

AOS-CX 10.07 Virtual Switching Framework (VSF) Guide

6200, 6300 Switch Series



a Hewlett Packard
Enterprise company

Part Number: 5200-7889
Published: April 2021
Edition: 1

Copyright Information

© Copyright 2021 Hewlett Packard Enterprise Development LP.

Open Source Code

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code 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, send a check or money order in the amount of US \$10.00 to:

Hewlett Packard Enterprise Company
6280 America Center Drive
San Jose, CA 95002
USA

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.

Contents	3
About this document	5
Applicable products	5
Latest version available online	5
Command syntax notation conventions	5
About the examples	6
Identifying switch ports and interfaces	6
Introduction	8
Terminology	8
Connection Topology	9
Feature overview	10
VSF Behavior	10
One Virtual Device	10
Interoperation	14
Link aggregation	14
Stack management	16
Consoles	16
Management interface	16
VSF configuration	16
Member number	16
Access to VSF members	16
VSF links	17
Member provisioning	17
Secondary member	17
Member remove	17
Automated image sync	18
Reboot	18
Member addition without auto-stacking	19
Member addition with auto-stacking	19
Member replacement without auto-stacking	19
Member replacement with auto-stacking	19
Stack and Port LED states	19
Configuring a VSF stack	21
VSF auto-stacking	21
Peer discovery	21
Auto-join eligibility	21
Force auto-join support	21
Designating master switch	22
Reserved interfaces for auto-stacking	23
Forming a four-member ring setup using auto-stacking command	23
Forming an eight-member ring setup manually using link configuration without auto-stacking ..	25
Misconfiguration recovery	28
Failure and recovery	30

Stack split	30
Management interface split detection	30
VSF recommendations and restrictions	33
VSF Commands	35
vsf member	35
member	36
type	37
link	38
vsf force-auto-join	40
vsf start-auto-stacking	40
vsf split-detect	41
vsf secondary-member	42
vsf renumber-to	43
vsf member reboot	44
interface	45
shutdown	45
show vsf	46
show vsf detail	46
show vsf link	48
show vsf link detail	49
show vsf link error-detail	50
show vsf link error-detail member	51
show vsf member	53
show vsf topology	53
Frequently asked questions	55
Frequently asked questions on Auto-stacking	61
Support and Other Resources	66
Accessing Aruba Support	66
Accessing Updates	66
Aruba Support Portal	66
My Networking	67
Warranty Information	67
Regulatory Information	67
Documentation Feedback	67

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

Applicable products

This document applies to the following products:

- Aruba 6200 Switch Series (JL724A, JL725A, JL726A, JL727A, JL728A)
- Aruba 6300 Switch Series (JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL762A)

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 ([]).
example-text	In code and screen examples, indicates text entered by a user.
Any of the following: <ul style="list-style-type: none">■ <code><example-text></code>■ <code><example-text></code>■ <i>example-text</i>■ <i>example-text</i>	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">■ For output formats where italic text cannot be displayed, variables are enclosed in angle brackets (< >). Substitute the text—including the enclosing angle brackets—with an actual value.■ 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.
	Vertical bar. A logical OR 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.

Convention	Usage
[]	Brackets. Indicates that the enclosed item or items are optional.
... or ...	Ellipsis: <ul style="list-style-type: none"> ■ In code and screen examples, a vertical or horizontal ellipsis indicates an omission of information. ■ 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.

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 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 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, defines a virtual switch, comprising multiple individual physical switches, interconnected through standard Ethernet links. These physical switches will operate with one control plane, visible to peers as a virtual switch stack. This composition simplifies management and provides the capability to scale the stack.

On-demand scalability in the access layer allows the user to increase the number of ports on a stack as per needs, without having to manage a new switch. The same stack can scale up or down to match the requirements.

- **6200F:** VSF allows stacks to be formed using any combination of SKUs of the 6200 family. Up to 8 member switches will be allowed. Connections between the switches must use 10G links.
- **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, or 50G links. All VSF links in a stack should operate at the same speed.

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 number 1.
Secondary	The configured secondary member.

Role	Definition
Master	Controls VSF administration and control plane.
Standby	Standby management and under control of the master; synchronizes control plane with the master.
Member	All devices in the stack other than the master and standby are called member switches. 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 master switch.

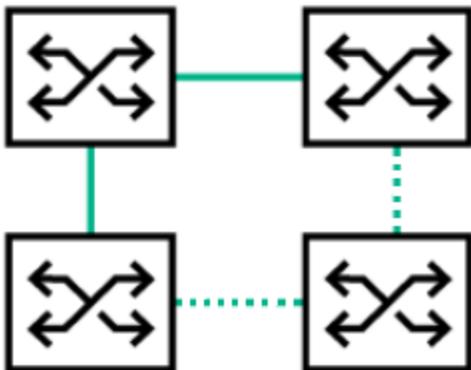
Connection Topology

VSF supports up to 8 member stacks (for 6200F devices) or 10 member stacks (for 6300 devices) in ring and chain topology.

Figure 1 Chain topology



Figure 2 Ring topology



Ring is the recommended deployment topology. It inherently provides resiliency against a single failure of a link or switch.

VSF is always enabled on supported switches. Within the stack, one switch is the Master 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.

VSF Behavior

Each stack member must have a unique member number. When deploying a stack, ensure that each member has a distinct number by renumbering the switches to the appropriate member numbers. Stack formation will fail if there is a member number conflict.

- The primary member will become Master and the secondary member will become Standby under normal circumstances.
- 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, and there is no default secondary member. 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 Master or Standby of the stack.

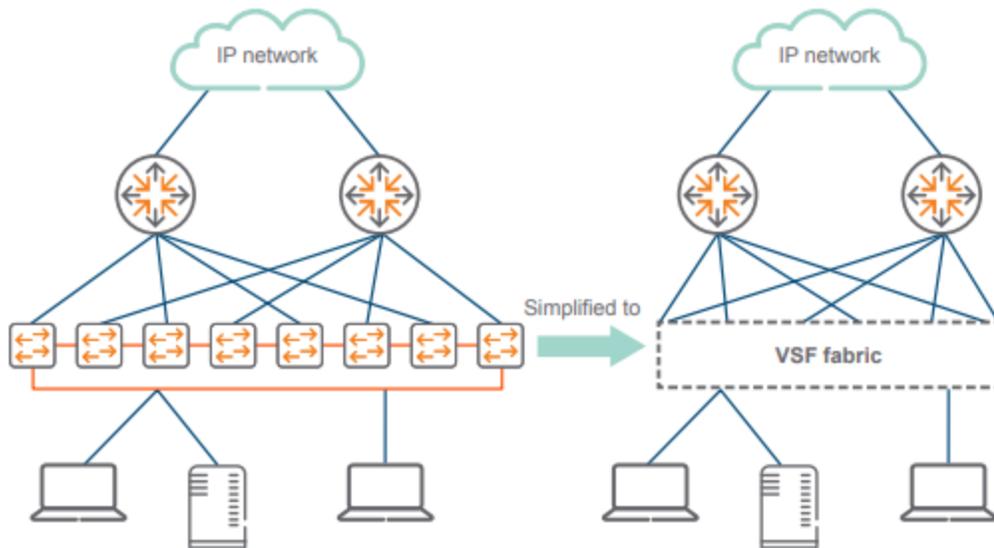


In a standard deployment, uplinks should be from primary and secondary. The management interface from primary and secondary members should be connected to the management network, providing management connectivity to the current master.

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 3 One virtual device example topology



```
switch# show vsf
```

```
Force Autojoin           : Disabled
Autojoin Eligibility Status: Not Eligible
MAC Address              : 08:97:34:b0:0e:00
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  Master
2  08:97:34:b1:43:00  JL665A  Standby
3  08:97:34:b7:cc:00  JL663A  Member
4                               JL662A  Not Present
```

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

Slot number is always fixed as 1. All interfaces are available for configuration.

```
switch# show interfaces brief
```

```
-----
Port      Native Mode  Type      Enabled Status Reason      Speed
```

VLAN							(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 ArubaOS-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 Master. Once a stack is deployed, the stack configuration is sticky. The user can safely remove all other configurations with the command `erase startup-configuration` without disturbing the stack configurations. To remove all configurations, including the stacking configurations, use the command `erase all zeroize`, where all members of the stack will be reset to factory defaults

Interoperation

The VSF stack supports either:

- 6200F devices, or
- 6300 devices (6300M or 6300F).



VSF stacking cannot be done with a mixed set of switches. The stack must be made up of *only* 6200 or *only* 6300 switches.



Aruba 6200F does not support modular units.

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

Link aggregation

Link aggregations (L-Agg) can span interfaces across multiple individual switches within the stack. Load balancing is performed on all interfaces of the L-Agg across the stack and is applicable to 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   State  System-ID      Pri   Key   State
```

```
-----  
1/1/18 lag1 up  
2/1/18 lag1 up  
1/1/23 lag2 up  
2/1/23 lag2 up
```

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
```

Consoles

The serial console of the Master 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 master and standby switch. This enables the stack master console to be reachable after a stack failover to the new Master.



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

Management interface

In a VSF stack, only the management interface on the Master 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 Master.

VSF configuration

The following aspects of VSF are user-configurable.

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 master switch via VSF link. If a renumbered member is unable to communicate with the master 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.

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.

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.

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 Master, the renumbered member will be removed from the stack and will reboot with factory defaults.

Secondary member

The stack will not have a standby member by default. A secondary member can be configured from available members and it will be assigned the role of stack standby.



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

When configured as secondary, a stack member that is already present in the stack will reboot and rejoin the stack as the standby.

A provisioned member can be configured as a secondary member. When the member joins the stack, it will boot up in the standby role, without any further reboot.

If a secondary member is already configured and physically present in the stack, removing the secondary will cause the secondary member to reboot and join as a member.

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 remove

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 master of the stack. If the master has to be removed, the recommendation is to switch over and wait for the old master 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 master. The removed member must be reset to factory-default once it boots to recovery.

Automated image sync

In a VSF environment, all stack members run the same software image. If the user upgrades the software on the Master 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 Master, the member will automatically update itself to the same version as the Master. The member will reboot itself to run the updated version before joining the stack.



Automated image sync is not applicable if the master 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 master. 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 Master, the command will trigger a failover and the Standby switch will take over as Master of the stack.
- If the Standby is unavailable at the time of master 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 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 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 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 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.

Stack and Port LED states

The following table describes the different states of stack **Stk** LED.

Table 3: 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 Port LED based on stack configurations and role of the members in the stack.

Table 4: 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 master and current member, all other ports LEDs glow half-bright green.
Slow flash green	Master 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 master. For example, if Stack Member 4 is the master, 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.

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 work-flow 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 given member of the stack. The new switch will automatically assigned with member ID and it will go for a 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 master using one of the methods described in [Designating master 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 denomination VSF port, sending a VSF peer discovery protocol packet. 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 master.

Auto-join eligibility

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 it moves out of factory default configuration state, it is not considered as auto-join eligible and cannot automatically join an existing stack. However, user can still manually configure the links, renumber the device to make it part of a stack. For more information, see [Forming an eight-member ring setup manually using link configuration without auto-stacking](#).

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.

Designating master switch

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

Following are different methods to designate the master:

- Using LED mode button: Physically connect the switches in a desired topology on the reserved VSF link ports and press the LED mode button until the mode changes to **Stk** on a factory defaults switch. This will automatically configure member 2 as the secondary member and VSF links. In addition to VSF secondary and link configurations, `ztp force-provision` will also be configured on the master 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](#).
- 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 master

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 master.

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 master 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 master 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 master through TFTP/ZTP, the physical connections between the switches should be made according to the downloaded configuration.

Reserved interfaces for auto-stacking

Based on the product type of a switch, the following two interfaces are reserved for the auto-stacking process:

- 24-port switch models: 25 and 26
- 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 following table shows the list of reserved auto-stacking interfaces based on the product type and platform:

Reserved Interfaces

Platform	SKU Part Number	Ports reserved for Auto-Stacking
6300	JL658A JL660A JL662A JL664A JL666A JL668A	25 and 26
6300	JL659A JL661A JL663A JL665A JL667A JL762A	49 and 50
6200	JL724A JL725A	25 and 26
6200	JL726A JL727A JL728A	49 and 50

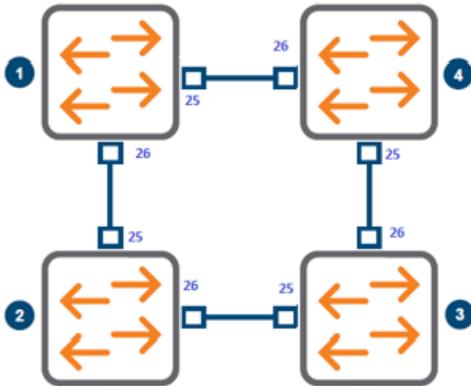
Forming a four-member ring setup using auto-stacking command

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 4 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 master. The links and secondary member will be automatically configured on the master.

```
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 ArubaOS-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 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
vsf member 3
type j1668a
link 1 3/1/26
link 2 3/1/25
vsf member 4
type j1668a
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
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	Master
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

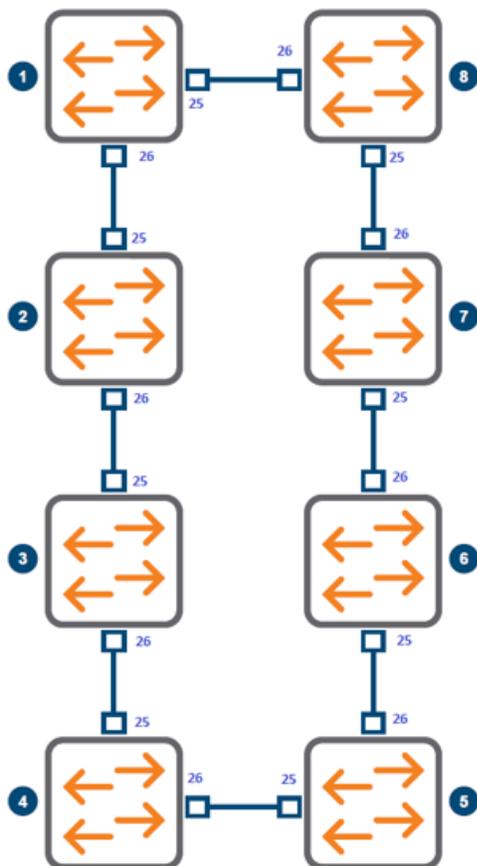


If full stack configuration is downloaded on to the master through TFTP/ZTP, the physical connections between the switches should be made according to the downloaded configuration.

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 5 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 master.
 - 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
Secondary                :
Topology                 : Ring
Status                   : Active Fragment
Split Detection Method   : None

Mbr Mac Address          type          Status
ID -----
1  38:21:c7:5d:d0:c0    JL668A      Master
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:

```
switch(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
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         Master
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.

Misconfiguration recovery

If a switch fails to join the stack because of 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 master, 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 `vsf-factory-reset` command.

```
^C
Login: admin
Password:

recovery# vsf-factory-reset
```

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 following section describes failure and recovery scenarios for stack split issues.

Stack split

Multiple link or member failures can result in a complete stack split, with the master and standby switch in different split fragments. In this case, both master and standby switches will become master switches. Each stack fragment will continue operating with the same configuration and state, and forward traffic between all stations that the fragment can reach.

The downside of this scenario is that each stack fragment will have the same MAC address and IP addresses. To avoid this scenario, configure split detection (described in the following section) which would bring down the interfaces on one fragment to prevent duplicate MAC/IPs.

Management interface split detection



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

VSF stack supports management split detection, which requires users to connect the management interfaces of the primary and secondary stack members to the same L2 network.

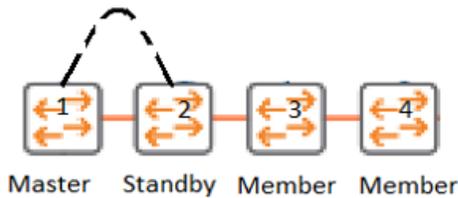
The Primary stack member is member "1", whereas the secondary member is the user-configured secondary switch. Once the stack is split, both of these switches become master of their respective fragments. The fragment with "1" as master is referred to as the primary fragment, and the fragment with the secondary switch as master is the secondary fragment.



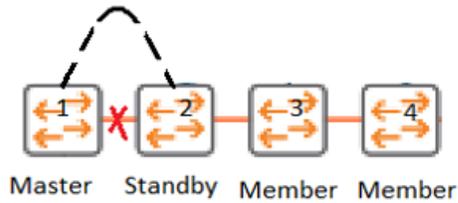
It is also possible to connect the management interfaces of primary and secondary to one another for split detection.

If the secondary fragment discovers that the primary fragment is operational, 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.

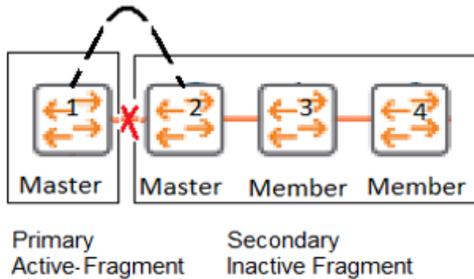
The interfaces will remain down until the stack is reconnected or the primary fragment goes down. The interfaces of the primary fragment will always remain operational.



A four member VSF with the management interfaces of primary and secondary connected and split-detection enabled.



A split happens between Member 1 and Member 2, with Master and Standby in different fragments.



The Primary fragment, that is the fragment with Member 1, becomes the 'Active Fragment'. The Secondary fragment with Members 2-4, will become 'Inactive Fragment' with Member 2 as Master. All the frontplane non-VSF interfaces in the Inactive Fragment will be brought down.

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
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      Master
2                               JL668A      In Other Fragment
3                               JL668A      In Other Fragment
4                               JL668A      In Other Fragment
switch#
switch# show vsf topology
Mstr
+---+
| 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
Secondary       : 2
Topology        : Chain
Status          : Inactive Fragment
Split Detection Method : mgmt
```

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

```
switch#
```

```
switch# show vsf topology
```

```
                Mstr
+---+   +---+   +---+
| 4 |1==2| 3 |1==2| 2 |
+---+   +---+   +---+
```

```
switch#
```

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 master fails and there is no standby, the stack will fail.
- If the master fails and there is a standby device, the standby becomes the new master and will take over stack management. When the old master 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 master. However, once recovered, it is not advisable to use the removed master 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 master 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 master 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 master switch with `vsf secondary <id>` configuration is unable to discover the standby switch. In such cases, the master 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 master'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 master when one or more members of the stack are booting.

vsf member

Syntax

```
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.

Command context

```
config
```

Parameters

<MEMBER-ID>

VSF member identifier.

- Range for 6200F devices: 1-8.
- Range for 6300 devices: 1-10.

Authority

Administrators or local user group members with execution rights for this command.

Examples

Configuring a VSF member:

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

Removing a non-master 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 master should be done with caution as it can make the stack unusable if there is no standby.

member

Syntax

```
member <MEMBER-ID>
```

Description

Connects to the specified member in a VSF environment.

Command context

Manager (#)

Parameters

<MEMBER-ID>

VSF member ID. Required.

- Range for 6200F devices: 1-8.
- Range for 6300 devices: 1-10.

Authority

Administrators or local user group members with execution rights for this command.

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
```

type

Syntax

```
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.

Command context

```
vsf-member-<ID>
```

Parameters

<TYPE>

The part number of the member being provisioned. Required.

Authority

Administrators or local user group members with execution rights for this command.

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

switch(vsf-member-2)# type j1662a

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

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

link

Syntax

```
link <LINK-ID> [<IFRANGE>]
```

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.

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



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

Command context

```
vsf-member-<ID>
```

Parameters

<LINK-ID>

The VSF link number. Required. Range: 1-2.

<IFRANGE>

The interface identifier range. Required.

Authority

Administrators or local user group members with execution rights for this command.

Examples

Creating and modifying VSF links:

```
switch(vsf-member-1) # link  
<1-2>                VSF Link number
```

```

switch(vsf-member-1)# link 1
  IFRANGE      Interface identifier range
  <cr>
switch(vsf-member-1)# link 1 1/1/49-1/1/50
  <cr>
switch(vsf-member-1)# link 2 1/1/52
  <cr>
switch(vsf-member-1)# link 1 1/1/51
  <cr>

switch(vsf-member-1)# show running-config
Current configuration:
!
!Version ArubaOS-CX SL.10.02.0020-741-g11104d6~dirty
!
!
!
!
ssh maximum-auth-attempts 6
!
!
!
!
vlan 1
interface 1/1/49
  no shutdown
interface 1/1/50
  no shutdown
interface 1/1/51
  no shutdown
interface 1/1/52
  no shutdown
vsf member 1
  type jl661a
  link 1 1/1/49-1/1/51
  link 2 1/1/52
exit

switch(vsf-member-1)# no link 1 1/1/47
Port 1/1/47 does not belong to link 1.

switch(vsf-member-1)# no link 1 1/1/48-1/1/49
Port 1/1/48 does not belong to link 1.

switch(vsf-member-1)# no link 2 1/1/49-1/1/51
Port 1/1/50 does not belong to link 2.

switch(vsf-member-1)# no link 1
  <cr>
switch(vsf-member-1)# no link 1
This will cause the stack to split.
Do you want to continue (y/n)? y

switch(vsf-member-1)# no link 2
This will cause the stack to split and the residual stack
fragment will become unusable.
Do you want to continue (y/n)? y

switch(vsf-member-1)# show running-config
Current configuration:
!

```

```
!Version ArubaOS-CX SL.10.02.0020-741-g11104d6~dirty
!  
!  
!  
!  
ssh maximum-auth-attempts 6  
!  
!  
!  
!  
!  
vlan 1  
interface 1/1/52  
    no shutdown  
vsf member 1  
    type j1661a  
exit
```

Before removing an individual interface from the VSF link using the `no vsf link <x> <interface>` command, ensure that the interface is admin shutdown at both local and peer ends.

vsf force-auto-join

Syntax

```
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.

Command context

```
config
```

Authority

Administrators or local user group members with execution rights for this command.

Examples

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

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

vsf start-auto-stacking

Syntax

```
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.

Command context

config

Authority

Administrators or local user group members with execution rights for this command.

Examples

Configuring a VSF secondary member and VSF link on master:

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

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
```

vsf split-detect

Syntax

```
vsf split-detect <MGMT-INTERFACE>
```

Description

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

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

Command context

config

Parameters

<MGMT-INTERFACE>

Configures mgmt-interface as the split detection method. Connect the management interfaces of the primary and secondary members to the same L2 network. Optionally, the management interfaces of primary and secondary can be directly connected to each other.

Authority

Administrators or local user group members with execution rights for this command.

Examples

Configuring mgmt-interface as the split detection method:

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

Removing split detection from the stack:

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

vsf secondary-member

Syntax

```
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.



Member 1 cannot be configured as the secondary member.

Command context

config

Parameters

<MEMBER-ID>

Secondary member number. Required.

- Range for 6200F devices: 2-8.
- Range for 6300 devices: 2-10.

Authority

Administrators or local user group members with execution rights for this command.

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
```

vsf renumber-to

Syntax

```
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.

Command context

```
config
```

Parameters

```
<MEMBER-ID>
```

Member number to which the member will be renumbered. Required.

- Range for 6200F devices: 2-8.
- Range for 6300 devices: 2-10.

Authority

Administrators or local user group members with execution rights for this command.

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
```

vsf member reboot

Syntax

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

Description

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

Command context

Manager (#)

Parameters

<MEMBER-ID>

Member number to be rebooted. Required.

- Range for 6200F devices: 1-8.
- Range for 6300 devices: 1-10.

Authority

Administrators or local user group members with execution rights for this command.

Examples

Rebooting the primary switch of the stack:

```
switch# vsf member 1 reboot
Rebooting the master 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 master switch will reboot and the standby will become the master.
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
```

interface

Syntax

```
interface <IFRANGE>
```

Description

Enters configuration context for one or more VSF link interfaces.

Command context

```
config
```

Parameters

```
<IFRANGE>
```

PORT identifier range. Required.



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

Authority

Administrators or local user group members with execution rights for this command.

Examples

Entering configuration context:

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

shutdown

Syntax

```
shutdown
```

Description

Shuts down one or more VSF link interfaces.

Command context

```
config-if-vsfc
```

Authority

Administrators or local user group members with execution rights for this command.

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.

show vsf

Syntax

```
show vsf
```

Description

Displays the list of provisioned VSF stack members.

Command context

Manager (#)

Authority

Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

Example

Showing the list of provisioned VSF stack members:

```
switch# show vsf  
  
Force Autojoin           : Disabled  
Autojoin Eligibility Status: Not Eligible  
MAC Address              : 08:97:34:b0:0e:00  
Secondary                : 2  
Topology                 : Chain  
Status                   : Active Fragment  
Split Detection Method   : mgmt  
  
Mbr MAC Address          Type           Status  
ID -----  
1  08:97:34:b0:0e:00    JL666A      Master  
2  08:97:34:b1:43:00    JL665A      In Other Fragment  
3  08:97:34:b7:cc:00    JL663A      Member  
4  JL662A               Not Present
```

show vsf detail

Syntax

```
show vsf detail
```

Description

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

Command context

Manager (#)

Authority

Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

Example

```
switch# show vsf detail

VSF Stack
  MAC Address           : ec:eb:b8:d0:80:40
  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                   : Aruba-VSF-6300F
  Contact                :
  Location               :

Member ID                : 1
  MAC Address           : ec:eb:b8:d0:80:40
  Type                  : JL666A
  Model                 : Aruba 6300F 24G PoE CLS 4 /4SFP56 Switch
  Status                : Master
  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                  : JL666A
  Model                 : Aruba 6300F 24G PoE CLS 4 /4SFP56 Switch
  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           : JL666A
Model          : Aruba 6300F 24G PoE CLS 4 /4SFP56 Switch
Status         : Not Present
ROM Version    :
Serial Number  :
Uptime        :
CPU Utilization :
Memory Utilization :
VSF link 1    :
VSF link 2    :

```

show vsf link

Syntax

```
show vsf link
```

Description

Displays the VSF link state for each member.

Command context

Manager (#)

Authority

Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

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

show vsf link detail

Syntax

```
show vsf link detail
```

Description

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

Command context

Manager (#)

Authority

Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

Example

```
switch# show vsf link detail

VSF Member: 1 Link 1

Port      State      Status Code Peer Interface Peer System MAC Peer
Product  Type
-----  -
1/1/27   up         S           2/1/27       38:21:c7:5c:e4:c0
JL668A

1/1/28   error      M           1/1/27       38:21:c7:5c:d7:40
JL668A

VSF Member: 2 Link 1

Port      State      Status Code Peer Interface Peer System MAC Peer
Product  Type
-----  -
2/1/27   up         S           1/1/27       38:21:c7:5c:99:80
JL668A

2/1/28   error      T

VSF Member: 2 Link 2

Port      State      Status Code Peer Interface Peer System MAC Peer
Product  Type
-----  -
```

```

-----
2/1/25   up           S           3/1/26     38:21:c7:5c:f0:00
JL668A

2/1/26   down          D

VSF Member: 3  Link 1

Port      State      Status Code Peer Interface Peer System MAC Peer
Product  Type
-----
-----
3/1/27   error      L           3/1/28     38:21:c7:5c:f0:00
JL668A

3/1/28   error      L           3/1/27     38:21:c7:5c:f0:00
JL668A

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     38:21:c7:5c:e4:c0
JL668A

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
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

```

show vsf link error-detail

Syntax

```
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.

Command context

Manager (#)

Authority

Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

Example

Showing error information of the interfaces about the loop detection:

```
switch# 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
free
                provisioned member entries available matching the peer
                switch with product type JL667A.

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

show vsf link error-detail member

Syntax

```
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.

Command context

Manager (#)

Parameters

<MEMBER-ID>

VSF member identifier. Required.

- Range for 6200F devices: 1-8.
- Range for 6300 devices: 1-10.

Authority

Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

Example

Showing error information and the suggestive action for member 1:

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

VVSF 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
link 1 is          : the same peer switch. Interface 1/1/52 of member 1
                    : 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
auto-join capable

this device via interfaces 25 or 26 - those are the
interfaces on this device.

show vsf member

Syntax

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

Description

Displays information about the specified VSF member.

Command context

Manager (#)

Parameters

<MEMBER-ID>

VSF member identifier. Required.

- Range for 6200F devices: 1-8.
- Range for 6300 devices: 1-10.

Authority

Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

Example

```
switch# show vsf member 1

Member ID           : 1
MAC Address         : ec:eb:b8:d0:80:40
Type                : JL557A
Model               : Aruba JL557A 2930F-48G-740W-PoE+-4SFP Switch
Status              : Master
ROM Version         : SL.10.02.0000-7755
Serial Number       : CN7ZK90012
Uptime              : 0d 0h 18m
CPU Utilization     : 0%
Memory Utilization  : 15%
VSF link 1          : Down
VSF link 2          : Down
```

show vsf topology

Syntax

```
show vsf topology
```

Description

Displays information about VSF stack member connections.

Command context

Manager (#)

Authority

Operators or Administrators or local user group members with execution rights for this command.
Operators can execute this command from the operator context (>) only.

Example

```
switch# show vsf topology
      Stby      Master
+----+ +----+ +----+
| 3 |1==2| 2 |1==1| 1 |
+----+ +----+ +----+
```

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 Aruba 6300F/M Switch Series supports VSF.

VSF can be formed with a combination of any of the Aruba 6300F/M Switch Series (JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL762A) or a combination of any of the Aruba 6200F Switch Series (JL724A, JL725A, JL726A, JL727A, JL728A).



Aruba 6200F Switch Series only supports fixed SKUs.

What port speeds do VSF links support?

For Aruba 6300F/M Switch Series: All uplink ports with 10G, 25G, and 50G speeds can be configured as VSF links.

For Aruba 6200F Switch Series: All uplink ports with 10G speed can be configured as VSF links.

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

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 master 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.

Aruba 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 master of the stack.

How are master and standby for a stack determined?

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

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

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

What is the role of other members in a stack?

All devices other than the master 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 master 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?

- 6200F platforms can stack up to 8 members with no modular SKU (only fixed SKU).
- 6300 F/M platforms can stack up to 10 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 master 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 master 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 ?

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

show run vsf	show vsf	show vsf topology
<pre>6300(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</pre>	<pre>6300(config)# show vsf Force Autojoin : Disabled Autojoin Eligibility Status: Not Eligible MAC Address : 90:20:c2:20:a2:80 Secondary : 2 Topology : Ring Status : No Split Split Detection Method : None Mbr Mac Address type Status ID ----- 1 90:20:c2:20:a2:80 JL668A Master 2 38:21:c7:5a:a5:40 JL668A Standby</pre>	<pre>6300(config)# show vsf topo Mstr Stdby +----+ +----+ 1 1==1 2 +----+ +----+ 2 +=====+</pre>

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 master, the master 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 master.

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

What happens when the VSF master switch goes down?

The standby switch, if present, will take the role of the master. The old master 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), master 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 master 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 master. 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 `vsf-factory-reset` 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 master management interface is active. The management interface can be assigned an IP address for device management purposes. When a master device fails, the standby becomes master and activates its management interface.

How does replacing the master 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 master 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 master role.
2. Once the stack is up with the new master, remove all physical connections from the old master 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 master and standby of the stack become part of two different fragments, the standby takes up the master role for its fragment. Network disruption can result because the two fragments are simultaneously active. Aruba highly recommends enabling VSF split-detection to gracefully handle split brain scenarios.
- If a stack splits and the master 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 master.

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?

VSF supports 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 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 master switch from the stack?

Aruba does not recommend removing a member that is master of a stack.

If the master switch has to be removed, complete a switchover and wait for:

- the standby to take up the master role, and
- the old master 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 master 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 master. Execute the same command on the master 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.

What is 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 user intervention touch points to form a VSF stack. You can add a new switch to the existing stack by physically connecting it to a member of the stack. A member ID is automatically assigned to the new switch and rebooted. After rebooting, the newly added member joins the stack.

Is it mandatory to connect the new switch in the direction of higher denomination master port only after configuring the VSF links on the master for auto-stacking?

Yes. Auto-stacking process always starts only in the direction of higher denomination port of the master. If no switches are connected to the higher denomination port, auto-stacking process will not start.

If a new switch added for stacking is connected in the direction of lower denomination port of the master, the master 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.

Can the size of the stack be extended in the direction of lower denomination port of the master?

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 master to bring up a stack using auto-stacking?

There are five different ways to designate the master and bring up the stack using auto-stacking. The different ways are:

1. Configuring the VSF links manually on the master switch.
2. Executing the `vsf start-auto-stacking` command using CLI on the master switch.
3. Pressing the **Stk** LED mode button on the master switch.
4. Downloading full stack configuration using ZTP.
5. Downloading full stack using TFTP, SFTP, NetEdit, or REST.

What happens when the master is designated manually by configuring lower denomination 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 becomes 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 ?

No. 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:

- 24-port switch models: 25 and 26
- 48-port switch models: 49 and 50

However, if the stack is formed by any of the other methods , you must ensure that the physical interconnections on the stack ports between different members of the topology is matching with the provisioned configuration on the master.

Can a stack be formed through auto-stacking when the master 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 master 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 master. So, press **Stk** LED mode button only on the switch which is supposed to be the master of the stack. There should be only one master 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 master. So, execute the command on the switch which is supposed to be the master of the stack. There should be only one master for a VSF stack.

What will happen if Stk mode button is pressed on the master 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 master 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 master 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 (10 for 6300 switch series, 8 for 6200 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?

Yes. Though it is supported , it is always recommended to have the entire stack with cable of same speed for VSF links.

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 master 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 master 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 master 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 master 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.

Can I use reserved auto-stacking interfaces as normal data ports?

Yes. The reserved auto-stacking interfaces can be used as normal data ports.

Accessing Aruba Support

Aruba Support Services	https://www.arubanetworks.com/support-services/
Aruba Support Portal	https://asp.arubanetworks.com/
North America telephone	1-800-943-4526 (US & Canada Toll-Free Number) +1-408-754-1200 (Primary - Toll Number) +1-650-385-6582 (Backup - Toll Number - Use only when all other numbers are not working)
International telephone	https://www.arubanetworks.com/support-services/contact-support/

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:

Airheads social forums and Knowledge Base	https://community.arubanetworks.com/
Software licensing	https://lms.arubanetworks.com/
End-of-Life information	https://www.arubanetworks.com/support-services/end-of-life/
Aruba software and documentation	https://asp.arubanetworks.com/downloads

Accessing Updates

You can access updates from the Aruba Support Portal or the HPE My Networking Website.

Aruba Support Portal

<https://asp.arubanetworks.com/downloads>

If you are unable to find your product in the Aruba Support Portal, you may need to search My Networking, where older networking products can be found:

My Networking

<https://www.hpe.com/networking/support>

To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

<https://support.hpe.com/portal/site/hpsc/aae/home/>

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HP Passport set up with relevant entitlements.

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://asp.arubanetworks.com/notifications/subscriptions> (requires an active Aruba Support Portal (ASP) account to manage subscriptions). Security notices are viewable without an ASP 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

Aruba 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

Aruba 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). 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.