a tagged member of VLAN 50

Rev. 4.11b
13

## Displaying trunk status



```
Edge_1(config)# show trunk
```

Load Balancing

Port	Name	Type	Group	Type
25	Core_1 trunk	100/1000T	Trk1	Trunk
26	Core_1 trunk	100/1000T	Trk1	Trunk

```
Edge_1(config)#
```

```
Core_1(config)# show trunk
```

Load Balancing

Port	Name	Type	Group	Type
c1	Edge_1 trunk	100/1000T	Trk1	Trunk
c2	Edge_1 trunk	100/1000T	Trk1	Trunk

```
Edge_1(config)#
```

Rev. 4.11b
14

## Dynamic Link Aggregation Control Protocol



### **Dynamic LACP requires the ports on one side of the trunk to be defined as 'lacp active'**

- The active side sends LACP-BPDUs (bridge protocol data units) across all of the links in the trunk
- LACP-BPDUs contain the switch's MAC address, a priority value, and a port number

### **'Passive' LACP is enabled by default for each port on switches that support LACP:**

- Each port listens for LACP-BPDUs from an active partner
- If two or more ports receive LACP-BPDUs that advertise the same switch MAC address, the switch recognizes a 'dynamic' trunk

Rev. 4.11b

15

## Configurable LACP states



- There are three interface configuration level options for LACP:
  - Passive
    - The default state for each port
    - Listens for LACP-BPDUs from an active LACP partner
  - Active
    - Transmits LACP-BPDUs
  - Disabled
    - Does not send or receive LACP-BPDUs
- HP recommends that you set the LACP state at 'disabled' for every port that will not be using dynamic LACP
- For example, to disable LACP on all ports on the 10/100 module in Slot A:

```
Edge_1(config)# no int a1-a24 lacp
```

Rev. 4.11b

16



## Static vs. dynamic link aggregation



### Static link aggregation:

- **Benefit:** Static trunks appear in the switch's configuration and therefore accept any configuration parameters that would be assigned to a single port
- **Limitation:** Static trunks can have a maximum of four members, with no backup links supported

### Dynamic link aggregation:

- **Benefit:** Dynamic trunks can recognize "standby" links
- **Limitation:** Dynamic trunks always take default interface parameters because they do not appear in the switch's configuration

Rev. 4.11b

17

## Link aggregation summary



### Benefits:

- Increases the capacity of links between switches and links between a switch and a server
- Very fast convergence
- On link failure, conversations assigned to the failed link will be distributed over the remaining links

### Limitations:

- Requires point-to-point coterminous links, that is, they must begin together and end together
- Requires links to operate at the same speed and have the same media type

Rev. 4.11b

18

## Link aggregation support in HP ProCurve switches



Switch Family	# of VLANs	IP routing	# of Trunks	Rapid STP	Mult. STP	Meshing	Traffic classification	PoE	IGMP
2500 series	30	No	one 4-port trunk	Yes	No	No	Preserves existing mark	No	Data driven
2600 series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	Yes	Data driven
2800 series	60	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	No	Data driven
4100gl series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	No	IGMP snooping
5300xl series	256	Yes; dynamic and 16 static	36 trunks; 8 ports per trunk	Yes	Yes	Yes	Full QoS – 802.1p and DiffServ	Yes	Data driven
6108	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	No	Data driven

Rev. 4.11b

19

## Layer 2 redundancy: STP and RSTP



✓ *Link aggregation*

### Layer 2 redundancy: STP and RSTP

- STP and RSTP similarities
- Setting Bridge Priority
- Spanning Tree and VLANs

*Multiple Spanning Tree*

*HP Switch Meshing*

Rev. 4.11b

20

## Layer 2 redundancy options



**All HP ProCurve switches support the following Layer 2 high-availability features:**

- Spanning Tree Protocol (STP)
- Rapid Reconfiguration Spanning Tree Protocol (RSTP)


**HP ProCurve Switch 5300xl series additionally support the following Layer 2 high-availability features:**

- Multiple Spanning Tree Protocol (MSTP)
- HP Switch Meshing

Rev. 4.11b

21

## Spanning Tree review



The Spanning Tree Protocol (STP) automatically:

- Elects one switch to be the root
- Detects loops in the topology
- Uses the lowest cost path to the root

Rev. 4.11b

22

## RSTP and STP similarities



- Rapid Reconfiguration Spanning Tree Protocol
  - IEEE 802.1w standard supersedes 802.1D
- Both STP and RSTP:
  - Use Bridge Priority to elect Root Bridge
  - Use BPDU messages to determine best path to Root Bridge
  - Specify default port costs based on link speed
- RSTP advantages over STP:
  - Faster convergence
  - Ports default to 'edge' state, allowing rapid transition to Forwarding State
- Spanning Tree is off by default on HP ProCurve switches

Rev. 4.11b

23

## Addresses and identifiers



### **Two identifiers play an important role in determining the active path through the bridged network**

- Bridge ID – 64-bit field comprising:
  - 16-bit user definable priority value
  - 48-bit bridge MAC address
- Port ID – 16-bit field comprising:
  - 8-bit user definable priority value
  - 8-bit port number

Rev. 4.11b

24

## Link costs and path costs




- When RSTP is enabled, all ports are assigned a default link cost:
  - Gigabit ports (fiber or copper, up or down) – 20,000
  - 10/100 ports
    - 100 Mbps (and 'down' ports) – 200,000
    - 10 Mbps – 2,000,000
- Root Path Cost
  - The cost of the shortest path between a switch and the Root Bridge
- Root Port
  - The port with the lowest cost path to the Root Bridge
  - If multiple ports on switch have the same lowest cost path, neighbors' Bridge IDs are used as a tie-breaker

Rev. 4.11b

25

## Setting Bridge Priority

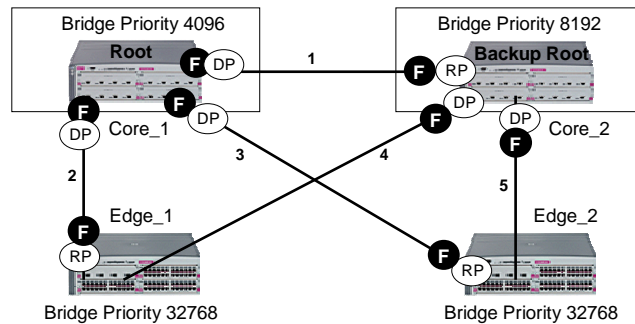


- Bridge Priority for RSTP switches is set in increments of 4096
- To set the Bridge Priority at 4096, enter the following command:  
`Core_1(config)# span priority 1`
- To set the Bridge Priority at 8192:  
`Core_2(config)# span priority 2`
- To return Bridge Priority to default setting of 32768:  
`Core_2(config)# span priority 8`

Rev. 4.11b

26

## The impact of Bridge Priority setting




- All ports on the Core\_1, the Root Bridge, are in Forwarding State
- Other switches select Root Port directly adjacent to Root Bridge
- One switch forwards traffic on behalf of each LAN – the 'Designated Bridge' – which is selected based on Bridge ID
  - Core\_1 is Designated Bridge for Links 1, 2, and 3
  - Core\_2 is Designated Bridge for Links 4 and 5 because it has better priority

Rev. 4.11b

27

## RSTP edge ports



- All ports are edge ports by default when RSTP is enabled, which causes rapid transition to Forwarding State
- Ports that are connected to other switches should NOT be treated as edge ports

```
Core_1(config)# no span b4,c1,d4 edge-port
Core_2(config)# no span a1,c1,c4 edge-port
Edge_1(config)# no span a1,a4 edge-port
Edge_2(config)# no span a1,a4 edge-port
```

Rev. 4.11b

28

## Configuring and enabling RSTP



- Spanning Tree can be configured before it is enabled

- Sample configuration commands:

```
Core_1(config)# span pri 1
Core_1(config)# no span b4, c1, d4 edge
Core_1(config)# span b4 path-cost 40000

Edge_1(config)# span c12 point-to-point-mac force-false
Edge_1(config)# span Trk1 pri 4
```

- Enable Spanning Tree when configuration is complete:

```
Core_1(config)# span
```

- Disabling Spanning Tree ('no span') does not remove existing settings from the configuration

Rev. 4.11b

29

## RSTP interoperability with STP



- When RSTP is enabled, it automatically interoperates with neighboring switches running STP (802.1D)
- If the Spanning Tree protocol version on an HP ProCurve switch should be changed to STP, issue the command:

```
Core_3(config)# span protocol-version stp
```

This command requires a reboot

- To disable rapid state transitions, issue the command:

```
Core_3(config)# span force-version stp-compatible
```

- To enable quicker recognition of STP neighbors, issue the following command for all ports known to be connected to STP switches:

```
Edge_1(config)# no span b18-b20 mcheck
```

- Set consistent costs for all links
  - STP maximum link cost is 65535
  - RSTP maximum link cost is 2 million

Rev. 4.11b

30

## Combining Spanning Tree and VLANs



### When Spanning Tree and VLANs are combined

- Both RSTP and STP standards specify a single Spanning Tree that resolves loops in a bridged network
- Regardless of the number of VLANs in the bridged network, BPDUs are sent untagged
- Redundant links between switches are blocked
  - Network designers must take care not to isolate VLANs



Rev. 4.11b

31

## Poor design can isolate VLANs



- Only three of the six links between the switches are required for full connectivity; three are blocked
- Are any hosts isolated from other hosts in the same VLAN?



Rev. 4.11b



32



## Assigning all VLANs to redundant links



- Assigning all three VLANs to switch-to-switch links assures that all VLANs will be reachable regardless of which links are blocked



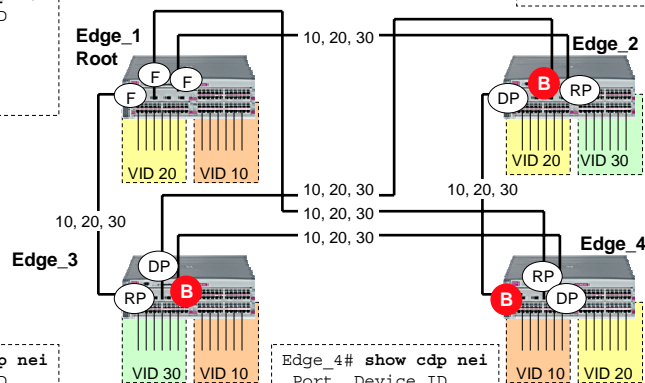
Rev. 4.11b

33

## Spanning Tree and CDP



```
Edge_1# show cdp nei
Port  Device ID
-----
A1    Edge_3
A3    Edge_4
A4    Edge_2
```



```
Edge_2# show cdp nei
Port  Device ID
-----
A3    Edge_3
A4    Edge_1
```

```
Edge_4# show cdp nei
Port  Device ID
-----
A1    Edge_1
A4    Edge_4
```

```
Edge_4# show cdp nei
Port  Device ID
-----
A1    Edge_2
A3    Edge_1
```

Rev. 4.11b

34

## Summary: STP and RSTP



- RSTP provides faster convergence than STP
- Bridge ID and Port ID are significant factors in determining the fastest path through a bridged network
  - Bridge Priority should be configured manually to ensure proper selection of the Root Bridge
  - Port Priority can be configured manually to affect active path selection
- RSTP and STP are interoperable by design
- Designers must ensure that Spanning Tree operation does not isolate VLAN members

Rev. 4.11b

35

## RSTP support on HP ProCurve switches



Switch Family	# of VLANs	IP routing	# of Trunks	Rapid STP	Mult. STP	Meshing	Traffic classification	PoE	IGMP
2500 series	30	No	one 4-port trunk	Yes	No	No	Preserves existing mark	No	Data driven
2600 series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	Yes	Data driven
2800 series	60	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	No	Data driven
4100gl series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	No	IGMP snooping
5300xl series	256	Yes; dynamic and 16 static	36 trunks; 8 ports per trunk	Yes	Yes	Yes	Full QoS – 802.1p and DiffServ	Yes	Data driven
6108	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	No	Port-based, preserves existing mark	No	Data driven

Rev. 4.11b

36

## Multiple Spanning Tree (MST)



- ✓ *Link aggregation*
- ✓ *Layer 2 redundancy: STP and RSTP*

### Multiple Spanning Tree (MST)

- **VLANs and Spanning Tree**
- **Interoperability with STP, RSTP**

*HP Switch Meshing*


Rev. 4.11b

37

## Multiple Spanning Tree



- The first IEEE standard (802.1s) to specify a method for using redundant links in a bridged network
- Allows VLANs to be assigned to different Spanning Tree instances that independently select the lowest cost path through the bridged network



Rev. 4.11b

38

## MST features



- **MST allows for multiple instances of a redundant path for a set of VLANs within the bridged network**
  - Each switch supports up to 16 Spanning Tree instances
  - Each Spanning Tree instance has its own Root Bridge
  - Traffic is balanced across redundant links
- **MST follows the same basic principles as STP and RSTP**
  - Compatible and interoperable with STP and RSTP
  - Emulates STP and RSTP behaviors of when encountering switches that don't support MST

Rev. 4.11b

39

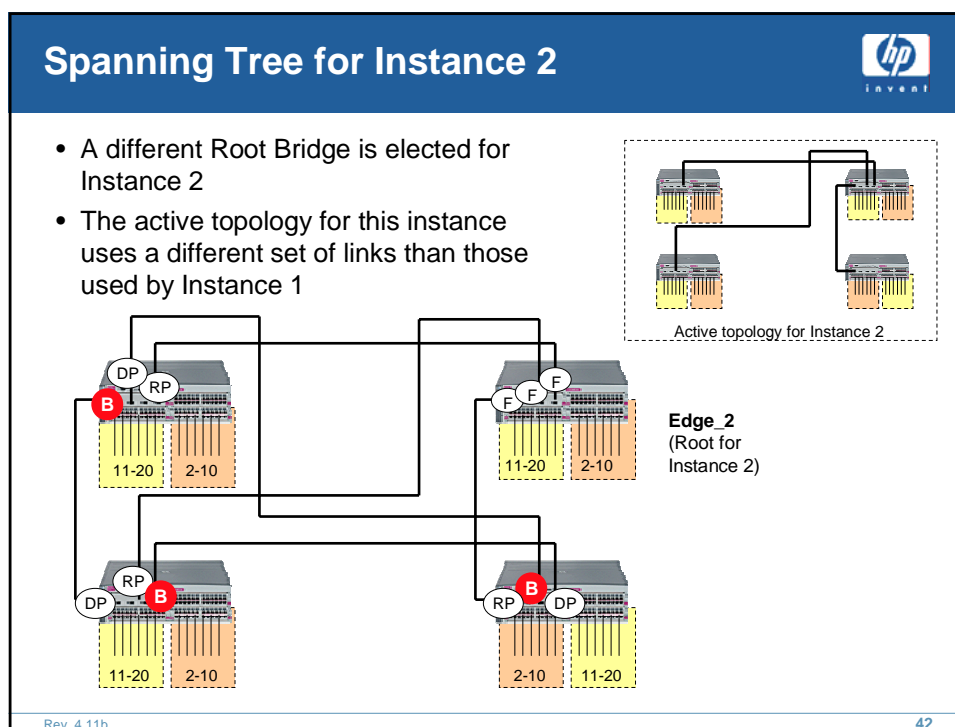
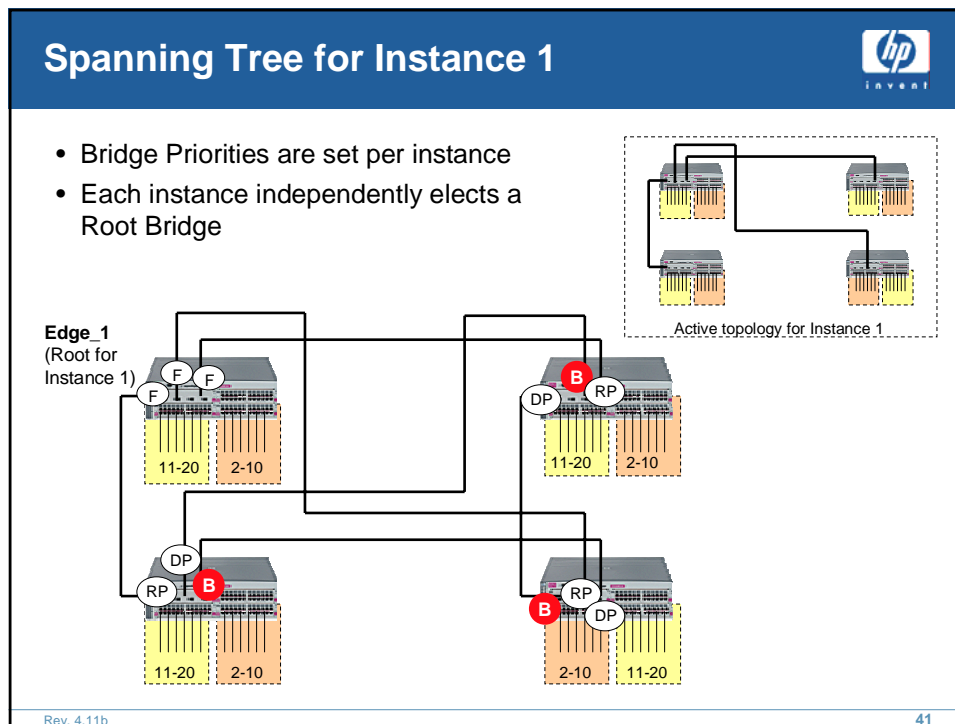
## Comparing RSTP, PVST, and MSTP




	<b>RSTP</b>	<b>PVST or PVST+</b>	<b>MSTP</b>
Relationship between VLAN IDs and Spanning Tree instances	All VIDs map to a common instance	Each VID maps to a separate instance	A configurable set of VIDs is mapped to each instance (range: 2-16)
Use of redundant links	Do not carry any traffic	Primary and backup links can be configured to carry traffic for a <i>different set of instances</i>	Primary and backup links can be configured to carry traffic for <i>different instances</i>
Planning and configuration overhead	Negligible; priorities and path costs are defined <i>once</i>	Priorities and path costs are defined for each VID so that each VID is mapped to a separate instance	Priorities and path costs are defined per instance so that multiple VIDs are mapped to each instance

Rev. 4.11b

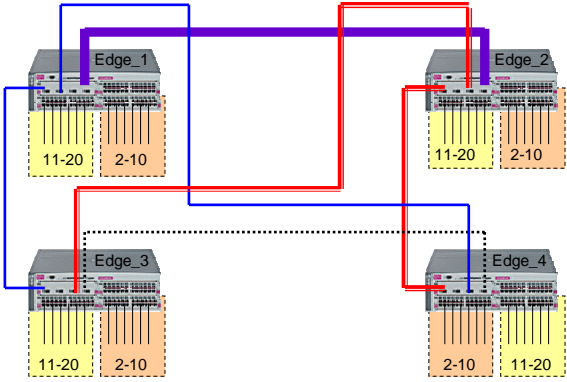
40



## High availability with increased capacity



- Defining multiple Spanning Tree instances results in more efficient network utilization
- The link between the two Root Bridges is active in both topologies



**Legend**

Active links for:

Instance 1      —


Instance 2      —

Both instances   —

Backup link:    .....

Rev. 4.11b 43

## MST regions



- A group of switches that collectively define multiple Spanning Tree instances is known as an 'MST region'
- Each switch can belong to only one region
- All switches in a region must have identical configuration attributes:
  - Alphanumeric configuration name
  - Configuration revision number
  - Associations between VLANs and Spanning-Tree instances
- A switch defines a region boundary if it receives BPDUs from:
  - A switch with different configuration attributes, or
  - An STP or RSTP switch

Rev. 4.11b 44

## MST instances within a single switch

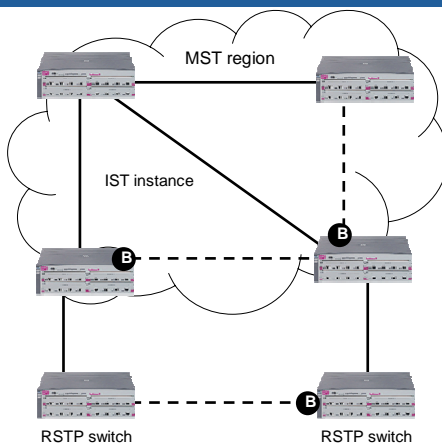


- When MST is initially enabled, the default conditions are as follows:
  - Each switch defines its MAC address as its configuration name and “0” as its configuration revision number
  - All of the VLANs defined on a switch belong to the Internal Spanning Tree (IST) instance
- To cause the switch to interact correctly with other switches in the MST region, you must define common configuration attributes
- Any VLAN not explicitly mapped to a user-defined instance remains associated with the IST. VLAN 1 is often associated with the IST

Rev. 4.11b

45

# The Common Spanning Tree



- Every MST switch supports the IST instance in addition to the user defined instances
- The IST instance within the MST region is a component within the Common Spanning Tree (CST)
- The CST interconnects the IST within MST regions and the Single Spanning Tree (SST) domain defined by STP and RSTP switches

Rev. 4.11b

46

## MST initial configuration



- When Spanning Tree is enabled, the default protocol version will be RSTP. To change the version to MSTP, enter the following command:

```
Edge_1(config)# span protocol-version mstp
Edge_1(config)# write mem
Edge_1(config)# reload
```

- Define an MST region identity for the switch

```
Edge_1(config)# span config-name blue
Edge_1(config)# span config-revision 1
```

- Associate user VLANs with MST instances (MSTI)

```
Edge_1(config)# span instance 1 vlan 2-10
Edge_1(config)# span instance 2 vlan 11-20
```

- Every VLAN belongs to one instance – either a user-defined MSTI or the Internal Spanning Tree (IST)

Rev. 4.11b

47

## Displaying MST configuration parameters



- All switches that *should* be in the same MST Region must have identical parameters
- To see a switch's MST configuration parameters, enter the following command:

```
Edge_1# show span mst-config

MST Configuration Identifier Information

MST Configuration Name: blue
MST Configuration Revision: 1
MST Configuration Digest : 0xE821CCEE7501115289B37C79A72E07C9

IST Mapped VLANs : 1

Instance ID Mapped VLANs
-----
1           2-10
2           11-20
```

- Switches with different configuration parameters will consider themselves to be in different MST Regions

Rev. 4.11b

48



## Additional MST configuration parameters



- Bridge Priority may be defined for each Spanning Tree instance. To define priorities for two MST bridges in the same region:

```
Edge_1(config)# span instance 1 pri 1
Edge_1(config)# span instance 2 pri 2
```

```
Edge_2(config)# span instance 2 pri 1
Edge_2(config)# span instance 1 pri 2
```

- Bridge Priority set for the IST instance:

```
Edge_1(config)# span pri 1
```

- Port priority and path-cost may be defined for each instance (including the IST):

```
Edge_3(config)# span instance 1 a3 pri 1
Edge_3(config)# span instance 2 a3 pri 2
Edge_3(config)# span a3 pri 1
Edge_3(config)# span instance 2 a3 path-cost 100000
```

- After defining parameters, enable spanning tree:

```
Edge_1(config)# span
```

Rev. 4.11b

49

## Implementing MST



### Strategy:

- Maximizing the use of bandwidth is one of the principal goals of MST
- Identify where redundant links exist
- On paper, plan the VLAN mappings for each group
  - Identify a Root Bridge for each instance
  - Determine which links should be in Forwarding State
  - Assign Bridge Priority for the Root Bridge and backup Root Bridge
  - Assign Port Priorities and/or Path Costs as necessary to accomplish your goal

Rev. 4.11b

50

## MST support on HP ProCurve switches



Switch Family	# of VLANs	IP routing	# of Trunks	Rapid STP	Mult. STP	Meshing	Traffic classification	PoE	IGMP
2500 series	30	No	one 4-port trunk	Yes	<b>No</b>	No	Preserves existing mark	No	Data driven
2600 series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	<b>No</b>	No	Port-based, preserves existing mark	Yes	Data driven
2800 series	60	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	<b>No</b>	No	Port-based, preserves existing mark	No	Data driven
4100gl series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	<b>No</b>	No	Port-based, preserves existing mark	No	IGMP snooping
5300xl series	256	Yes; dynamic and 16 static	36 trunks; 8 ports per trunk	Yes	<b>Yes</b>	Yes	Full QoS – 802.1p and DiffServ	Yes	Data driven
6108	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	<b>No</b>	No	Port-based, preserves existing mark	No	Data driven

Rev. 4.11b

51

## Summary: Multiple Spanning Tree (MST)



- MST enables the configuration of VLAN-aware Spanning Trees
  - Under MST, each VLAN is mapped to a single MST instance
  - The state and role of ports in the Spanning Tree can vary according to instance
- MST is compatible and interoperable with STP and RSTP
- The HP ProCurve Switch 5300xl series supports MST

Rev. 4.11b

52

## HP Switch Meshing



- ✓ *Link aggregation*
- ✓ *Layer 2 redundancy: STP and RSTP*
- ✓ *Multiple Spanning Tree (MST)*

### HP Switch Meshing

- Enabling
- HP Switch Meshing domain
- Load balancing
- Design guidelines

Rev. 4.11b


53

## HP Switch Meshing



### HP Switch Meshing is another option for providing Layer 2 redundancy

- Available on HP ProCurve Switch 5300xl series
- Enables redundant links to be used rather than blocked



Rev. 4.11b

54


## Enable HP Switch Meshing



- Switch meshing is enabled on a per-port basis – **only** on point-to-point links with other meshed ports. For example:

```
Edge_3(config)# mesh a1,a2
```

- The links in a mesh can be made up of any combination of speeds and media types
- When meshing is added to or removed from ports, switches must be rebooted




Rev. 4.11b

55

## Initialize HP Switch Mesh ports



- When switch meshing is enabled on a port, the switch initially places the port in a Blocking State and sends Switch Meshing Protocol packets
- If the port connected to a meshed port does not respond with Switch Meshing Protocol packets, the port remains blocked
- If the connected port responds with Switch Meshing Protocol packets, the two begin negotiating a peer relationship



```
Edge_3(config)# show mesh
```

Status and Counters - Switch Mesh Information

Port	State	Adjacent Switch	Peer Port
A1	Established	001083-cb2100	001083-cb31cb
A2	Established	c001e7-932800	c001e7-9318c9

```
Edge_3(config)#
```


Rev. 4.11b

56

## Switch mesh domain



- A group of meshed switch ports exchanging meshing protocol packets is called a "switch mesh domain"
- A switch mesh domain can contain up to 12 switches. Each switch can have up to 24 meshed ports
- An "edge" switch has some mesh ports and some non-meshed ports. Switches 1-5 are edge switches



Rev. 4.11b

57

## Conversation-based load balancing



### Determining lowest cost path

- When the mesh is fully initialized, each path through the mesh is assigned a cost based on link speed, outbound and inbound queue depths, and packet drop counts
- Costs are recalculated every 30 seconds
  - At any given moment, one path is considered the lowest cost path


### Forwarding decisions

- Frames that are part of a new conversation are forwarded over the current lowest cost path
- Frames that are part of an established conversation are forwarded through the same port as the first frame in that conversation

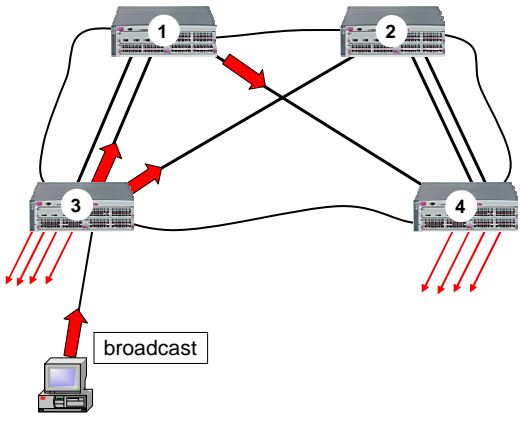
Rev. 4.11b

58

## Flooded traffic within the mesh domain




- Edge switches prevent broadcast storms by defining a "broadcast path"
  - This assures that only one copy of every broadcast frame reaches the other edge switches
- Each edge switch independently defines a loop-free broadcast path through the mesh
- Rather than flooding frames destined for *unknown unicast* addresses, mesh edge switches use a query/response mechanism to learn their locations

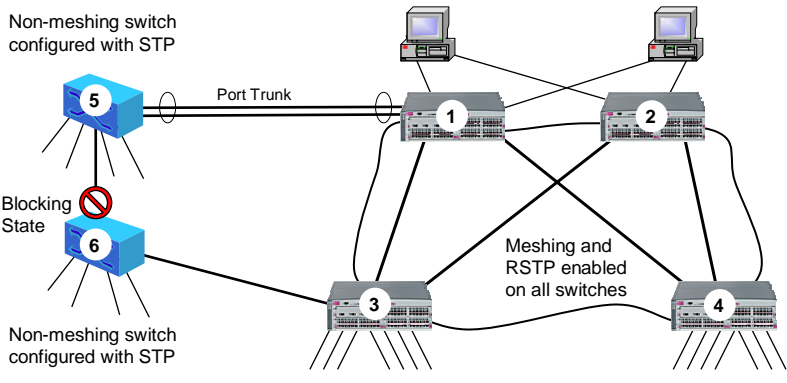


Rev. 4.11b 59

## Switch Meshing compatibility with STP and RSTP



- To interoperate with non-meshing switches within the Layer 2 domain, enable STP or RSTP on meshed switches
- The mesh appears to non-meshing switches STP/RSTP switches as a single switch

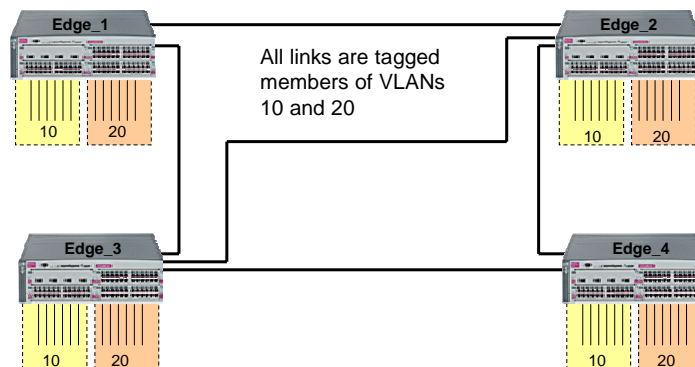


Rev. 4.11b 60

## Meshing and VLANs



- Ports on which meshing has been configured are collectively referred to in the configuration as 'Mesh'
- The Mesh is automatically made a tagged member of all user-defined VLANs on the switch, immediately enabling the included links to carry traffic for all VLANs



Rev. 4.11b


61

## Meshing and IP routing at the Medical Center




- Although the Mesh is automatically made a member of all VLANs, a meshed switch *cannot* perform IP forwarding between VLANs
- The switch that performs IP routing cannot be meshed


**IP routing enabled**  
**RSTP enabled**  
No meshing allowed




**Meshing enabled**  
**RSTP enabled**



**Meshing enabled**  
**RSTP enabled**



**Meshing enabled**  
**RSTP enabled**



Meshed link

Meshed link

Each meshed link is a tagged member of all user VLANs

Meshed link

Rev. 4.11b

62

## HP Switch Meshing design guidelines



- A mesh consists of up to 12 HP ProCurve 5300xl series switches
  - Each switch can belong to only one mesh
  - A switch can have up to 24 meshed ports using any combination of media types and link speeds
  - Meshing and IP routing cannot simultaneously be enabled on the same switch
- Meshing is enabled per port
  - Enable only on ports that directly connect to other meshed ports
  - A port cannot simultaneously be configured for trunking and meshing
- The maximum hop count through the mesh is 5 switches
- HP Switch Meshing supports full mesh and partial mesh topologies

Rev. 4.11b

63

## Summary: HP Switch Meshing



- HP Switch Meshing can be used to improve availability while increasing capacity within a Layer 2 switched network
- HP Switch Meshing is similar to the Spanning Tree Protocol in that it allows designers to create topologies that contain redundant paths
  - HP Switch Meshing deals with redundant links in a more intelligent way than STP or RSTP
  - Instead of placing redundant links in the Blocking state, switches using HP Switch Meshing can use all available links to forward traffic
- The operation of HP Switch Meshing is transparent to non-meshing devices

Rev. 4.11b

64



## HP Switch Meshing support on HP ProCurve switches



Switch Family	# of VLANs	IP routing	# of Trunks	Rapid STP	Mult. STP	Meshing	Traffic classification	PoE	IGMP
2500 series	30	No	one 4-port trunk	Yes	No	<b>No</b>	Preserves existing mark	No	Data driven
2600 series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	<b>No</b>	Port-based, preserves existing mark	Yes	Data driven
2800 series	60	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	<b>No</b>	Port-based, preserves existing mark	No	Data driven
4100gl series	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	<b>No</b>	Port-based, preserves existing mark	No	IGMP snooping
5300xl series	256	Yes; dynamic and 16 static	36 trunks; 8 ports per trunk	Yes	Yes	<b>Yes</b>	Full QoS – 802.1p and DiffServ	Yes	Data driven
6108	30	Yes; 16 static routes	6 trunks; 4 ports per trunk	Yes	No	<b>No</b>	Port-based, preserves existing mark	No	Data driven

Rev. 4.11b

65

## Summary



**Several technologies exist to increase the availability and capacity of switched networks, including:**

- Link aggregation
  - HP Port Trunking
  - Fast EtherChannel (FEC)
  - Link Aggregation Control Protocol (LACP)
- Spanning Tree Protocol and Rapid Reconfiguration Spanning Tree Protocol
- Multiple Spanning Tree Protocol
- HP Switch Meshing

Rev. 4.11b

66



**Questions**  
**?**

