

TECHNICAL WHITE PAPER

BACKPLANE STACKING AND VSF BEST PRACTICES ON AOS-SWITCH

AOS-SWITCH VERSION 16.08

CONTENTS

Revision History	2
Purpose	3
Backplane Stack Deployment	3
Automatic provisioning	3
Strict provisioning	4
Stack member priority	4
Split stack mitigation	4
Global OoBM address assignment	5
Virtual Switching Framework (VSF) Deployment	5
Semi-automatic provisioning	5
Strict provisioning	6
VSF member priority	6
VSF split mitigation	6
Global OoBM address assignment (5400R only)	7
Backplane Stacking Maintenance and Troubleshooting	7
Split stack operations	7
Forcing a Commander failover	8
Removing a stack member	8
Renumbering stack members	8
Replacing a stack member	9
Moving switch from one stack to another	10
Software upgrades	10
Erasing stack configuration	10
VSF Maintenance and Troubleshooting	10
Split stack operations	10
Forcing a Commander failover	11
Removing a VSF member	11
Renumbering VSF members	12
Replacing a VSF member	12
VSF Fast Software Upgrade (5400R only)	13
Software upgrades (2930F and 5400R)	14
Appendix: Timings	15
3810M backplane stack (4 members)	15
2930M backplane stack (4 members)	15
5400R VSF (2 members)	15
2930F VSF (4 members)	15
Supported Platforms	16

REVISION HISTORY

Document Version	Reason for Change	Revision Date
Initial Release		Jan 2019
1.1	Formatting updates; updated software platform name to AOS-Switch; added clarification on software version mismatch behavior on new VSF/stack members	Jun 2019

PURPOSE


Combining multiple switch chassis into a single logical device can provide increased capacity and improved redundancy for access and distribution workloads, while providing a single point of management access. Several Aruba switch families — 2930F, 2930M, 3810M, and 5400R — provide either dedicated backplane stacking hardware or frontplane stacking capability (Virtual Switching Framework, or VSF). This document contains guidance and best practices for deploying and maintaining switch stacks using the latest release of AOS-Switch.

BACKPLANE STACK DEPLOYMENT

The 2930M and 3810M switch series support dedicated backplane stacking modules and cables that provide high-bandwidth links for traffic forwarding and state synchronization between stack members. Note that stacking modules are not hot-swappable; the switch must be powered down before a module may be installed or removed.


The recommended stack topology for the 2930M is a *ring*, in which every stack member has two available paths back to the commander (for 2-member stacks, both stack ports are used on each member; for 3+ members, each stack member is connected to two other members such that there is a continuous path connecting the entire stack). This ensures that the failure of any individual stacking link or member will not result in a stack split.

A stack operating without the additional link between the first and last members is described as a *chain*; the failure of a stack link or member in a chain could result in a stack split, the effects of which are described in a later section.



For the 3810M, the recommended topology for stacks of up to 5 members is a *mesh*; each stack member is connected to every other member to maximize reliability by providing multiple paths for rerouting traffic in the event of a link or chassis failure. For stacks of greater than 5 members, the ring topology, as described above, is recommended.

A backplane stack may contain a mixture of models from the same switch series — any 2930M can stack with any other 2930M, and any 3810M may stack with any other 3810M, as long as they are running the same switch software version. Stacks cannot be formed between members of different switch series, even if they share a common stacking cable type (i.e., a 3810M cannot stack with a 3800, and a 2930M cannot stack with a 2920).



Automatic provisioning

To provision a new backplane stack, it is recommended that all prospective stack members start powered down in a factory default state. Install stacking modules in all members and connect all stacking cables. Power up stack members in the desired stack ID order (i.e., power up stack member 1 first); allow each member to complete the boot sequence and, for members 2 and onward, complete synchronization with the Commander (which is, by default, the first stack member to boot) before powering up the next member in sequence. If the new member is running a different software version than the Commander, note that additional reboots may occur as the software images on the new member are replaced with those on the Commander. While Standby synchronization is in progress, no configuration changes can be made. When the last member has been added to the stack, ensure remaining stacking cables are

connected to either close the ring or complete the mesh, as appropriate.

Note that no pre-configuration is required to deploy a backplane stack in this fashion; a switch powering up in a factory default state with a stacking module installed and stacking cables connected will automatically boot in stacking mode.

Note: In some cases, a new member may be running a version of the backplane stacking protocol that is incompatible with that running on the Commander; this typically occurs when the new member is running a software image that is several major releases older than the Commander (e.g., the Commander is running 16.08 and the new member is running 16.02). If this occurs, the software image on the new member will need to be manually updated to match the Commander's running software and then factory reset before it can be joined to the stack, using these commands as an example:

```
switch# copy sftp flash sftpuser@10.10.10.1 KB_16_08_0003.swi primary
The primary image will be deleted.
```

```
Continue (y/n)? y
switch# erase all
The system will be rebooted and all management module files except
software images will be erased.
```

```
Continue (y/n)? y
```

Strict provisioning

As an alternative, once the first stack member has been booted, the remaining members of the stack can be pre-provisioned to ensure that they are assigned specific member IDs regardless of the order in which they are powered up and connected to the Commander.

To do so, once the initial Commander is up and running with stacking enabled, use the following commands to provision additional members by type and MAC address (replace the listed values as needed):

```
switch(config)# stacking member 2 type R0M67A mac-address 308d99-000002
switch(config)# stacking member 3 type R0M68A mac-address 308d99-000003
switch(config)# stacking member 4 type R0M68A mac-address 308d99-000004
```

Once members have been provisioned, member-specific interface, VLAN, and other settings can also be configured prior to adding the new members to the stack.

Stack member priority

You can specify a designated stack Commander and/or Standby by setting priority values for each member; when a stack commander election occurs, the member with the highest priority is selected as Commander, and the second highest is selected as Standby. If all members have the same priority, the member with the lowest MAC address is selected as Commander. To manually set priority values for stack members, use the following examples as a guideline (changes will be applied after a stack reboot):

```
switch(config)# stacking member 1 priority 255
switch(config)# stacking member 2 priority 192
```

Split stack mitigation

In order to mitigate the effects of a split-stack scenario, the out-of-band management (OoBM) ports on all members should be connected to a common upstream device, which will allow stack member discovery to take place in the event of a hardware failure that splits the stack.

Global OoBM address assignment

Another purpose of having all stack OoBM ports connected is the ability to configure a global OoBM management address and default gateway for the entire stack, allowing an administrator to open a stack management session using a common IP address regardless of which member is acting as the commander. To do so, use the following commands, replacing listed values as needed:

```
switch(config)# oobm
switch(oobm)# ip address 10.100.10.10/24
switch(oobm)# ip default-gateway 10.100.10.1
```

VIRTUAL SWITCHING FRAMEWORK (VSF) DEPLOYMENT

The 2930F and 5400R switch series support the ability to use front panel switch ports to link switches in a ‘virtual’ stack using the VSF feature, with these links being used for traffic forwarding and state synchronization between VSF members. Multiple ports can be aggregated to increase VSF link bandwidth; on the 5400R, up to eight 10G or 40G ports can be used in a VSF link, while on the 2930F, up to eight 1G ports or four 10G ports can be used per link.

The 5400R supports up to two VSF members, while the 2930F supports up to eight members. The recommended topology for the 2930F is a ring, which is functionally identical to the topology described for the 2930M in the previous section. Where possible, use at least two ports per VSF link at the highest speed available, both for redundancy and to maximize bandwidth for intra-fabric traffic forwarding.

As with backplane stacking, a 2930F VSF fabric may contain a mixture of models from the same switch series — any 2930F may stack with any other 2930F, while a pair of 5400R chassis (either two 5406Rs or two 5412Rs, not a mix of both) may form a fabric with any combination of interface modules, as long as they are running the same switch software version and support the configured VSF link speed. A VSF fabric cannot be formed between different switch series, even if the VSF feature is supported on both.

Semi-automatic provisioning

For this provisioning method, start with the initial Commander powered up with a management console open (SSH, Telnet, or serial port), and leave the remaining members powered down in a factory default state. Use the following command to enable VSF and specify a unique domain ID (substitute the number 1 with your desired value):

```
switch(config)# vsf enable domain 1
This will save the current configuration and reboot the switch.
Continue (y/n)? y
```

Once the switch reboots into VSF mode and assumes the Commander role, configure **one or more** ports as the initial VSF link:

```
switch(config)# vsf member 1 link 1 1/49,1/50
```

If the fabric will contain **more than two members**, configure a second VSF link as well:

```
switch(config)# vsf member 1 link 2 1/51,1/52
```

Connect the VSF link cable(s) **to the second member and power it on**. Once the switch brings up ports and detects the active VSF link, it will reboot and join the fabric automatically. If the new member is running a different software version than the Commander, additional reboots may occur as the software images on the new member are replaced with those on the Commander. For 2930F stacks of greater **than two members**, **configure each member’s second VSF link once it has joined the fabric** and completed synchronization with the Commander:

```
switch(config)# vsf member 2 link 2 2/49,2/50
```

When the last member has been added to the fabric, remember to connect its second VSF link to the second link on member 1 to close

the ring.

Note: In some cases, a new member may be running a version of the VSF protocol that is incompatible with that running on the Commander; this typically occurs when the new member is running a software image that is several major releases older than the Commander (e.g., the Commander is running 16.08 and the new member is running 16.02). If this occurs, the software image on the new member will need to be manually updated to match the Commander's running software and then factory reset before it can be joined to the fabric, using these commands as an example:

```
switch# copy sftp flash sftpuser@10.10.10.1 WC_16_08_0003.swi primary
The primary image will be deleted.
```

```
Continue (y/n)? y
switch# erase all
The system will be rebooted and all management module files except
software images will be erased.
```

```
Continue (y/n)? y
```

Strict provisioning

The strict provisioning process is similar to that used for backplane stacking, with the caveat that VSF links must be configured in addition to the type and MAC address for each member. Use the following commands, replacing the listed values as needed:

```
switch(config)# vsf member 2 type j1256a mac-address e0071b-000002
switch(config)# vsf member 2 link 1 2/49,2/50
switch(config)# vsf member 2 link 2 2/51,2/52
switch(config)# vsf member 3 type j1256a mac-address e0071b-000003
switch(config)# vsf member 3 link 1 3/49,3/50
switch(config)# vsf member 3 link 2 3/51,3/52
```

VSF member priority

You can specify a designated VSF Commander and/or Standby by setting priority values for each member; when a VSF Commander election occurs, the member with the highest priority is selected as Commander, and the second highest is selected as Standby. If all members have the same priority, the member with the lowest MAC address is selected as Commander.

To set priority for VSF members, use the following commands (replace listed values as appropriate):

```
switch(config)# vsf member 1 priority 255
switch(config)# vsf member 2 priority 192
```

The same logic used for selecting a backplane stack Commander is used for VSF Commander elections.

VSF split mitigation

In order to mitigate the effects of a VSF split, a Multi-Active Detection (MAD) method should be configured and utilized.

For the 5400R, OoBM ports on both members should be connected to a common second device, as with backplane stacks, and OoBM-MAD should be enabled using the following command:

```
switch(config)# vsf oobm-mad
```

On the 2930F, the simplest method is VLAN-MAD, which requires designating a specific VLAN with at least one port assigned per VSF member and connected to a common second device, allowing discovery of other members in the event of a VSF split.

To specify the VLAN-MAD VLAN and assign it to member ports, use the following command (replacing VLAN ID and port numbers as needed):

```
switch(config)# vlan 999 name VLAN-MAD
switch(config)# vlan 999 untag 1/47,2/47,3/47,4/47,5/47,6/47
switch(config)# vsf vlan-mad 999
```

Once configuration is complete, connect the VLAN-MAD ports to the second device; doing so before issuing the VLAN-MAD command may cause a broadcast storm, impacting fabric performance.

Global OoBM address assignment (5400R only)

As with a backplane stack, the 5400R management module OoBM ports can be assigned a global management IP address and gateway for use in managing the fabric from a single point of access. Use the following commands, replacing values as appropriate:

```
switch(config)# oobm
switch(oobm)# ip address 10.100.20.10/24
switch(oobm)# ip default-gateway 10.100.20.1
```

BACKPLANE STACKING MAINTENANCE AND TROUBLESHOOTING

Split stack operations

When the stack is configured in a chain topology, in the event of a stacking link or chassis failure that results in a stack split, the stack manager software component will automatically designate one of the two resulting stack fragments as *inactive*, shutting down all switch ports on members that are part of that fragment, while the other fragment remains active with traffic being forwarded normally.

The behaviors described in the following scenarios assume that all member OoBM ports are connected as described in the deployment section.

If a stacking link fails, but all stack members remain operational, one of the following occurs:

- If both fragments are the same size, the fragment containing the current Commander remains active and the other fragment becomes inactive; a new Standby is selected in the active fragment if not already present and at least one other active member is available.
- If the fragments are not of equal size, the *larger* fragment remains (or becomes) active and the *smaller* fragment becomes inactive.
 - If either the Commander or Standby are in the larger fragment, they retain or assume the Commander role; a new Standby is selected, if not already present.
 - If both the Commander and Standby are in the smaller fragment, that fragment remains active until all members of the larger fragment have rebooted and a new Commander and Standby are selected.

If a stack member fails and causes a stack split, one of the following occurs:

- If the remaining fragments are the same size, the fragment containing the Commander remains active.
 - If the Commander failed, the fragment containing the Standby remains active, the Standby assumes the Commander role, and a new Standby is selected in the active fragment if at least one other active member is present.
 - If the Standby failed, the fragment containing the Commander remains active, and a new Standby is selected if there is at least one other active member present.
- If the remaining fragments are not of equal size, the larger fragment remains (or becomes) active.

- If the Commander is in the larger fragment, it retains the Commander role; a new Standby is selected if not already present.
- If the Commander failed, and the Standby is in the larger fragment, the Standby assumes the Commander role and a new Standby is selected.
- If neither the Commander nor Standby are in the larger fragment, the smaller fragment will remain active, either with the existing Commander or with the Standby assuming the Commander role, until all members of the larger fragment have rebooted and a new Commander and Standby are selected.

Once the failure has been resolved and connectivity between members has been restored, a stack merge will occur and all members of the inactive fragment will reboot. The Commander in the active fragment will retain its role; a new Standby will be selected if one was not present.

Forcing a Commander failover

There are a few situations in which you might wish to manually trigger a Commander failover to the Standby: the Commander is to be removed from the stack and replaced or repurposed; testing for network behavior and potential performance issues during initial deployment; or having the primary Commander reassume the role after a previous failover.

The preferred method for initiating a manual failover is to use the following command:

```
switch# redundancy switchover
```

This will trigger a reboot of the Commander, resulting in the Standby assuming the Commander role. If there is at least one other operating member available, a new Standby will be selected. Once the former Commander reboots, it will rejoin the stack as either the Standby (if part of a 2-member stack) or a member (if stack is 3 or more members).

Note: the `reload` command functions identically to `redundancy switchover` if executed in a stack with at least 2 members.

Removing a stack member

To remove a stack member without replacement (i.e., you are repurposing the switch for use elsewhere), use the following command:

```
switch(config)# stacking member 5 remove
```

```
The specified stack member will be removed from the stack and
its configuration will be erased. The resulting configuration
will be saved. The stack member will be shutdown. Continue [y/n]? y
```

This command will cause the member to shut down, and it will be removed from both the running and startup configuration. It can then be disconnected and repurposed elsewhere.

The stack Commander cannot be removed with this command; a failover to the Standby must first be performed. If the Standby is removed, a new Standby will be selected if at least one other operating member is present. This command also cannot be used if removing the member would result in a split stack; ensure that the stack is operating in a mesh or ring topology before attempting to remove the member.

Renumbering stack members

There is no native capability to renumber or reorder stack members; doing so requires removing the members in question from the stack, then manually re-provisioning them in the desired order. This process is disruptive, as each member removed from the stack will be shut down and must be manually powered back up once renumbering is complete. (Shutting down each member, rather than forcing them to reboot, prevents them from immediately rejoining the stack with an automatically-assigned member ID.)

Before performing these steps, back up the stack configuration for use when reapplying member-specific settings after renumbering is

complete.

In this example, the switches acting as stack members 3 (MAC address 308d99-000004) and 4 (MAC address 308d99-000003) were powered up in the wrong order during initial stack provisioning, and need to be renumbered to reflect their physical positions in their rack.

Start by removing and shutting down the stack members to be renumbered:

```
switch(config)# stacking member 3 remove
switch(config)# stacking member 4 remove
```

Next, re-add both members to the stack in the desired order:

```
switch(config)# stacking member 3 type ROM68A mac-address 308d99-000003
switch(config)# stacking member 4 type ROM68A mac-address 308d99-000004
```

Power-cycle the two stack members and verify that they join the stack in the desired order. Finally, re-apply any member-specific configuration settings that were cleared when the members were removed.

Replacing a stack member

A number of scenarios may require the replacement of an individual member of the stack — a partial or complete failure of switch hardware, replacing a non-PoE model with a PoE model, or replacing a 24-port model with a 48-port model to increase capacity. Note that replacing a stack member with a different model requires the *removal* of the member from the stack configuration (see **Removing a stack member**) before the replacement member is added, and any member-specific configuration must be re-applied after the replacement has been added.

To replace a stack member, it must not currently be operating as part of the stack; an operating member must be shut down or disconnected from the stack. If the stack is operating using a ring or mesh topology, shut down the member using the following command (replace the 2 with the desired member ID):

```
switch# stacking member 2 shutdown
The specified stack member will be shut down. Continue [y/n]? y
```

The *shutdown* command cannot be used if shutting down the specified member would result in a stack split:

Shutting down this member is not allowed since it would result in a stack split.

If possible, add sufficient stacking links to place the stack in a ring or mesh topology prior to removing the original member. If this is not feasible, and a temporary stack split condition is considered acceptable, simply power down the member to be replaced.

Once the switch has been shut down, disconnected, or removed from the stack configuration (if required), use the following command to replace it in the stack configuration (replace the type and MAC address with the desired values):

```
switch(config)# stacking member 2 type ROM68A mac-address 308d99-000001
```

If attempting to replace the original member with a different model, but the original member has not yet been removed from the stack configuration, the following error message will be received:

```
switch(config)# stacking member 2 type JL074A mac-address 308d99-000001
This will save the current configuration. Continue [y/n]? y
The switch type does not match the configured
switch type for this memberId.
```

Install a stacking module in the replacement member, connect the stacking cables, and connect it to power. Verify that the member joins

the stack as expected and that the resulting port and VLAN configuration is correct; re-apply member-specific configuration if required.

Moving switch from one stack to another

When a switch is disconnected from a stack without being properly removed, and is intended for use in another stack at some point in the future, the existing stacking configuration must be removed before the switch can be utilized in the new stack (as the switch retains the stack ID and configuration from its previous stack association).

To do so, use the following command after the switch has been disconnected from the original stack, but before it has been added to the new stack:

```
switch# stacking factory-reset  
The current configuration will be deleted and the device rebooted.  
  
Continue (y/n)? y
```

Once the reboot is complete, the switch will be in a factory default state, and can be connected to another stack.

Software upgrades

In order to upgrade software on a 2930M or 3810M backplane stack, the new software image must be loaded into either the primary or secondary flash on the Commander, the new image must be automatically synchronized to the same flash storage on all other stack members (this occurs in the background), then the stack must be rebooted into the new software version (there is no capability to perform an in-service software upgrade on either platform).

To load the software image from an SFTP server, use the following command (replace server IP, username/password, and software image filename as appropriate):

```
switch# copy sftp flash sftpuser@10.10.10.1 WC_16_08_0003.swi primary
```

To reboot the stack into the new software version, use this command:

```
switch# boot system flash primary
```

The entire stack will reboot; once the boot process is complete, a Commander election will occur as normal.

Erasing stack configuration

Because backplane stacking provides a persistent connection between members and some stacking state information is stored in a stack data file in flash separate from the startup and running configuration, the stack's startup configuration can be erased while preserving the stack topology and member IDs. Use the following command:

```
switch# erase startup-configuration  
The current configuration will be deleted, existing login passwords removed, and the device rebooted.  
  
Continue (y/n)? y
```

Note that using the command `erase all` or `erase all zeroize` will instead erase *all* files in flash storage on all stack members, including the stack data file.

VSF MAINTENANCE AND TROUBLESHOOTING

Split stack operations

As with backplane stacks, a VSF fabric operating as a chain may experience a split stack condition due to the failure of a VSF link or fabric

member. The logic applied by the VSF manager software component, however, differs somewhat from its counterpart.

The behaviors described in the following scenarios presume that either OoBM-MAD, LLDP-MAD, or VLAN-MAD are configured and functioning.

If a VSF link fails, but all fabric members remain operational, the following occurs:

- The fragment containing the Commander remains active, and a new Standby is selected if not already present and another active member is present.
 - If the existing Standby ended up in the inactive fragment, it assumes the Commander role, and a replacement Standby is selected if another member is present; otherwise, all members of the fragment reboot and a new Commander and Standby are selected, though all ports (except configured VSF link ports) are disabled.

If a VSF member fails and causes a split, the following occurs:

- The fragment containing the Commander remains active, and a new Standby is selected if not already present and at least one active member is available.
 - If the existing Standby ended up in the inactive fragment, it assumes the Commander role, and a replacement Standby is selected if another member is present; otherwise, all members of the fragment reboot and a new Commander and Standby are selected, though all ports (except configured VSF link ports) are disabled.
- If the split was caused by the failure of the Commander, the fragment containing the Standby remains active and the Standby assumes the Commander role; a new Standby is selected if at least one other active member is available. All members of the inactive fragment reboot and a new Commander and Standby are selected, though all ports (except VSF link ports) are disabled.

Forcing a Commander failover

There are a few situations in which you might wish to manually trigger a Commander failover to the Standby: the Commander is to be removed from the stack and replaced or repurposed; testing for network behavior and potential performance issues during initial deployment; or having the primary Commander reassume the role after a previous failover.

The preferred method for initiating a manual failover is to use the following command:

```
switch# redundancy switchover
```

This will trigger a reboot of the Commander, resulting in the Standby assuming the Commander role. If there is at least one other operating member available, a new Standby will be selected. Once the former Commander reboots, it will rejoin the stack as either the Standby (if part of a 2-member fabric) or a member (if stack is 3 or more members).

Note: the `reload` command functions identically to `redundancy switchover` if executed in a VSF fabric with at least 2 members.

Removing a VSF member

To remove a VSF fabric member, use the following command (replace the member ID with the desired value):

```
switch(config)# vsf member 4 remove
```

```
The specified VSF virtual chassis member will be removed and
its configuration will be erased. The resulting configuration
will be saved. The VSF member will be shutdown. Continue (y/n)? y
```

This command will erase the member's saved configuration and immediately shut it down; this prevents it from rebooting and automatically rejoining the fabric.

The operating Commander cannot be removed using this command, nor can a member be removed if doing so would cause a VSF split. Ensure that the VSF fabric is operating in a ring topology, if needed, and execute a manual failover if the Commander is to be removed.

Renumbering VSF members

As with a backplane stack, renumbering VSF members requires removing the members in question from the fabric, then manually re-provisioning them in the desired order. This process is disruptive, as each member removed from the fabric will be shut down and must be manually powered back up once renumbering is complete.

Before performing these steps, back up the VSF configuration for use when reapplying member-specific settings after renumbering is complete.

In this example, the switches acting as stack members 2, (MAC address ecebb8-000004), 3 (MAC address ecebb8-000002), and 4 (MAC address ecebb8-000003) were connected to the fabric in the wrong order during initial provisioning, and need to be renumbered to reflect their physical positions in their rack.

Start by removing and shutting down the stack members to be renumbered; in this example, members 3 and 4 will have their member IDs swapped:

```
switch(config)# stacking member 2 remove
switch(config)# stacking member 3 remove
switch(config)# stacking member 4 remove
```

Next, re-add both members to the stack in the desired order:

```
switch(config)# stacking member 2 type JL557A mac-address ecebb8-000002
switch(config)# stacking member 3 type JL557A mac-address ecebb8-000003
switch(config)# stacking member 4 type JL557A mac-address ecebb8-000004
```

Power-cycle the three fabric members and verify that they join the stack in the desired order. Finally, re-apply any member-specific configuration settings that were cleared when the members were removed.

Replacing a VSF member

The process for replacing a VSF member is similar to replacing a backplane stack member, in that the member to be replaced must be shut down, disconnected, or removed from the VSF configuration prior to replacement. If the member to be replaced is currently in operation, and the VSF topology is a ring (or the member is at either end of a chain), use the following command to perform a controlled shutdown (replace the 6 with the desired member ID):

```
switch# vsf member 6 shutdown
The specified VSF virtual chassis member will be shut down. Continue (y/n)? y
```

As with the backplane stack equivalent, the command will not allow shutdown of a member if it would result in a VSF split:

```
Shutting down this member is not allowed since it would result in a stack split.
```

If possible, add VSF links as required to place the fabric in a ring topology prior to attempting to remove the member. If this is not feasible, and a temporary VSF split condition is considered acceptable, simply power down the member to be replaced before continuing.

If the member is to be replaced with a different model, the member must first be removed from the configuration (see **Removing a VSF Member**) before the replacement is added. Failure to do so will result in the following error message being displayed when attempting to replace the original member's configuration:

```
switch(config)# vsf member 6 type j1255a mac-address e0071b-000001
This will save the current configuration. Continue (y/n)? y
The switch type does not match the configured switch type for this memberId.
```

Keep in mind that, due to the nature of VSF utilizing front panel ports for inter-switch links and that the number and types of ports differ between switch models, when replacing a switch with a different model, remember to update VSF port assignments when configuring the replacement (e.g., when replacing a 24-port switch with a 48-port switch and 10G SFP+ ports are being used for VSF links, use port numbers from 49-52 instead of 25-28).

Once the switch has been shut down, disconnected, or removed from the VSF configuration (if required), use the following command to replace it in the VSF configuration and re-add VSF link port assignments (replace the type and MAC address with your desired values):

```
switch(config)# vsf member 6 type j1256a mac-address e0071b-000001
switch(config)# vsf member 6 link 1 6/49,6/50
switch(config)# vsf member 6 link 2 6/51,6/52
```

If the original member was removed from the VSF configuration, you will need to re-apply any member-specific configuration to the replacement member.

If the switch being replaced is still operational, and is intended to be reused as part of another VSF fabric or in standalone mode, be sure to disable VSF to ensure that the previous VSF configuration and state are removed; use this command on the former member after all links have been disconnected:

```
switch(config)# vsf disable
This will remove VSF-related configuration and reboot the switch.
Continue (y/n)? y
```

VSF Fast Software Upgrade (5400R only)

The 5400R, when used in a VSF pair, has the capability to perform software upgrades using a sequenced reboot method, which minimizes downtime when used in conjunction with redundant uplinks and downlinks.

The process of loading new software images on a 5400R VSF pair is identical to any other standalone or stacked switch; use the following command, replacing server IP, username/password, and software image filename as appropriate:

```
switch# copy sftp flash sftpuser@10.10.10.1 KB_16_07_0002.swi primary
```

Once the image has been loaded and automatically synchronized between Commander and Standby, use the following command to initiate a sequenced reboot:

```
switch# vsf sequenced-reboot primary
The standby will reboot from the primary image and will become commander.
Do you want to continue [y/n]? y
```

When executed, the Standby will reboot into the new software first, while the Commander remains operational and forwarding traffic normally. Once the Standby has completed the boot process, it will assume the Commander role and trigger a reboot of the previous Commander, which will assume the Standby role once it completes the boot process.

For dual-homed devices utilizing link aggregation on both VSF members, this should result in minimal packet loss as traffic fails over to the active link, with expected interruption from 2-3 seconds for static routed traffic to up to 90 seconds for flows utilizing IGMP. For single-homed devices, traffic interruption is expected for 2-3 minutes, the time required for the member to reboot and ports to come back up.

Software upgrades (2930F and 5400R)

The process for upgrading a 2930F VSF fabric, as well as a 5400R VSF pair when a sequenced reboot is deemed unnecessary, is identical to that of any other stacked or standalone switch. Use the following command to load the new software version on the fabric:

```
switch# copy sftp flash sftpuser@10.10.10.1 WC_16_07_0002.swi primary
```

Once the image has been loaded and automatically synchronized to all VSF members, reboot the fabric into the new software version using this command:

```
switch# boot system flash primary
```

The entire VSF fabric will reboot into the new software version; once the boot process is complete, a Commander election will occur.

APPENDIX: TIMINGS

Values are in minutes and seconds; all timings are approximate and may vary with software version, features configured, hardware variations, etc.

3810M backplane stack (4 members)

Operation	Estimated time to complete
Cold boot (all members)	3:10
Member addition	1:50 to 11:30, depending on required software and firmware updates

2930M backplane stack (4 members)

Operation	Estimated time to complete
Cold boot (all members)	3:00
Member addition	1:50 to 11:30, depending on required software and firmware updates

5400R VSF (2 members)

Operation	Estimated time to complete
Cold boot	2:05
Failover with Standby re-sync	2:30
Member addition	4:45
Sequenced reboot	2:00 to 3:00 per member, with brief traffic interruption on aggregated links when second member reboots

2930F VSF (4 members)

Operation	Estimated time to complete
Cold boot	1:35
Member addition	3:10 to 5:00, depending on required software and firmware updates

SUPPORTED PLATFORMS

Backplane stacking is supported on the following switch platforms:

- Aruba 2930M Switch Series
- Aruba 3810M Switch Series

Virtual Switching Framework (VSF) is supported on the following switch platforms:

- Aruba 2930F Switch Series
- Aruba 5400R Switch Series