

Small and Medium-sized Data Center Stack+BGP Configuration Guide

Models: Full L3 Networking

Contents

1. Solution Introduction	1
1.1 Background.....	1
1.2 Networking Introduction.....	1
2. Network Topology	1
3. Device Connection Planning Table	2
4. Configuring Ideas	2
5. Configuring Steps	3
5.1 Spine and Leaf Configuring Stacking.....	3
5.2 Basic Configurations of Spine and Leaf.....	4
5.3 Configuring BFD of Spine and Leaf.....	5
5.4 Interconnection port & IP address of Spine, Leaf.....	6
5.5 BGP Configuring Routing Protocol BGP.....	7
5.6 Configuring Downlink Business Port of Leaf.....	8
6. Configuring Review	9
6.1 Stack Status Review.....	9
6.2 BGP neighbor, routing table configuration check.....	10

1. Solution Introduction

1.1 Background

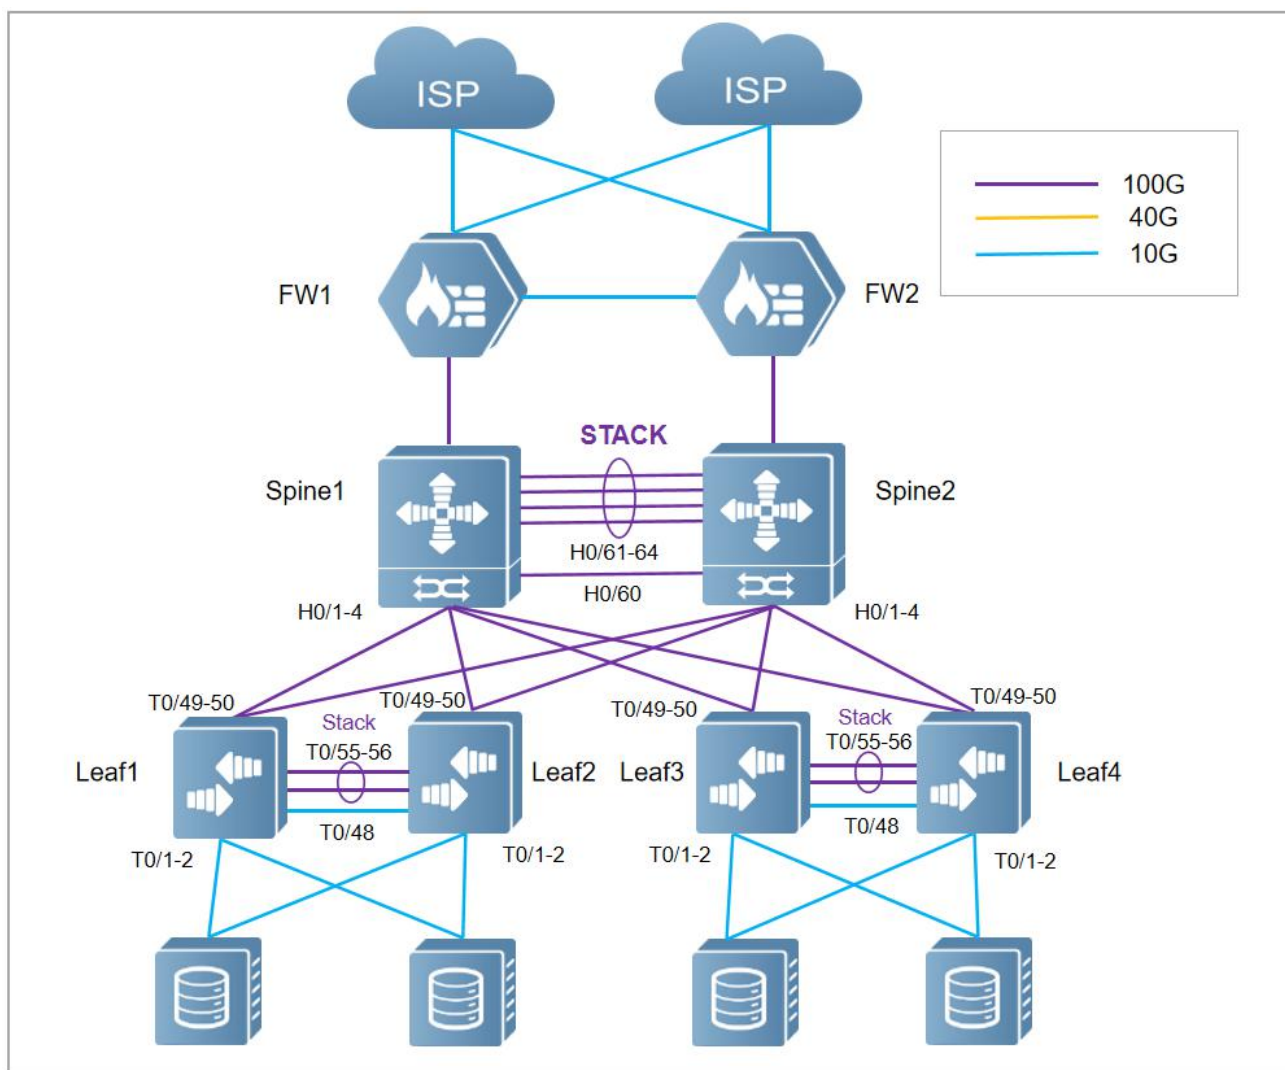
The next-generation data center and cloud computing have more urgent requirements for the non-blocking transmission of large-flow data. The network is evolving toward the trend of "east and west", fast switching- failover time for link failure is within 50 to 200ms, devices virtualization, Spine-Leaf full L3 network architecture, etc, will be the standard requirements of data center switches.

1.2 Networking Introduction

Considering the reliability and stability of small and medium-sized data centers, the design ideas are as follows:

- The Spine layer adopts 2 units N8560-64C stacking, enable redundancy by the active-active uplinks; Leaf layer adopts 2 units N5860-48SC stacking.
- Adopts Virtual Switch Unit(stack) and BGP technology to build a L3 full-mesh network.
- The spine-leaf interconnection, to improve the link reliability and bandwidth utilization, realizing high-speed switching of Intranet.

2. Network Topology



3. Device Connection Planning Table

a. Device connection

Local Device	Local Interface	interface IP address	Peer Device	Peer interface	Peer IP address
spine1	H0/61	VSL	spine2	H0/61	VSL
spine1	H0/62	VSL	spine2	H0/62	VSL
spine1	H0/63	VSL	spine2	H0/63	VSL
spine1	H0/64	VSL	spine2	H0/64	VSL
spine1	H0/60	BFD	spine2	H0/60	BFD
spine1	H0/1	10.1.1.1/30	Leaf1	H0/1	10.1.1.2/30
Spine1	H0/2	10.1.2.1/30	Leaf1	H0/2	10.1.2.2/30
Spine1	H0/3	10.1.3.1/30	Leaf2	H0/1	10.1.3.2/30
Spine1	H0/4	10.1.4.1/30	Leaf2	H0/2	10.1.4.2/30
Spine2	H0/1	10.2.1.1/30	Leaf3	H0/1	10.2.1.2/30
Spine2	H0/2	10.2.2.1/30	Leaf3	H0/2	10.2.2.2/30
Spine2	H0/3	10.2.3.1/30	Leaf4	H0/1	10.2.3.2/30
Spine2	H0/4	10.2.4.1/30	Leaf4	H0/2	10.2.4.2/30
Leaf1	H0/55	VSL	Leaf2	H0/55	VSL
Leaf1	H0/56	VSL	Leaf2	H0/56	VSL
Leaf1	T0/48	BFD	Leaf2	T0/48	BFD
Leaf3	H0/55	VSL	Leaf4	H0/55	VSL
Leaf3	H0/56	VSL	Leaf4	H0/56	VSL
Leaf3	T0/48	BFD	Leaf4	T0/48	BFD
Leaf1&2	T0/1-2	Lacp	Server	T0/1-2	LACP
Leaf1&2	T0/1-2	Lacp	Server	T0/1-2	LACP

b. Management and business address planning

Device Name	IP address	Remark
stack-Spine	10.200.0.200/24	MGMT IP
stack-Leaf1&2	10.200.0.1/24	MGMT IP
stack-Leaf3&4	10.200.0.2/24	MGMT IP
stack-Leaf1&2	10.100.1.1/24	Business Vlan
stack-Leaf3&4	10.100.2.1/24	Business Vlan

4. Configuring Ideas

- Spine, Leaf form stack hot stack to provide device redundancy while increasing forwarding capacity.
- Fullmesh interconnection between Spine and leaf, with 4 ecmps on each leaf.
- Interconnected using BGP routing protocol, providing rich routing control and strong scalability.

5. Configuring Steps

5.1 Spine and Leaf Configuring Stacking

- a. Configure stacking domain and priority of Spine, meanwhile, add the interface "HundredGigabitEthernet 0/61-64" to VSL (Virtual Switching Link)

```
FS# configure terminal
FS(config)#hostname Spine1
Spine1(config)# switch virtual domain 100
Spine1(config-vs-domain)# switch 1 priority 200
Spine1(config-vs-domain)# switch 1 description stack-Spine1
Spine1(config-vs-domain)# exit
Spine1(config)# vsl-port
Spine1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/61
Spine1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/62
Spine1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/63
Spine1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/64
Spine1(config-vsl-port)#end
Spine1# switch convert mode virtual
Convert mode will backup and delete config file, and reload the switch. Are you sure to continue[yes/no]:yes
```

```
FS# configure terminal
FS(config)#hostname Spine2
Spine2(config)# switch virtual domain 100
Spine2(config-vs-domain)# switch 2 priority 100
Spine2(config-vs-domain)# switch 2 description stack-Spine2
Spine2(config-vs-domain)# exit
Spine2(config)# vsl-port
Spine2(config-vsl-port)#port-member interface HundredGigabitEthernet 0/61
Spine2(config-vsl-port)#port-member interface HundredGigabitEthernet 0/62
Spine1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/63
Spine1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/64
Spine2(config-vsl-port)#end
Spine2# switch convert mode virtual
Convert mode will backup and delete config file, and reload the switch. Are you sure to continue[yes/no]:yes
```

- b. Configure stacking domain and priority of Leaf, meanwhile, add the interface "T0/19" and "T0/20" to VSL (Virtual Switching Link)

```
FS>enable
FS# configure terminal
FS(config)#hostname Leaf1
Leaf1(config)# switch virtual domain 100
Leaf1(config-vs-domain)# switch 1 priority 200
Leaf1(config-vs-domain)# switch 1 description stack-Leaf1
Leaf1(config-vs-domain)# exit
```

```

Leaf1(config)# vsl-port
Leaf1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/55
Leaf1(config-vsl-port)#port-member interface HundredGigabitEthernet 0/56
Leaf1(config-vsl-port)#end
Leaf1# switch convert mode virtual
Convert mode will backup and delete config file, and reload the switch. Are you sure to continue[yes/no]:yes

FS>enable
FS# configure terminal
FS(config)#hostname Leaf2
Leaf2(config)# switch virtual domain 100
Leaf2(config-vs-domain)# switch 2 priority 100
Leaf2(config-vs-domain)# switch 2 description stack-Leaf2
Leaf2(config-vs-domain)# exit
Leaf2(config)# vsl-port
Leaf2(config-vsl-port)#port-member interface HundredGigabitEthernet 0/55
Leaf2(config-vsl-port)#port-member interface HundredGigabitEthernet 0/56
Leaf2(config-vsl-port)#end
Leaf2# switch convert mode virtual
Convert mode will backup and delete config file, and reload the switch. Are you sure to continue[yes/no]:yes
  
```

c. Stacking Verification

```

FS#show switch virtual
  
```

Switch_id	Domain_id	Priority	Position	Status	Role	Description
1(1)	100(100)	200(200)	LOCAL	OK	ACTIVE	leaf1
2(2)	100(100)	100(100)	REMOTE	OK	STANDBY	leaf2

5.2 Basic Configurations of Spine and Leaf

a. Basic Configurations of Spine (hostname, ssh, log, snmp, MGMT address)

```

FS# configure terminal
FS(config)#hostname stack-Spine
stack-Spine(config)#username admin privilege 15 password admin
stack-Spine(config)#line vty 0 4
stack-Spine(config-line)#login local
stack-Spine(config-line)#exit
stack-Spine(config)#logging userinfo command-log
stack-Spine(config)#logging buffered 1048576
stack-Spine(config)#logging file flash:syslog
stack-Spine(config)#snmp-server community FS-snmp ro
stack-Spine(config)#enable service ssh-server
stack-Spine(config)#inter mgmt 0
stack-Spine(config-if-Mgmt 0)#ip add 10.200.0.200 24
  
```

```
stack-Spine(config-if-Mgmt 0)#gateway 10.200.0.254
```

b. Basic Configurations of Leaf 1&2 (hostname, ssh, log, snmp, MGMT address), the configs of Leaf3 &4 are similar to this, will not repeat

```
FS#config
FS(config)#hostname stack-Leaf1&2
stack-Leaf1&2(config)#hostname stack-Leaf1&2
stack-Leaf1&2(config)#username admin privilege 15 password admin
stack-Leaf1&2(config)#line vty 0 4
stack-Leaf1&2(config-line)#login local
stack-Leaf1&2(config-line)#exit
stack-Leaf1&2(config)#logging userinfo command-log
stack-Leaf1&2(config)#logging buffered 1048576
stack-Leaf1&2(config)#logging file flash:syslog
stack-Leaf1&2(config)#snmp-server community FS-snmp ro
stack-Leaf1&2(config)#enable service ssh-server
stack-Leaf1&2(config)#interface Ten1/0/48
stack-Leaf1&2 (config)#inter mgmt 0
stack-Leaf1&2 (config-if-Mgmt 0)#ip add 10.200.0.1 24
stack-Leaf1&2 (config-if-Mgmt 0)#gateway 10.200.0.254
```

5.3 Configuring BFD of Spine and Leaf

a. Configuring BFD of Spine

```
stack-Spine(config)#interface Hun1/0/60
stack-Spine(config-if-HundredGigabitEthernet 1/0/60)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 1/0/60)#exit
stack-Spine(config)#interface Hun2/0/60
stack-Spine(config-if-HundredGigabitEthernet 2/0/60)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 2/0/60)#exit
stack-Spine(config)#switch virtual domain 100
stack-Spine(config-vs-domain)#dual-active detection bfd
stack-Spine(config-vs-domain)#dual-active bfd interface Hun 1/0/60
stack-Spine(config-vs-domain)#dual-active bfd interface Hun 2/0/60
stack-Spine(config-vs-domain)#Exit
stack-Spine(config)#
```

b. Configuring BFD of Leaf 1&2, the configs of Leaf3 &4 are similar to this, will not repeat

```
stack-Leaf1&2(config)#interface Ten1/0/48
stack-Leaf1&2(config-if-TenGigabitEthernet 1/0/48)#no switchport
stack-Leaf1&2(config-if-TenGigabitEthernet 1/0/48)#exit
stack-Leaf1&2(config)#interface Ten2/0/48
stack-Leaf1&2(config-if-TenGigabitEthernet 2/0/48)#no switchport
stack-Leaf1&2(config-if-TenGigabitEthernet 2/0/48)#exit
stack-Leaf1&2(config)#switch virtual domain 100
```

```

stack-Leaf1&2(config-vs-domain)#dual-active detection bfd
stack-Leaf1&2(config-vs-domain)#dual-active bfd interface TenGigabitEthernet 1/0/48
stack-Leaf1&2(config-vs-domain)#dual-active bfd interface TenGigabitEthernet 2/0/48
stack-Leaf1&2(config-vs-domain)#exit
  
```

c. Check the configuration status of the dual main chassis

```

FS# show switch virtualdual-active summary
BFDDual-active detection enabled: No
Aggregateportdual-active detection enabled: Yes
Interfaces excluded from shutdown in recovery mode:
In dual-active recovery mode: NO
  
```

d. Check the configuration status of BFD dual main chassis

```

FS# show switch virtualdual-active bfd
BFD dual-active detection enabled: Yes
BFD dual-active interface configured:
  TenGigabitEthernet 1/0/48: UP
  TenGigabitEthernet 2/0/48: UP
  
```

5.4 Interconnection port & IP address of Spine, Leaf

a. Configuring interconnection address of Spine

```

stack-Spine#configure terminal
stack-Spine(config)#interface hundredGigabitEthernet 1/0/1
stack-Spine(config-if-HundredGigabitEthernet 1/0/1)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 1/0/1)#ip address 10.1.1.2 30
stack-Spine(config-if-HundredGigabitEthernet 1/0/1)#exit
stack-Spine(config)#interface hundredGigabitEthernet 1/0/2
stack-Spine(config-if-HundredGigabitEthernet 2/0/1)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 2/0/1)#ip address 10.1.2.2 30
stack-Spine(config-if-HundredGigabitEthernet 2/0/1)#exit
stack-Spine(config)#interface hundredGigabitEthernet 1/0/3
stack-Spine(config-if-HundredGigabitEthernet 1/0/2)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 1/0/2)#ip address 10.1.3.2 30
stack-Spine(config-if-HundredGigabitEthernet 1/0/2)#exit
stack-Spine(config)#interface hundredGigabitEthernet 1/0/4
stack-Spine(config-if-HundredGigabitEthernet 2/0/2)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 2/0/2)#ip address 10.1.4.2 30
stack-Spine(config-if-HundredGigabitEthernet 2/0/2)#exit
stack-Spine(config)#interface hundredGigabitEthernet 2/0/1
stack-Spine(config-if-HundredGigabitEthernet 1/0/1)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 1/0/1)#ip address 10.2.1.2 30
stack-Spine(config-if-HundredGigabitEthernet 1/0/1)#exit
stack-Spine(config)#interface hundredGigabitEthernet 2/0/2
  
```



```

stack-Spine(config-if-HundredGigabitEthernet 2/0/1)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 2/0/1)#ip address 10.2.2.2 30
stack-Spine(config-if-HundredGigabitEthernet 2/0/1)#exit
stack-Spine(config)#interface hundredGigabitEthernet 2/0/3
stack-Spine(config-if-HundredGigabitEthernet 1/0/2)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 1/0/2)#ip address 10.2.3.2 30
stack-Spine(config-if-HundredGigabitEthernet 1/0/2)#exit
stack-Spine(config)#interface hundredGigabitEthernet 2/0/4
stack-Spine(config-if-HundredGigabitEthernet 2/0/2)#no switchport
stack-Spine(config-if-HundredGigabitEthernet 2/0/2)#ip address 10.2.4.2 30
stack-Spine(config-if-HundredGigabitEthernet 2/0/2)#exit
  
```

b. Configuring interconnection address of Leaf 1&2, the configs of Leaf3 &4 are similar to this, will not repeat

```

stack-Leaf1&2#configure terminal
stack-Leaf1&2(config)#interface hundredGigabitEthernet 1/0/49
stack-Leaf1&2(config-if-HundredGigabitEthernet 1/0/49)#no switchport
stack-Leaf1&2(config-if-HundredGigabitEthernet 1/0/49)#ip address 10.1.1.1 30
stack-Leaf1&2(config)#interface hundredGigabitEthernet 1/0/50
stack-Leaf1&2(config-if-HundredGigabitEthernet 1/0/50)#no switchport
stack-Leaf1&2(config-if-HundredGigabitEthernet 1/0/50)#ip address 10.1.2.1 30
stack-Leaf1&2(config)#interface hundredGigabitEthernet 2/0/49
stack-Leaf1&2(config-if-HundredGigabitEthernet 2/0/49)#no switchport
stack-Leaf1&2(config-if-HundredGigabitEthernet 2/0/49)#ip address 10.1.3.1 30
stack-Leaf1&2(config)#interface hundredGigabitEthernet 2/0/50
stack-Leaf1&2(config-if-HundredGigabitEthernet 2/0/50)#no switchport
stack-Leaf1&2(config-if-HundredGigabitEthernet 2/0/50)#ip address 10.1.4.1 30
  
```

5.5 BGP Configuring Routing Protocol BGP

a. Configuring BGP of Spine and Leaf

```

stack-Spine#configure terminal
stack-Spine(config)#router bgp 100
stack-Spine(config-router)#bgp log-neighbor-changes
stack-Spine(config-router)#bgp graceful-restart restart-time 120
stack-Spine(config-router)#bgp graceful-restart stalepath-time 360
stack-Spine(config-router)#bgp graceful-restart
stack-Spine(config-router)#maximum-paths ebgp 32
stack-Spine(config-router)#bgp nsr
stack-Spine(config-router)#neighbor 10.1.1.1 remote-as 1001
stack-Spine(config-router)#neighbor 10.1.2.1 remote-as 1001
stack-Spine(config-router)#neighbor 10.1.3.1 remote-as 1001
stack-Spine(config-router)#neighbor 10.1.4.1 remote-as 1001
  
```

b. Configuring BGP of Leaf1&2, the configs of Leaf3 &4 are similar to this, will not repeat

```

stack-Leaf1&2#configure terminal
stack-Leaf1&2(config)#router bgp 1001
stack-Leaf1&2(config-router)#bgp log-neighbor-changes
stack-Leaf1&2(config-router)#bgp graceful-restart restart-time 120
stack-Leaf1&2(config-router)#bgp graceful-restart stalepath-time 360
stack-Leaf1&2(config-router)#bgp graceful-restart
stack-Leaf1&2(config-router)#maximum-paths ebgp 32
stack-Leaf1&2(config-router)#bgp nsr
stack-Leaf1&2(config-router)#neighbor 10.1.1.2 remote-as 100
stack-Leaf1&2(config-router)#neighbor 10.1.2.2 remote-as 100
stack-Leaf1&2(config-router)#neighbor 10.1.3.2 remote-as 100
stack-Leaf1&2(config-router)#neighbor 10.1.4.2 remote-as 100
stack-Leaf1&2(config-router)#network 10.100.1.0 mask 255.255.255.0
  
```

c. Configuring Spine Load Balance

```

stack-Spine
stack-Spine(config)#load-balance-profile ap&ecmp
stack-Spine(config-load-balance-profile)#ipv4 field src-ip dst-ip protocol l4-src-port l4-dst-port
stack-Spine(config-load-balance-profile)#exit
stack-Spine(config)#aggregateport load-balance enhanced profile ap&ecmp
  
```

d. Configuring Leaf1&2 Load Balance, the configs of Leaf3 &4 are similar to this, will not repeat

```

stack-Leaf1&2
stack-Leaf1&2(config)#load-balance-profile ap&ecmp
stack-Leaf1&2(config-load-balance-profile)#ipv4 field src-ip dst-ip protocol l4-src-port l4-dst-port
stack-Leaf1&2(config-load-balance-profile)#exit
stack-Leaf1&2(config)#aggregateport load-balance enhanced profile ap&ecmp
  
```

5.6 Configuring Downlink Business Port of Leaf

a. Configuring MGMT address of leaf1&2, the configs of Leaf3 &4 are similar to this, will not repeat

```

stack-Leaf1&2
stack-Leaf1&2(config)#interface VLAN 2
stack-Leaf1&2(config-if-VLAN 2)#description PM-access
stack-Leaf1&2(config-if-VLAN 2)#carrier-delay 0
stack-Leaf1&2(config-if-VLAN 2)#ip address 10.100.1.1 24
  
```

b. Configuring link aggregation of Leaf1&2, the configs of Leaf3 &4 are similar to this, will not repeat

```

stack-Leaf1&2(config)#interface TenGigabitEthernet 1/0/1
stack-Leaf1&2(config-if-TenGigabitEthernet 1/0/1)#port-group 1 mode active
stack-Leaf1&2(config-if-TenGigabitEthernet 1/0/1)#lacp short-timeout
stack-Leaf1&2(config)interface TenGigabitEthernet 2/0/1
stack-Leaf1&2(config-if-TenGigabitEthernet 2/0/1)#port-group 1 mode active
  
```

```
stack-Leaf1&2(config-if-TenGigabitEthernet 2/0/1)#lACP short-timeout
stack-Leaf1&2(config)#interface TenGigabitEthernet 1/0/2
stack-Leaf1&2(config-if-TenGigabitEthernet 1/0/1)#port-group 2 mode active
stack-Leaf1&2(config-if-TenGigabitEthernet 1/0/1)#lACP short-timeout
stack-Leaf1&2(config)interface TenGigabitEthernet 2/0/2
stack-Leaf1&2(config-if-TenGigabitEthernet 2/0/1)#port-group 2 mode active
stack-Leaf1&2(config-if-TenGigabitEthernet 2/0/1)#lACP short-timeout
```

c. Configuring security of leaf1&2, the configs of Leaf3 &4 are similar to this, will not repeat

```
stack-Leaf1&2(config)#interface aggregatePort 1
stack-Leaf1&2(config-if-AggregatePort 1)#storm-control broadcast 10240
stack-Leaf1&2(config-if-AggregatePort 1)#storm-control multicast level 1
stack-Leaf1&2(config-if-AggregatePort 1)#storm-control unicast 10240
stack-Leaf1&2(config-if-AggregatePort 1)#description server-1
stack-Leaf1&2(config)#interface aggregatePort 2
stack-Leaf1&2(config-if-AggregatePort 1)#storm-control broadcast 10240
stack-Leaf1&2(config-if-AggregatePort 1)#storm-control multicast level 1
stack-Leaf1&2(config-if-AggregatePort 1)#storm-control unicast 10240
stack-Leaf1&2(config-if-AggregatePort 1)#description server-1
stack-Leaf1&2(config-if-AggregatePort 1)#carrier-delay 0
stack-Leaf1&2(config-if-AggregatePort 1)#switchport mode access
stack-Leaf1&2(config-if-AggregatePort 1)#switchport access vlan 2
stack-Leaf1&2(config-if-AggregatePort 1)#spanning-tree portfast
```

6. Configuring Review

6.1 Stack Status Review

Function	Command
stack Info	Show switch virtual[topology config role]
Current configs info of dual main chassis	Show switch virtual dual-active
VSL-AP Operation Info	Show switch virtual link[port]
Switch ID	Show switch id

6.2 BGP neighbor, routing table configuration check

Function	Command
Review single wavelength routing of BGP IPv4	show ip bgp
View BGP stats	show bgp statistics
Clear BGP IPv4 unicast routes	clear ip bgp
Turn on all BGP debug switches	debug ip bgp all
Turn on the BGP event handling debugging switch	debug ip bgp event
Turn on the BGP Neighbor Keep Alive debug switch.	debug ip bgp keepalives



 <https://www.fs.com>



All statements, technical information, and recommendations related to the products here are based upon information believed to be reliable or accurate. However, the accuracy or completeness thereof is not guaranteed, and no responsibility is assumed for any inaccuracies. Please contact FS for more information.