

# **AOS-CX 10.16.xxx Virtual Switching Framework (VSF) Guide**

**4100i, 6100, 6200, 6300 Switch Series**



**Hewlett Packard**  
Enterprise

Published: October 2025

Version: 1

## Copyright Information

© Copyright 2025 Hewlett Packard Enterprise Development LP.

This product includes code licensed under certain open source licenses which require source compliance. The corresponding source for these components is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, please check if the code is available in the HPE Software Center at <https://myenterpriselicense.hpe.com/cwp-ui/software> but, if not, send a written request for specific software version and product for which you want the open source code. Along with the request, please send a check or money order in the amount of US \$10.00 to:

Hewlett Packard Enterprise Company  
Attn: General Counsel  
WW Corporate Headquarters  
1701 E Mossy Oaks Rd Spring, TX 77389  
United States of America.



## Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.



---

<b>About this document</b> .....	<b>8</b>
Applicable products .....	8
Latest version available online .....	8
Command syntax notation conventions .....	8
About the examples .....	9
Identifying switch ports and interfaces .....	10
<b>Protocol and feature details</b> .....	<b>11</b>
Terminology .....	11
Connection Topology .....	12
Ring topology .....	12
Chain topology .....	13
VSF Behavior .....	13
One Virtual Device .....	14
VSF auto-stacking .....	17
Peer discovery .....	17
Auto-join Eligibility .....	17
Reserved interfaces for auto-stacking .....	17
Force auto-join support .....	19
Interoperation .....	20
Link aggregation .....	20
VSF port shaping .....	21
Supported VSF port speeds .....	22
VSF Support on SmartRate Ports .....	23
<b>Configuration task list</b> .....	<b>24</b>
Stack deployment using auto-stacking .....	24
Designating conductor switch .....	24
Auto-stacking using LED mode button .....	25
Auto-stacking using CLI command .....	26
Auto-stacking using zero-touch provisioning (ZTP) .....	27
Auto-stacking using HPE Aruba Networking CX mobile app .....	27
Manual configuration .....	28
VSF links .....	29
Secondary member .....	29
Member number .....	29
Member provisioning .....	29
Access to VSF members .....	30
Stack management .....	30
Consoles .....	30
Management interface .....	30
Split detection .....	30
Automated image sync .....	33
Reboot .....	33
Member addition with auto-stacking .....	33
Member addition without auto-stacking .....	33
Member replacement with auto-stacking .....	34
Member replacement without auto-stacking .....	34
Member removal .....	34

Stack and Port LED states .....	34
<b>Use cases .....</b>	<b>36</b>
Forming a four-member ring setup using auto-stacking command .....	36
Forming a four-member chain setup using link configuration command with auto-join .....	38
Forming an eight-member ring setup manually using link configuration without auto-stacking .....	39
<b>VSF In-Service Software Upgrade .....</b>	<b>43</b>
Upgrade scope .....	43
VSF ISSU operation .....	43
Step 1: Image download and validation .....	43
Step 2: Install the image in the standby & members .....	43
Step 3: Freeze the system .....	44
Step 4: Perform the upgrade for the line modules .....	44
Step 5: Recover the previous state .....	44
Step 6: Failover occurs between the conductor and the standby .....	44
Step 7: The previous Conductor will upgrade its image .....	44
VSF ISSU considerations .....	44
VSF ISSU limitations .....	45
Hardware faults .....	45
Prohibited operations .....	45
Recommendations .....	45
Scheduled jobs .....	45
Information not updated during ISSU .....	45
Show tech .....	45
Features that don't support uninterrupted operation during ISSU .....	45
AAA .....	46
PIM with MSDP .....	46
Dynautz ISSU .....	46
BFD ISSU .....	46
NAE .....	46
Feature readiness check .....	46
Restrictions .....	46
VSF ISSU errors and recovery .....	46
<b>VSF commands .....</b>	<b>48</b>
clear vsf split-detection statistics .....	48
description .....	48
interface .....	49
issu rollback-timer .....	50
issu update-software .....	51
link .....	56
member .....	57
show issu .....	58
show vsf .....	63
show vsf detail .....	66
show vsf member .....	69
show vsf link .....	69
show vsf link detail .....	70
show vsf link error-detail .....	72
show vsf link error-detail member .....	74
show vsf split-detection .....	76
show vsf topology .....	77
shutdown .....	78
type .....	79

---

vsf egress-shape .....	80
vsf force-auto-join .....	82
vsf member .....	83
vsf member reboot .....	84
vsf renumber-to .....	85
vsf secondary-member .....	86
vsf split-detect mgmt .....	88
vsf split-detect vlan .....	89
vsf start-auto-stacking .....	91
<b>Configuration conflict finder recommendations .....</b>	<b>93</b>
<b>Considerations and best practices .....</b>	<b>96</b>
<b>Debugging and troubleshooting .....</b>	<b>98</b>
Stack split .....	98
Step 1: Verify split has occurred .....	98
Step 2: Identify failed link or member .....	99
Step 3: Recover or replace failed link or member .....	100
Step 4: Verify proper stack operation .....	100
Misconfiguration recovery .....	100
VSF switchover behaviors .....	101
<b>Frequently asked questions .....</b>	<b>102</b>
General .....	102
Auto-stacking .....	111
<b>Support and Other Resources .....</b>	<b>118</b>
Accessing HPE Aruba Networking Support .....	118
Accessing Updates .....	119
Warranty Information .....	119
Regulatory Information .....	119
Documentation Feedback .....	119

This document describes features of the AOS-CX network operating system. It is intended for administrators responsible for installing, configuring, and managing HPE Aruba Networking switches on a network.

## Applicable products

This document applies to the following products:

- HPE Aruba Networking 4100i Switch Series (JL817A, JL818A)
- HPE Aruba Networking 6100 Switch Series (JL675A, JL676A, JL677A, JL678A, JL679A, R9Y04A)
- HPE Aruba Networking 6200 Switch Series (JL724A, JL725A, JL726A, JL727A, JL728A, R8Q67A, R8Q68A, R8Q69A, R8Q70A, R8Q71A, R8V08A, R8V09A, R8V10A, R8V11A, R8V12A, R8Q72A, JL724B, JL725B, JL726B, JL727B, JL728B, S0M81A, S0M82A, S0M83A, S0M84A, S0M85A, S0M86A, S0M87A, S0M88A, S0M89A, S0M90A, S0G13A, S0G14A, S0G15A, S0G16A, S0G17A, R8V13A)
- HPE Aruba Networking 6300 Switch Series (JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL762A, R8S89A, R8S90A, R8S91A, R8S92A, S0E91A, S0X44A, S3L75A, S3L76A, S3L77A, S4P41A, S4P42A, S4P43A, S4P44A, S4P45A, S4P46A, S4P47A, S4P49A, S0G03A, S0G05A, S0F99A, S0G01A, S0G96A, S0G98A, S4P48A, S0G04A, S0G06A, S0G00A, S0G95A, S0G97A, S0G02A)

## Latest version available online

Updates to this document can occur after initial publication. For the latest versions of product documentation, see the links provided in [Support and Other Resources](#).

## Command syntax notation conventions

Convention	Usage
<code>example-text</code>	Identifies commands and their options and operands, code examples, filenames, pathnames, and output displayed in a command window. Items that appear like the example text in the previous column are to be entered exactly as shown and are required unless enclosed in brackets ( <b>[ ]</b> ).
<b>example-text</b>	In code and screen examples, indicates text entered by a user.
Any of the following: <ul style="list-style-type: none"><li>▪ <code>&lt;example-text&gt;</code></li><li>▪ <code>&lt;example-text&gt;</code></li><li>▪ <i>example-text</i></li><li>▪ <i>example-text</i></li></ul>	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: <ul style="list-style-type: none"><li>▪ For output formats where italic text cannot be displayed, variables are enclosed in angle brackets (&lt; &gt;). Substitute the text—including the enclosing angle brackets—with an actual value.</li></ul>

Convention	Usage
	<ul style="list-style-type: none"> <li>For output formats where italic text can be displayed, variables might or might not be enclosed in angle brackets. Substitute the text including the enclosing angle brackets, if any, with an actual value.</li> </ul>
	Vertical bar. A logical <b>OR</b> that separates multiple items from which you can choose only one. Any spaces that are on either side of the vertical bar are included for readability and are not a required part of the command syntax.
{ }	Braces. Indicates that at least one of the enclosed items is required.
[ ]	Brackets. Indicates that the enclosed item or items are optional.
... or ...	Ellipsis: <ul style="list-style-type: none"> <li>In code and screen examples, a vertical or horizontal ellipsis indicates an omission of information.</li> <li>In syntax using brackets and braces, an ellipsis indicates items that can be repeated. When an item followed by ellipses is enclosed in brackets, zero or more items can be specified.</li> </ul>

## About the examples

Examples in this document are representative and might not match your particular switch or environment.

The slot and port numbers in this document are for illustration only and might be unavailable on your switch.

### Understanding the CLI prompts

When illustrating the prompts in the command line interface (CLI), this document uses the generic term **switch**, instead of the host name of the switch. For example:

```
switch>
```

The CLI prompt indicates the current command context. For example:

```
switch>
```

Indicates the operator command context.

```
switch#
```

Indicates the manager command context.

```
switch (CONTEXT-NAME) #
```

Indicates the configuration context for a feature. For example:

```
switch (config-if) #
```

Identifies the **interface** context.

### Variable information in CLI prompts

In certain configuration contexts, the prompt may include variable information. For example, when in the VLAN configuration context, a VLAN number appears in the prompt:

```
switch (config-vlan-100) #
```

When referring to this context, this document uses the syntax:

```
switch (config-vlan-<VLAN-ID>) #
```

Where **<VLAN-ID>** is a variable representing the VLAN number.

# Identifying switch ports and interfaces

Physical ports on the switch and their corresponding logical software interfaces are identified using the format: *member/slot/port*.

## On the HPE Aruba Networking 4100i Switch Series

- *member*: The member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 4. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on the switch.

## On the HPE Aruba Networking 6100 Switch Series

- *member*: The member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 6. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on the switch.

## On the HPE Aruba Networking 6200 Switch Series

- *member*: Member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 8. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 in slot 1 on member 1.

## On the HPE Aruba Networking 6300 Switch Series

- *member*: Member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 10. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on member 1.

Virtual Switching Framework, or VSF, allows network administrators to stack multiple individual switches into a single logical device using standard Ethernet links. VSF stacks provide increased network capacity and improved redundancy, allowing administrators to scale stack size with user and device requirements while simplifying configuration complexity and providing a single point of management access with a shared control plane across all stack members.

- **4100i:** VSF allows stacks to be formed using any combination of SKUs of the 4100i family. Up to 4 member switches will be allowed. Connections between the switches must use 1G SFP/SFP+ either with 1G copper or 10G ports. All VSF links in a stack should operate at the same speed.
- **6100:** VSF allows stacks to be formed using any combination of SKUs of the 6100 family except 12-port SKU. A 12-port switch will form a VSF stack only with other 12-port switches. Up to 6 member switches will be allowed. Connections between the switches must use 1G SFP/SFP+ either with 1G copper or 10G ports. All VSF links in a stack should operate at the same speed.
- **6200:** VSF allows stacks to be formed using any combination of 24 and 48-port SKUs of the 6200 family. (A 6200F 12-port switch will form a VSF stack only with other 6200F 12-port switches.) Up to 8 member switches will be allowed. Connections between the switches must use 1G SFP/SFP+ either with 1G copper downlink ports, 10G links, or SmartRate ports.
- **6300:** VSF allows stacks to be formed using any combination of SKUs of the 6300 family. Up to 10 member switches will be allowed. Connections between the switches must use 10G, 25G, 50G links, or SmartRate ports. All VSF links in a stack should operate at the same speed.  
The VSF stack, containing only the S0E91A and S0X44A SKUs, supports 10G, 25G, and 100G speed links. If a VSF stack includes a mix of S0E91A/S0X44A SKU and other SKUs, then it supports 10G, and 25G VSF links only.

---

4100 and 6100 switch series do not support using ports with different speeds for VSF links.

---

In 6100 and 6200 switch series, 12-port switches cannot form a stack with 24-port or 48-port switches.

---

In 6300 switch series, S3L75A, S3L76A and S3L77A SKUs can only stack with other S3L75A, S3L76A and S3L77A SKUs.

---



VSF is enabled by default on all supported switch models and cannot be disabled. Within the stack, one switch (normally the primary, member 1) is the Conductor that runs all control plane software and manages the ASICs of all stack members. Any switch apart from primary can be configured as Standby switch, which maintains a synchronized copy of the Conductor's configuration database and is capable of assuming the Conductor role in the event of a failure of, or loss of connectivity to, the Conductor.

## Terminology

**Table 1: Acronyms used in this book**

Term	Definition
VSF	Virtual Switching Framework
L2	Layer 2 of the OSI 7-layer model
L3	Layer 3 of the OSI 7-layer model
SKU	Stock Keeping Unit
FRU	Field Replaceable Unit
ASIC	Application-Specific-Integrated Circuit
L-Agg	Link Aggregation
CLI	Command Line Interface

**Table 2: Role types**

Role	Definition
Primary	The primary member is member ID 1; normally operates as the stack Conductor.
Secondary	User-configurable using any valid member ID other than 1; normally operates as the stack Standby.
Conductor	The Conductor maintains the VSF stack configuration, software images, and control plane.
Standby	The Standby maintains a synchronized copy of the VSF stack configuration from the Conductor; automatically assumes the Conductor role if connectivity is lost to the existing Conductor due to hardware or link failures.
Member	The member switch does not run any networking protocols and has no states. The interfaces on this switch are directly controlled and programmed by the conductor switch.

## Connection Topology

VSF supports up to 4 member stacks (for 4100i devices), 6 member stacks (for 6100 devices), 8 member stacks (for 6200F devices) and 10 member stacks (for 6300 devices) in ring and chain topologies.

### Ring topology




A two-member ring topology is not supported on 4100 and 6100 series switches. For link-level redundancy in two-member stacks, it is recommended to bundle multiple interfaces into a single VSF link using a chain topology.

In a ring topology, each stack member has a VSF link connection to two other members, providing resiliency against link and member switch hardware failures as any single failure does not isolate

remaining stack members from each other. HPE Aruba Networking strongly recommends deploying VSF stacks using ring topologies whenever feasible.

**Figure 1** *Ring topology*



## Chain topology

In a chain topology, there is only one path between any two stack members. A VSF link or hardware failure in a chain topology may cause a stack split and result in network disruption; VSF split detection may be used to mitigate this scenario.

**Figure 2** *Chain topology*



## VSF Behavior

Each stack member must have a unique member ID number. Auto-stacking automatically assigns the lowest available member ID when adding a new member to the stack; if deploying or expanding a stack manually, ensure that each new member is assigned a valid member ID not already in use by an existing stack member, as a member ID conflict will result in the new member failing to join the stack.

- During normal stack operation, the primary member will assume the Conductor role and the secondary member will assume the Standby role during normal operation.
- The primary member is member number 1. This setting is not configurable and 1 is the default. A factory-default switch boots up as a VSF-enabled switch with a member number of 1.
- The secondary member number is user configurable; when auto-stacking is used via the push-button or CLI methods, member 2 is automatically assigned as the secondary. It is recommended that the customer configures a secondary member in the stack, since a stack with a standby offers resiliency and high availability.
- No members other than primary and secondary members can become Conductor or Standby of the stack.



---


In a standard deployment, uplinks should be from primary and secondary. The management interface from primary and secondary members should be connected to the management network, providing management connectivity to the current conductor.

---

# One Virtual Device

Once the VSF stack is formed, all interconnected switches operate as a single virtual switch with a single control plane. All interfaces of all switches in the stack are available for configuration and management.

**Figure 1** One virtual device example topology



```
switch# show vsf
```

```
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address               : 08:97:34:b0:0e:00
Egress Shape              : Enabled
Egress Shape Rate         : None
Secondary                 : 2
Topology                  : Chain
Status                    : No Split
Split Detection Method    : None
```

Mbr ID	MAC Address	Type	Status
1	38:21:c7:5c:62:40	JL668A	Conductor
2	18:7a:3b:1b:68:c0	R8S90A	Standby
3	38:21:c7:5c:57:c0	JL668A	Member
4	18:7a:3b:1b:66:40	R8S89A	Member

Interfaces will be numbered as noted in the following table.

Name	Member Number	Slot	Port
1/1/1	1	1	1
2/1/14	2	1	14
8/1/12	8	1	12

For VSF-capable switches, the slot number is always 1. All interfaces except for those assigned to VSF links are available for normal configuration.

```
switch# show interfaces brief
```

Port	Native VLAN	Mode	Type	Enabled	Status	Reason	Speed (Mb/s)
1/1/1	10	access	SFP+DA3	yes	up		10000
1/1/2	--	routed	--	no	down	No XCVR installed	--
1/1/3	--	routed	--	no	down	No XCVR installed	--
1/1/4	--	routed	--	no	down	No XCVR installed	--
1/1/5	--	routed	--	no	down	No XCVR installed	--
1/1/6	--	routed	--	no	down	No XCVR installed	--
...							
1/1/33	--	routed	--	no	down	No XCVR installed	--
1/1/34	--	routed	--	no	down	No XCVR installed	--
1/1/35	--	routed	--	no	down	No XCVR installed	--
1/1/36	--	routed	--	no	down	No XCVR installed	--
2/1/1	10	access	SFP+DA3	yes	up		10000
2/1/2	--	routed	--	no	down	No XCVR installed	--
2/1/3	--	routed	--	no	down	No XCVR installed	--
2/1/4	--	routed	--	no	down	No XCVR installed	--
...							
2/1/35	--	routed	--	no	down	No XCVR installed	--
2/1/36	--	routed	--	no	down	No XCVR installed	--

A single control plane operates for the entire VSF stack.

```
6300(config)# show run
Current configuration:
!
!Version AOS-CX FL.10.07.xxxx
!export-password: default
cli-session
  timeout 0
!
!
!
!
!
!
ssh server vrf default
ssh server vrf mgmt
vsf secondary-member 2
vsf member 1
  type j1666a
  link 1 1/1/26
  link 2 1/1/25
vsf member 2
  type j1666a
  link 1 2/1/25
  link 2 2/1/26
vlan 1
spanning-tree
interface mgmt
  no shutdown
  ip dhcp
interface 1/1/1
  no shutdown
  no routing
```

```

    vlan access 1
interface 1/1/2
    no shutdown
    no routing
    vlan access 1
interface 1/1/3
    no shutdown
    no routing
    vlan access 1
interface 1/1/4
    no shutdown
    no routing
    vlan access 1
interface 1/1/5
    no shutdown
    no routing
    vlan access 1
interface 1/1/6
    no shutdown
    no routing
    vlan access 1

...
...
interface 2/1/1
    no shutdown
    no routing
    vlan access 1
interface 2/1/2
    no shutdown
    no routing
    vlan access 1
interface 2/1/3
    no shutdown
    no routing
    vlan access 1
interface 2/1/4
    no shutdown
    no routing
    vlan access 1
interface 2/1/5
    no shutdown
    no routing
    vlan access 1
interface 2/1/6
    no shutdown
    no routing
    vlan access 1

...
...

!
!
!
!
!
https-server vrf default
https-server vrf mgmt
switch(config)#

```

As shown in this configuration, interfaces of all member switches can be configured from the Conductor.

Once a stack is deployed, the stack configuration is persistent and stored separate from the startup configuration. The user can safely remove the startup configuration with the command `erase startup-configuration` without disturbing the stacking topology. To remove all configurations, including the stacking topology, use the command `erase all zeroize` which will automatically reboot and zeroize all stack members, restoring them to a factory default state.

## VSF auto-stacking

VSF auto-stacking feature provides a mechanism to automatically form a stack when the stack members are physically connected in a desired topology. This reduces the number of user intervention touch points to form a VSF stack.

A manual stack formation procedure generally requires the user to explicitly log in to each of the switch, configure the links, renumber it, and then make the physical connection to form a VSF stack of desired size and topology. This is error prone since there are multiple touch-points involved in the whole workflow for each member. The auto-stacking feature eases this problem by reducing the number of touch-points involved to simple physical connections of the links. A new factory default switch can be added into an existing stack by physically connecting it to a VSF link port on an existing stack member. The new switch will automatically be assigned the lowest available member ID and will automatically reboot. After reboot, the newly added member will join the stack.

There are two major components to the auto-stacking solution:

- Peer discovery—Initiated from the conductor using one of the methods described in [Designating conductor switch](#).
- Auto-join Eligibility—Determined by the configuration state of each stack member. A switch with a factory default configuration is auto-join eligible.

### Peer discovery

Auto-stacking peer discovery is a uni-directional process. It starts with the VSF link containing the higher-numbered VSF port, sending a VSF peer discovery protocol packet to a connected peer switch. The peer receives the packet, determines if it is valid, and sends a response with information including its auto-join eligibility, MAC address, and part number. If the peer is auto-join eligible, the VSF member and link configurations are automatically added to the running configuration of the conductor.

### Auto-join Eligibility

Auto-join eligibility determines whether a switch will join a VSF stack if connected to a configured VSF link port on an existing stack member. A switch in its factory defaults configuration state is considered to be auto-join eligible. If the auto-join eligible switch is connected to existing stack, it will automatically reboot and join the stack. Once a switch is no longer using a factory default configuration, it is no longer auto-join eligible and will not automatically join an existing stack. A CLI command is available to override this behavior; see [Force auto-join support](#) for details. However, user can still manually configure the links, renumber the device to make it part of a stack.

### Reserved interfaces for auto-stacking

The following two interfaces are reserved for the auto-stacking process for most switches:

- 12-port switch models: 13 and 14
- 24-port switch models: 25 and 26

- 40-port switch models: 41 and 42
- 48-port switch models: 49 and 50

Users can physically connect the switch to an existing stack on one of these reserved auto-stacking interfaces.



The 6300 Switch series models R8S89A, R8S90A, R8S91A and R8S92A reserve the last pair of ports for other features, and reserve the first page of ports for stacking @ 50G speeds.

The following table shows the list of reserved auto-stacking interfaces based on the product type and platform:

*Reserved Interfaces*

<b>Platform</b>	<b>SKU Part Number</b>	<b>Ports reserved for Auto-Stacking</b>
4100i	JL817A	13, 14
4100i	JL818A	25, 26
6100	JL679A	13, 14
6100	JL677A, JL678A	25, 26
6100	JL675A, JL676A, R9Y04A	49, 50
6200	R8Q72A, R8V13A	13, 14
6200	JL724B, JL725B, S0M81A, S0M82A, S0M86A, S0M87A, S0G13A, S0G14A, JL724A, JL725A, R8Q67A, R8Q68A, R8V08A, R8V09A	25, 26
6200	JL726A, JL727A, JL728A, R8Q69A, R8Q70A, R8Q71A, R8V10A, R8V11A, R8V12A, JL726B, JL727B, JL728B, S0M83A, S0M84A, S0M85A, S0M88A, S0M89A, S0M90A, S0G15A, S0G16A, S0G17A	49, 50
6300	JL658A, JL660A, JL662A, JL664A, JL666A, JL668A, S0G03A, S0G05A, S0F99A, S0G01A, S0G96A, S0G98A, R8S89A, R8S92A, S4P44A, S4P48A	25, 26
6300	S4P42A, S4P46A	41, 42
6300	JL659A, JL661A, JL663A, JL665A, JL667A, JL762A, S0G04A, S0G06A, S0G00A, S0G95A, S0G97A, S0G02A, R8S90A, R8S91A, S0E91A, S0X44A, S4P41A, S4P43A, S4P45A, S4P47A	49, 50

Platform	SKU Part Number	Ports reserved for Auto-Stacking
6300L	S3L75A	25, 26
6300L	S3L76A, S3L77A	49, 50

## Force auto-join support

Only a switch with factory default configuration is considered to be auto-join eligible. In order to support factory express deployments where the user wants to add a switch which is in its non-factory default configuration, the force auto-join configuration support is provided. Use the `vsf force-auto-join` command to force the switch to join the stack automatically. Once the user sets force auto-join in the switch configuration, the switch will be considered as auto-join eligible and will join the stack even though the switch does not have the factory default configuration.



---

Force auto-join will work only if the switch does not have any pre-existing VSF configurations.

---

# Interoperation

A VSF stack supports any combination of models within a supported switch family:

- 6200F switches, or
- 6300 switches (6300L, 6300M, and 6300F).



---

Stacking is not supported across switch families. For example, a VSF stack can include either 4100i, 6100, 6200 or 6300 switches, but not a mix. Additionally, 6100 switch series JL679A SKUs can only stack with other JL679A SKUs, 6200 switch series models R8Q72A and R8V13A SKUs can only stack with other R8Q72A and R8V13A SKUs, and 6300L switch series S3L75A, S3L76A and S3L77A SKUs can only stack with other S3L75A, S3L76A and S3L77A SKUs.

---



---

Firmware versions prior to AOS-CX 10.07 are not interoperable with 10.07 or later versions.

---

## Link aggregation

Link aggregations (L-Agg) may span interfaces across multiple stack members. Load balancing is performed on all interfaces of the L-Agg across the stack and is applicable to both L2 and L3 L-Aggs.

```
interface lag 1
  no shutdown
  no routing
  vlan access 1
  loop-protect
interface lag 2
  no shutdown
  bfd min-transmit-interval 1000
  ip address 192.168.12.7/24
interface 1/1/18
  no shutdown
  lag 1
interface 2/1/18
  no shutdown
  lag 1
interface 1/1/23
  no shutdown
  lag 2
interface 2/1/23
  no shutdown
  lag 2
```

```
switch# show lacp interfaces
```

```
State abbreviations :
A - Active           P - Passive           F - Aggregable   I - Individual
S - Short-timeout   L - Long-timeout   N - InSync       O - OutofSync
C - Collecting      D - Distributing
X - State m/c expired   E - Default neighbor state
```

```
Actor details of all interfaces:
```

```

-----
Intf      Aggr      Port  Port  State  System-ID      System Aggr  Forwarding
          Name      Id    Pri
          -----
1/1/18   lag1
2/1/18   lag1
1/1/23   lag2
2/1/23   lag2
-----

```

Partner details of all interfaces:

```

-----
Intf      Aggr      Port  Port  State  System-ID      System Aggr
          Name      Id    Pri
          -----
1/1/18   lag1
2/1/18   lag1
1/1/23   lag2
2/1/23   lag2
-----

```

## VSF port shaping

The VSF port shaping feature ensures seamless stacking between S0E91A/S0X44A SKU and other 6300 SKUs. The S0E91A and S0X44A SKU have a speed limit of 100G ports, while the speed limit of other 6300 SKUs are limited up to 50G ports. To ensure that all VSF interfaces run at the lowest common speed supported by the entire VSF stack, port shaping is used. This feature is applicable only in the VSF stack containing at least one S0E91A or S0X44A SKU member. If all the S0E91A/S0X44A SKU members are removed from the stack, port shaping will not work.

VSF port shaping configuration is enabled by default. Use the command 'no vsf egress-shape' to disable port shaping. See [vsf egress-shape](#) for details.




---

This feature is applicable only for 6300 switch series (except for S3L75A, S3L76A and S3L77A).

---

Port shaping is triggered for recalculating the speed in the following scenarios:

- New member addition to the existing stack.
- Existing Member removal from the stack.
- Link state changes from up to down or down to up.
- VSF configuration changes.
- Topology transitions from ring-to-chain and chain-to-ring.




---

Port shaping will be applied only to the stack members that are up and operational.

---

### Limitations

The following are limitations associated with the VSF port shaping feature:

- Port shaping is not supported on 4100i, 6100 and 6200 switch series.
- In the 6300 switch series, this feature is applicable only in a VSF stack containing at least one S0E91A or S0X44A SKU member.
- In the case of two-member stack, port shaping will work only with ring topology.
- Port shaping is not applicable for the standalone conductor device.

## Supported VSF port speeds

### For 4100i and 6100 switch series

4100i and 6100 stacks support only homogeneous port speeds across the stack which means only one port speed is supported across all VSF ports in the stack. The port speed is internally set to the highest speed capable by the first VSF interface that is configured.

If the user wants to change the port speed, then all VSF links must be unconfigured and the desired port can be added as the first VSF interface, which will set the port speed to the desired value.



---

In the case of transceiver ports, the speed is determined by the bay type. An SFP+ bay, if configured as the 1st VSF interface, will set the port speed to 10G.

---

## VSF Support on SmartRate Ports

AOS-CX supports VSF on SmartRate ports. The following table lists the platforms and the supported SmartRate speed.

**Table 1:** *Supported SmartRate Speed*

Platform	Product Type	Supported SmartRate Speed
6300	R8S89A, S0E91A, S0X44A	10G SmartRate
6300	R8S90A, R8S91A, JL659A, JL660A, S0G04A, S0G05A	5G SmartRate
6300L	S3L75A	10G SmartRate
6300L	S3L76A, S3L77A	5G SmartRate
6200	R8Q71A, R8V12A	5G SmartRate
6300	S4P41A, S4P42A, S4P45A, S4P46A	5G/10G SmartRate

---

SKUs S4P41A, S4P42A, S4P45A and S4P46A support SmartRate ports with dual-mode and can operate either on 5G or 10G.

---



Dynamic speed changes of dual-mode SR ports between 5G and 10G, are not supported when stacking interfaces are in use. If the stack is up and operating with dual-mode SR ports at 5G speed, and the user wishes to change the speed to 10G (or vice versa), using the command `system interface-group <GROUP> member <MEMBER-ID> speed <SPEED>`, the stack may split.

---

The following sections describe the prerequisites and procedures to configure a VSF stack.

### Stack deployment using auto-stacking

Utilize the VSF auto-stacking feature to quickly deploy pre-cabled stacks with minimal configuration required.

#### Designating conductor switch

Auto-Stacking feature requires the conductor of the stack to be configured with VSF links. Optionally, a secondary or standby device can also be configured.

The following are different methods to designate the conductor:

- Using LED mode button: Physically connect the switches in the desired topology on the reserved VSF link ports and press the LED mode button until the mode changes to **Stk** on a factory default switch. This will automatically configure member 2 as the secondary member and configure the reserved VSF link ports as VSF links 1 and 2. In addition to VSF secondary and link configurations, `ztp force-provision` will also be configured on the conductor switch. The status of stack formation can be verified using Stk LED and Port LEDs states. For more information on LED states, see [Stack and Port LED states](#).



---

Auto-Stacking configures the higher-numbered reserved port on member 1 (14, 26, 42 or 50, depending on the SKU) as VSF link 1, and the lower-numbered reserved port (13, 25, 41 or 49, depending on the SKU) as VSF link 2, unless explicitly defined in a configuration downloaded by ZTP.

---

- Using start auto-stacking CLI: Physically connect the switches in a desired topology on the reserved VSF link port and execute the `vsf start-auto-stacking` command to automatically configure links. The command also configures member 2 as secondary. For more information, see [Forming a four-member ring setup using auto-stacking command](#)

Example:

```
switch(config)# vsf start-auto-stacking
This will configure links and secondary on conductor

Do you want to continue (y/n)? y
```

For information on interfaces that should be configured as VSF links, refer to the [Reserved interfaces for auto-stacking](#) section.



---

To use this command, the switch must be in the factory default configuration.

---

- Using link configuration CLI: Execute the `vsf member` command to configure VSF links on the conductor.

Example:

```
6300 (vsf-member-1) # link 1 1/1/26
6300 (vsf-member-1) # link 2 1/1/25
```



---

To form an ordered stack, it is recommended to configure higher denomination interface first into VSF link .

---

- TFTP download: Full stack configuration can be downloaded into the conductor of the stack. The recommendation is to first download the configuration to the startup and then move the startup to the running configuration.
- ZTP download: Full stack configuration can be downloaded into the conductor of the stack from TFTP server using ZTP. Once the configuration has been downloaded and applied, auto-stacking peer discovery proceeds and forms a stack.

For more information on ZTP, refer to the *Zero Touch Provisioning* chapter in the *Fundamentals Guide*.



---

If full stack configuration is downloaded into the conductor through TFTP/ZTP, the physical connections between the switches should be made according to the downloaded configuration.

---

## Auto-stacking using LED mode button

### Prerequisites

- All switches must be in the factory default configuration.
- All stack members must be connected in a ring topology on the reserved VSF link ports. For more information, see [Reserved interfaces for auto-stacking](#).
- Operator must be connected to the conductor's USB-C console port or a terminal server connected to the USB-A port via a supported adapter or cable.
- This feature is not supported on 4100i or 6100 switch series.

### Procedure

1. Physically connect the switches in a desired topology on the reserved VSF link ports.
2. Press the LED mode button on the conductor until the mode changes to "Stk". The stack members reboot one after another and join the stack.

During stacking operation, the port LEDs are displayed in three different states:

- Flashing green—Indicates that the member is the conductor.
- Flashing orange—Indicates that the member is rebooting to join the stack or offline due to error condition.
- Solid green—Indicates that the member joined the stack and is operational.

For more information on stacking LED states, refer to the Monitoring Guide.

3. Issue a "show vsf" command to ensure that the stack has successfully formed. Alternatively, you can also verify the stack formation using LED states.

```
switch (config)# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 70:72:cf:ef:b7:f2
Egress Shape            : Enabled
Egress Shape Rate       : None
Secondary                : 2
Topology                 : Chain
Status                   : No Split
Split Detection Method   : None
```

Member ID	Mac Address	type	Status
1	70:72:cf:ef:b7:f2	JL664A	Conductor
2	90:20:c2:23:67:40	JL664A	Standby
3	90:20:c2:24:71:c0	JL667A	Member
4	38:21:c7:5a:33:40	JL668A	Member

## Auto-stacking using CLI command

### Prerequisites

- All switches must be in the factory default configuration.
- All stack members must be connected in a ring topology on the reserved VSF link ports. For more information, see [Reserved interfaces for auto-stacking](#).
- Operator must be connected to the conductor's USB-C console port or a terminal server connected to the USB-A port via a supported adapter or cable.

### Procedure

1. Physically connect the switches in a desired topology on the reserved VSF link ports.
2. Connect to the switch console and log in using the admin user; set a password when prompted.
3. Issue the `vsf start-auto-stacking` from the configuration context to start auto-stacking. The stack members reboot one after another and join the stack.

4.

```
switch(config)# vsf start-auto-stacking
```

5. Issue a "show vsf" command to ensure that the stack has successfully formed. Alternatively, you can also verify the stack formation using LED states.

```
switch(config)# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 70:72:cf:ef:b7:f2
Egress Shape            : Enabled
Egress Shape Rate       : None
Secondary                : 2
Topology                 : Chain
Status                   : No Split
Split Detection Method   : None
```

Member ID	Mac Address	type	Status
1	70:72:cf:ef:b7:f2	JL664A	Conductor
2	90:20:c2:23:67:40	JL664A	Standby
3	90:20:c2:24:71:c0	JL667A	Member
4	38:21:c7:5a:33:40	JL668A	Member

## Auto-stacking using zero-touch provisioning (ZTP)

### Prerequisites

- All switches must be in factory default configuration.
- The conductor switch must be connected to the management network.
- All stack must be connected in a ring topology using SFP using SFP uplink ports with a minimum VSF link speed of 10Gbps.
- A supported ZTP method, such as a TFTP server must be defined by DHCP options.

### ZTP auto-stacking via DHCP

AOS-CX uses the following DHCP options to specify network locations and filenames for automatic configuration and software image downloads:

- Option 60: Vendor Class Identifier (VCI).

VCI provided in DHCP request from switch is matched to an Option 43 vendor class defined on the DHCP server to provide configuration filename on TFTP server.

- Option 66: TFTP server name (IPv4 address)
- Option 43: Vendor-specific information
  - suboption 144: Name of the configuration file
  - Sub-option 145: Name of the firmware image file

Use the `show dhcp client vendor-class-identifier` command from the conductor switch to display the VCI string needed to configure the DHCP Option 43 vendor class on the DHCP server

On the DHCP server, configure predefined option 144 for each vendor class that will be used to specify the filename with the ZTP configuration to be downloaded by the stack member.

Optionally, you can also configure option 145 with the filename for an AOS-CX software image that will be downloaded by the conductor for automatic image upgrades.

ZTP auto-stacking is initiated by powering up stack members, with the conductor (member 1) receiving a DHCP address with TFTP config download parameters in DHCP suboption 144. The conductor downloads the ZTP configuration from the TFTP server, which includes configuration for all members. If VSF links are defined and connected, they are used for peer discovery; otherwise, the reserved VSF link ports are configured and used instead.

Once the configuration has been downloaded and applied, auto-stacking peer discovery proceeds and forms a stack.

For more information on ZTP, refer to the Zero Touch Provisioning chapter in the Fundamentals Guide.

## Auto-stacking using HPE Aruba Networking CX mobile app

### Prerequisites

- All switches must be in the factory default configuration.
- All stack members must be connected in a ring topology.
- Bluetooth adapters are installed in all stack member USB-A ports.
- Supported iOS, iPadOS, or Android mobile device is running the HPE Aruba Networking CX mobile app version 2.0 or later.

## Procedure




---

Note: Spanning Tree is enabled by default on 4100i, 6100, 6200, and 6300 switch families, which will prevent a loop from forming when VSF link cables are connected prior to the stack being fully provisioned.

---

1. On the mobile device with the HPE Aruba Networking CX mobile app installed, first use Bluetooth to discover and connect to the switch that will be the stack primary; each switch should show up in the device list using the format 6x00-SERIAL\_NUMBER.
2. Once the device is connected to the switch, launch the HPE Aruba Networking CX mobile app. Within a few seconds, the app should display an active Bluetooth connection to the switch with the message Login Required. Tap the Initial Config button to start the stack configuration.
3. On the Device Login screen, leave the Connection Type as Bluetooth. Enter the username admin and leave the password field blank, then tap the Log In button at the bottom of the screen.
4. On the Initial Config screen, tap the Stack button to begin stack setup. The app will automatically discover all switches that are connected to the primary via VSF link cables, connect to them via Bluetooth, and will display them in the topology view on the screen.
5. To change member IDs or assign a member as the stack secondary, tap the switch in the topology view. Enabling the LED switch for each switch causes the blue UID LED on its front panel to flash; use them to verify that the physical stack layout matches the displayed topology. Select **Configure Members** to apply member IDs and secondary configuration to all switches in the stack, which will cause each member other than the primary to reboot.
6. Once all switches have rebooted and joined the stack, the message Stack Set Up Successful! will be displayed. Select **Configure Stack** to continue.
7. If NetEdit will be used to manage the stack, enter the NetEdit server address, username, and password, then tap **Log In**; if not, then tap Skip.
8. Choose the desired switch management interface from the dropdown menu; configure the stack hostname, admin password, and management IP interface (static or DHCP); Alternately, you may deploy a custom configuration template (saved on your device or available on a connected file sharing service such as OneDrive, Dropbox, or iCloud Drive) by tapping the interface dropdown menu, and selecting **Import Custom Template...** Once the desired configuration has been selected, tap **Next** to continue.
9. Review the configuration generated by the app or imported from a template; then, tap **Deploy** to apply the configuration to the stack.
10. Once the configuration has been successfully deployed, the connection between the device and the switch (now the stack primary) will turn green, and the message **Device Deployment Successful!** will be displayed. Tap **Done** to return to the app's main page.

## Manual configuration

In cases where auto-stacking or the HPE Aruba Networking CX mobile app cannot be used to provision a stack, stack members can be configured manually.

## VSF links

The user can specify the interfaces which comprise the VSF links. Refer to [link](#) for information about specifying interfaces.

When the interface is configured, any existing configuration is removed, including VLAN memberships, ACL/Quality of Service rules and any speed/duplex/MTU configuration.

Once the interface becomes part of a VSF link, no protocol or feature will be allowed to run on it as it is now part of the fabric.



---

A VSF link will be a routed interface.

---

## Secondary member

When auto-stacking is used to provision a stack, member 2 is automatically designated as the secondary. A secondary member can be designated from the Conductor using any valid stack member ID other than 1 before or after that member has actually joined the stack.



---

Member number 1 can never be configured as a secondary member.

---

An existing non-secondary stack member that is designated as the secondary will reboot and rejoin the stack to assume the Standby role. If a member not present in the stack is designated as the secondary, that member will automatically assume the Standby role without an additional reboot when it subsequently joins the stack.

If a secondary member is already configured and present in the stack, removing the secondary designation will cause that member to reboot and rejoin the stack with the Member role.

Refer to `vsf secondary-member` for information about configuring a secondary member.

In the case of auto-stacking, member 2 is automatically configured as secondary member through LED button press or `vsf start-auto-stacking` command.

## Member number

To add a device to a VSF stack, the device must be renumbered to the corresponding member ID. The user can specify the member number of the switch. The default member number is 1.

- For the 6200F device, the default number can be changed to any value from 2 through 8. (The device supports up to 8 members.)
- For 6300 devices, the default number can be changed to any value from 2 through 10. (The device supports up to 10 members.)

Refer to [vsf renumber-to](#) and [Misconfiguration recovery](#) for information about renumbering a member.



---

Changing the member number causes the switch to reboot and all configuration on the switch is removed.

---

A switch with a member number other than 1 cannot boot completely unless it has reachability to a VSF conductor switch via VSF link. If a renumbered member is unable to communicate with the conductor switch and is waiting in booting state, the user can:

- Go to a recovery console with a `ctrl+c` sequence and collect the diagnostic information, or
- Reset the VSF configuration.

## Member provisioning

VSF allows the user to provision or pre-configure any member before the member is physically added to the stack. Provisioning the member allows the user to complete the required configuration as if the member is present in the stack. When the member eventually joins the stack, it will boot up with the configuration made on the pre-provisioned interfaces.

To provision a member, the part number of the member must be specified. Refer to [type](#) for information about provisioning a member.



---

If a member tries to join the stack with a different part number to the one provisioned on the Conductor, the renumbered member will be removed from the stack and will reboot with factory defaults.

---

## Access to VSF members

In addition to serial console connections, any stack member can be accessed from any other member using the `member` command.

Refer to [member](#) for information about console connection to a member switch.

## Stack management

### Consoles

The serial console of the Conductor switch provides a full CLI configuration interface for a user with valid credentials. The serial console of the other stack members, including the Standby, provides a reduced CLI configuration interface, with only a limited set of commands for troubleshooting the stack.

In a standard deployment, connect to the console interface of the conductor and standby switch. This enables the stack conductor console to be reachable after a stack failover to the new Conductor.



---

Any switch configuration or monitoring must be performed from the console of the stack Conductor switch only.

---

### Management interface

In a VSF stack, only the management interface on the Conductor switch will be assigned an IP address (configured or assigned by DHCP). The stack allows connectivity to management protocols and Console through the management interface on the Conductor.



---

The management interface is applicable only for 6200 and 6300 devices.

---

### Split detection

#### 4100i and 6100 Switch Series

4100i and 6100 switch series which do not have OOBM ports, use front-plane ports to detect active stack fragments upon split. These ports should be part of a VSF Multi-Active Detection (MAD) VLAN. Only one MAD VLAN can be configured. One front-plane port from primary and one front-plane port from secondary should be part of a MAD VLAN as access ports and interconnected either directly or over the same Layer 2 (L2) network.

#### Limitations:

- Any other protocol which may alter the functionality of split detection should not be enabled on MAD VLAN and front-plane ports used for split detection.
- Only one port from primary and secondary can be assigned to the MAD VLAN.
- Trunk and Aggregation ports cannot be assigned to the MAD VLAN.
- The default VLAN cannot also be the MAD VLAN.



For more information, refer to [vsf split-detect vlan](#).

### 6200, 6300 and 6300L Switch Series

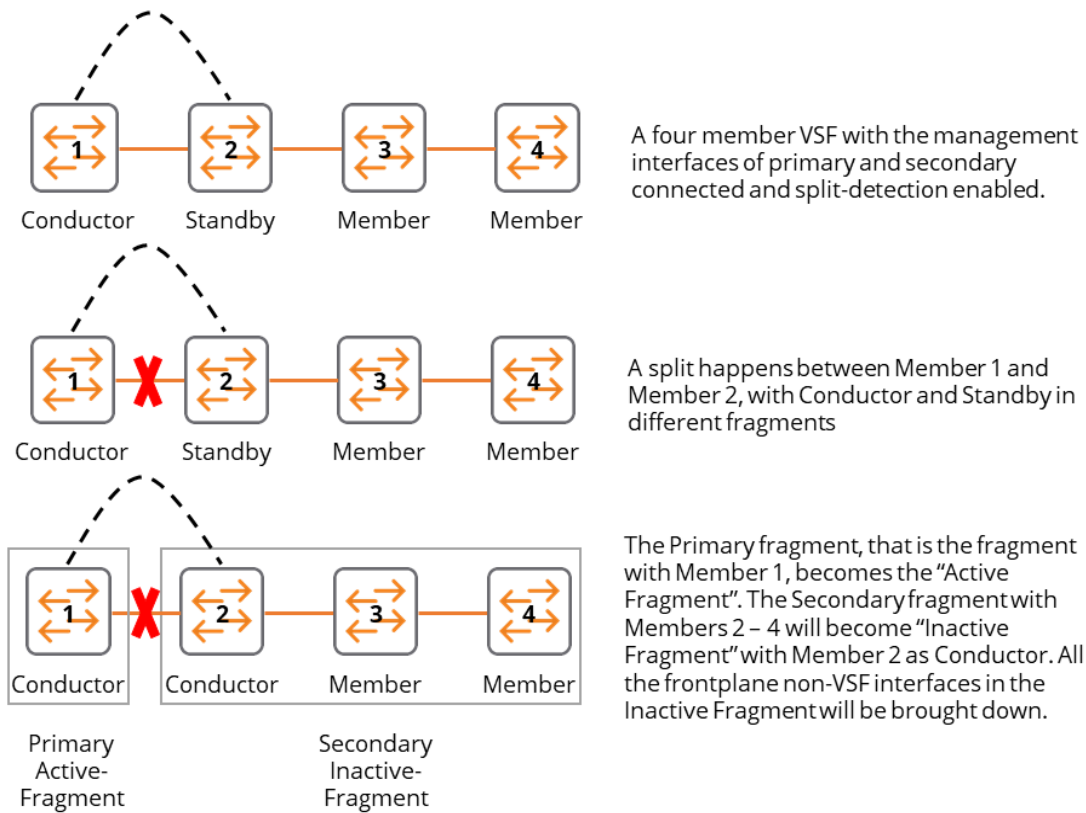
On 6200, 6300 and 6300L switch series VSF stack supports split detection utilizing the management interfaces, which requires users to connect the management interfaces of the primary and secondary stack members to the same L2 network.



It is also possible to connect the management interfaces of primary and secondary directly to one another for split detection.

In the event of a stack split where the primary and secondary members are on opposite sides of the split, if the secondary fragment discovers that the primary fragment is operational via the management port connection, it will bring down all front-plane non-VSF interfaces on the secondary fragment to minimize network disruption due to duplicate MAC or IP addresses.

The interfaces will remain down until the stack is reconnected or the primary fragment goes down. The interfaces of the primary fragment will always remain operational.



The `show vsf` output in the Primary fragment will look like this:

```

switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 08:97:34:b0:0e:00
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                 : Chain
Status                   : Active Fragment
Split Detection Method    : mgmt

Mbr Mac Address          type          Status
ID
-----
1  38:21:c7:5c:f4:c0     JL668A          Conductor
2                               JL668A          In Other Fragment
3                               JL668A          In Other Fragment
4                               JL668A          In Other Fragment
switch#
switch# show vsf topology
Conductor
+----+
| 1 |
+----+
switch#

```

The show vsf output in the secondary fragment will look like this:

```

switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 08:97:34:b0:0e:00
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                 : Chain
Status                   : Inactive Fragment
Split Detection Method    : mgmt

Mbr Mac Address          type          Status
ID
-----
1                               JL668A          In Other Fragment
2  38:21:c7:5c:77:40     JL668A          Conductor
3  38:21:c7:5a:a5:80     JL668A          Member
4  38:21:c7:5c:b3:00     JL668A          Member
switch#
switch# show vsf topology
Conductor
+----+ +----+ +----+
| 4 |1==2| 3 |1==2| 2 |
+----+ +----+ +----+
switch#

```



For more information, refer to [vsf split-detect mgmt.](#)

## Automated image sync

In a VSF environment, all stack members run the same software image. If the user upgrades the software on the Conductor by downloading a new software image using SFTP/TFTP, all members of the stack will simultaneously upgrade.

When forming a stack, if the software version on a member is different from the version of the Conductor, the member will automatically update itself to the same version as the Conductor. The member will reboot itself to run the updated version before joining the stack.



---

Automated image sync is not applicable if the conductor is running the firmware version 10.07 or later and the member is booted with firmware version 10.06 or earlier versions and vice-versa.

---

## Reboot

An individual stack member can be rebooted from a CLI command.

- The member will reboot and re-join the stack, with the same role that it had prior to the reboot.
- If the stack topology is a ring, no traffic disruption is expected on any other stack members when a single member is rebooted.
- If the stack topology is a chain, rebooting a member may cause a stack split, resulting in members being unreachable from the conductor. This result can cause significant disruption of the stack, so use this option with caution.
- If the member is the stack Standby, there will be no Standby in the stack until the member reboots and re-joins the stack. At this point, the member will again have the role of Standby.
- If the member is the stack Conductor, the command will trigger a failover and the Standby switch will take over as Conductor of the stack.
- If the Standby is unavailable at the time of conductor reboot, the whole stack will reboot.

The whole stack can also be rebooted by using the `boot system` command.

- All members will reboot and the stack will re-form.
- Traffic will be disrupted for the duration of the reboot.

Refer to [vsf member reboot](#) for information about rebooting a member.

## Member addition with auto-stacking

A new factory default switch can be added into an existing stack by physically connecting it to a given member of the stack on the auto-stacking reserved interfaces. The newly added member will be automatically assigned with member ID and go for a reboot. After reboot, the newly added member will join the stack.

For more information auto-stacking reserved interfaces, [Reserved interfaces for auto-stacking](#).

## Member addition without auto-stacking

A member can be added to the stack to augment an existing stack. The member being added can be a factory-default switch or a switch with pre-existing configuration.

1. Configure interfaces to VSF links on the member being added.
2. Renumber the member being added.



---

The member will not join the stack if there is a member number conflict.

---

3. Renumbering will cause a reboot of the switch.
4. Connect the configured VSF links to a previously configured VSF link on the stack.
5. The member joins the stack, with default configuration on its interfaces. Any previous configuration on the member will be lost.

## Member replacement with auto-stacking

Disconnect all the physical connections of the member that will be replaced and connect the new replacement member to the same interfaces as the switch being replaced. The new member joins the stack, with the same configuration as the member it is replacing.



---

The replacement member must be of the same part number as the switch being replaced.

---

## Member replacement without auto-stacking

The replacement member must be of the same part number as the switch being replaced.

1. Power off or disconnect all physical connections of the member that will be replaced.
2. Configure interfaces to VSF links on the replacement member. These interfaces **must** match the interfaces configured on the switch being replaced.
3. Renumber the replacement member to the same number as the switch being replaced.
4. Renumbering will cause a reboot of the switch.
5. Connect the replacement member to the stack.
6. The member joins the stack, with the same configuration as the member it is replacing.

## Member removal

A member can be removed from a running stack. All configuration associated with the member will be removed.

If the member is physically present in the stack at the time it is removed, all VSF configurations on that member will be erased and it will lose its identity as a member of the stack from which it was removed. The member will come back as member 1 with factory default configuration.



---

It is not advisable to remove the member that is the conductor of the stack. If the conductor has to be removed, the recommendation is to switch over and wait for the old conductor to come up as standby before removing it.

---

Refer to the `vsf member` command for information about removing a member.

Though it is not recommended as it can cause traffic outages, if an active member needs to be removed from a stack, member must be physically removed after issuing `no vsf member` command. Else, the member will join the stack back through auto-stacking. Alternatively, the links can be disabled first and the member can be removed from the conductor. The removed member must be reset to factory-default once it boots to recovery.

## Stack and Port LED states

The following table describes the different states of stack **Stk** LED for 6200 and 6300 switch series.

**Table 1:** *Stk LED States*

State	Meaning
On - Green	Stacking Mode is selected.
On - Amber	Stacking Mode is selected and stacking-related error has occurred.
Slow flash Amber	Stacking Mode is not selected but still stacking-related error has occurred.
Off	Stacking mode is not selected.

The following table describes the different states of speed **Spd** LED for 4100i and 6100 switch series.

**Table 2:** *Spd LED States*

State	Meaning
On - Green	Speed Mode is selected.
Slow flash Green	Stacking Mode is selected.
Off	Speed and Stacking mode are not selected.

The following table describes the different states of Port LED based on stack configurations and role of the members in the stack.

**Table 3:** *Port LED States*

State	Meaning
On - Green	Current Stack Member and is operational. For example, if Port 3 is on green, this indicates that the current chassis is member 3 in the stack.
Half-bright green	Total members in the Stack. Except port LEDs indicating the conductor and current member, all other ports LEDs glow half-bright green.
Slow flash green	Conductor of the Stack. In a six-member stack, one of the six port LEDs glows slow flashing green indicating that unit in the stack is the conductor. For example, if Stack Member 4 is the conductor, Port 4 LED glow slow flashing green.
On- Amber	The stack member is not reachable or in booting condition. When the member is fully booted and joined the stack successfully, then LED glows solid green.
Slow flash amber	The stack member is in a known fault condition. Only the global Status LED of faulted member glows slow flash amber.
Off	The stack Member does not exist in the stack.

The following sections describe the prerequisites and procedures to configure a VSF stack.


## Forming a four-member ring setup using auto-stacking command

### Prerequisites

- All switches must be in the factory default configuration.
- All stack members must be connected in a desired topology on the reserved VSF link ports. For more information, see [Reserved interfaces for auto-stacking](#).

In the following procedure, the `vsf start-auto-stacking` command is used to form a four-member stack with the ring topology:

**Figure 1** *Four-member ring setup*



### Procedure

1. Rack up all the four switches and physically connect them on the reserved auto-stacking interfaces in a ring setup. For more information on reserved interfaces, see



Alternatively, you can also add members one after another.

2. Designate the first switch of the rack using the `vsf start-auto-stacking` command as the conductor. The links and secondary member will be automatically configured on the conductor.

```
switch(config)# vsf start-auto-stacking
```

Since the switches are already physically connected, starting with the second switch, each switch in the stack reboots automatically and join the stack one after another automatically. The running configuration will appear as shown below:

```

6300# show run
Current configuration:
!
!Version AOS-CX FL.10.07.xxxx
!export-password: default
cli-session
    timeout 0
!
!
!
!
!
ssh server vrf default
ssh server vrf mgmt
vsf secondary-member 2
vsf member 1
    type jl668a
    link 1 1/1/26
    link 2 1/1/25
vsf member 2
    type jl668a
    link 1 2/1/25
    link 2 2/1/26
vsf member 3
    type jl668a
    link 1 3/1/26
    link 2 3/1/25
vsf member 4
    type jl668a
    link 1 4/1/25
    link 2 4/1/26

```

3. Issue a "show vsf" command to ensure that the ring has successfully formed. You can also verify stack formation using different LED states. For more information on LED states, [Stack and Port LED states](#).

```

6300(config)# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 70:72:cf:ef:b7:f2
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                 : Ring
Status                   : No Split
Split Detection Method    : None

```

Mbr ID	Mac Address	type	Status
1	70:72:cf:ef:b7:f2	JL668A	Conductor
2	90:20:c2:23:67:40	JL668A	Standby
3	90:20:c2:24:71:c0	JL668A	Member
4	38:21:c7:5a:33:40	JL668A	Member




If full stack configuration is downloaded on to the conductor through TFTP/ZTP, the physical connections between the switches should be made according to the downloaded configuration.

## Forming a four-member chain setup using link configuration command with auto-join

Manual configuration of a VSF stack requires the user to individually configure each switch in the stack.

**Figure 1** *Four-member chain setup*



### Procedure

1. Execute the following command on the member 1:

```
switch# configure
switch(config)# vsf member 1
switch(vsf-member-1)# link 1 1/1/26
switch(config)# vsf member 2
```

2. Physically connect member 2 to member 1 on the auto-stacking reserved interfaces. For example, if part number of the switch is JL659A, then reserved auto-stacking interfaces are 25 and 26. Once the physical connections are made, member 2 will reboot automatically and join the stack as standby switch with member 1 as the conductor. The running configuration on the conductor when member 2 join the stack will appear as shown below:

```

vsf secondary-member 2
vsf member 1
  type jl659a
  link 1 1/1/26
vsf member 2
  type jl661a
  link 1 2/1/25

```



The conductor's VSF link must be connected to interface of higher value. For example, when you are connecting member 2 to member 1, you must connect interface 26 of member 1 to interface 25 of member 2. Otherwise, the member 2 will not join in auto-stacking.

3. Physically connect member 3 to member 2 on the auto-stacking reserved interfaces. The member 3 will reboot automatically and join the stack as member.
4. Repeat the step 3 for stack member 4. Once the auto-stacking process is complete, member 4 will reboot automatically and join the stack as member 4. This forms a chain topology.
5. Issue a "show vsf" command to ensure that the ring has successfully formed. You can also verify the stack formation using LED states. For more information on LED states, [Stack and Port LED states](#).

```

switch(config)# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 70:72:cf:ef:b7:f2
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                  : Chain
Status                    : No Split
Split Detection Method    : None

```

Mrbr ID	Mac Address	type	Status
1	70:72:cf:ef:b7:f2	JL659A	Conductor
2	90:20:c2:23:67:40	JL661A	Standby
3	90:20:c2:24:71:c0	JL668A	Member
4	38:21:c7:5a:33:40	JL668A	Member

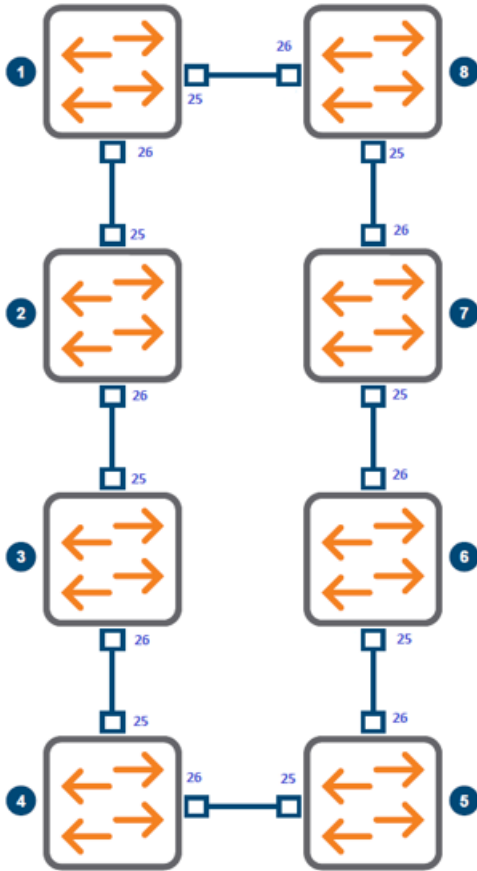


If full stack configuration is downloaded into the conductor through TFTP/ZTP, the physical connections between the switches should be made according to the downloaded configuration.

## Forming an eight-member ring setup manually using link configuration without auto-stacking

Manual configuration of a VSF stack requires the user to individually configure each switch in the stack. This process provides the best control for the user to configure VSF member number and links.

**Figure 1** *Eight-member ring setup*



## Procedure

To form an eight-member ring setup as shown, do not make the connections initially. Connect the ports only after each device is fully configured.

1. Log in to the first device, numbered 1.
  - a. The default member number is 1, so no member number change is required.
  - b. At the prompt, enter the following commands:

```
switch# configure
switch(config)# vsf member 1
switch(vsf-member-1)# link 1 1/1/25
switch(vsf-member-1)# link 2 1/1/26
```

- c. The preceding sequence of commands will configure the links for member 1.
  - d. Ports 25 and 26 are configured as link 1 and 2 respectively.
2. Log in to the second device, numbered 2.
  - a. Execute the following commands.

```
switch# configure
switch (config)# vsf member 1
switch(vsf-member-1)# link 1 1/1/25
switch(vsf-member-1)# link 2 1/1/26
switch(vsf-member-1)# exit
switch(config)# vsf renumber-to 2
```

```
This will save the VSF configuration and reboot the switch.
Do you want to continue (y/n)? y
```

- b. The preceding sequence of commands will configure the links on member 2.
  - c. The default member number is "1". The command "vsf renumber-to" changes this member number.
  - d. Links are configured **before** renumbering, and the member identifier in the interface name is "1" at this point.
  - e. The switch will reboot after executing the renumber command.
3. Physically connect member 2 to member 1 as shown in the figure.
    - a. This action will cause member 2 to join the stack, with member 1 as the conductor.
    - b. This result can be verified by executing "show vsf" on member 1.
  4. Repeat steps 2 and 3, for each stack member 3 through 8.
    - a. Be sure to specify the member number correctly on each member.
    - b. If a member number conflict is detected, the member will NOT join the stack.
  5. Once member 8 has successfully joined the stack, connect member 8 link 2 to member 1 link 1, to complete the ring.

Issue a `show vsf` command to ensure that the ring has successfully formed.

```
switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 38:21:c7:5d:d0:c0
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                 :
Topology                  : Ring
Status                    : Active Fragment
Split Detection Method    : None

Mbr Mac Address          type          Status
ID -----
1 38:21:c7:5d:d0:c0     JL668A      Conductor
2 38:21:c7:6a:10:c0     JL668A      Member
3 38:21:c7:5c:15:80     JL668A      Member
4 38:21:c7:5a:61:40     JL668A      Member
5 38:21:c7:62:66:00     JL668A      Member
6 38:21:c7:58:22:40     JL668A      Member
7 38:21:c7:5a:9c:00     JL668A      Member
8 38:21:c7:63:a5:00     JL668A      Member
```

6. The preceding steps will form an eight-member stack without a standby. To make any member the standby (for example, member 8), use the secondary command:
  - a. From the primary VSF member, configure member 8 as VSF secondary member:

```
swtich(config)# vsf secondary-member 8
This will save the configuration and reboot the specified switch.
Do you want to continue (y/n)? y
```

```
switch(config)#
```

- b. This action will reboot member 8 and it will rejoin as standby.

```
switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 38:21:c7:5d:d0:c0
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 8
Topology                  : Ring
Status                    : Active Fragment
Split Detection Method    : None

Mbr Mac Address          type          Status
ID
-----
1  38:21:c7:5d:d0:c0    JL668A    Conductor
2  38:21:c7:6a:10:c0    JL668A    Member
3  38:21:c7:5c:15:80    JL668A    Member
4  38:21:c7:5a:61:40    JL668A    Member
5  38:21:c7:62:66:00    JL668A    Member
6  38:21:c7:58:22:40    JL668A    Member
7  38:21:c7:5a:9c:00    JL668A    Member
8  38:21:c7:63:a5:00    JL668A    Standby
```

7. Alternatively, before adding member 8 to the stack, pre-configure the secondary as 8 and then renumber device 8. This action will ensure that device 8 will join the stack directly as standby.

## VSF In-Service Software Upgrade

---

In VSF In-Service Software Upgrade (ISSU), the firmware image of the stack can be updated dynamically with minimal downtime. The conductor, standby and members of the stack are upgraded dynamically without having the necessity for them to reboot at all. This means that the traffic won't be interrupted during the upgrade and if an error happens and traffic is interrupted to keep the interruption as short as possible. HPE Aruba Networking's switch architecture decouples the control plane and the data plane, so it is possible to implement an ISSU mechanism by keeping the data plane up while the control plane manages the upgrade.



---

VSF ISSU is supported *only* on the HPE Aruba Networking 6300 Switch Series (except for S3L75A,S3L76A and S3L77A).

---



---

VSF ISSU is supported for upgrades between minor releases.

---



---

The VSF ISSU feature uses ISSU commands for its operations (configuration & show commands). Refer to the [VSF commands](#) section for command details.

---

### Upgrade scope

The current ISSU feature is intended for upgrades between minor releases. This is suitable for stable networks that need critical bug fixes or security fixes instead of full version upgrades to get new functionality.

### VSF ISSU operation

ISSU in the 6300 consists of a series of steps designed to keep the system in a stable state while the firmware is upgraded and to recover this state after such upgrade.

#### Step 1: Image download and validation

The first step of the process is to download the image with the new version of the firmware to the alternate flash. Verification is performed by the switch that it complies with the internal rules and the security checks (e.g. valid signed image).

Upon successful verification, the new image will be installed in the bank that doesn't store the image currently running on the switch. The old image will be kept in the alternative bank in case a fail back process needs to be triggered due to an error.

The ISSU validation triggers after the ISSU operation is started. User can perform the validation without triggering ISSU as well by using the CLI. For more information, refer the [issu update-software](#) command.

#### Step 2: Install the image in the standby & members

Once the new image is verified, the standby and members of the stack will be upgraded with this image without having them to reboot.

### **Step 3: Freeze the system**

Once the standby and members of the stack are successfully upgraded, the system's state will be frozen. This means that no further configuration will be allowed until the ISSU is complete. In fact, no configuration via CLI, REST, or SNMP is allowed once the ISSU process starts (Step 1) in order to make sure that the system doesn't become unstable once ISSU has started. Informational queries are still allowed. Any change in the network such as MAC learning will not be reflected until ISSU is complete. Please refer to [Information not updated during ISSU](#) for additional information.

### **Step 4: Perform the upgrade for the line modules**

The Line Card (LC) modules services on standby and all the members of the stack will be upgraded at this time, however they will not be rebooted.

### **Step 5: Recover the previous state**

The previous state will be recovered and synchronized with the new firmware running in the standby. Some features can't recover the previous state at this point and therefore do not support uninterrupted operation during ISSU. Please refer to [Features that don't support uninterrupted operation during ISSU](#) for more information.

### **Step 6: Failover occurs between the conductor and the standby**

The Standby will take over and becomes the new conductor running the new firmware version. All sessions will be closed and the user will be required to log in again.

### **Step 7: The previous Conductor will upgrade its image**

The previous conductor will upgrade its image and take the role of Standby and joins the stack without reboot. The ISSU is complete at this point.

## **VSF ISSU considerations**

The first requirement for ISSU is to ensure that traffic continues to pass through the stack uninterrupted during the upgrade operation. In the event that something goes wrong traffic must be restored as soon as possible. This could mean falling back to the previous image, or in some cases to fall forward to the new one.

The upgrade process can take minutes to complete. During that time as many control plane protocols as possible will remain operational, however some of them need to stop processing packets or avoid reacting to events. Please refer to [VSF ISSU limitations](#) for additional information. The amount of time when they are not operational will be kept as small as possible.

In order to start ISSU, the stack topology must be a Ring. The Standby and all members must be operational and running. If any of them is in another state (e.g., booting) then the ISSU process will not start.

## VSF ISSU limitations

The ISSU process currently has the following limitations:

### Hardware faults

It is recommended that hardware faults are resolved prior to performing ISSU. Fan faults and abnormal temperature statuses should be addressed. Fans may increase in speed during the ISSU process.

### Prohibited operations

During ISSU the following operations are not permitted:

- Any hardware hot swap
- Pressing any button

Performing any of these operations during an ISSU could put the system in an error state that could adversely affect the ISSU.

The system may not be able to address failures during the ISSU. For example, if network protocols require blocking a port, flush MAC, or routing tables, etc. the switch would not be able to react to those tasks until the ISSU is complete.

### Recommendations

Before initiating ISSU, it is recommended to enable spanning-tree root-guard on the supposed designated ports of the VSF stack. This ensures and enforces the root bridge placement in the network.

### Scheduled jobs

Since no configuration is allowed during the ISSU, it is possible that some scheduled jobs will fail if they are attempted while the ISSU is in progress. It is recommended to not execute any scheduled job during ISSU.

### Information not updated during ISSU

During ISSU the following information may not be current until after ISSU completes:

- Link state. Transitions may not be reflected until ISSU is done.
- Temperatures.
- Power consumption values.
- System state as indicated by LEDs.
- MAC Learning.
- Change in MAC for a neighbor via Gratuitous ARP.
- Some counters will stop increasing their value during ISSU.

### Show tech

The show tech command is enabled during ISSU using the CLI, however it isn't possible to use it with the REST API.

### Features that don't support uninterrupted operation during ISSU

During ISSU the following features may be interrupted until after ISSU completes:

## AAA

1. During ISSU new client authentication will occur on other members except the VSF conductor that is rebooting. Clients will not be able to access to the network until ISSU is complete.
2. Session timeout/inactivity timeout applied to clients will be delayed. Clients will remain longer than the configured timer.

## PIM with MSDP

Currently, PIM is not supported in ISSU on RP routers, therefore traffic loss is expected during ISSU on RP routers where PIM and MSDP are enabled.

## Dynauthz ISSU

All Dynamic Authorization requests will be rejected with NAK during ISSU operation. Captive portal workflow will be affected due to NAK received for Dynauthz requests during ISSU.

## BFD ISSU

It's important to note that this mechanism is only recommended for sessions that have a failure detection time higher than 6 seconds (e.g. Tx=Rx=3000 and Detect Multiplier=2, or Tx=Rx=2000 and Detect Multiplier=3). The setup should have less than 64 sessions, otherwise the ISSU failover might not be hitless for all sessions.

## NAE

NAE will have a momentary disruption during ISSU events due to the failover event. See the Network Analytics Engine Guide for more information.

## Feature readiness check

ISSU and feature readiness check will fail if the feature ready check does not pass. The feature ready check validates that none of the following features are in a failed state:

- ACLs, object groups, and their applications
- Classifier policies, classes, and their applications
- Port-access polices, classes, and their applications
- Port-access group-based policies, classes, and their applications
- Active PBR actions and their applications

In general, if any of the above is in a failed state it is recommended to fix it and then upgrade to the newer operating system image using the system reboot method instead of ISSU. Some failed states will remain even after it is fixed on a running system and a reboot is good best practice for clearing it.

## Restrictions

The following caveats apply:

1. VSF ISSU is not supported on 4100i, 6100 or 6200 platforms.
2. VSF ISSU is not supported on S3L75A, S3L76A, and S3L77A SKUs.

## VSF ISSU errors and recovery

If any VSF stack member encounters an error during the "Initiate ISSU" or "Validate System Readiness" phases then the ISSU process will be aborted.

If any VSF stack member encounters an error during the "Upgrade Standby and Member Modules" or "Upgrade Line Module Services" phases, then the stack will be unable to recover. The entire stack will reboot with the upgraded image directly.

If any stack member encounters an error during failover, the whole stack will be rebooted. The system will boot with the new image. The error will be handled as a fail forward.

### clear vsf split-detection statistics

```
clear vsf split-detection statistic
```

#### Description

Resets the VSF split-detection statistics counters to zero.

#### Example

The following example runs the **clear vsf split-detection statistics** command to clear VSF split-detection statistic counters, then uses the optional **show vsf split-detection statistics** command to verify that these statistics were correctly reset back to zero:

```
switch# clear vsf split-detection statistics
switch#
switch# show vsf split-detection statistics
VSF Split Detection Statistics:
=====
Total Packets Transmitted           = 0
Total Packets Received              = 0
Total Packets Received And Discarded = 0
```

#### Command History

Release	Modification
10.16.1000	Command introduced

#### Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

### description

```
description
no description
```

#### Description

Adds a description for one or more VSF link interfaces.

The **no** form of this command removes the interface description.

## Examples

Adding a description for VSF link interface **1/1/25**:

```
switch(config)# interface 1/1/25  
switch(config-if-vsfc)# description mem 1 intf 1/1/25
```

Removing the description from interface 1/1/25

```
switch(config)# int 1/1/25  
switch(config-if-vsfc)# no description
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.10	Command introduced.

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config-if-vsfc	Administrators or local user group members with execution rights for this command.

## interface

interface <IFRANGE>

### Description

Enters configuration context for one or more VSF link interfaces.



VSF link interfaces cannot be included in a range with other interfaces.

Parameter	Description
<IFRANGE>	Port identifier range. Required.

## Examples

Entering configuration context:

```
switch(config)# interface 1/1/1
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config	Administrators or local user group members with execution rights for this command.

## issu rollback-timer

```
issu rollback-timer [wait-time <TIME>]
```

### Description

Enables the ISSU rollback timer on the system. The rollback timer automatically rolls the system back to the configuration and OS image used before starting the ISSU, unless the upgrade is confirmed with **issu update-software confirm**. Changing the rollback-timer will not affect an active timer and will apply on the next ISSU. To cancel the active timer, confirm the previous ISSU with **issu update-software confirm**. Disabled by default.

The **no** form of this command disables the rollback timer on the system.

Parameter	Description
<TIME>	Specifies how many minutes the system will wait for confirmation that the last ISSU is accepted before triggering a system reboot and roll back to the previous configuration and OS version. This change will not affect an active timer and will apply on the next ISSU. Range: 30-1440.

### Examples

Enabling the ISSU rollback timer:

```
switch(config)# issu rollback-timer
```

Disabling the ISSU rollback timer:

```
switch(config)# no issu rollback-timer
```

Disabling the ISSU rollback timer on the system while a previous ISSU's timer is active:

```
switch(config)# no issu rollback-timer
The ISSU rollback timer is active. This change will apply on the next ISSU
operation
```

```
To cancel the active timer, confirm the previous ISSU with 'issu update-software confirm'
```

Setting the ISSU rollback timer wait time to 80 minutes:

```
switch# issu rollback-timer wait-time 80
```

Setting the ISSU rollback timer wait time to 81 minutes while a previous ISSU's timer is active:

```
switch# issu rollback-timer wait-time 81  
The ISSU rollback timer is active. This change will apply on the next ISSU operation
```

Resetting the ISSU rollback timer wait time to default:

```
switch(config)# no issu rollback-timer wait-time
```

## Command History

Release	Modification
10.11	Command introduced

## Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A and S3L77A	config	Administrators or local user group members with execution rights for this command.

# issu update-software

```
issu update-software [validate|confirm]
```

## Description

Initiates ISSU to the alternate boot location. The newer operating system image must be downloaded to the alternate boot location prior to running this command. Additionally, the current running operating system version must match the version that is stored in the current boot location or ISSU will not be allowed.



This operation is disruptive and will result in the management interface being disconnected during the process.

Starting an ISSU will cause the running configuration to be saved in case an error is encountered that requires a reboot to recover the switch. In addition, a special configuration checkpoint will be stored to disk and is used to roll back to the pre-ISSU firmware. During the ISSU process all management

methods (CLI, REST, WebUI) will be blocked from making configuration changes to the switch. The configuration block is active from the time ISSU starts until the time ISSU switchover is complete. After the ISSU switchover is completed, switch configuration can resume.

Parameter	Description
<code>validate</code>	Runs all pre-ISSU validations without executing the actual upgrade. The validation runs in the background, however its results will be displayed in real time for approximately the next three minutes. If the validation is not finished within that time frame or if the display is aborted with Control+C or Control+Z, the results can be queried using the <code>show issu validation</code> command.
<code>confirm</code>	Confirms the software update and cancels the configured rollback timer. If the rollback timer is configured then this command has to be executed after an ISSU before the timer expires. Else, the pre-ISSU checkpoint is copied to the startup configuration and the system is rebooted to the image used before ISSU. <b>Note: To perform an intentional system rollback before the timer expires, a manual downgrade must be executed through the following steps:</b> <ol style="list-style-type: none"><li>1. Copy the pre-ISSU checkpoint to the startup configuration using the <code>copy checkpoint pre-issu-startup-config startup-config</code> command.</li><li>2. Boot to the previous image using the <code>boot system</code> command.</li></ol>

## Usage

Note the following points before running this command:

- The newer operating system image must be downloaded to the alternate boot location prior to running this command.
- The current running operating system version must match the version that is stored in the current boot location or ISSU will not be allowed.
- This operation is disruptive and will result in the management interface being disconnected during the process.
- The running configuration will be stored in case an error is encountered that requires a reboot to recover the switch.
- During the ISSU process all management methods (CLI, REST, WebUI) are blocked from making configuration changes to the switch. The configuration changes are not allowed from the time ISSU starts until ISSU switchover is complete. After the ISSU switchover is completed, switch configuration can resume.
- In case of ISSU, conductor gets transitioned to standby role without reboot.
- The stack topology must be a ring before initiating ISSU. ISSU is not supported in chain topologies and the process is aborted if the ISSU is initiated.
- During the ISSU process, the **show core-dump all**, **show tech all**, and **copy support-files** commands all may fail to run or display correct output.

## Examples

Initiating an ISSU:

```
switch# issu update-software

This command will perform an in-service software upgrade
using pre-staged secondary operating system image
FL.10.13.1000M

This will save the current running configuration

WARNING:

The rollback timer is enabled and configured to 30 minutes.

After the upgrade is done, execute "issu update-software confirm"
to confirm the new image works as expected. If the command is not
entered, the system will be rebooted to the previous version.

Continue (y/n)? y

Starting in-service software upgrade.

Use "show issu" to monitor status and progress.

Use "show events -c issu" to view event notifications.
```

Initiating an ISSU, but stopping it without confirming the upgrade:

```
switch(config)# issu update-software

This command will perform an in-service software upgrade
using pre-staged secondary operating system image
FL.10.13.1000M
This will save the current running configuration

WARNING:

The rollback timer is enabled and configured to 30 minutes.
After the upgrade is done, execute "issu update-software confirm"
to confirm the new image works as expected. If the command is not
entered, the system will be rebooted to the previous version.

Continue (y/n)? n
In-service software upgrade aborted. No changes were made.
```

Confirming the ISSU configuration when the rollback timer has been configured and started:

```
switch# issu update-software confirm

The ISSU has been confirmed and the rollback timer has been cancelled.
```

Confirming the ISSU configuration when the rollback timer has not started:

```
switch# issu update-software confirm

No rollback timer has been started, no action was done.
```

Executing an ISSU "dry run" where all pre-ISSU validations are run without executing the actual upgrade:

```
switch# issu update-software validate

ISSU Validation

=====
Condition                               Status
-----
Current Image Valid                     ---
Target Image Valid                       ---
Target Version Compatible                 ---
Management Modules Ready                 ---
Line Modules Ready                       ---
Features Ready                           ---

In Progress[/]
```

Executing an ISSU "dry run" where all pre-ISSU validations are run and the user aborts the ISSU validation on screen display without user confirmation:

```
switch# issu update-software validate

ISSU Validation

=====
Condition                               Status
-----
Current Image Valid                     Pass
Target Image Valid                       Pass
Target Version Compatible                 Failed
Management Modules Ready                 ---
Line Modules Ready                       ---
Features Ready                           ---

In Progress[\]
To view the validation progress and results, execute "show issu validation"
```

Executing an ISSU "dry run" where all pre-ISSU validations are run without executing the actual upgrade and the validation progress has finished successfully:

```
switch# issu update-software validate

ISSU Validation

=====
Condition                               Status
-----
Current Image Valid                     Pass
Target Image Valid                       Pass
```

```

Target Version Compatible    Pass
Management Modules Ready    Pass
Line Modules Ready          Pass
Features Ready               Pass

ISSU Validation has completed

```

Executing an ISSU validation while a previous ISSU is unconfirmed, i.e. the rollback timer is still running:

```

switch# issu update-software validate
The previous ISSU has not been confirmed. Please confirm it with
'issu update-software confirm' before starting a new ISSU or running a validation.

```

Executing an ISSU "dry run" when the validations are taking more than three minutes to complete, then checking the result of the validation afterwards:

```

switch# issu update-software validate
ISSU validation is taking longer than expected. Check the final result with 'show
issu validation'

switch# show issu validation

ISSU Validation
=====

Condition                               Status
-----
Current Image Valid                      Pass
Target Image Valid                       Pass
Target Version Compatible                 Pass
Management Modules Ready                  Pass
Line Modules Ready                        Pass
Features Ready                            Pass

```

Following example shows ISSU performed with chain topology:

```

switch# issu update-software
Stack topology is not a ring. ISSU upgrade aborted.

```

## Command History

Release	Modification
10.11	Validate and confirm parameters added.
10.10	Command introduced.

## Command Information

Platforms	Command context	Authority
6300, except	config	Administrators or local user group members with execution rights for this command.

for S3L75A,  
S3L76A and  
S3L77A

## link

```
link <LINK-ID> [<IFRANGE>][description <DESCRIPTION>]
no link <LINK-ID> [<IFRANGE>][description <DESCRIPTION>]
```

### Description

Creates or modifies a VSF link. The user can specify the physical interfaces that make up the VSF link. Once an interface is part of a VSF link, all existing configuration on the interface is removed and the interface will operate as a VSF interface. At least one interface must be specified for the creation of a VSF link. VSF interfaces carry VSF traffic and can only be connected to other VSF interfaces. Before removing an individual interface from the VSF link using the `no vsf link <x> <interface>` command, ensure that the interface is administratively shutdown at both local and peer ends.



Interface(s) configured with MACsec cannot be added as VSF links. You have to remove the MACsec configuration before adding an interface to a VSF link.

The **no** form of the command can be used to remove interfaces from a link or remove the link completely.



When configuration is removed from a link, it may cause the stack to split.

Parameter	Description
<LINK-ID>	The VSF link number. Range: 1 to 2.
<IFRANGE>	The interface identifier range.
<DESCRIPTION>	Adds a description for the link. Range: 1 to 64 printable ASCII characters.

### Examples

Creating a VSF link called **link 1** with an interface range of **1/1/51** and a description, and a VSF link called **link 2** with an interface range of **1/1/52**:

```
switch(vsf-member-1)# link 1 1/1/51
switch(vsf-member-1)# link 1 description link 1 connected to member 2
switch(vsf-member-1)# link 2 1/1/52
```

Removing VSF **link 1** and **link 2** completely:

```
switch(vsf-member-1)# no link 1
switch(vsf-member-1)# no link 2
```

Removing an assigned interface **1/1/51** from VSF **link 1**:

```
switch(vsf-member-1)# no link 1 1/1/51
```

Attempting to add an interface configured with MACsec to a VSF link:

```
switch(vsf-member-1)# link 1 1/1/51
VSF link cannot be configured on an interface with MACsec policy enabled.
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.10	Added the <b>description</b> parameter.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	vsf-member- <i>&lt;ID&gt;</i>	Administrators or local user group members with execution rights for this command.

## member

member *<MEMBER-ID>*

### Description

Connects to the specified member in a VSF environment.

Parameter	Description
<i>&lt;MEMBER-ID&gt;</i>	VSF member ID. <ul style="list-style-type: none"><li>Range for 4100i devices: 1-4</li><li>Range for 6100 devices: 1-6</li><li>Range for 6200F devices: 1-8.</li><li>Range for 6300 devices: 1-10.</li></ul>

### Examples

VSF stack is formed with two members:

```
switch# member 2
admin@172.17.17.2's password:

Last login: 2019-09-30 11:42:17 from the console
User "admin" has logged in 1 time in the past 30 days
```

```
member-2#
```

Member to self:

```
switch# member 1  
Already on member id 1
```

VSF stack is not formed and member not available:

```
switch# member 2  
No stack role for member id 2
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	Administrators or local user group members with execution rights for this command.

## show issu

```
show issu [brief|history|validation]
```

### Description

Shows information about the current state of ISSU. If no ISSU is currently in progress, the command displays the progress details of the last ISSU.

The command with the `brief` parameter displays a short summary of the ISSU state and indicates if the system is ready to accept an ISSU command and whether or not an ISSU is in progress. If the **brief** parameter is not included, then more details about an in progress ISSU or the last ISSU are displayed.

Parameter	Description
brief	Shows a short summary of the ISSU state.
history	Shows details of ISSU software update history.
validation	Shows information about the current state of an ISSU validation.

### Examples

Showing detailed ISSU status with an ISSU in progress for the first time:

```
switch# show issu
ISSU Summary
=====
ISSU Status      : In Progress
Current Version  : FL.10.13.0005K      Upgrade Version : FL.10.13.1000M
Upgrade Image    : secondary          Start Date      : 2023-11-08 07:01:47
Last ISSU Result: --
Rollback timer   : Not started

ISSU Progress
=====
Upgrade Operation      Status      Start Date
-----
Initiate ISSU          Complete    2023-11-08 07:01:47
Validate System Readiness In Progress 2023-11-08 07:01:47
Upgrade Standby and Member Modules Pending     --
Upgrade Line Module Services Pending     --
Prepare for Switchover Pending     --
Finalize Upgrade      Pending     --
ISSU Complete         Pending     --
```

Showing detailed status for VSF ISSU:

```
switch# show issu
ISSU Summary
=====
ISSU Status      : In Progress
Current Version  : FL.10.11.0001      Upgrade Version : FL.10.11.1000BD
Upgrade Image    : secondary          Start Date      : 2023-02-02 14:22:31
Last ISSU Result: --
Rollback timer   : Not started

ISSU Progress
=====
Upgrade Operation      Status      Start Date
-----
Initiate ISSU          Complete    2023-02-02 14:22:31
Validate System Readiness Complete    2023-02-02 14:22:31
Upgrade Standby and Member Modules In Progress 2023-02-02 14:22:54
Upgrade Line Module Services Pending     --
Prepare for Switchover Pending     --
Finalize Upgrade      Pending     --
ISSU Complete         Pending     --
```

Showing detailed ISSU status with ISSU in progress after successfully completing a previous ISSU:

```
switch# show issu
ISSU Summary
=====
ISSU Status      : In progress
Current Version  : FL.10.10.0001      Upgrade Version : FL.10.10.0002
Upgrade Image    : secondary          Start Date      : 2021-10-15 08:37:49
Last ISSU Result: Completed (Without errors)

ISSU Progress
=====
Upgrade Operation      Status      Start Date
```

```

-----
Initiate ISSU                               Complete      2021-10-13 23:05:41
Validate System Readiness                   Complete      2021-10-13 23:05:41
Upgrade Standby Management Module           Complete      2021-10-13 23:05:41
Upgrade Line Modules                        In Progress   2021-10-13 23:07:07
Prepare for Switchover                     Pending       --
Finalize Upgrade                           Pending       --
ISSU Complete                              Pending       --

```

Showing detailed ISSU status with ISSU in progress after aborting the previous ISSU:

```

switch# show issu
ISSU Summary
=====
ISSU Status      : In progress
Current Version  : FL.10.10.0001      Upgrade Version : FL.10.10.0002
Upgrade Image    : secondary          Start Date      : 2021-10-15 08:37:49
Last ISSU Result: Aborted (One or more line modules are not ready to start ISSU)

ISSU Progress
=====
Upgrade Operation      Status      Start Date
-----
Initiate ISSU         Complete    2021-10-13 23:05:41
Validate System Readiness Complete    2021-10-13 23:05:41
Upgrade Standby Management Module Complete    2021-10-13 23:05:41
Upgrade Line Modules  In Progress 2021-10-13 23:07:07
Prepare for Switchover Pending     --
Finalize Upgrade      Pending     --
ISSU Complete         Pending     --

```

Showing detailed ISSU status with no ISSU in progress and no previous ISSU performed:

```

switch# show issu
ISSU Summary
=====
ISSU Status      : Ready
Current Version  : FL.10.10.0001      Upgrade Version : --
Upgrade Image    : --                  Start Date      : --
Last ISSU Result: -- (--)

ISSU Progress
=====
Upgrade Operation      Status      Start Date
-----
Initiate ISSU         --          --
Validate System Readiness --          --
Upgrade Standby Management Module --          --
Upgrade Line Modules  --          --
Prepare for Switchover --          --
Perform Switchover   --          --
Finalize Upgrade      --          --
ISSU Complete         --          --

```

Showing detailed ISSU status after completion and before system is ready to start a new ISSU:

```

switch# show issu
ISSU Summary

```

```

=====
ISSU Status      : Not ready
Current Version  : FL.10.10.0001      Upgrade Version : --
Upgrade Image   : --                  Start Date      : --
Last ISSU Result: Completed (Without errors)

ISSU Progress
=====
Upgrade Operation      Status      Start Date
-----
Initiate ISSU          Complete   2021-10-13 23:05:41
Validate System Readiness Complete   2021-10-13 23:05:41
Upgrade Standby Management Module Complete   2021-10-13 23:05:41
Upgrade Line Modules   Complete   2021-10-13 23:07:07
Prepare for Switchover Complete   2021-10-13 23:07:50
Finalize Upgrade       Complete   2021-10-13 23:07:53
ISSU Complete          Complete   2021-10-13 23:08:10

```

Showing detailed ISSU status after an error occurred and the process is aborted:

```

switch# show issu
ISSU Summary
=====
ISSU Status      : Aborted
Current Version  : FL.10.10.0001      Upgrade Version : FL.10.10.0002
Upgrade Image   : secondary           Start Date      : 2021-12-09 19:17:15
Last ISSU Result: Aborted (System failed to prepare for ISSU)

ISSU Progress
=====
Upgrade Operation      Status      Start Date
-----
Initiate ISSU          Complete   2021-12-09 19:17:15
Validate System Readiness Complete   2021-12-09 19:17:15
Upgrade Standby Management Module Complete   2021-12-09 19:17:15
Upgrade Line Modules   Error      2021-12-09 19:19:22
Prepare for Switchover Aborted    --
Finalize Upgrade       Aborted    --
ISSU Complete          Aborted    --

```

Showing summary of ISSU status with no ISSU in progress where system is ready to start a new ISSU and with no previous ISSU performed:

```

switch# show issu brief
ISSU Summary
=====
ISSU Status      : Ready
Current Version  : FL.10.10.0001      Upgrade Version : --
Upgrade Image   : --                  Start Date      : --
Last ISSU Result: -- (--)

```

Showing summary of ISSU status with no ISSU in progress where system is ready to start a new ISSU and after successfully completing a previous ISSU:

```

switch# show issu brief
ISSU Summary
=====

```

```
ISSU Status      : Ready
Current Version  : FL.10.10.0001      Upgrade Version : --
Upgrade Image    : --                  Start Date      : --
Last ISSU Result: Completed (Without errors)
```

Showing a summary of ISSU status with no ISSU in progress where system is ready to start a new ISSU and after aborting the previous ISSU:

```
switch# show issu brief
ISSU Summary
=====
ISSU Status      : Ready
Current Version  : FL.10.10.0001      Upgrade Version : --
Upgrade Image    : --                  Start Date      : --
Last ISSU Result: Aborted (One or more line modules are not ready to start ISSU)
```

Showing a summary of ISSU status with ISSU in progress:

```
switch# show issu brief
ISSU Summary
=====
ISSU Status      : In progress
Current Version  : FL.10.10.0001      Upgrade Version : FL.10.10.0002
Upgrade Image    : secondary          Start Date      : 2021-10-15 08:37:49
Last ISSU Result: Completed (Without errors)
```

Showing a summary of ISSU status with no ISSU in progress where the system is not ready to start a new ISSU:

```
switch# show issu brief
ISSU Summary
=====
ISSU Status      : Not ready
Current Version  : FL.10.10.0001      Upgrade Version : FL.10.10.0002
Upgrade Image    : secondary          Start Date      : 2021-10-15 08:37:49
Last ISSU Result: Completed (Without errors)
```

Showing summary of ISSU status after an error occurred and the process is aborted:

```
switch# show issu brief
ISSU Summary
=====
ISSU Status      : Aborted
Current Version  : FL.10.10.0001      Upgrade Version : FL.10.10.0002
Upgrade Image    : secondary          Start Date      : 2021-12-09 19:17:15
Last ISSU Result: Aborted (System failed to prepare for ISSU)
```

Showing ISSU validation status per condition/validation:

```
switch# show issu validation

ISSU Validation
```

```

=====
Condition                               Status
-----
Current Image Valid                     ---
Target Image Valid                      ---
Target Version Compatible                ---
Management Modules Ready                ---
Line Modules Ready                      ---
Features Ready                          ---

```

Showing ISSU history:

```

switch# show issu history
Upgrade: 1
  From Version : FL.10.11.0001
  To Version   : FL.10.11.1000
  Start Time   : 2022-09-14 15:37:33
  End Time     : 2022-09-14 15:40:45
  Status       : Completed

```

## Command History

Release	Modification
10.11.1000	History parameter introduced.
10.11	Support for 6300 Switch Series added.
10.10	Command introduced

## Command Information

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A and S3L77A	Manager (#)	Administrators or local user group members with execution rights for this command.

## show vsf

```
show vsf
```

### Description

Displays the information about the configuration and status of a VSF stack and its members.

### Example

Showing the information about the configuration and status of a VSF stack and its members. Applies only to HPE Aruba Networking 4100i, 6100 Switch series.

```

switch# show vsf

Force Autojoin           : Enabled

```

```

Autojoin Eligibility Status: Eligible
MAC Address                : 90:20:c2:05:30:c0
Secondary                  : 2
Topology                   : Chain
Status                     : No Split
Split Detection Method     : vlan
Split Detection VLAN      : 10

```

Member ID	Mac Address	type	Status
1	90:20:c2:05:30:c0	JL675A	Conductor
2	90:20:c2:06:30:40	JL676A	Standby

Showing the information about the configuration and status of a VSF stack and its members :  
(Applies only to HPE Aruba Networking 6200, 6300 and 6300L Switch Series)

```

switch# show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 08:97:34:b0:0e:00
Egress Shape                : Enabled
Egress Shape Rate          : 10000000 kbps
Secondary                   : 2
Topology                    : Chain
Status                     : Active Fragment
Split Detection Method      : mgmt

```

Member ID	Mac Address	type	Status
1	08:97:34:b0:0e:00	JL666A	Conductor
2	08:97:34:b1:43:00	JL665A	In Other Fragment
3	08:97:34:b7:cc:00	S0E91A	Member
4		JL662A	Not Present

Showing the information about the configuration and status of a VSF stack and its stack members (with S0E91A or S0X44A SKU member):  
(Applies only to HPE Aruba Networking 6300 Switch Series)

```

switch# show vsf

Force Autojoin              : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address                 : 08:97:34:b0:0e:00
Egress Shape                : Enabled
Egress Shape Rate          : 10000000 kbps
Secondary                   : 2
Topology                    : Chain
Status                     : No Split
Split Detection Method      : None

```

Member ID	MAC Address	Type	Status
1	08:97:34:b0:0e:00	JL666A	Conductor
2	08:97:34:b1:43:00	JL665A	Standby

3	08:97:34:b7:cc:00	SOE91A	Member
4		JL662A	Not Present

Showing the information about the configuration and status of a VSF stack and its stack members with egress shape rate is populated with **Not Applied** indicating that port shaping failed to apply on one or more active VSF interfaces:

(Applies only to HPE Aruba Networking 6300 Switch Series)

```
switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 08:97:34:b0:0e:00
Egress Shape             : Enabled
Egress Shape Rate        : Not Applied
Secondary                : 2
Topology                 : Chain
Status                   : No Split
Split Detection Method    : None

Mbr MAC Address          Type          Status
ID
-----
1  08:97:34:b0:0e:00     JL666A      Conductor
2  08:97:34:b1:43:00     JL665A      Standby
3  08:97:34:b7:cc:00     SOE91A      Member
```



The **Egress Shape** displays the status of VSF port shaping(Enabled/Disabled). The **Egress Shape Rate** displays the operational speed of the stack when the VSF port shaping is applied. An error message is displayed if port shaping fails to apply to the interface. The purpose of the VSF port shaping feature is to ensure that all VSF interface operate at a lowest common port speed across the stack. This feature is supported only in the HPE Aruba Networking 6300 Switch series.

Showing VSF information when the port speed has been modified on a 4100i or 6100 Switch series using the **vsf port-speed** command:

```
6100# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 90:20:c2:05:30:c0
Operational Port Speed   : 1G
Secondary                : 2
Topology                 : Chain
Status                   : No Split
Split Detection Method    : vlan
Split Detection VLAN     : 100
Mbr Mac Address          type          Status
ID
-----
1  90:20:c2:05:30:c0     JL675A      Conductor
2  90:20:c2:06:30:40     JL676A

6100(config)# vsf port-speed 10g
This command will reboot the stack and reset the VSF port speed to configured
```

```

value
Do you want to continue (y/n)? y

6100# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 90:20:c2:05:30:c0
Operational Port Speed   : 10G
Secondary                : 2
Topology                 : Chain
Status                   : No Split
Split Detection Method    : vlan
Split Detection VLAN     : 100
Mbr Mac Address          type           Status
ID
-----
1   90:20:c2:05:30:c0   JL675A   Conductor
2   90:20:c2:06:30:40   JL676A   Standby

```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.15	Command updated to display <b>Egress Shape</b> . Applicable only for 6300 Switch Series.
10.13.1000	Command updated to display <b>Egress Shape Rate</b> . Applicable only for 6300 Switch Series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

## show vsf detail

```
show vsf detail
```

### Description

Displays detailed information related to the current state of the VSF stack and the stack members.

### Example

The output below shows example information for a 6300 Switch series. The information in the **type** and fields will vary, depending upon the individual switch model type.

```

switch# show vsf detail

VSF Stack
  MAC Address          : ec:eb:b8:d0:80:40
  Egress Shape        : Enabled
  Egress Shape Rate   : Not Applied
  Secondary           : 2
  Topology            : Chain
  Status              : No Split
  Uptime              : 0d 0h 23m
  Split Detection Method : None
  Software Version    : SL.10.02.0000-7755
  Force Autojoin      : Disabled
  Autojoin Eligibility Status : Not Eligible
  Autojoin Ineligibility Reason: Configuration changes detected
  Name                : HPE-ANW-VSF-<model>
  Contact             :
  Location            :

Member ID              : 1
  MAC Address          : ec:eb:b8:d0:80:40
  Type                 : <part number>
  Model                : <model>
  Status               : Conductor
  ROM Version         : SL.10.02.0000-7755
  Serial Number       : CN7ZK90012
  Uptime              : 0d 0h 23m
  CPU Utilization     : 0%
  Memory Utilization  : 20%
  VSF link 1         : Up, connected to peer member 2, link 1
  VSF link 2         : Down

Member ID              : 2
  MAC Address          : eb:ec:d8:e0:50:60
  Type                 : <part number>
  Model                : <model>
  Status               : Standby
  ROM Version         : SL.10.02.0000-7755
  Serial Number       : CN7ZK90012
  Uptime              : 0d 0h 23m
  CPU Utilization     : 0%
  Memory Utilization  : 15%
  VSF link 1         : Up, connected to peer member 1, link 1
  VSF link 2         : Down

Member ID              : 3
  MAC Address          :
  Type                 : <part number>
  Model                : <model>
  Status               : Not Present
  ROM Version         :
  Serial Number       :
  Uptime              :
  CPU Utilization     :
  Memory Utilization  :
  VSF link 1         :
  VSF link 2         :

```

The output below shows example information for a 4100, 6100 or 6200 Switch series. The information in the **type** and fields will vary, depending upon the individual switch model type.

```

switch(config)# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : bc:d7:a5:03:39:c0
Secondary                :
Topology                 : Standalone
Status                   : No Split
Split Detection Method   : None
Mbr Mac Address          type                Status
ID
-----
1   bc:c7:a5:23:39:c0   <part number>          Conductor

switch(config)# show vsf detail
VSF Stack
MAC Address              : bc:d7:a5:03:39:c0
Secondary                :
Topology                 : standalone
Status                   : No Split
Split Detection Method   : None
Software Version         : RL.10.16.1000G
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
Autojoin Ineligibility Reason: Configuration changes detected
Name                     : HPE-ANW-VSF-<model>
Contact                  :
Location

Member ID                : 1
MAC Address              : bc:c7:a5:23:39:c0
Type                     : <part number>
Model                    : <model>
Status                   : Conductor
ROM Version              : <XX>.01.17.0001
Serial Number            : TW42KYC01D
Uptime                   : 10 weeks, 1 day, 23 hours, 9 minutes
CPU Utilization          : 5%
Memory Utilization       : 11%
VSF Link 1               :
VSF Link 2               :

```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

# show vsf member

show vsf member <MEMBER-ID>

## Description

Displays information about the specified VSF member.

Parameter	Description
<MEMBER-ID>	VSF member ID. <ul style="list-style-type: none"><li>▪ Range for 4100i devices: 1-4</li><li>▪ Range for 6100 devices: 1-6</li><li>▪ Range for 6200F devices: 1-8.</li><li>▪ Range for 6300 devices: 1-10.</li></ul>

## Example

```
switch# show vsf member 1

Member ID           : 1
MAC Address         : ec:eb:b8:d0:80:40
Type                : <part number>
Model               : <model>
Status              : Conductor
ROM Version         : <xx>.10.02.0000-7755
Serial Number       : CN7ZK903012
Uptime              : 0d 0h 18m
CPU Utilization     : 0%
Memory Utilization  : 15%
VSF link 1          : Down
VSF link 2          : Down
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

# show vsf link

show vsf link

## Description

Displays the VSF link state for each member.

## Example

```
switch# show vsf link

VSF Member 1

  Link      Peer      Peer
Link State Member Link  Interfaces
-----
1    up      2         1     1/1/50
2    up      10        2     1/1/49

VSF Member 2

  Link      Peer      Peer
Link State Member Link  Interfaces
-----
1    up      1         1     2/1/49
2    up      3         1     2/1/50

VSF Member 3

  Link      Peer      Peer
Link State Member Link  Interfaces
-----
1    up      2         2     3/1/25
2    up      4         1     3/1/26

VSF Member 4

  Link      Peer      Peer
Link State Member Link  Interfaces
-----
1    up      3         2     4/1/25
2    up      5         1     4/1/26
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

## show vsf link detail

```
show vsf link detail
```

## Description

Shows detailed information of the interfaces configured on links of all stack members.

## Example

```
switch# show vsf link detail

VSF Member: 1 Link 1 Description: link 1 connected to member 2

Port      State      Status Code Peer Interface Peer System MAC Peer Product
Type
-----
---
1/1/27    up         S           2/1/27         38:21:f7:4c:a4:c0 <model>
1/1/28    error      M           1/1/27         38:21:d7:4e:d7:40 <model>

VSF Member: 2 Link 1 Description: link 1 connected to member 1

Port      State      Status Code Peer Interface Peer System MAC Peer Product
Type
-----
---
2/1/27    up         S           1/1/27         38:21:b7:4a:99:80 <model>
2/1/28    error      T

VSF Member: 2 Link 2 Description: link 2 connected to member 3

Port      State      Status Code Peer Interface Peer System MAC Peer Product
Type
-----
---
2/1/25    up         S           3/1/26         38:20:c7:4a:f0:00 <model>
2/1/26    down      D

VSF Member: 3 Link 1 Description: link 1 in loop

Port      State      Status Code Peer Interface Peer System MAC Peer Product
Type
-----
---
3/1/27    error      L           3/1/28         38:21:c7:4c:f0:00 <model>
3/1/28    error      L           3/1/27         38:21:c7:4c:f0:00 <model>

VSF Member: 3 Link 2

Port      State      Status Code Peer Interface Peer System MAC Peer Product
Type
-----
---
3/1/25    down      D
3/1/26    up         S           2/1/25         38:22:c7:4c:e4:c0 <model>

Flag abbreviation:
S - Success          D - Interface physically down      T - Peer timed out
L - Loop detected on the interface  AP - Peer autojoin in
progress
P - Peer with incompatible product type  ANE - Peer is not autojoin
eligible
```

```

SV - Peer with incompatible software version   AF - Peer autojoin validations
failed
M - Peer with inconsistent system MAC address
ILC - Peer with inconsistent VSF link configuration
AMS - Peer autojoin failed as it has MACsec configuration
AMI - Peer with multiple VSF interfaces attempting to autojoin
ACM - Peer attempting to autojoin on non-provisioned interface
AND - Peer with non-default VSF interface attempting to autojoin
AID - Peer autojoin failed as it is connected in incorrect direction
AFN - Peer autojoin failed as there is no free member number available

```

Showing output for a 6100 or 4100i switch series configured with an Port Speed Mismatch (PSM) status due to misconfigured port speed:

```

6100# show vsf link detail
VSF Member: 1 Link 1
Port      State   Status Code Peer Interface Peer System MAC Peer Product
Type
-----
1/1/28   up      S           2/1/25         f8:60:f0:c9:70:40 JL678A

VSF Member: 1 Link 2
Port      State   Status Code Peer Interface Peer System MAC Peer Product Type
-----
1/1/19   error  PSM       2/1/2          f8:60:f0:c9:70:40 JL678A

VSF Member: 2 Link 1
Port      State   Status Code Peer Interface Peer System MAC Peer Product Type
-----
2/1/25   up      S           1/1/28         ec:50:aa:2e:2f:80 JL678A

VSF Member: 2 Link 2
Port      State   Status Code Peer Interface Peer System MAC Peer Product Type
-----
2/1/2    error  PSM       1/1/19         ec:50:aa:2e:2f:80 JL678A

```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

## show vsf link error-detail

```
show vsf link error-detail
```

## Description

Shows detailed error information of the interfaces configured on links of all stack members. Also, the corrective action is also recommended to recover from the error.

## Example

Showing error information for the interfaces regarding loop detection:

```
switch(config)# show vsf link error-detail

VSF Member: 2 Link 1

Port                : 2/1/27
Status Code         : L - `Loop detected on the interface`

Error Description   : There is a loop detected between interfaces 2/1/27
and                                                         2/1/28 of member 2 indicating wrong cabling.

Suggested Corrective Action : VSF interfaces 2/1/27 and 2/1/28 are connected back
to                                                         back - please fix the cabling.

VSF Member: 2 Link 1

Port                : 2/1/28
Status code         : L - `Loop detected on the interface`

Error Description   : There is a loop detected between interfaces 2/1/28
and                                                         2/1/27 of member 2 indicating wrong cabling.

Suggested Corrective Action : VSF interfaces 2/1/28 and 2/1/27 are connected back
to                                                         back - please fix the cabling.

VSF Member: 10 Link 1

Port                : 10/1/26
Status Code         : AFN - `Peer autojoin failed as there is no free
member number available`

Error Description   : Maximum stack size has been reached or there are no
the                                                         free provisioned member entries available matching
peer switch with product type JL667A.

Suggested Corrective Action : Remove a member using "no vsf member x" CLI and then
physically disconnect and reconnect the new switch
with product type JL667A for adding it into the
stack.
```

Showing error information when peer member is connected to VSF link via its MACsec-configured interface for autojoin:

```
switch(config)# show vsf link error-detail

VSF Member: 2 Link 2
```

```

Port                : 2/1/26
Status Code         : AMS - `Peer autojoin failed as it has MACsec
                    configuration`

Error Description    : Autojoin failed as interface 2/1/26 is connected to
                    peer with MAC 38:21:c7:5c:d4:00 on interface 1/1/27
                    which has MACsec configuration.

Suggested Corrective Action : MACsec configuration should be removed from the
                    peer with MAC 38:21:c7:5c:d4:00 on interface 1/1/27.

```

Showing information about a 4100i or 6100 switch series configured with an incorrect port speed:

```

6100(config)# show vsf link error-detail
VSF Member: 1 Link 1
Port                : 1/1/49
Status Code         : IPS - `Peer switch with inconsistent VSF port speed`
Error Description    : Peer switch is operating at a different
                    vsf port speed of 1Gbps.
Suggested Corrective Action : Change vsf port speed to 1Gbps on this switch
                    or change vsf port speed to 10Gbps on the peer switch.

```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

## show vsf link error-detail member

```
show vsf link error-detail member <MEMBER-ID>
```

### Description

Shows error information and the suggestive action to resolve the error of the interfaces configured on links of a particular stack member.

Parameter	Description
<MEMBER-ID>	VSF member identifier. Required. <ul style="list-style-type: none"> <li>▪ Range for 4100i devices: 1-4.</li> <li>▪ Range for 6100 devices: 1-6.</li> <li>▪ Range for 6200F devices: 1-8.</li> <li>▪ Range for 6300 devices: 1-10.</li> </ul>

## Example

Showing error information and the suggestive action for member 1:

```
switch# show vsf link error-detail member 1

VSF Member: 1 Link 1

Port           : 1/1/52
Status Code    : M - `Peer with inconsistent system MAC address`

Error Description : All interfaces within a single VSF link must
terminate
                into the same peer switch. Interface 1/1/52 of
                member 1 link 1 is connected to a wrong peer with
                MAC 38:21:c7:5c:26:40.

Suggested Corrective Action : Multiple VSF neighbors detected on this VSF link 1.
Interface 1/1/52 is connected to device MAC
38:21:c7:5c:26:40. Please make sure the VSF
interfaces
                of link 1 terminate on the same peer device.
```

Showing error information and the suggestive action for member 4:

```
switch# show vsf link error-detail member 4

VSF Member: 4 Link 1

Port           : 4/1/27
Status Code    : AND - `Peer with non-default VSF interface
attempting
                to autojoin`

Error Description : Switch with MAC 38:21:c7:5c:a0:c0 is connected on
port
                1/1/27 which is a non default autojoin VSF
                interface.

Suggested Corrective Action : Auto-join failed on device with MAC
38:21:c7:5c:a0:c0.
                Please connect this device via interfaces 25 or 26 -
                those are the auto-join capable interfaces on this
                device.
```

Showing error information when the peer member is connected to member 2's VSF link via its MACsec-configured interface for autojoin:

```

switch(config)# show vsf link error-detail member 2

VSF Member: 2 Link 2

Port                : 2/1/26
Status Code         : AMS - `Peer autojoin failed as it has MACsec
                    : configuration`

Error Description    : Autojoin failed as interface 2/1/26 is connected to
                    : peer with MAC 38:21:c7:5c:d4:00 on interface 1/1/27
                    : which has MACsec configuration.

Suggested Corrective Action : MACsec configuration should be removed from the
                    : peer with MAC 38:21:c7:5c:d4:00 on interface 1/1/27

VSF Member: 2 Link 2

Port                : 2/1/26
Status Code         : AMS - `Peer autojoin failed as it has MACsec
                    : configuration`

Error Description    : Autojoin failed as interface 2/1/26 is connected to
                    : peer with MAC 38:21:c7:5c:d4:00 on interface 1/1/27
                    : which has MACsec configuration.

Suggested Corrective Action : MACsec configuration should be removed from the
                    : peer with MAC 38:21:c7:5c:d4:00 on interface 1/1/27

```

Showing output when there is no error-detail for a particular member:

```

switch# show vsf link error-detail member 2
No Error found in member 2

```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

## show vsf split-detection

```
show vsf split-detection [status|statistics]
```

### Description

Displays the VSF split detection status or statistics.

Parameter	Description
status	Displays whether the split detection operational status is up or down.
statistics	Displays information related to transmitted and received packets.

## Examples

```
switch# show vsf split-detection status
Split Detection Method      : mgmt
Split Detection Operational Status : Down (Mgmt interface admin state is down)

switch# show vsf split-detection status
Split Detection Method      : vlan
Split Detection VLAN       : 9
Split Detection VLAN Interfaces :
Split Detection Operational status : Down (VLAN 9 is not configured)

switch# show vsf split-detection statistics

VSF Split Detection Statistics:
=====
Total Packets Transmitted      = 10
Total Packets Received        = 20
Total Packets Received And Discarded = 20
```

## Command History

Release	Modification
10.16.1000	Command introduced.

## Command Information

Platforms	Command context	Authority
4100i 6100 6200 6300	Manager (#)	

## show vsf topology

```
show vsf topology
```

### Description

Displays information about VSF stack member connections.

### Example

```
switch# show vsf topology

      Stby      Conductor
+----+ +----+ +----+
| 3 |1==2| 2 |1==1| 1 |
+----+ +----+ +----+
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	

# shutdown

```
shutdown
no shutdown
```

## Description

Shuts down one or more VSF link interfaces.

The **no** form of this command turns on one or more VSF link interfaces.

## Examples

Shutting down a VSF link interface:

```
switch(config)# interface 1/1/1-1/1/2
switch(config-if-vsfc-1/1/1-1/1/2)# shutdown
```



Shutdown configuration for VSF interfaces is not persistent across reboots.

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config-if-vsfc	Administrators or local user group members with execution rights for this command.

## type

```
type <TYPE>
no type <TYPE>
```

### Description

Configures the part number of the VSF member being provisioned. After provisioning, the interfaces of the member are available for configuration.

When the member eventually joins the stack, it will boot up with the configuration made on the pre-provisioned interfaces.

To provision a member, the member number and the part number of the member must be specified.

The **no** form of this command removes the configuration for the part number of the VSF member provisioned.

Parameter	Description
<TYPE>	The part number of the member being provisioned. Required.

### Examples

Configuring the part number of a VSF member:

```
switch(vsf-member-2)#
  type The part number of the member being provisioned

switch(vsf-member-2)# type ?
j1658a 6300M 24SFP+ /4SFP56 Switch
j1659a 6300M 48SR PoE CLS 6 /4SFP56 Switch
j1660a 6300M 24SR PoE CLS 6 /4SFP56 Switch
j1661a 6300M 48G PoE CLS 4 /4SFP56 Switch
j1662a 6300M 24G PoE CLS 4 /4SFP56 Switch
j1663a 6300M 48G /4SFP56 Switch
j1664a 6300M 24G /4SFP56 Switch
j1665a 6300F 48G PoE CLS 4 /4SFP56 Switch
j1666a 6300F 24G PoE CLS 4 /4SFP56 Switch
j1667a 6300F 48G /4SFP56 Switch
j1668a 6300F 24G /4SFP56 Switch
j1762a 6300M 48G 4SFP56 Pwr2Prt Switch
s3175a HPE ANW 6300L 24SRX CL6 2L 2Y L2 Switch
s3176a HPE ANW 6300L 48SR5 CL8 2L 2Y L2 Switch
s3177a HPE ANW 6300L 48SR5 2L 2P LRM L2 Switch

switch(vsf-member-2)# type j1662a

switch(vsf-member-2)# show running-config
Current configuration:
!
!Version AOS-CX
```

```

!
!
!
!
ssh maximum-auth-attempts 6
!
!
!
!
!
vlan 1
vsf member 1
    type j1661a
exit
vsf member 2
    type j1662a
exit

```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	<code>vsf-member-&lt;ID&gt;</code>	Administrators or local user group members with execution rights for this command.

## vsf egress-shape

```

vsf egress-shape
no egress-shape

```

### Description

When the VSF links are of different interface speeds, the command ensures that all VSF interfaces in a stack operate at a chosen common denominator interface speed when there is at least one active stack member of type S0E91A or S0X44A in the stack.

The **no** form of this command removes the common denominator speed from all VSF links and ensure that all the VSF links operate at their original speed.



VSF egress-shape is enabled by default on the switch.

This feature is applicable only for 6300 switch series (except for S3L75A, S3L76A and S3L77A).

Parameter	Description
<code>egress-shape</code>	Enables egress shaping on all VSF interfaces.

## Examples

### Example 1:

```
switch(config)#vsf egress-shape
switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 38:21:c7:5c:62:40
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                 : 2
Topology                  : Ring
Status                    : No Split
Split Detection Method    : None
```

Mbr ID	Mac Address	type	Status
1	38:21:c7:5c:62:40	JL668A	Conductor
2	18:7a:3b:1b:68:c0	R8S90A	Standby
3	38:21:c7:5c:57:c0	JL668A	Member
4	18:7a:3b:1b:66:40	R8S89A	Member

### Example 2:

```
switch(config)#vsf egress-shape
switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 38:21:c7:5c:62:40
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                 : 2
Topology                  : Ring
Status                    : No Split
Split Detection Method    : None
```

Mbr ID	Mac Address	type	Status
1	38:21:c7:5c:62:40	JL668A	Conductor
2	18:7a:3b:1b:68:c0	R8S90A	Standby
3	38:21:c7:5c:57:c0	JL668A	Member
4	18:7a:3b:1b:66:40	S0E91A	Member

### Example 3:

```
switch(config)#vsf egress-shape
switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 38:21:c7:5c:62:40
Egress Shape             : Enabled
```

```
Egress Shape Rate      : 10000000 kbps
Secondary              : 2
Topology              : Ring
Status                : No Split
Split Detection Method : None
```

Id	Mac Address	Type	Status
1	38:21:c7:5c:62:40	JL668A	Conductor
2	18:7a:3b:1b:68:c0	R8S90A	Standby
3	38:21:c7:5c:57:c0	JL668A	Member
4	18:7a:3b:1b:66:40	S0E91A	Member

**Example 4:**

```
switch(config)#no vsf egress-shape
switch# show vsf

Force Autojoin          : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address             : 38:21:c7:5c:62:40
Egress Shape           : Disabled
Egress Shape Rate      : None
Secondary              : 2
Topology              : Ring
Status                : No Split
Split Detection Method : None
```

Id	Mac Address	Type	Status
1	38:21:c7:5c:62:40	JL668A	Conductor
2	18:7a:3b:1b:68:c0	R8S90A	Standby
3	38:21:c7:5c:57:c0	JL668A	Member
4	18:7a:3b:1b:66:40	S0E91A	Member

**Command History**

Release	Modification
10.15	Command introduced.

**Command Information**

Platforms	Command context	Authority
6300, except for S3L75A, S3L76A, S3L77A.	config	Administrators or local user group members with execution rights for this command.

**vsf force-auto-join**

vsf force-auto-join

## Description

Forces the switch with non-factory default configuration to join a stack. The switch should not have any existing VSF configurations for force auto-join to work. If VSF configurations are made after force auto-join is enabled, the switch will no longer be eligible for auto-join.

## Examples

Forcing a switch with non-factory default configuration to join a stack:

```
switch(config)# vsf force-auto-join
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config	Administrators or local user group members with execution rights for this command.

## vsf member

```
vsf member <MEMBER-ID>  
no vsf member <MEMBER-ID>
```

## Description

Creates VSF member context in the switch for the specified member.

The **no** form of this command removes the specified member from the stack. All configuration associated with the member, as well as the subsystems and interfaces of the member will also be removed.

If the member is physically present in the stack at the time it is removed, it will reboot with the default configuration and lose its identity as a member of the stack from which it was removed.



When a physically present member is removed, it may cause the stack to split.

Parameter	Description
<MEMBER-ID>	VSF member identifier. <ul style="list-style-type: none"><li>Range for 4100i devices: 1 to 4.</li><li>Range for 6100 devices: 1 to 6.</li><li>Range for 6200F devices: 1 to 8.</li><li>Range for 6300 devices: 1 to 10.</li></ul>

## Examples

Configuring a VSF member:

```
switch(config)# vsf member 2
switch(vsf-member-2)#
```

Removing a non-conductor member from the stack:

```
switch(config)# no vsf member 2
The specified switch will be unconfigured and rebooted
Do you want to continue (y/n)? y
```



Removing the running conductor should be done with caution as it can make the stack unusable if there is no standby.

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config	Administrators or local user group members with execution rights for this command.

## vsf member reboot

```
vsf member <MEMBER-ID> reboot
```

### Description

Reboots the specified VSF member. Upon reboot, if the conductor is reachable, the member will rejoin the stack.

Parameter	Description
<MEMBER-ID>	Member number to be rebooted. Required. <ul style="list-style-type: none"><li>Range for 4100i devices: 1-4.</li><li>Range for 6100 devices: 1-6.</li><li>Range for 6200F devices: 1-8.</li><li>Range for 6300 devices: 1-10.</li></ul>

## Examples

Rebooting the primary switch of the stack:

```
switch# vsf member 1 reboot
Rebooting the conductor switch of the stack without a standby
will make the stack unusable.
Do you want to continue (y/n)? y

switch# vsf member 1 reboot
The conductor switch will reboot and the standby will become the conductor.
Do you want to continue (y/n)? y

switch# vsf member 2 reboot
This will reboot the specified switch.
Do you want to continue (y/n)? y
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	Manager (#)	Administrators or local user group members with execution rights for this command.

## vsf renumber-to

vsf renumber-to <MEMBER-ID>

### Description

Renumbers VSF member 1 to a value from 2 through 10 (for 6300 devices) and 2 through 8 (for the 6200F device). Changing the member number causes the switch to reboot with the new member number. Only member 1 can be renumbered.



VSF links must be configured before renumbering a switch. Renumbering will be disallowed if no links are configured or there are provisioned/physically present members.

Parameter	Description
<MEMBER-ID>	Member number to which the member will be renumbered. Required. <ul style="list-style-type: none"><li>▪ Range for 4100i devices: 2-4.</li><li>▪ Range for 6100 devices: 2-6.</li><li>▪ Range for 6200F devices: 2-8.</li><li>▪ Range for 6300 devices: 2-10.</li></ul>

## Examples

Renumbering primary VSF member from 1 to 2:

```
switch(config)# vsf renumber-to 2
Member 1 cannot be renumbered until all other members are removed.

switch(config)# vsf renumber-to 2
Member 1 cannot be renumbered until a VSF link is configured.

switch(config)# vsf renumber-to 2
This will save the VSF configuration and reboot the switch.
Do you want to continue (y/n)? y
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config	Administrators or local user group members with execution rights for this command.

## vsf secondary-member

```
vsf secondary-member <MEMBER-ID>
no vsf secondary-member <MEMBER-ID>
```

### Description

Configures a secondary member from the available members. The secondary member will normally operate as the Standby member of the stack.

The **no** form of this command removes the configuration of the secondary member.



---

Member 1 cannot be configured as the secondary member.

---

Parameter	Description
<MEMBER-ID>	Secondary member number. Required. <ul style="list-style-type: none"><li>▪ Range for 4100i devices: 2-4.</li><li>▪ Range for 6100 devices: 2-6.</li><li>▪ Range for 6200F devices: 2-8.</li><li>▪ Range for 6300 devices: 2-10.</li></ul>

## Examples

## Configuring and un-configuring a secondary member:

```
switch(config)# vsf secondary-member 3
This will save the configuration and reboot the specified switch.
Do you want to continue (y/n)? y

switch(config)# no vsf secondary-member
The secondary member will go for a reboot.
Do you want to continue (y/n)? y
```

## Configuring a secondary member when secondary member is already configured:

```
switch(config)# vsf secondary-member 3
This will save the configuration and reboot the specified switch.
Do you want to continue (y/n)? y

switch (config)# vsf secondary-member 4
A secondary member is already configured. Existing secondary member
will be unconfigured and rebooted to join the stack as a member. The
specified switch is then rebooted and will join the stack as the new
standby.
Do you want to continue (y/n)? y
```

## Configuring a secondary member when one or more members are booting:

```
switch(config)# vsf secondary-member 3
One or more members are currently booting. Allowing this configuration
may cause stack to split leading to traffic disruption.
Do you want to continue (y/n)? y
This will save the configuration and reboot the specified switch.
Do you want to continue (y/n)? y

switch(config)#no vsf secondary-member
One or more members are currently booting. Allowing this configuration
may cause stack to split leading to traffic disruption.
Do you want to continue (y/n)? y
The secondary member will go for a reboot.
Do you want to continue (y/n)? y
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config	Administrators or local user group members with execution rights for this command.

# vsf split-detect mgmt

```
vsf split-detect
  mgmt
  no ...
```

## Description

Configures the VSF split detection method that specifies the mechanism used for stack fragment discovery when there is a VSF stack split.

Once the stack fragments are discovered, the fragment having the primary member always takes precedence. All non-VSF interfaces on the secondary stack fragment will be brought down to minimize network disruption due to duplicate MAC/IP addresses.

Parameter	Description
mgmt	Configures a management interface as the split detection method on 6200 and 6300 Switch series. The management interfaces of the primary and secondary members should be either connected to the same L2 network or directly connected to each other.
no ...	Negates any configured parameter and removes the VSF split detection configuration.

## Examples

Configuring the **mgmt** interface as the split detection method:

```
switch(config)# vsf split-detect mgmt
```

```
switch# sh vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 4c:d5:87:30:26:80
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                 : Ring
Status                   : No Split
Split Detection Method    : mgmt
```

```
 Mbr Mac Address          type          Status
ID -----
1  4c:d5:87:30:26:80      JL660A      Conductor
2  e4:de:40:9b:10:00      JL659A      Standby
3  14:ab:ec:eb:81:80      R8S91A      Member
4  14:ab:ec:eb:53:80      R8S91A      Member
```

Removing split detection from the stack:

```
switch(config)# no vsf split-detect
```

```

switch# sh vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 4c:d5:87:30:26:80
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                 : Ring
Status                   : No Split
Split Detection Method    : None

```

```

Mbr Mac Address      type      Status
ID
-----
1 4c:d5:87:30:26:80  JL660A  Conductor
2 e4:de:40:9b:10:00  JL659A  Standby
3 14:ab:ec:eb:81:80  R8S91A  Member
4 14:ab:ec:eb:53:80  R8S91A  Member

```

## Command History

Release	Modification
10.16.1000	Command introduced.

## Command Information

Platforms	Command context	Authority
6200 6300	config	Administrators or local user group members with execution rights for this command.

## vsf split-detect vlan

```

vsf split-detect
vlan <VLAN_ID>
no ...

```

### Description

Configures the VSF split detection method that specifies the mechanism used for stack fragment discovery when there is a VSF stack split.

Once the stack fragments are discovered, the fragment having the primary member always takes precedence. All non-VSF interfaces on the secondary stack fragment will be brought down to minimize network disruption due to duplicate MAC/IP addresses.

Parameter	Description
vlan <VLAN_ID>	Configures a VLAN as the split detection method on supported switches. This method uses two access ports, one from the primary member and one from the secondary member. These ports should be connected on the same L2 network, therefore, should not be part of another VLAN. The VLAN identifier should

Parameter	Description
	be reserved and not be used to carry any data traffic. Range: 2-4094.
no ...	Negates any configured parameter and removes the VSF split detection configuration.

## Example

```

6100(config)# vlan 100
6100(config-vlan-100)# exit
6300(config)# interface 1/1/1,2/1/1
6300(config-if-<1/1/1,2/1/1>)# vlan access 100
6100(config)# vsf split-detect vlan 100
6100# show vsf
Force Autojoin                : Disabled
Autojoin Eligibility Status   : Not Eligible
MAC Address                    : 90:20:c2:05:30:c0
Secondary                      : 2
Topology                      : Chain
Status                        : No Split
Split Detection Method        : vlan
Split Detection VLAN          : 100

Mbr Mac Address               type                Status
ID                            -----
1   90:20:c2:05:30:c0       JL675A           Conductor
2   90:20:c2:06:30:40       JL676A           Standby

switch# show vsf split-detection status
Split Detection Method        : vlan
Split Detection VLAN         : 100
Split Detection VLAN interfaces : 1/1/1,2/1/1
Split Detection Operational status : Up

```

The user is advised to connect only one front-plane port between the Conductor and Standby. Only these ports should be part of a dedicated split detection VLAN(MAD VLAN) as Access ports. However, if more than one port on the Conductor or Standby is configured with MAD VLAN, the port selection process will depend on the following conditions:

- If the command `vsf split detect vlan <VLAN-ID>` is issued after multiple ports have been configured with MAD VLAN, the port with the lowest number will be selected on both the Conductor and Standby.
- If the command `vsf split detect vlan <VLAN-ID>` is issued before configuring multiple ports with MAD VLAN, the ports are selected based on the order in which they were added to the split detection VLAN.

---

Split detection VLAN configuration is applicable only on access ports and is not supported on trunk or aggregation interfaces.

---

Split detection ports are automatically enabled for `spanning-tree bpdu-filter` whenever a port is configured as part of MAD VLAN. This configuration is displayed in the `show running-config` CLI output, even if the end user has not explicitly configured it. When the port is removed from MAD VLAN, the `spanning-tree bpdu-filter` configuration is also removed from that port.

---

Once the MAD VLAN is enabled on a port, it is recommended not to configure any other features on the same port. Similarly, `spanning-tree bpdu-filter` should not be removed from the port which is part of MAD VLAN.

---

The split detection port will be in a blocked state whenever a port is configured as part of the MAD VLAN, and no dataplane traffic will be allowed. The port will be unblocked when it is removed from the MAD VLAN configuration.

---

## Command Information

Platforms	Command context	Authority
4100i 6100	<code>config</code>	Administrators or local user group members with execution rights for this command.

---

## vsf start-auto-stacking

vsf start-auto-stacking

### Description

Configures the secondary member and VSF links automatically. To use this command, the switch must be in the factory default configuration.

This command is applicable only on the primary switch. The primary switch must be in factory default condition and must not have any VSF configuration.

---

### Examples

Configuring a VSF secondary member and VSF link on conductor:

```
switch(config)# vsf start-auto-stacking
This will configure links and secondary on conductor

Do you want to continue (y/n)? y
```

Running the configuration on non-factory default switch:

```
switch(config)# vsf start-auto-stacking
The switch is having non-factory default running configuration.
Command is not applicable
```

Running the configuration on non-primary switch:

```
switch(config)# vsf start-auto-stacking
The command is applicable only on Primary switch
```

## Command History

Release	Modification
10.16.1000	Command introduced on 4100 and 6100 Switch series.
10.07 or earlier	--

## Command Information

Platforms	Command context	Authority
4100 6100 6200 6300	config	Administrators or local user group members with execution rights for this command.




---

Configuration Conflict Finder (CCF) is supported only on 6200 and 6300 platforms.

---

Configuration Conflict Finder (CCF) is a configuration troubleshooting solution that allows admins and support to automatically determine switch configuration anomalies using a set of feature configuration templates. This ensures specific features are error-free and validates the switch's configuration. CCF is an especially useful tool for CX users, support personnel, and ERT.

CCF detects misconfigurations including:

- Incomplete or inter-dependent configurations
- Mutually exclusive configurations

CCF provides the following recommendations for VSF environments:

### Secondary is not configured in a stack of size larger than one

If the VSF switch is not standalone, then the recommendation is to configure the secondary for redundancy. If the secondary is not configured then a `Secondary member is recommended for vsf stack (vsf secondary-member <mem-id>)` recommendation is displayed.

```
switch# sh running-config | line
 1 Current configuration:
 2 !
 3 !Version ArubaOS-CX FL.10.10.0000T-62-gb03e61e2b788a
 4 !export-password: default
 5 cli-session
 6     timeout 0
 7 !
 8 !
 9 !
10 !
11 !
12 !
13 ssh server vrf default
14 ssh server vrf mgmt
15 vsf member 1
16     type r8s92a
17 vsf member 2
18     type jl668a
19 vlan 1
20 spanning-tree
21 interface mgmt
22     no shutdown
23     ip dhcp
switch# switch config-validator feature vsf
Line number 15: Secondary member is recommended for vsf stack (vsf secondary-
member <mem-id>)
```

### Secondary is configured and stack split mode is not configured

Whenever the secondary is configured, it is recommended to configure split detection to avoid multiple stack fragments being active at the same time. The `Split detect (MAD)` is recommended for vsf stack (`vsf split-detect mgmt`) recommendation is displayed if split detection is not configured.

```
switch# sh running-config | line
 1 Current configuration:
 2 !
 3 !Version ArubaOS-CX FL.10.10.0000T-62-gb03e61e2b788a
 4 !export-password: default
 5 cli-session
 6     timeout 0
 7 !
 8 !
 9 !
10 !
11 !
12 !
13 ssh server vrf default
14 ssh server vrf mgmt
15 vsf secondary-member 2
16 vsf member 1
17     type r8s92a
18 vsf member 2
19     type j1668a
20 vlan 1
21 spanning-tree
22 interface mgmt
23     no shutdown
24     ip dhcp
switch# switch config-validator feature vsf
Line number 15: Split detect (MAD) is recommended for vsf stack (vsf split-detect mgmt)
```

## VSF links are administratively shutdown

Administratively shutting down the VSF links will affect the connectivity to other members of the stack and may result in splitting of the stack into multiple fragments. A `VSF interface should be administratively up` recommendation is displayed when any VSF links are administratively shutdown.

```
switch# sh running-config | line
 1 Current configuration:
 2 !
 3 !Version ArubaOS-CX FL.10.10.0000T-62-gb03e61e2b788a
 4 !export-password: default
 5 cli-session
 6     timeout 0
 7 !
 8 !
 9 !
10 !
11 !
12 !
13 ssh server vrf default
14 ssh server vrf mgmt
15 vsf split-detect mgmt
16 vsf secondary-member 2
17 vsf member 1
18     type r8s92a
19     link 1 1/1/25
20 vsf member 2
```

```
21     type j1668a
22     vlan 1
23     spanning-tree
24     interface mgmt.

.
.
.
.
.

123    interface 1/1/25
124        shutdown
125    interface 1/1/26
126        no shutdown
127        no routing

switch# switch config-validator feature vsf
Line number 19: VSF interface should be administratively up
```

## Considerations and best practices

---

The following recommendations and restrictions apply to VSF.

- Before applying a configuration on a stack through checkpoint restore or TFTP/SFTP/USB download, make sure that current VSF-specific configurations and the intended configurations match exactly. In other words, the VSF stack and the intended configuration must have the same:
  - Total number of members
  - Member types
  - Member number/ID
  - VSF link configurations
- A functional stack must be configured with a standby for redundancy purposes. If the conductor fails and there is no standby, the stack will fail.
- If the conductor fails and there is a standby device, the standby becomes the new conductor and will take over stack management. When the old conductor device is replaced, it seamlessly becomes the standby device for the stack and there no disruption.

The MAC address of the stack will remain the same until the entire stack is rebooted, after which the stack MAC address will be the MAC address of the new conductor. However, once recovered, it is not advisable to use the removed conductor elsewhere in the same network until the stack is rebooted to avoid MAC address conflicts.

- After downloading firmware to a stack, the stack must be rebooted to complete the upgrade process. Adding or rebooting individual members before the upgrade process is completed can cause the individual member to fail while joining the stack. A member with 10.07 software version cannot join a stack running on earlier versions.
- If there is a discrepancy between a VSF member link configuration on the conductor and the VSF member link configuration on the member, the link configuration on the member is used.
- If there is a split, failure in the connectivity between management interfaces of the conductor and standby might result in two active fragments. This issue can occur even if management split-detect is enabled.
- Replacing member 1 in a stack without a standby with a new switch booted as member 1 will reset all configurations on the stack.
- Do not connect a renumbered member to multiple primary devices through VSF links.
- Before removing an individual interface from VSF link using the command `no link <x> <interface>`, ensure that the interface is admin shutdown at both local and peer ends. For example: Interface 1/1/25 on member 1 link 1 is connected to 2/1/25 on member 2 link 2. The user intends to remove 1/1/25 from link 1 of member 1. Both the interfaces 1/1/25 and 2/1/25 have to be admin shutdown before actually removing them from the link configuration. To delete the link completely using the `no link <x>` command, all individual interfaces in the VSF link have to be admin shutdown both at local and peer ends.
- There may be instances in which a conductor switch with `vsf secondary <id>` configuration is unable to discover the standby switch. In such cases, the conductor switch will wait for up to 6 minutes to detect the standby switch.

- When applying a configuration on a stack from Central/NetEdit/ZTP/TFTP or through a checkpoint restore to remove members from the stack, consider the following recommendations:
  - a. If you are removing members that are physically present, it is recommended to remove one member at a time. In the case of ring topology, once a single member is removed, the topology will get transitioned to a chain topology. After that, members must be removed starting from the farthest end.
  - b. If you are removing provisioned members, then you can remove multiple members at the same time.



---

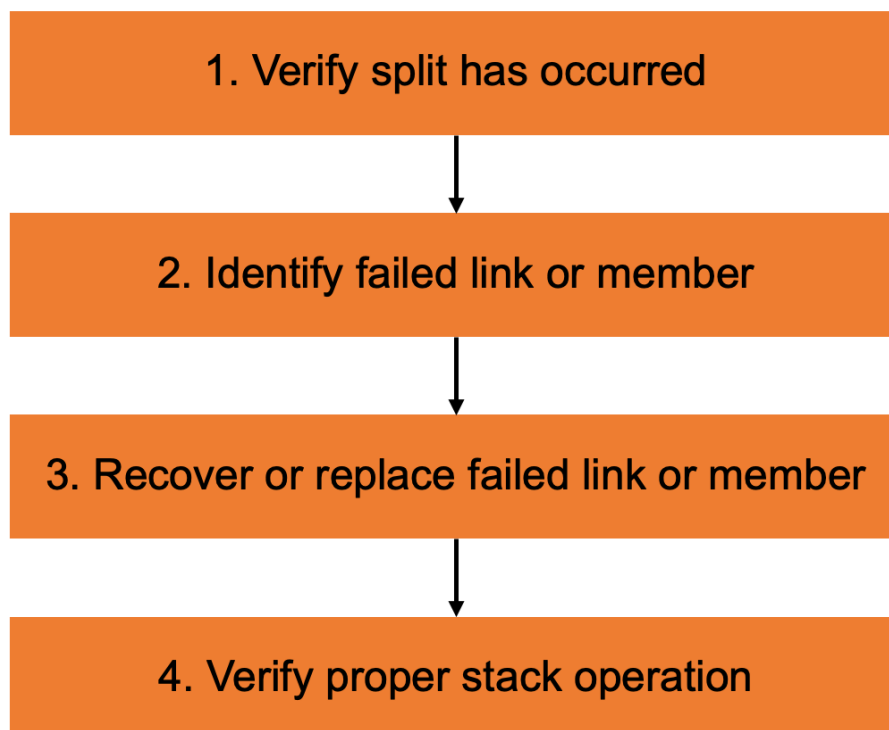
Removing more than one member at a time through configuration restoration (ZTP/Central/Checkpoint) might result in non-deterministic behavior. This might cause the members to reboot and drop to the console.

---

- If the entire stack configuration needs to be provisioned manually using CLIs, ensure that the conductor's VSF link configuration is done at the last.
- For TFTP download, the recommended work-flow is to copy the configuration to startup first, and then copy to running-configuration. The direct download of TFTP to running-configuration is not recommended.
- It is not recommended to change the VSF configurations (links & secondary) on the conductor when one or more members of the stack are booting.
- On 4100i and 6100 platforms, it is recommended to use VSF ports with the same speed across the stack.
- For SKUs S4P41A, S4P42A, S4P45A, and S4P46A, which support dual-mode SmartRate VSF ports, it is recommended not to change the port speed while the stack is up.
- With split detection MAD VLAN, it is recommended to use only one VLAN ID and one front-plane port from both the Primary and Secondary switches. Additionally, it is recommended not to enable any other features on this VLAN ID and front-plane ports used with split-detection.

The following section describes failure and recovery scenarios for VSF stacks.

## Stack split



### Step 1: Verify split has occurred

Use the `show vsf` command from the primary and secondary members to determine whether or not a split has occurred.

Output from the primary member will display a stack status of Active Fragment and member status of any members on the other side of the split as In Other Fragment:

```
switch# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 08:97:34:b0:0e:00
Egress Shape            : Enabled
Egress Shape Rate       : None
Secondary               : 2
Topology                : Chain
Status                  : Active Fragment
Split Detection Method   : mgmt
```

Mbr ID	Mac Address	type	Status
1	38:21:c7:5c:f4:c0	JL668A	Conductor
2		JL668A	In Other Fragment
3		JL668A	In Other Fragment
4		JL668A	In Other Fragment

Output from the secondary member will display a stack status of Inactive Fragment, with members on the other side of the split listed as In Other Fragment:

```
switch# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 08:97:34:b0:0e:00
Egress Shape            : Enabled
Egress Shape Rate       : None
Secondary                : 2
Topology                 : Chain
Status                   : Inactive Fragment
Split Detection Method   : mgmt
```

Mbr ID	Mac Address	type	Status
1		JL668A	In Other Fragment
2	38:21:c7:5c:77:40	JL668A	Conductor
3	38:21:c7:5a:a5:80	JL668A	Member
4	38:21:c7:5c:b3:00	JL668A	Member

## Step 2: Identify failed link or member

Symptoms may include:

- No link lights for VSF link ports between stack members
- No power to one or more stack members
- Event logs indicating loss of connectivity to one or more members and/or VSF links going down

Utilize show commands, event logs, and physical inspection of stack members and associated cabling to determine which link(s) or member(s) have failed to cause the split.

```
6300# show events -r -d vsfd
-----
Event logs from current boot
-----
2021-11-23T20:08:01.173123+00:00 6300 vsfd[732]: Event|9927|LOG_
INFO|CDTR|1|Fragment with conductor 1 is Active
2021-11-23T20:07:59.400936+00:00 6300 vsfd[732]: Event|9924|LOG_INFO|CDTR|1|VSF
link 1 is down
2021-11-23T20:07:59.400841+00:00 6300 vsfd[732]: Event|9913|LOG_WARN|CDTR|1|Lost
member 2 with Loss of communication
2021-11-23T20:07:59.400733+00:00 6300 vsfd[732]: Event|9908|LOG_
INFO|CDTR|1|Topology is Standalone
2021-11-23T20:07:58.534186+00:00 6300 vsfd[732]: Event|9924|LOG_INFO|CDTR|1|VSF
```

```
link 2 is down
```

### Step 3: Recover or replace failed link or member

If the split was caused by the failure of a stack member or VSF link cable, replace the affected hardware.

If the split was caused by a misconfiguration, such as inadvertently disabling one or more VSF links or otherwise modifying the stack configuration, revert the applicable configuration changes. If the configuration change resulted in removing members from the stack, re-add those members as appropriate.

### Step 4: Verify proper stack operation

Once the cause of the split has been identified and corrected, verify that the stack is operating normally.

```
switch(config)# show vsf
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 38:21:c7:5c:f4:c0
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                 : Ring
Status                   : No Split
Split Detection Method    : mgmt
```

Id	Mac Address	Type	Status
1	38:21:c7:5c:f4:c0	JL668A	Conductor
2	38:21:c7:5c:77:40	JL668A	Standby
3	38:21:c7:5a:a5:80	JL668A	Member
4	38:21:c7:5c:b3:00	JL668A	Member

## Misconfiguration recovery

If a switch fails to join the stack, or fails to rejoin after a reboot, due to misconfiguration, use the following procedure to restore the switch back to a factory default configuration.



---

The user must have management connectivity to the failed member for support files from the member in recovery mode.

---

1. Press **Ctrl+C** on the switch console.



---

If the member is not able to reach conductor, it will go to recovery console after 10 minutes. You can press **Ctrl+C** to redirect the switch to the recovery console immediately.

---

2. Log in using administrator credentials.

3. At the prompt, issue the `erase all zeroize` command.

```
^C
Login: admin
Password:

recovery# erase all zeroize
```

4. This resets the member to factory-default settings and the switch will come up with a default member ID of 1.
5. Now the user can reconfigure the VSF link and renumber it to the preferred member ID.



---

The `no vsf member <id>` command can also be used for switch recovery, if support files from lost members must be preserved in a local copy.

---

## VSF switchover behaviors

The following behaviors are expected during a VSF switchover event:

- The count of console login attempts is cleared (reset to 0).
- The count of login attempts for the **aaa authentication limit-login-attempts** feature is cleared (reset to 0).
- The output of the command **show authentication locked-out-users list** is cleared of users locked out via the console.
- The output of the command **show authentication locked-out-users list** is cleared of users locked out via SSH, TELNET, or REST (verified on an SSH channel.)

## General

### What is VSF?

Virtual Switching Framework, or VSF, defines a single virtual switch comprised of multiple individual physical switches that are interconnected through standard Ethernet links. These links are referred to as VSF links.

These physical switches will function as one device with a unified control and management plane. Multiport VSF links are supported.

### What are the supported platforms for VSF?

The HPE Aruba Networking 4100, 6100, 6200F and 6300F/M/L Switch Series support VSF.

VSF can be formed with the following switch models:

- HPE Aruba Networking 6300F/M/L Switch Series (JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL762A, R8S89A, R8S90A, R8S91A, R8S92A, S0E91A, S0X44A, S0G03A, S0G05A, S0F99A, S0G01A, S0G96A, S0G98A, S4P48A, S0G04A, S0G06A, S0G00A, S0G95A, S0G97A, S0G02A, S3L75A, S3L76A, S3L77A, S4P41A, S4P42A, S4P43A, S4P44A, S4P45A, S4P46A, S4P47A, S4P49A)
- HPE Aruba Networking 6200F Switch Series (JL724A, JL725A, JL726A, JL727A, JL728A, R8Q67A, R8Q68A, R8Q69A, R8Q70A, R8Q71A, R8V08A, R8V09A, R8V10A, R8V11A, R8V12A, R8V13A, R8Q72A, JL724B, JL725B, JL726B, JL727B, JL728B, S0M81A, S0M82A, S0M83A, S0M84A, S0M85A, S0M86A, S0M87A, S0M88A, S0M89A, S0M90A, S0G13A, S0G14A, S0G15A, S0G16A, S0G17A)
- HPE Aruba Networking 4100i Switch series (JL817A, JL818A)
- HPE Aruba Networking 6100 Switch series (JL675A, JL676A, JL677A, JL678A, JL679A, R9Y04A)



---

HPE Aruba Networking 6200F Switch Series only supports fixed SKUs.

---

### Can I create a stack with a mix of different switch models?

No, a VSF stack cannot be created with a mix of different switch families. The stack must be made up of only a single switch model type. For example, a VSF stack can include either 4100i, 6100, 6200 or 6300 switches, but not a mix. Additionally:

- 6200 switch series models with SKUs R8Q72A and R8V13A can only stack with other R8Q72A and R8V13A SKUs
- 6100 Switch series models with SKU JL679A can stack on with other JL679A SKUs
- 6300L switch series models with SKUs S3L75A, S3L76A, and S3L77A can only stack with other 6300L S3L75A, S3L76A and S3L77A SKUs.

### **Can I form a stack with a mix of 6200 switch series such as R8Q72A and any other SKU other than R8Q72A or R8V13A?**

No, an R8Q72A SKU switch can stack only with switches of type RQ72A or R8V13A SKU.

### **Can I form a stack with a mix of 6100 switch series JL679A and any other 6100 SKU, other than JL679A?**

No, a 6100 JL679A SKU switch can stack only with switches of type JL679A .

### **Can I form a stack with S3L75A, S3L76A, S3L77A and any other 6300 switches?**

No, S3L75A, S3L76A, and S3L77A SKUs can only stack with other S3L75A, S3L76A, and S3L77A SKUs and cannot stack with other 6300 switches.

### **What port speeds do VSF links support?**

- **4100:** Connections between the switches must use 1G or 10G links.
- **6100:** Connections between the switches must use 1G or 10G links.
- **6200F:** Connections between the switches must use 1G SFP/SFP+ either with 1G copper downlink ports, 10G links, or SmartRate ports.
- **6300:** Connections between the switches must use 10G, 25G, or 50G links.  
The VSF stack, containing only the S0E91A and S0X44A SKUs, supports 10G, 25G, and 100G speed links. If a VSF stack includes a mix of S0E91A/S0X44A SKU and other SKUs, then it supports 10G, 25G VSF links, SmartRate ports and only.

HPE Aruba Networking recommends that all VSF links be configured to run at the same speed.

### **On S4P41A, S4P42A, S4P45A, and S4P46A SKUs, can I change the speed of dual mode smartRate ports when the stack is up with smartRate ports?**

No, changing the speed of dual mode smartRate ports when the stack is up can cause the stack to split and lead to network instability.

### **On S4P41A, S4P42A, S4P45A, and S4P46A SKUs, how do I change the speed of dual mode smartRate VSF ports when the stack is up?**

To change the speed of dual mode smartRate VSF ports on S4P41A, S4P42A, S4P45A, and S4P46A SKUs when the stack is up, follow these steps:

1. The stack must first be dismantled using `erase all zeroize` command. This will reset all stack members and bring them up as standalone units.
2. After dismantling, reconfigure the port group corresponding to the VSF links with desired port speed using `system interface-group <GROUP> member <MEMBER-ID> speed <SPEED>` command.
3. Once reconfiguration is complete, bring the stack up again with the new port speed settings.

### **On S4P41A, S4P42A, S4P45A, and S4P46A SKUs, how do I change the speed of dual mode smartRate VSF ports interconnected between 2 members when the stack is up?**

To change the speed of dual mode smartRate VSF ports interconnected between 2 members when the stack is up, follow these steps:

1. Unconfigure the desired members using `no vsf member <member_id>` command. This will reset the desired members and bring them up as a standalone units.

2. >Once the units are up as a standalone, reconfigure the port group corresponding to the VSF links with desired port speed using `system interface-group <GROUP> member <MEMBER-ID> speed <SPEED>` command.
3. Once reconfiguration is complete, configure the members to join the stack back.

### **Can VSF be disabled?**

Users cannot disable VSF. A factory default switch boots up as a VSF-enabled device with its Member ID set to 1.

### **What is a primary switch in VSF stack? Is it configurable?**

Only the switch with a Member ID of 1 will be the primary switch in a VSF stack. This switch will function as the stack conductor and will drive the control and management plane for the stack.

### **What is a secondary switch in a VSF stack? Is it configurable?**

The secondary switch will function as the standby in a stack. In the case of auto-stacking, secondary member is automatically configured through button press or `vsf start-auto-stacking` command. In addition, any member other than Member 1 can be configured manually as the secondary switch using the `vsf secondary-member <MEMBER-ID>` command.

HPE Aruba Networking strongly recommends that you configure a secondary member (standby) for stack high-availability.

### **How many secondary member switches are configurable in a VSF stack?**

A VSF stack can be configured with one secondary member only.

### **Once it is configured, is it possible to change the secondary member?**

Yes, a new secondary member can be configured using the `vsf secondary-member <MEMBER-ID>` command. The old standby device will boot first and join the stack with the member role. Then, the newly configured secondary member will go for boot and join the stack with the standby role. The secondary member configuration can only be changed when Member 1 is conductor of the stack.

### **How are conductor and standby for a stack determined?**

By default, the primary member (Member 1) becomes the conductor of the stack and the user-configured secondary member becomes the standby.

The secondary member synchronizes all its states with the conductor. If the current conductor (Member 1) fails, the standby (secondary member) will seamlessly transition to the conductor role. In this state, if Member 1 comes back up, it will take the standby role.

Only primary and secondary members can take up conductor and standby roles in a stack.

### **What is the role of other members in a stack?**

All devices other than the conductor and standby are called members. These devices do not have any network, control, or management plane functions. Their interfaces are directly controlled and managed by the conductor switch.

### **Is there any restriction in the order of VSF member numbering?**

There is no restriction on the order in which VSF members can be numbered. Each member, however, must have a unique number in the range of 1-10 (for 6300 switches) or 1-8 (for 6200F switches).

### **What is the supported stack height and topology?**

- 4100 platforms can stack up to four members.
- 6100 platforms can stack up to six members.
- 6200F platforms can stack up to eight members with no modular SKU (only fixed SKU).
- 6300 F/M and 6300L platforms can stack up to ten members in a chain or ring topology.

Ring is the recommended topology. This topology requires that each member is configured with two VSF links, interconnecting each member with two other members in the stack.

### Can features be configured on a VSF link?

Once an interface becomes part of a VSF link, no standard network layer protocol or feature can run on that interface because it is part of the VSF stack fabric.

### Will configurations in an individual member switch be retained after joining a stack?

Individual member device configurations are not retained after the switch is renumbered and becomes part of a stack.

### How do the consoles of each member in a stack work?

The console of the conductor switch provides a full CLI that can be used to manage the stack. Consoles of other stack members, including the standby, have a limited set of CLI commands that are useful for troubleshooting the device from a stacking functionality standpoint.

### How does an image upgrade for a stack work?

To upgrade a stack to a new firmware image, use the `copy <TFTP/SFTP> image` command to download the image to the device. The image will be downloaded to the stack conductor first and then synced with other members of the stack automatically.

After downloading the firmware, reboot the stack using the `boot system <PRIMARY/SECONDARY>` command. This action completes the upgrade process.

Adding or rebooting individual members before the upgrade process is complete can cause the individual member to fail while joining the stack.

### Is two-member ring supported ?




---

A two-member ring topology is not supported on 4100 and 6100 series switches. For link-level redundancy in two-member stacks, it is recommended to bundle multiple interfaces into a single VSF link using a chain topology.

---

Yes. It is supported from AOS-CX 10.07 onwards.

**show run vsf:**

```
switch(config)# show run vsf
vsf secondary-member 2
vsf member 1
  type j1668a
  link 1 1/1/26
  link 2 1/1/25
vsf member 2
  type j1668a
  link 1 2/1/25
  link 2 2/1/26
```

**show vsf:**

```
switch(config)# show vsf

Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 90:20:c2:20:a2:80
Egress Shape             : Enabled
Egress Shape Rate        : None
Secondary                : 2
Topology                 : Ring
Status                   : No Split
Split Detection Method    : None
```

Member ID	Mac Address	type	Status
1	90:20:c2:20:a2:80	JL668A	Conductor
2	38:21:c7:5a:a5:40	JL668A	Standby

**show vsf topology:**

```
switch# sh vsf topology
Conductor      Standby
+-----+      +-----+
|  1  |1==1|  2  |
+-----+      +-----+
      2              2
      +=====+
```

### Can I add a member to the VSF stack when the member is running an image with a different version than the stack?

When a device joins a stack and its firmware version is different from the version on the conductor, the conductor will push its firmware copy to the device. Once the device receives a copy of the firmware, it will reboot and rejoin the stack, now running the same version as the conductor.

This is not supported if either members or the conductor running on firmware prior to 10.07 version.

### What happens when the VSF conductor switch goes down?

The standby switch, if present, will take the role of the conductor. The old conductor switch will boot and join the stack as the standby switch. This transition will be seamless with limited network impact.

In the absence of a standby (no secondary member configuration), conductor device failure causes the remaining VSF members to reboot and come back up. At this point, members will enter a state in which they are waiting for the conductor to come back up. During this time, front plane ports of the members will be down.

### How do I recover a device that has not joined a stack due to misconfiguration?

The **vsf renumber-to** command is used to trigger a device to take up its new member number and light up its VSF links. This command causes the device to reboot, come back up and wait for messages from the stack conductor. If the VSF link is configured incorrectly or the member number is wrong, the device could be waiting in this state indefinitely.

To recover a device in this state, execute the following commands:

1. Execute the **ctrl+c** command on the device console. This action launches the recovery console.
2. Execute the **erase all zeroize** command on the recovery console.

This action resets the device to factory-default.

- The device will come back up as member ID 1 with no VSF link configuration.
- The device can be configured with the correct member number and VSF links.
- The **vsf renumber-to** command will trigger this configuration to take effect.

The recovery console also has commands that allow the user to copy support files to an external server. This functionality is useful for troubleshooting stacking-related issues.

### **How do the management ports of each member in a stack work?**

In a stack, only the conductor management interface is active. The management interface can be assigned an IP address for device management purposes. When a conductor device fails, the standby becomes conductor and activates its management interface.

### **How does replacing the conductor switch in a stack work?**

The replacement device must be of the same part number as the switch being replaced. You must also have a standby switch configured for replacing the conductor of a stack without losing configuration.

Complete the following steps:

1. Execute the **vsf switchover** command to trigger the standby switch to take over the conductor role.
2. Once the stack is up with the new conductor, remove all physical connections from the old conductor switch that is being replaced.
3. Configure VSF interfaces/links on the new device. It is of critical importance to match the interfaces configured on the switch being replaced.
4. Physically connect the new device to the stack through configured VSF links.
5. The new switch will join the stack and take up the role of standby.

### **What is the workflow for replacing a standby or member switch?**

The replacement device must be of the same part number as the switch being replaced.

Complete the following steps:

1. Configure VSF interfaces/links on the new device. It is of critical importance to match the interfaces configured on the switch being replaced.
2. Renumber the new device to match the switch being replaced.
3. Physically connect the new device to the stack through configured VSF links.
4. The new switch will join the stack and take up the standby or a member role based on the secondary configuration for the stack.

### **What happens if a VSF link fails?**

- If the stack topology is a ring, it will degenerate to a chain when a VSF link in the stack fails.
- If the topology is a chain, a VSF link failure will result in a stack being split into two independent stack fragments.

- When a stack splits and the conductor and standby of the stack become part of two different fragments, the standby takes up the conductor role for its fragment. Network disruption can result because the two fragments are simultaneously active. HPE Aruba Networking highly recommends enabling VSF split-detection to gracefully handle split brain scenarios.
- If a stack splits and the conductor and standby are in the same fragment with the other members on a different fragment, the members-only fragment will:
  - Reboot.
  - Come back up.
  - Wait for communication from the stack conductor.

## What is VSF split-detect?

When a stack splits, the split-detect feature provides a mechanism for the fragments to discover each other.

Once the two stack fragments are discovered, the fragment that has the primary member becomes the active fragment and keeps its front plane (non-VSF) interfaces up and running. The other fragment becomes inactive and all non-VSF interfaces on the inactive fragment are brought down to avoid network disruption.

## How do I configure split-detect?

6200 and 6300 switch series support split-detection through the management interface.

Connect the management interfaces of the primary and secondary members to the same management VLAN/network or connect them directly to one another.

- The CLI command to enable split detection is **vsf split-detect mgmt**.

## How do I configure split-detect using MAD VLAN?

The 4100i and 6100 platforms support split-detection using a dedicated Multi-Active Detection (MAD) VLAN and front-plane ports.

To enable split-detection:

- **Physical Connection:** Connect one of the front-plane ports of the primary member to one of the front-plane ports of the secondary member. Ensure both ports are either: Connected to the same management VLAN/network, or Directly connected to each other.
- **VLAN Configuration:** Reserve a dedicated VLAN (e.g., VLAN ID 100) for split-detection.

## CLI Commands

Use the following commands to enable split-detection:

```
6100(config)# vlan 100
6100(config-vlan-100)# exit
6300(config)# interface 1/1/1,2/1/1
6300(config-if-<1/1/1,2/1/1>)# vlan access 100
6100(config)# vsf split-detect vlan 100
6100# show vsf
Force Autojoin                : Disabled
Autojoin Eligibility Status   : Not Eligible
MAC Address                    : 90:20:c2:05:30:c0
Secondary                      : 2
Topology                       : Chain
Status                         : No Split
```

```

Split Detection Method      : vlan
Split Detection VLAN       : 100

Mbr Mac Address           type           Status
ID
-----
1  90:20:c2:05:30:c0      JL675A      Conductor
2  90:20:c2:06:30:40      JL676A      Standby

switch# show vsf split-detection status
Split Detection Method      : vlan
Split Detection VLAN       : 100
Split Detection VLAN interfaces : 1/1/1,2/1/1
Split Detection Operational status : Up

```




---

Replace interface numbers and VLAN ID as appropriate for your deployment.

---

### How do I remove the non-VSF configurations in a stack?

Use the `erase startup-config` command on the VSF stack. This action will remove all non-VSF related configurations from the startup-config. Then reboot the stack.

### Can a VSF member be removed from a stack?

Yes, remove a member from the stack using the `no vsf member <MEMBER-ID>` command. All configurations associated with the member will also be removed. The member will boot and come back up with the factory default configuration.

### How do I remove the conductor switch from the stack?

HPE Aruba Networking does not recommend removing a member that is conductor of a stack. If the conductor switch has to be removed, complete a switchover and wait for:

- the standby to take up the conductor role, and
- the old conductor to reboot and join the stack as standby.

Then use a member remove command to remove the device from the stack.

### How can I boot the whole VSF stack and individual members using CLI?

The `boot system` command can be used to boot the whole stack.

To boot an individual member, use the `vsf member <MEMBER-ID> reboot` command.

### Is modifying the VSF-specific configuration using Checkpoint restore or TFTP/SFTP/USB download supported?

This functionality is not supported. Before applying a configuration on a stack through Checkpoint restore or TFTP/SFTP/USB download, you must ensure that the following configurations match exactly:

- The current stack VSF configurations.
- The VSF configurations that are part of the configuration file that is being restored or downloaded from the server.

Specifically, the current VSF stack and the Checkpoint/downloaded configuration that will be applied on the stack must have the same:

- Number of members
- Member part number (J#)
- Member number
- VSF link configurations
- Secondary member configuration
- Split-detect configuration

### **How can I dismantle a stack?**

A VSF stack can be dismantled by using the `erase all zeroize` command.

This action will cause each member to reboot, come back up with factory defaults, and function as individual/standalone devices.

### **How do I collect support files for a stacked device?**

The `copy support-files all` command executed on the conductor console will collect support and troubleshooting information from all members that are part of the stack.

If a member is not part of the stack, you must run the same command from the recovery console of the respective member.

If a stack has split into two fragments, both fragments will have a conductor. Execute the same command on the conductor console of both fragments.

### **Can VSF configurations be changed when some of the members are booting?**

No. It is recommended to change VSF configurations only when the stack is in steady state.

To ensure that there is no stack split, it is recommended to form the VSF stack in the ring topology before changing the VSF configurations. This might result in reboot of some of the members.

### **Is there a way to troubleshoot if the members did not join the stack?**

Yes. Use the `show vsf link error-details` command to check if any of the VSF links are down due to error scenarios. For most of the error scenarios, corrective action is also recommended to resolve the issue.

### **Does the command '[no] vsf egress-shape' have any effect on a stack that has no stack members of type S0E91A & S0X44A?**

No. VSF port shaping feature ensures seamless stacking between S0E91A/S0X44A SKU and other 6300 SKUs. The S0E91A and S0X44A SKU have a speed limit of 100G ports, while the speed limit of other 6300 SKUs are limited up to 50G ports. To ensure that all VSF interfaces run at the lowest common speed supported by the entire VSF stack, port shaping is used when the stack has at least one member of type S0E91A or S0X44A.

### **How does VSF port shaping apply when VSF link has multiple interfaces?**

VSF port shaping will ensure VSF links to run at the lowest common speed irrespective of the number of interfaces configured on a VSF link given that the stack consist one or more stack member of type S0E91A or S0X44A. VSF port shaping will be calculated based on all the active VSF interfaces which include multiple interfaces configured on a VSF link.

### **What happens to VSF port shaping when I remove a stack member of type S0E91A or S0X44A from the stack?**

When the stack consist of only one stack member of type S0E91A or S0X44A, the VSF port shaping will be removed from all the active VSF interfaces and the VSF stack operates at the default speed of the VSF interfaces. Whereas when the stack has more that one stack member of type S0E91A or S0X44A, VSF port shaping will be applied on the active VSF interfaces based on the new lowest common speed calculated.


## Auto-stacking

### Is it mandatory to connect the new switch in the direction of the higher-numbered conductor port after configuring the VSF links on the conductor for auto-stacking?

Yes. Auto-stacking process always starts only in the direction of the higher-numbered VSF link port on the conductor. If no switches are connected to the end of the stack connected to the higher-numbered port, the auto-stacking process will not start.

If a new switch being added to the stack is connected in the direction of the lower-numbered port on the conductor, the conductor will show it as an error. Use the `show vsf link error-detail` command to see the error and its recommendation to fix the error.

In this example, Switch3 will join the stack only when it is connected in the direction of the higher-numbered port on the conductor (i.e. to port 1/1/50 on Switch2) as shown in the following figure:



### Can the size of the stack be extended in the direction of lower denomination port of the conductor?

No. You can still renumber manually and add the members to the stack. But the newly added member will not join the stack automatically through auto-stacking.

### What are the different methods to designate the conductor to bring up a stack using auto-stacking?

There are five different ways to designate the conductor and bring up the stack using auto-stacking. The different ways are:

1. Configuring the VSF links manually on the conductor switch.
2. Executing the `vsf start-auto-stacking` command using CLI on the conductor switch.
3. Pressing the **Stk** LED mode button on the conductor switch.
4. Downloading full stack configuration using ZTP.
5. Downloading full stack using TFTP, SFTP, NetEdit, or REST.

### Do all platforms support auto-stacking by pressing the Stk LED mode button on the conductor switch?

No, Auto-stacking by pressing the Stk LED mode button on the conductor switch is supported only on 6200 and 6300 platforms. This feature is not available on other VSF supported platforms.

### What happens when the conductor is designated manually by configuring the lower-numbered port as VSF port first?

This can potentially lead to formation of out-of-order stack since auto stacking happens only in the direction of highest denomination port. If physical connections are already made, the newly added switch might not join the stack.

### **What is the eligibility criteria for a switch to be connected to an existing stack through auto-stacking?**

For a switch to connect to an existing stack, it must be in the auto-join eligible state. A switch in its factory default state is considered to be auto-join eligible.

### **When will a switch become auto-join eligible? Is there a way to make a switch auto-join eligible again to take part in the auto-join process to form the stack?**

If a switch moves out of factory default configuration state, then the switch cannot join an existing stack. In this case, use the `vsf force-auto-join` command to make a non-factory default switch to auto-join eligible again. Once the user sets force auto-join in the switch configuration, the switch will be considered as auto-join eligible and will join the stack even though the switch does not have factory default configuration.

`vsf force-auto-join` command will only work if the switch does not have any pre-existing VSF configurations such as secondary or VSF links. If the switch has VSF configurations already, then the recommendation is to unconfigure and reconfigure `vsf force-auto-join` once all VSF configurations are removed from the switch.

### **Is it mandatory to use only the internally reserved ports to bring up a stack through auto-stacking ?**

If you need to form a stack using **vsf start-auto-stacking** command or by pressing **Stk** LED mode button, then it is mandatory to use the internally reserved VSF ports.

Based on the product type of a switch, the following two interfaces are reserved for the auto-stacking process:

- 12-port switch models: 13 and 14
- 24-port switch models: 25 and 26
- 40-port switch models: 41 and 42
- 48-port switch models: 49 and 50

If auto-stacking via zero-touch provisioning is being used to build the stack, any of the four SFP ports on each member can be used for VSF links.

### **Can a stack be formed through auto-stacking when the conductor is running on 10.07 firmware version and the newly added member is running on firmware version prior to 10.07?**

No. It is mandatory to have all the switches running on 10.7 or later releases to form a stack through auto-stacking .

You cannot form a stack through auto-stacking if either conductor or the stack members running on different firmware versions prior to 10.07.

### **Will a stack be formed if the Stk mode button is pressed on all the members before physically connecting the cables?**

No. Pressing the **Stk** mode button on all the members will configure VSF links and secondary on the switches which will make the members not eligible for auto-join. The members will join with the stack only when it becomes auto-join eligible again.

Pressing **Stk** LED mode button is to designate the conductor. So, press **Stk** LED mode button only on the switch which is supposed to be the conductor of the stack. There should be only one conductor for a VSF stack.

### **Will a stack be formed if *vsf start-auto-stacking* is executed on all the members before connecting the cables physically?**

No. Executing the `vsf start-auto-stacking` command will configure VSF links and secondary on a switch which will make the switch not eligible for auto-join, The members will join with the stack only when they become auto-join eligible again.

Executing `vsf start-auto-stacking` is to designate the conductor. So, execute the command on the switch which is supposed to be the conductor of the stack. There should be only one conductor for a VSF stack.

### **What will happen if **Stk** mode button is pressed on the conductor of an active stack?**

Since the VSF configurations are already present , pressing the **Stk** mode button will not have any effect on the stack configuration. But the LEDs of the stack will now glow to depict the state of the stack.

For more details on the LED states, [Stack and Port LED states](#)

### **What will happen if the *vsf start-auto-stacking* command is executed on the conductor of an active stack?**

Since the VSF configurations are already present , configuring `vsf start-auto-stacking` will not have any effect on the stack configuration. An error message also will be displayed to show that the switch does not have factory default configuration.

### **After downloading the VSF stack configuration to the conductor through TFTP/ZTP/NetEdit, what happens if a new member added has a different SKU than the one provisioned for that particular member-id through auto-stacking?**

If the existing stack size configuration is less than maximum size supported (4 for 4100i switch series, 6 for 6100 switch series, 8 for 6200 switch series and 10 for 6300 switch series), the newly added member will join the stack with the least member-id available, but not with the provisioned member id.

If the existing stack size configuration is already the maximum size supported, then the newly added member will go for a reboot , but will not join the stack. This member will again come up with the factory default configuration as there is a SKU mismatch.

### **What happens if the **Stk** mode button is pressed when the cables are not connected properly on the reserved interfaces, later connected correctly on the reserved interfaces?**

Members switches go for a reboot and join the stack when the cables are connected on the reserved interfaces correctly.

### **Is multi-port VSF configuration supported to bring up a VSF stack through auto-stacking?**

Forming a stack using auto-stacking with multi-ports can be done only when the configuration of all the members are fully pre-provisioned on Member 1.

By default, **Stk** mode button press or `vsf start-auto-stacking` command configures only one port per VSF link. So even if multiple ports were connected physically, stack will come up with single port per VSF link only.

## Can cables of different speed be connected to the members to form a stack through auto-stacking?

Cables of different speed are only supported on 6200 and 6300 switch series.

It is always recommended to have the entire stack with cables of the same speed for VSF links.

## Do 4100i and 6100 platforms support different port speeds for VSF links?

No, the 4100 and 6100 platforms do not support different port speeds for VSF links. All VSF member links must operate at the same speed on these platforms.

## Do 4100i and 6100 platforms support using 1G Copper port and 10G SFP+ port for VSF links?

No, 4100i and 6100 platforms do not support using a 1G Copper port and a 10G SFP+ port for VSF links to form a stack. Stack can be formed by either using 1G Copper port or 10G SFP+ for VSF links across the stack.

## On 4100i and 6100 platforms, how to recover a member that has entered recovery mode due to using a VSF port with a different speed?

To recover a member that has entered recovery mode due to using a VSF port with a different speed, follow these steps:

1. Use the `erase all zeroize` command from the recovery console of the affected member. This will reset the device and bring it up as a standalone unit.
2. Once the device is up as a standalone, configure the VSF ports with desired port speed and then bring up the member.

## On 4100i and 6100 platforms, how do you change the speed of the VSF links from 1G to 10G (or vice versa) when the stack is up?

To change the speed of the VSF links from 1G to 10G (or vice versa) when stack is up, follow these steps:

1. The stack must first be dismantled using the `erase all zeroize` command. This will reset all the members in the stack and bring them up as standalone units.
2. After dismantling, configure VSF links with desired port speed and bring the stack up again with the new port speed.

## What happens when non-reserved ports of the newly added switch is connected to the auto-stacking reserved ports of member 1?

Newly added member will not go for a reboot unless there is a provisioned configuration of the member matching with non-reserved ports on member 1.

To use **Stk** mode button or `vsf start-auto-stacking` command, the cables must be connected to the reserved interfaces on the new switch to start the auto-stacking process. For more information on reserved interfaces, see [Reserved interfaces for auto-stacking](#)

## Can the Stk mode button or vsf start-auto-stacking command be used on a switch which has some VSF configuration already?

No. To use **Stk** mode button or `vsf start-auto-stacking` command, the switch must be in factory-default configuration.

## Is there a way to download configuration automatically after forming a stack via Stk mode button press?

Yes. Once the stack is formed , `ztp force-provision` will be automatically enabled on the stack. But you must have an uplink connectivity to DHCP Server (which can provide the ZTP options) and TFTP server to download the firmware and configuration files.

### **Can I download the full stack configuration via TFTP to the running configuration directly?**

Full stack configuration can be downloaded into the conductor of the stack. The recommendation is to first download the configuration to the start-up and then move the startup to the running configuration. Copying the configuration first to startup will help in detecting errors in the deployed configuration. Once the configuration is copied to startup without any errors, then the startup configuration can be applied to the running configuration. This will also ensure that auto-stacking process did not start prematurely. The stacking process might start prematurely if the configuration is applied directly to the running configuration.

### **How to check whether the switch is auto-join eligible or not?**

Executing the `show vsf` or `show vsf details` command shows an entry called `Autojoin Eligibility Status` which shows whether the switch is eligible or not eligible .

### **Is it possible to change the secondary member of the stack which is formed through `vsf start-auto-stacking` or `stk mode` button press?**

Yes. You can execute the `vsf secondary-member <member-id>` on the conductor to change the secondary member. This will reboot the member 2 (default secondary) and make it to join the stack as member. Then the newly configured secondary member will go for a reboot and joins the stack as standby. Changing the secondary member can only be done from the primary switch (member 1).

### **Will the switch stay as auto-join eligible if any new VSF configurations are made after executing the `vsf force-auto-join` command?**

No. `vsf force-auto-join` is a command used to make an auto-join ineligible switch to auto-join eligible again. If the VSF configuration of the switch gets changed again after executing the `vsf force-auto-join` command, the switch will become auto-join ineligible again. When the VSF configurations are removed , switch will automatically become auto-join eligible.

### **What happens if a member configuration is removed from the conductor of a stack using the `no vsf member <id>` command?**

If the member is part of the stack (not in "Not Present" state) and its ports are connected through reserved auto-stacking ports, then the removed switch will join back the stack after it comes up as standalone. This is because the auto-stacking starts on the reserved ports. So, after member removal, make sure you immediately disconnect the cables physically as well.

### **Will both auto-stacking and ZTP process start simultaneously if the conductor is designated using the configuration download through ZTP,?**

No. The auto-stacking process will start only after the completion of ZTP process.


### **Is there any difference between forming the stack using the `vsf start-auto-stacking` command and `Stk mode` button?**

There is no difference in forming the stack. But the `ztp force-provision` configuration will be added in addition to the VSF related configurations only when the stack is formed using the **Stk** mode button press. If the stack has an uplink connectivity to DHCP server, then the configuration and firmware for the stack can be downloaded from a TFTP server through ZTP.


## Does auto-stacking support ring topology without the need for any configuration changes?

Yes. By default, auto-stacking feature configures two links on each of the VSF members. In case, if there is a need to change the stack from chain to ring topology, connect the first member with the last member of the stack with a cable on the auto-stacking reserved ports.

For example, consider the three-member stack in a chain topology as shown in the following figure:



To change this into ring topology, connect Switch3 and Switch1 as shown in the following figure:



The three-member VSF Stack in chain topology can also be converted into a four-member VSF stack in Ring topology by connecting the Switch4 to port 1/1/50 of Switch3 and port 1/1/49 of Switch1 as shown in the following figure:



## Can I use reserved auto-stacking interfaces as normal data ports?

Yes. The reserved auto-stacking interfaces can be used as normal data ports.

## Can I form a 10G dual-mode SmartRate VSF using ZTP provisioning on the conductor?

No. When the conductor is provisioned with a configuration specifying 10G speed, it tries to establish a 10G link over the VSF ports. However, the member switch is not preconfigured for 10G and defaults to its native speed of 5G. As a result, the VSF link operates at 5G instead of the intended 10G. To achieve the desired 10G link, the member switch must be preconfigured with the appropriate speed settings before provisioning.

## What happens to the VLAN MAD status when performing a dynamic standby change?

When the standby role is assigned to a new member, the VLAN MAD configuration on the previous standby is removed, causing the VLAN MAD interface to go down. To restore functionality, the VLAN MAD interface must be reconfigured on the new standby member.

## Accessing HPE Aruba Networking Support

HPE Aruba Networking Support Services	<a href="https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html">https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html</a>
AOS-CX Switch Software Documentation Portal	<a href="https://arubanetworking.hpe.com/techdocs/AOS-CX/help_portal/Content/home.htm">https://arubanetworking.hpe.com/techdocs/AOS-CX/help_portal/Content/home.htm</a>
HPE Aruba Networking Support Portal	<a href="https://networkingsupport.hpe.com/home">https://networkingsupport.hpe.com/home</a>
North America telephone	1-800-943-4526 (US & Canada Toll-Free Number) +1-650-750-0350 (Backup—Toll Number)
International telephone	<a href="https://www.hpe.com/psnow/doc/a50011948enw">https://www.hpe.com/psnow/doc/a50011948enw</a>

Be sure to collect the following information before contacting Support:

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

### Other useful sites

Other websites that can be used to find information:

HPE Aruba Networking Developer Hub	<a href="https://developer.arubanetworks.com/hpe-aruba-networking-aoscx/docs/about">https://developer.arubanetworks.com/hpe-aruba-networking-aoscx/docs/about</a>
Airheads social forums and Knowledge Base	<a href="https://community.arubanetworks.com/">https://community.arubanetworks.com/</a>
AOS-CX Software Technical Update channel on YouTube.	Videos on new features introduced in this release: <a href="https://www.youtube.com/playlist?list=PLsYGHuNuBZcbWPEjjHuVMqP-Q_UL3Csks">https://www.youtube.com/playlist?list=PLsYGHuNuBZcbWPEjjHuVMqP-Q_UL3Csks</a>

HPE Aruba Networking Hardware Documentation and Translations Portal	<a href="https://arubanetworking.hpe.com/techdocs/hardware/DocumentationPortal/Content/home.htm">https://arubanetworking.hpe.com/techdocs/hardware/DocumentationPortal/Content/home.htm</a>
HPE Aruba Networking software	<a href="https://networkingsupport.hpe.com/downloads">https://networkingsupport.hpe.com/downloads</a>
Software licensing and Feature Packs	<a href="https://licensemanagement.hpe.com/">https://licensemanagement.hpe.com/</a>
End-of-Life information	<a href="https://networkingsupport.hpe.com/end-of-life">https://networkingsupport.hpe.com/end-of-life</a>

## Accessing Updates

You can access updates from the HPE Aruba Networking Support Portal at <https://networkingsupport.hpe.com>.

Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

To subscribe to eNewsletters and alerts:

<https://networkingsupport.hpe.com/notifications/subscriptions> (requires an active HPE Aruba Networking Support Portal account to manage subscriptions). Security notices are viewable without an HPE Aruba Networking Support Portal account.

## Warranty Information

To view warranty information for your product, go to <https://www.arubanetworks.com/support-services/product-warranties/>.

## Regulatory Information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at <https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

### Additional regulatory information

HPE Aruba Networking is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements, environmental data (company programs, product recycling, energy efficiency), and safety information and compliance data, (RoHS and WEEE). For more information, see <https://www.arubanetworks.com/company/about-us/environmental-citizenship/>.

## Documentation Feedback

HPE Aruba Networking is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback

[docsfeedback-switching@hpe.com](mailto:docsfeedback-switching@hpe.com)). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.