

Cisco Data Center Core: Configuring Networking

APPLYING SWITCHING PROTOCOLS



Sean Douglas

DATA CENTER ENGINEER

@ocdlearning



Overview



Data centers require high availability, redundancy, and load balancing

Examine layer 2 protocols that provide loop-free redundant links


- STP
- Port channels
- vPCs



STP



Rapid PVST+



Each VLAN has a separate instance. Enabled by default on VLAN 1



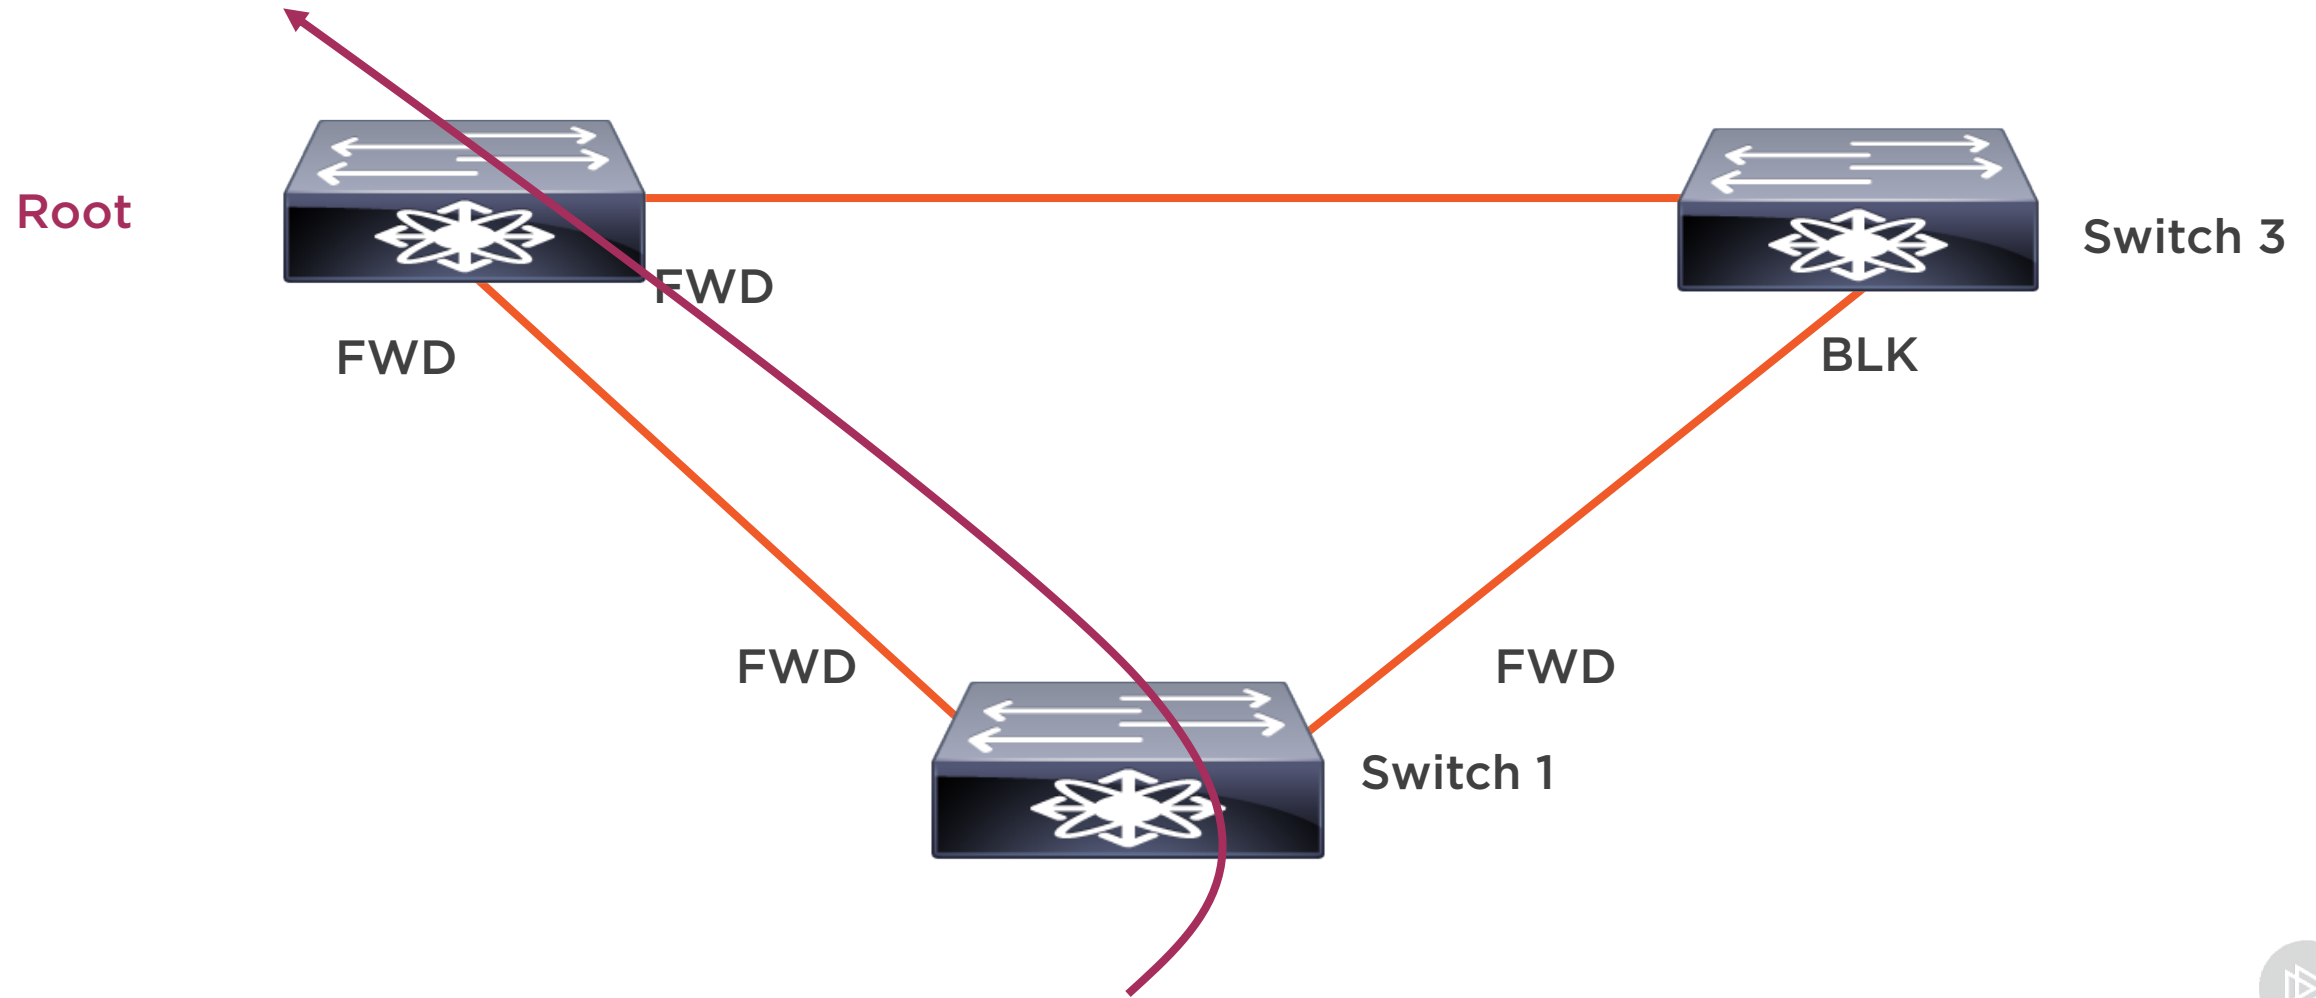
Evolution of the IEEE 801.d standard. Rapid PVST+ is the IEEE 802.1w standard. Offers fast convergence




Default STP on Cisco Nexus Switches




Rapid PVST+



Rapid PVST+



Rapid PVST+



IEEE 802.1D




**Nexus switches
do not run
802.1D**



IEEE 802.1D




**Nexus switches
do not run
802.1D**



**Nexus switches
are 802.1D
compatible**

Interoperability with 802.1d



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

```
spanning-tree vlan 1,2,4,6,8 root primary
```

```
spanning-tree vlan 3,5,7,9,11 root secondary
```

Rapid PVST+

Enabled by default. If modified, to replace Rapid-PVST+ use the following command



STP Enhancements



STP Edge Port



Access ports only

Bypasses listening and learning

Forwards immediately

Deployed on L2 switches connected to server, PC, or phone

No received BPDUs expected

No topology changes when link state changes

```
Interface Ethernet 2/1
```


```
spanning-tree port type edge
```

STP Edge Port

Configure STP edge port on an interface



STP Edge Port with BPDU Guard



```
interface Ethernet 2/1  
spanning-tree bpduguard enable
```

BPDU Guard

Enable BPDU guard on an interface



**BPDUs are sent out all
ports**

Used in special cases

Edge ports only

Workaround

Use cautiously

BPDUs Filter



```
interface Ethernet 2/1
```

```
spanning-tree bpdupfilter enable
```

BPDU Filter

Enable BPDU filter on an interface



STP Root Guard



If bridge ID is lowered, an access switch could become the root bridge

Older switch added to network could become the root bridge


Access switches should never be the root

Root guard limits ports which the root may exist

Superior BPDUs cause root-inconsistent state



STP Root Guard



```
interface Ethernet 3/1  
spanning-tree guard root
```

Root Guard

Enable root guard on an interface



STP Loop Guard



When a port no longer receives BPDUs, STP considers that the segment is loop-free


Blocking port transitions to forwarding

Can create a loop

Loop guard checks

Not receiving any BPDUs on protected port will cause loop-inconsistent state

STP Loop Guard



```
interface Ethernet 3/1  
spanning-tree guard loop
```

Loop Guard

Enable loop guard on an interface



LACP



Port Channels



Bundling ports adds more bandwidth and creates load balancing capability



Layer 2 - STP blocks redundant ports to prevent loops




Layer 3 - Routing protocols requires a routing adjacency to be formed for every link, which increases routing protocol overhead




Port channels aggregate multiple physical links into a single logical link between two switches

STP Prevents Redundant Links

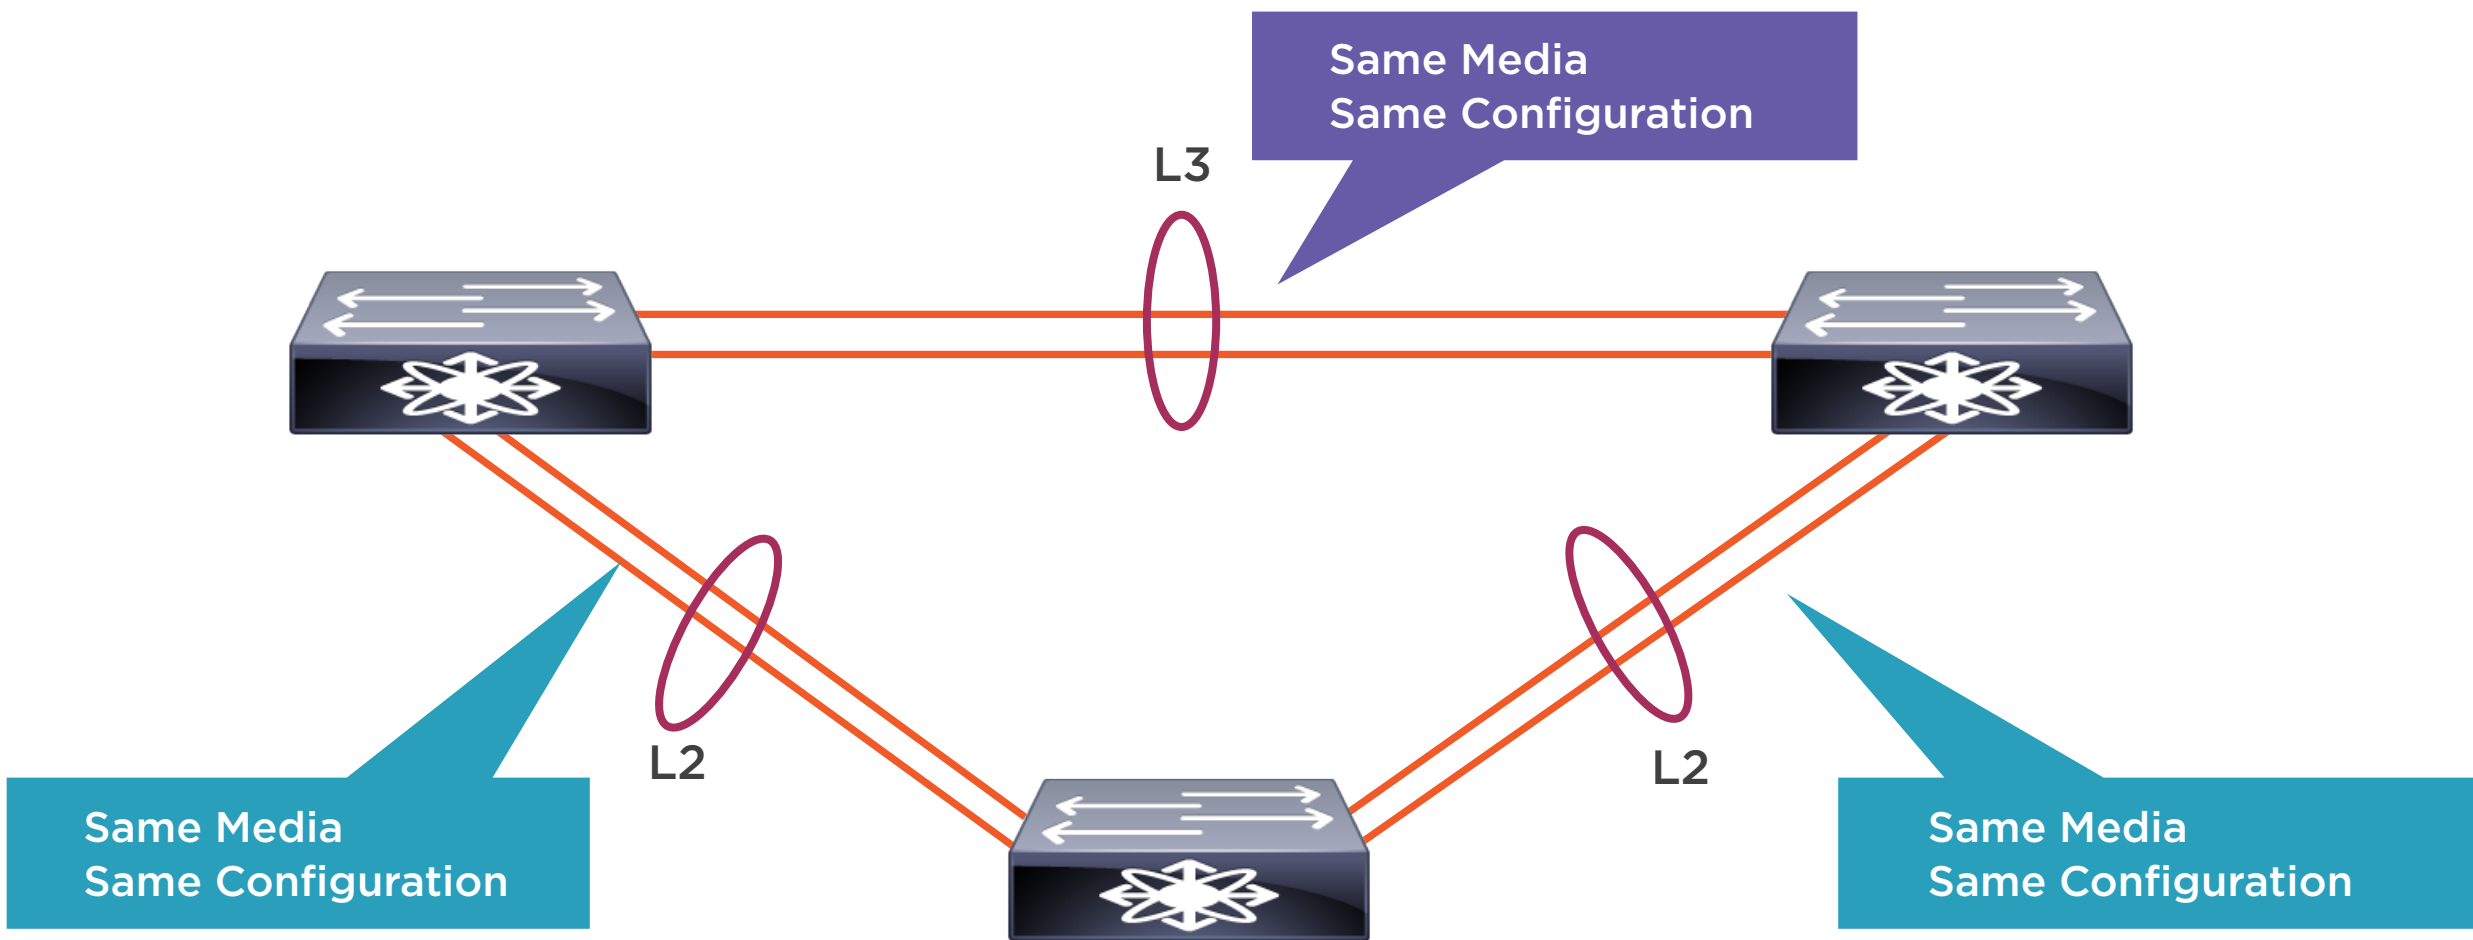
STP blocks additional ports to prevent loops



Port Channels



EtherChannel - L2 and L3




EtherChannel Views



Physical View

Multiple ports are defined as being part of an EtherChannel group



Logical View

Subsystems running on the switch see only once logical link per EtherChannel

```
feature lacp
!  
interface type ethernet 2/1  
channel-group number mode { active | on | passive }
```

Enabling LACP

Enable LACP globally, enable LACP for each channel by setting the channel mode for each interface to **active** or **passive**

Configure either channel mode for individual links in the LACP channel group



Examine Port Channel

```
show port-channel summary
```

Flags: D - Down P - Up in port-channel (members)

I - Individual H - Hot-standby (LACP only)

s - Suspended r - Module-removed

S - Switched R - Routed


U - Up (port-channel)

M - Not in use. Min-links not met

Group	Port-Channel -	Type	Protocol	Member Ports
1	Po1(SU)	Eth	LACP	Eth2/1(P) Eth2/2(P)



Port Channel Load Balancing



Source and/or
destination MAC
address at Layer 2

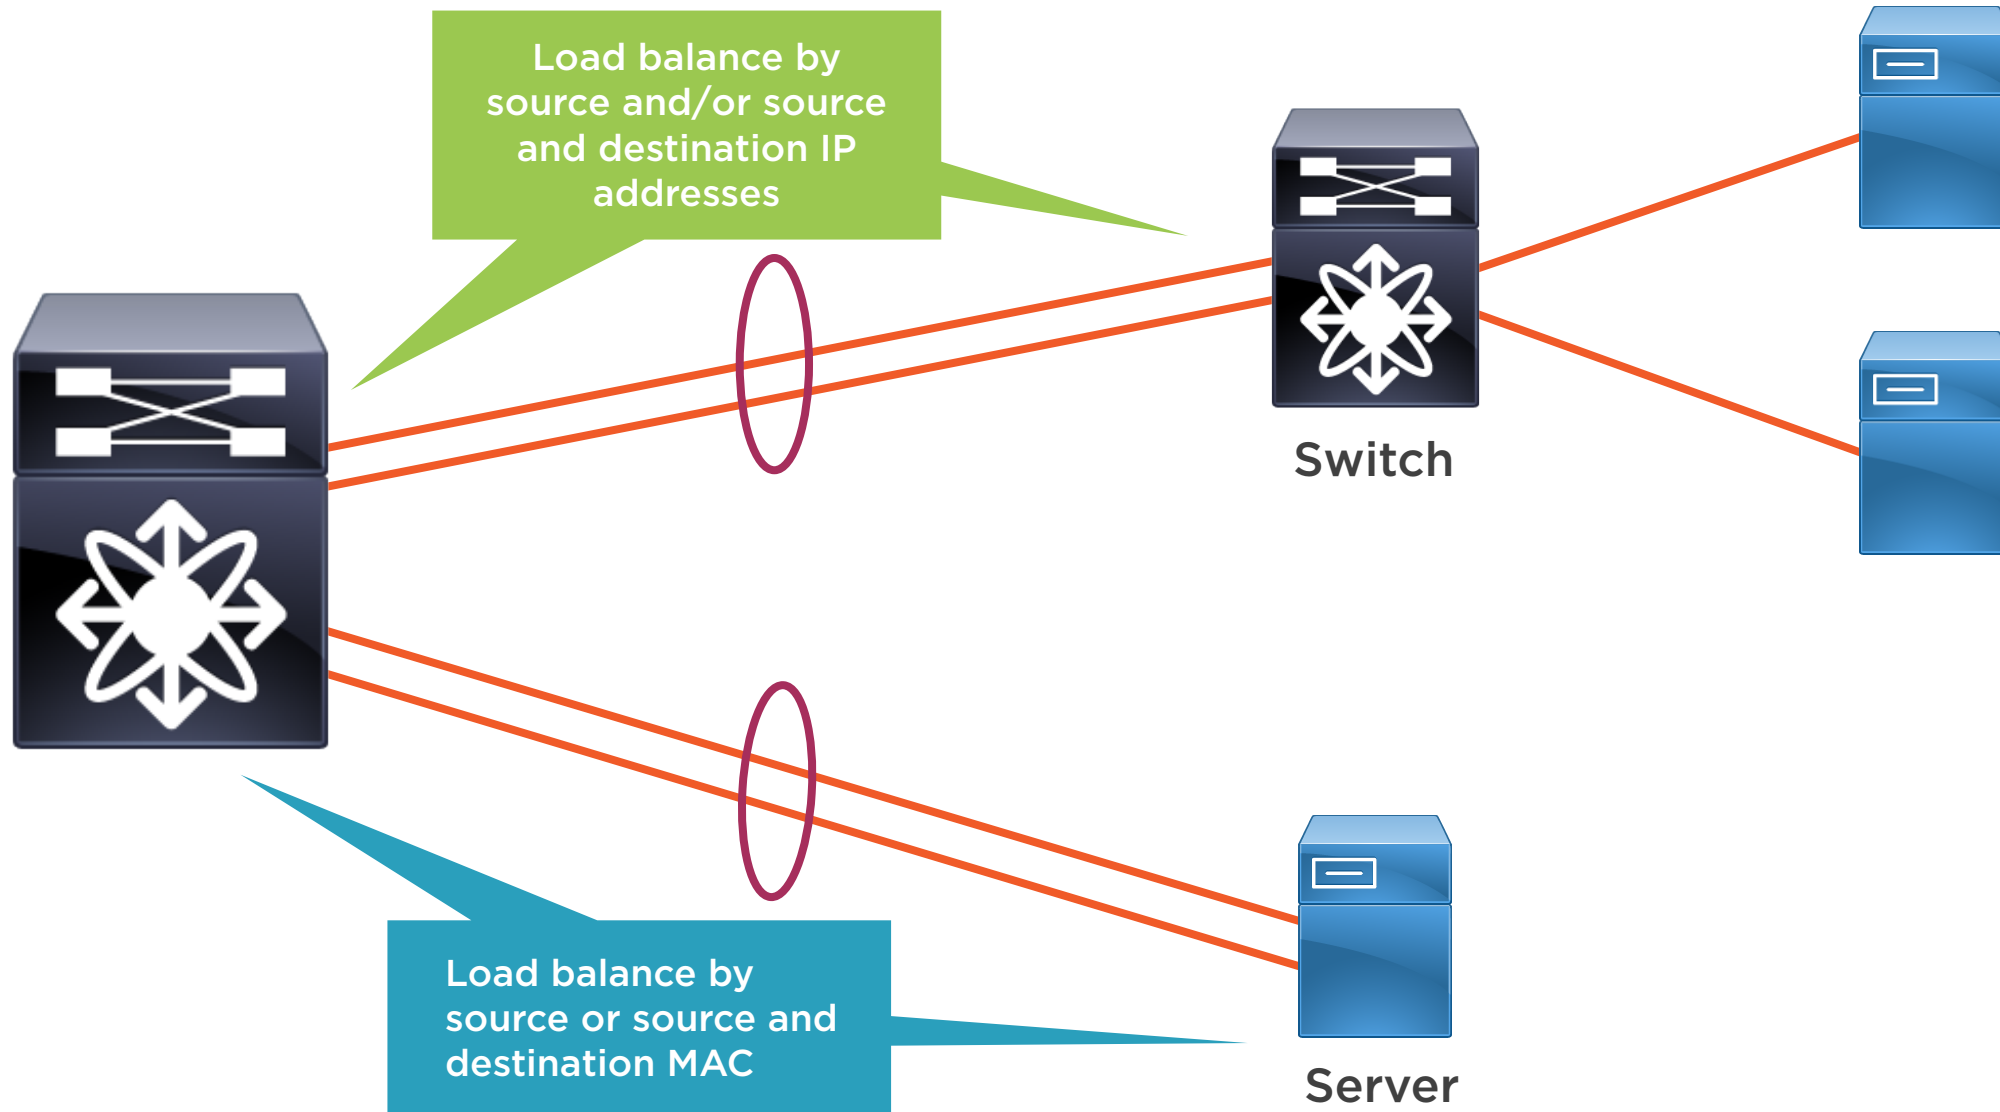
Source and/or
destination IP
addresses at Layer 3

Source or source and
destination TCP or
UDP ports at Layer 4

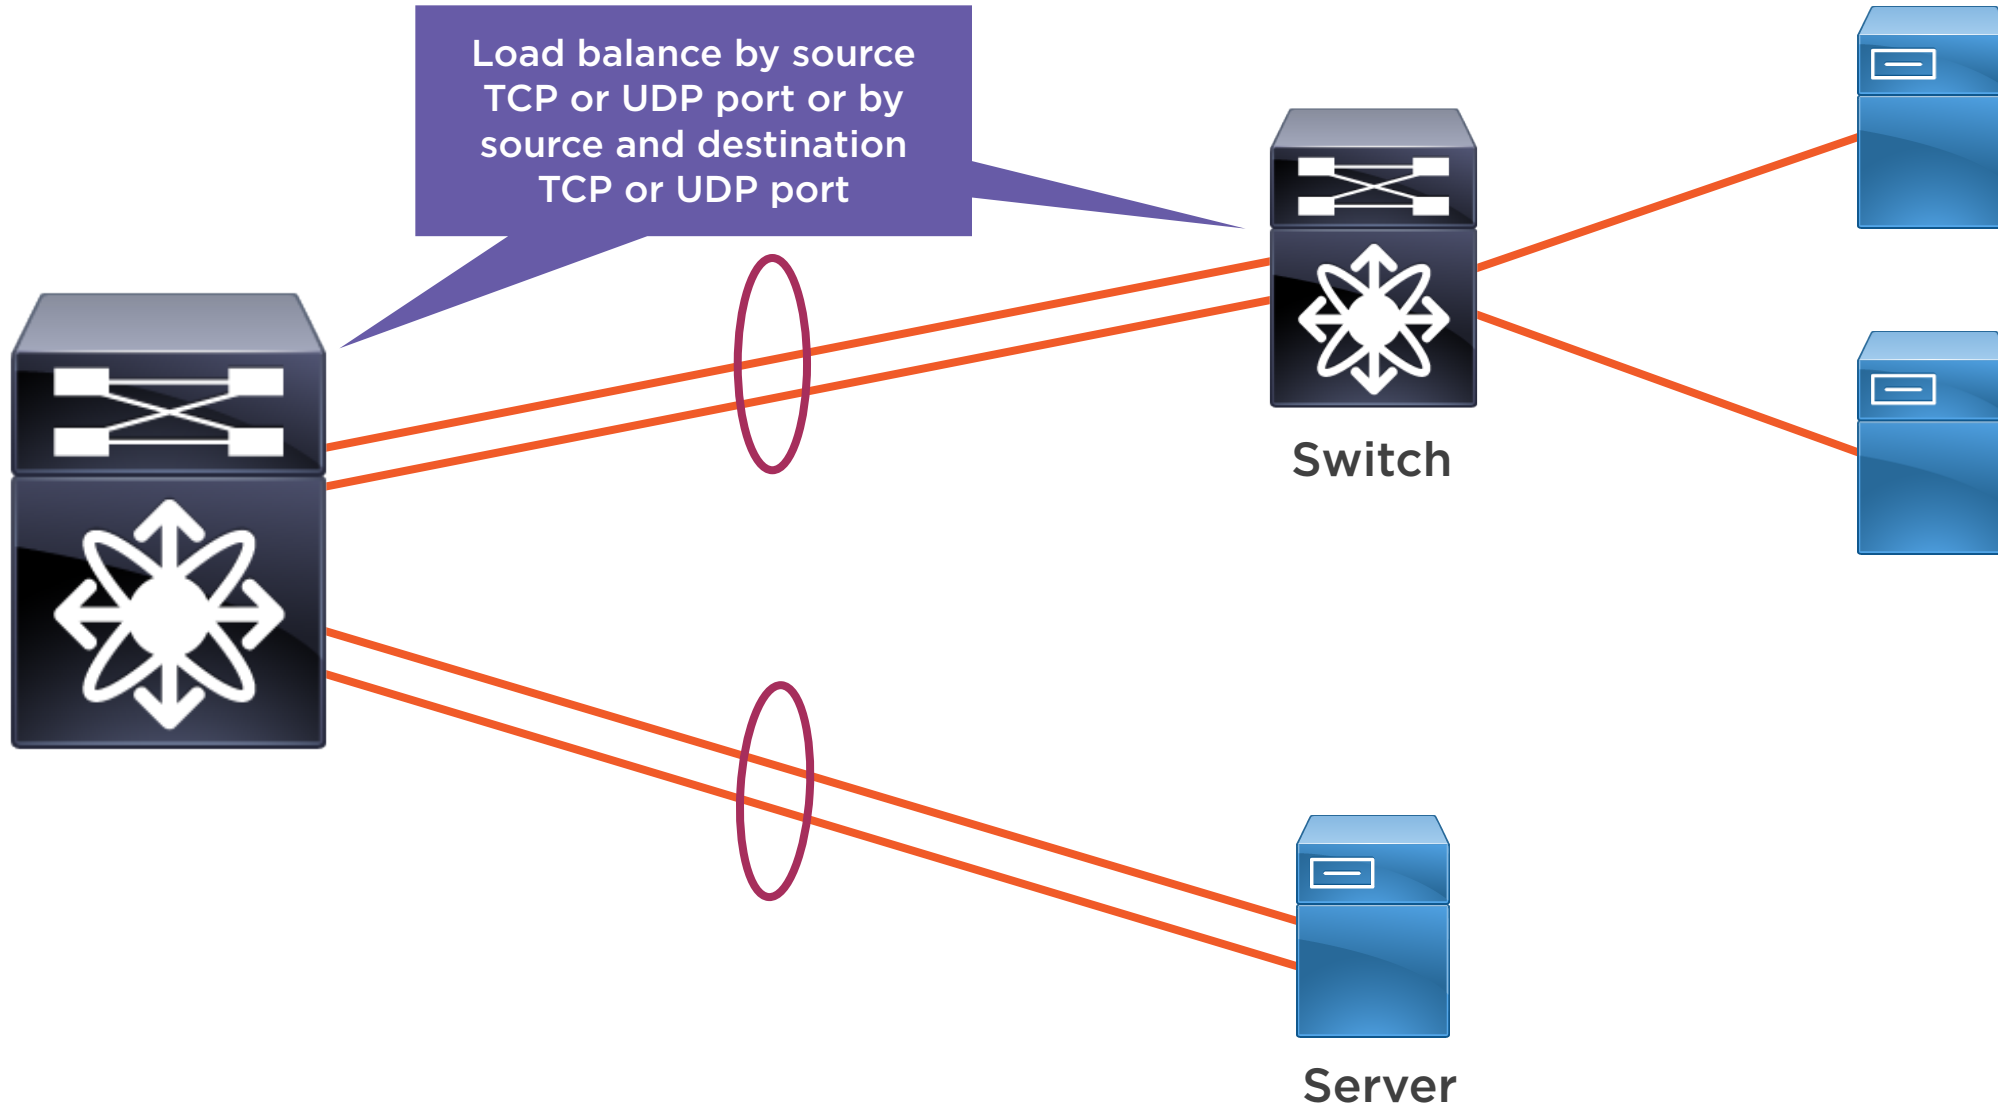
Load Balance Options



Port Channel Load Balancing



Port Channel Load Balancing



```
show port-channel load-balance
```

System config:

Non-IP: src-dst mac

IP: src-dst ip

Port Channel Load Balance


Examine the switch load balancing options enabled for the port channel




vPC



vPC Domain




vPC Domain




vPC Switches

Primary vPC generates BDPUs
using its Bridge ID

Secondary vPC relays BDPUs,
doesn't generate them



vPC Domain



vPC Advantages



Allows a single device to use a port channel across two upstream devices



Loop-free topology—both uplinks are active. No STP blocked ports



Uses all available uplink bandwidth



Downstream device experiences vPC peers as if they are a single device



Ensures high-availability and link-level resiliency



Cisco Fabric Services and vPC



CFS over Ethernet synchronizes MAC address table entries



CFS over Ethernet synchronizes IGMP snooping entries



CFS ensures configuration consistency between vPC peer switches




CFS tracks the vPC status on the peer



CFS synchronizes ARP tables



vPC Domain



vPC Configuration

```
feature vpc
```

```
!
```

```
vpc domain 10
```

```
peer-keepalive destination 10.10.10.11 source 10.10.10.10 vrf vPC-VRF
```

```
!
```

```
interface port-channel 1
```

```
vpc peer-link
```

```
!
```

```
interface port-channel 2
```


```
vpc 10
```



vPC at Multiple Layers

Nexus
5000 / 9000

Nexus
5000 / 9000



vPC Guidelines



Switch type must be the same platform (not 5600 to 9300 vPC)



Only two switches in a vPC and only a single vPC domain per switch




Avoid running vPC keepalive over vPC peer link




The vPC peer link requires a minimum 10 Gbps Ethernet port



vPC Peer-Switch



STP Logical Topology



```
Switch(config-vpc-domain)# peer-switch
```

vPC Peer-Switch

Activate the vPC peer-switch capability on both vPC peer devices



Summary



Data centers require high availability, redundancy, and load balancing

Examine layer 2 protocols that provide loop-free redundant links

- STP
- Port channels
- vPCs

