

Virtual Chassis Stacking (VCStack™)

Feature Overview and Configuration Guide

Introduction

Virtual Chassis Stacking—VCStack™—is the name given to two or more Allied Telesis switches that are configured to operate as a single switch. From a configuration and management point of view, it is as though the switches are just one device with a seamless transition from the ports of one stack member to the ports of the next.

When configuring a VCStack, there are no limitations on how the ports of one stack member can interact with the ports of another stack member—they can belong to VLANs together, they can belong to port aggregations together, they can mirror to each other, and port-range configuration can span multiple stack members. The stack member ports truly operate as though they all belong to one virtual switch.

The same applies with Layer 2 and Layer 3 switching (both unicast and multicast). The stack operates as a single device and is not perceived by end users, or the traffic itself, to be any more than a single network node.

There are some limitations to the seamlessness of virtual chassis stacking; for example, the file systems on the individual stack members remain discrete.

This guide explains the physical creation of a VCStack, the configuration required on stack members, and how to monitor the operation of the VCStack. It also provides an understanding of how the stack behaves when a stack member stops responding.

Products and software version that apply to this guide

This guide applies to AlliedWare Plus™ products that support VCStack, running version **5.4.7** or later.

To see whether your product supports VCStack, see the following documents:

- the [product's Datasheet](#)
- the product's [Command Reference](#)

These documents are available from the above links on our website at alliedtelesis.com.

Feature support may change in later software versions. For the latest information, see the above documents.

Content

Introduction.....	1
Products and software version that apply to this guide.....	2
Benefits of Virtual Chassis Stacking	4
High Availability	4
Simplified network management.....	5
The Forms of Virtual Chassis Stacking	6
Stacking Capability for Each Product Series	7
Connecting Switches into a Stack	8
Front-port stacking on DC2552XS/L3 switches.....	8
Back-port stacking on SwitchBlade x908 switches.....	9
Front-port and back-port stacking on x930 Series switches	10
Front-port stacking using XEM-STKs on x900 Series switches.....	12
Back-port stacking on x610 Series switches	14
AT-StackXG slide-in modules on x600 Series switches	16
Stacking on the x510 Series switches.....	17
Front-port stacking on CentreCOM XS900MX Series switches.....	19
Front-port stacking on CentreCOM GS900MX/MPX Series switches	20
Front-port stacking on CentreCOM FS980M Series switches	21
Front-port stacking on the x310 switches.....	22
General Stacking Restrictions by Product.....	23
How the Stack Communicates	26
The Roles of each Switch in a Stack	27
Selecting the active master	27
Identifying each Switch with Stack IDs.....	28
Displaying the stack IDs	29
Assigning stack IDs	30
Caution with stack ID renumbering	33
Steps to Set up a VCStack	33
Steps to Replace a Stack Member	37
Provisioning.....	38
Provisioning a bay	39
Provisioning a switch.....	42
Re-provisioning	44
Configuring the Stack	44
Port numbering.....	45
VLAN and IP subnet restrictions.....	45
Quality of Service (QoS) restriction.....	46
Stacking triggers	46

Licensing	46
Software and Configuration File Synchronisation	47
Software release auto-synchronisation	47
Shared running configuration	48
Shared startup configuration	49
Scripts.....	49
Rolling Reboot.....	50
Failover and Recovery.....	51
The resiliency link feature	51
Virtual MAC.....	53
Repairing a broken stack.....	54
Disabled master.....	54
Disabled Master Monitoring (DMM).....	54
Replacing a stack member	56
Executing Commands and Managing Files on Stack Members.....	56
Executing commands.....	56
Managing files.....	58
Remote login.....	59
VCStack Plus – Stacking on the SBx8100.....	62
Introduction.....	62
Enabling stacking	62
CFC operation	65
LIF operation.....	65
File synchronization	65
Stack failover	66
CFC failover	66
Chassis failover.....	66
Troubleshooting	67
Monitoring and Troubleshooting	68
Checking stack status	68
Stack debug output.....	73
Counters	79
Converting Stacking Ports to Network Ports.....	83
x610 Series: Reconfiguring AT-x6EM/XS2 stacking module ports as network ports	83
x930 Series: Reconfiguring AT-StackQS stacking module ports as network ports...	83
DC2552XS/L3 Series: Reconfiguring the front QSFP+ ports as network ports	84

Benefits of Virtual Chassis Stacking

VCStack provides a highly available solution for uninterrupted network access, simplified network management options and a high degree of flexibility.

This section is an overview of the beneficial features of VCStacking.

High Availability

Link aggregation across stack members

Aggregated links from access switches to the VCStack can terminate on different stack members. If a link in the aggregation is removed or fails, there is little network disruption.

Rolling reboot

A major benefit of the VCStack solution is that it provides unit resiliency - if one unit in the stack goes down, the other members of the stack continue to forward data. It is desirable for this continuity of service to persist even when the stack is being rebooted. Rolling reboot maintains continuous service by rebooting the stack in a rolling sequence, so that there is at least one unit actively forwarding data at all times during the stack reboot sequence.

Resiliency link

The dedicated stacking link is backed up by a further resiliency link. If the stacking link fails, communications between the stack members is maintained to enable graceful reconfiguration, and to avoid the problems caused by 'split brain'.

Fast failover

In a VCStack environment, one of the stack members acts as the master switch, and provides decision making for the virtual chassis. All of the other VCStack members are in active standby, also having learnt routing and forwarding information for the network to ensure that if the master were to fail, another member is able to seamlessly assume control of the virtual chassis with absolutely minimal network downtime. Synchronization of hardware forwarding tables, VLAN state tables, and port state tables is maintained across stack members for the following protocols:

- Spanning tree
- EPSR
- 802.1x
- ARP
- IPv6 Neighbor Discovery
- RIP
- RIPng
- OSPFv2
- OSPFv3
- BGP
- BGP4+
- IGMP
- MLD
- PIM
- PIMv6

Coupled with link aggregation across stack members, to ensure rapid recovery of Layer 2 forwarding, this full synchronization of IPv4 and IPv6 unicast and multicast forwarding tables ensures almost no packet loss upon failover.

Moreover, non-stop forwarding techniques like Graceful restart and protocol packet replication are used to ensure that the software routing processes return rapidly to an operational state, ready to handle neighbors' updates, after a failover.

Additionally, state information for other features is shared across stack members, to enable a seamless transition upon failover.

This applies to the following features:

- AMF
- DHCP
- DHCP Snooping
- LLDP
- VRRP
- Power over Ethernet

Long distance stacking

Most of the Allied Telesis range of stackable devices support long-distance stacking, so that stack members can be kilometers apart. This provides the location resiliency required for effective disaster recovery. Long distance stacking provides a genuine distributed virtual network core. The complete distributed virtual chassis provides a solution with no single point of failure, and a single management entity.

Virtual MAC

When a VCStack is central to network design, this virtual chassis uses a virtual MAC address for communication with other devices. As this single virtual MAC address is used for the complete VCStack, there is no change of MAC address if a new stack member is required to become master. Therefore, there is no need for other devices in the network to learn a new MAC address into their MAC or ARP tables.

Simplified network management

Configuration synchronization

If the startup-config on the master switch is updated, the new startup-config is automatically saved to Flash memory on all stack members. Similarly, it is automatically copied to any unit that subsequently joins the stack. This ensures seamless master failover and zero-touch stack unit replacement.

Provisioning for pre-configuration of network devices

Provisioning provides the ability to pre-configure the switch ports of devices that are not currently physically present. This allows a network administrator to configure the ports of an additional VCStack member before it is actually added to the stack. On the physical addition of the unit, the configuration is automatically applied, minimizing network disruption. A VCStack will also retain interface configuration for a device that is removed, facilitating effortless replacement of units.

License distribution

When a feature license is installed on the master switch, it will be copied across to the other stack members.

Management as a single unit

Management of a VCStack is simplified for network administrators, as the stack acts as a single virtual chassis, and can be managed as though it were a single unit.

Occasionally however, it can be desirable to login to a specific member of the stack. Remote login facilitates this by allowing a user on the master switch to log into the CLI of another stack member. Consequently, the management interface into the VCStack provides the best of both worlds - the VCStack can be managed as a single unit, or as individual units, depending on which is more convenient for the task at hand.

Extensive statistics

Extensive statistics available from a VCStack virtual chassis provide a wealth of information about data throughput on a per-port, per-resource, or traffic type basis.

The Forms of Virtual Chassis Stacking

There are three forms that Virtual Chassis Stacking can take. The three operate internally in essentially the same way, and have essentially the same user management interface. However, each has a specific physical form.


1. Standard Virtual Chassis Stacking – **VCStack** (sometimes referred to as **local** stacking)

The basic form of virtual chassis stacking consists of a set of Pizza Boxes or mini-chassis stacked together in close physical proximity (typically less than a metre apart). VCStack provides a highly available system where network resources are spread out across stacked units, reducing the impact if one of the units fails. Aggregating switch ports on different units across the stack provides excellent network resiliency.

2. Long Distance Virtual Chassis Stacking – **LD-VCStack**

Long-distance stacking allows a VCStack to be created over longer distances, perfect for a distributed network environment. The increased distance provided by fiber stacking connectivity means that members of the virtual chassis do not have to be collocated, but can be kilometres apart. All of the benefits and powerful features of VCStack remain exactly the same – Allied Telesis long distance stacking provides a genuine distributed virtual network core.

There is no restriction on the distance over which long distance stacking can operate. The only limitation on the distance between members of a long distance stack is the distance over which the SFP+ modules can operate.



The network diagram above shows a campus where the VCStack network core is distributed across two separate buildings. By having two stack members in each location, the benefits of using link aggregation between the core and edge switches remain. The complete distributed virtual chassis provides a solution with no single point of failure, and a single management entity.

3. Virtual Chassis Stacking Plus – **VCStack Plus**

The very high level of resiliency provided by stacking together two SBx8100 chassis is termed VCStack Plus. In this form of stacking, there is component redundancy within each chassis, as well as resiliency between the chassis.

Stacking Capability for Each Product Series

VCStack is implemented on a number of different series of Allied Telesis switch products. The implementations on different product series have different capabilities. The table below provides a summary of the capabilities available on each product series:

PRODUCT SERIES	MAX UNITS IN STACK	LONG DISTANCE SUPPORT	STACKING BANDWIDTH
DC2552XS/L3	2	Yes	160Gbps
SBx8100	2	Yes	160Gbps
SBx908	2	No	160Gbps
x930	8	Yes	160Gbps or 40Gbps
x900	2	No	60Gbps
x610	8	Yes	48Gbps
x600	4	No	48Gbps
x510	4	Yes	40Gbps
x310	4	No	1Gbps
CentreCOM GS900MX/MPX	4	No	40Gbps
CentreCOM XS900MX	2	No	40Gbps
CentreCOM FS980M	4	No	4 Gbps

Connecting Switches into a Stack

Different x-Series switches use different cables and connectors to connect to a stack. The types of cables and connections available are dependent on the type of x-Series switches you are stacking.

The stacking options are:

- ["Front-port stacking on DC2552XS/L3 switches" on page 9](#)
- ["Back-port stacking on SwitchBlade x908 switches" on page 10](#)
- ["Front-port and back-port stacking on x930 Series switches" on page 11](#)
- ["Front-port stacking using XEM-STKs on x900 Series switches" on page 13](#)
- ["Back-port stacking on x610 Series switches" on page 15](#)
- ["AT-StackXG slide-in modules on x600 Series switches" on page 17](#)
- ["Stacking on the x510 Series switches" on page 18](#)
- ["Front-port stacking on CentreCOM XS900MX Series switches" on page 20](#)
- ["Front-port stacking on CentreCOM GS900MX/MPX Series switches" on page 21](#)
- ["Front-port stacking on the x310 switches" on page 23](#)
- ["VCStack Plus – Stacking on the SBx8100" on page 63](#)

Front-port stacking on DC2552XS/L3 switches


The DC255XS/L3 Series switch provides four QSFP+ 40G ports which are **stacking** ports by default. For stacking purposes, the DC2552XS/L3 Series switch allows two units to be stacked using front-port stacking. There are two stacking options available:

1. Use all four QSFP+ ports, providing 320Gbps connectivity.
2. Use two QSFP+ ports, to provide 160Gbps connectivity and still provide redundancy.

Note: When VCStack is configured on the DC255XS/L3 Series switch, **all** four QSFP+ ports become **stacking** ports, and none are available for network connections, even if only one or two ports are used for stacking.

Unlike the SwitchBlade and x-Series switches, the stacking cables on the DC2552XS/L3 must connect ports of the same number. A stack will form regardless of which port numbers are used for linking, as long as the same numbers are linked together.


For example, connecting port 1.0.57 on switch 1 to port 2.0.57 on switch 2 and port 1.0.61 on switch 1 to port 2.0.61 on switch 2, as shown below, is a valid option:




Whereas, connecting port 1.0.57 on switch 1 to port 2.0.61 on switch 2 and port 1.0.57 on switch 1 to port 2.0.61 on switch 2, would not work.

Back-port stacking on SwitchBlade x908 switches

On the rear (lower center) of the SwitchBlade x908 chassis, there is a pair of fixed stacking ports.



Back port stacking requires a specific cable type. The product code is AT-HS-STK-CBL650.



You can stack two SwitchBlade x908 switches together using these ports and cables.




Note that the cables are crossed over—port 1 of the top switch is connected to port 2 of the bottom switch, and vice versa.

Front-port and back-port stacking on x930 Series switches

The x930 Series allow up to 8 units to be stacked using either front or back port stacking. You cannot combine front and back port stacking.

Front-port stacking

Front-port stacking on the x930 Series provides 40Gbps connectivity. The x930 Series have 4 x 10GbE SFP+ ports, two of which may be used for stacking instead of network connectivity. The stacking cables must form a ring, as shown in the diagram below.



If the switch does not have a StackQS module installed, then the front-panel ports S1 and S2 will be set up as the stacking ports by default when VCStacking is enabled by the **stack enable** command. If a StackQS module is installed, and you still want to use the

front ports for stacking, then you need to explicitly set the front ports as the stacking ports by using the command **stack enable builtin-ports**.

If the S1 and S2 ports are not configured as stacking ports, then they operate as normal network ports. The numbers assigned to these ports are simply the sequential numbers beyond those of the other ports on the switch, i.e. ports 1.0.27 and 1.0.28 on a 28-port switch and ports 1.0.51 and 1.0.52 on a 52-port switch.

Stacking cables


Cables that can be used for x930 Series front-port stacking:

- AT-SP10TW1—1 meter SFP+ direct attach cable
- AT-SP10TW3—3 meter SFP+ direct attach cable
- AT-SP10TW7—7 meter SFP+ direct attach cable

In addition, Allied Telesis SFP+ modules can be inserted and connected by cables of whatever length the SFP+ modules can support.


Rear-port stacking

Rear-port stacking on the x930 Series provides 160Gbps connectivity using the StackQS module and direct attached cables. The stacking cables must form a ring, in the same manner as when using front-port stacking.



Stacking modules and cables

- AT-StackQS: 2 x QSFP+ stacking module
- AT-QSFP1CU: 1 meter QSFP+ direct attach stacking cable
- AT-QSFPSR Transceiver: 150 meter, 12-strand OM4 fiber optic cable
- AT-QSFPLR4 Transceiver (has a duplex LC connector): 2 meter to 10 km single mode fiber (SMF).




By default, if a StackQS module is installed in the switch, then its ports will be chosen as the stacking ports if stacking is enabled.

If the configuration has been changed so that the front-panel ports are set up as the stacking ports, the role of stacking ports can be changed back to the StackQS ports by using the command **stack enable expansion-ports**.

If the ports on the AT-StackQS card are not configured as stacking ports, then they operate as 40Gbps network switch ports. The ports are numbered port1.1.1 and port1.1.5.

Front-port stacking using XEM-STKs on x900 Series switches

You can fit the XEM bays on x900 Series switches with a specialized stacking XEM called the XEM-STK.




The LED number on the front of the XEM-STK shows the stack ID of the switch that the XEM-STK is installed in. See the section ["Identifying each Switch with Stack IDs"](#) on [page 29](#) for more information about stack IDs.

The specific cable type that connects these XEMs are purchased individually as either 350mm or 2 meter long cables. The product codes are:

- AT-XEM-STK-CBL350
- AT-XEM-STK-CBL2.0


With these XEMs and cables, you can create a stack of two x900 Series switches.



Connecting the cables to the switches

Stacked switches are connected in a ring, with port 1 of one switch connected to port 2 of the next switch. The last switch in the stack then connects to the first switch using port 2 of the first switch and port 1 of the last switch.

In two switch stacks, this means that the connection consists of just a crossed pair of cables. You may also use only one cable to connect the switches, but you will halve the bandwidth between stack members.



Back-port stacking on x610 Series switches

With a choice of 24-port and 48-port versions and optional 10 Gigabit uplink units, the x610 Series can connect anything from a small workgroup to a large business. A fully resilient solution is created where up to eight units can form a single virtual chassis, with dual connections to key servers and access switches. Expansion modules are available for local and long-distance stacking or can be configured to provide two additional 10G ports.

The stacking cables must form a ring, as shown in the diagram below.



Expansion modules

The x610 Series switches use the following two expansion modules:

- AT-x6EM/XS2-00 —Expansion module (2 x SFP+) for long-distance stacking or for use as two additional 10GbE ports.



If the AT-x6EM/XS2 module is installed in the switch, and stacking is enabled, with the command **stack enable** (enabled by default), then the two ports on the AT-x6EM/XS2 module will operate as stacking ports.

If stacking is disabled, using the command **no stack enable**, (followed by saving the configuration and rebooting the switch), then the two ports on the AT-x6EM/XS2 module are re-purposed as normal 10Gbps network ports. The port numbers assigned to the ports are 1.1.1 and 1.1.2.

- AT-StackXG-00—Expansion module with two full-duplex, 12 Gbps stacking ports, one AT-StackXG/0.5-00 cable included. This is exactly the same module as is used in x600 Series switches.



Stacking cables

A variety of stacking cables are available for use with the x610 Series switches.

Cables that can be used with AT-StackXG-00:

- AT-StackXG/0.5-00—0.5 meter cable for stacking
- AT-StackXG/1-00—1 meter cable for stacking

Cables that can be used with AT-x6EM/XS2-00

- AT-SP10TW1—1 meter SFP+ direct attach cable
- AT-SP10TW3—3 meter SFP+ direct attach cable
- AT-SP10TW7—7 meter SFP+ direct attach cable

In addition, standard SFP+ units can be inserted into the x6EM/XS2-00, and connected by cables of whatever length the SFP+ modules can support.

AT-StackXG slide-in modules on x600 Series switches

On the rear of the x600 chassis you can insert a slide-in module, the AT-StackXG-00. This connector can also be used in x610 Series switches, see ["Expansion modules" on page 15](#).




The specific cable type that connects the AT-StackXG are purchased as either 0.5 or 1 meter long cables. The product codes are:

- AT-STACKXG/0.5-00—0.5 meter cable for stacking
- AT-STACKXG/1-00—1 meter cable for stacking

You can stack up to four x600 switches using the AT-StackXG cables.



Like the other stacking methods, the connections are crossed-over—port 1 on one switch is connected to port 2 on its neighbor and the switches are connected in a ring. The following figure shows how to connect a 4 switch stack of x600 Series switches. Once again, you could connect the switches without one of the cables, but you would halve the bandwidth between stack members.



Stacking on the x510 Series switches

This section specifies which stacking components to use with x510 Series switches to correctly form Virtual Chassis Stacking. The x510 Series Switches come with two stacking ports. These are the last two SFP+ slots on the switches and are labeled **S1/27** and **S2/28** on the 28-port switches.

The ports have two functions. You may use them as stacking ports when VCStack is enabled by the **stack enable** command (the default state).

Or on the other hand, by disabling the VCStack feature, using the **no stack enable** command, you may also use them with regular SFP or SFP+ transceivers as additional networking ports.

The names of the ports depend on the status of the VCStack feature on the switch. They are presented by the software as S1 and S2 when the VCStack feature is enabled. When you disable the VCStack feature to use the ports with regular SFP or SFP+ transceivers on a stand-alone switch, they are referred to as port1.0.27 and port1.0.28 on the 28-port switches and port1.0.51 and port1.0.52 on the 52-port switches.

Local stacking using direct attach cables

Requirements:


- x510 switches running AlliedWare Plus 5.4.3 GA or later.

For switches running AlliedWare Plus **prior** to 5.4.5, use AT-StackXS/1.0 (1m) stacking cables (one per switch). A StackXS cable is similar to a TwinAx direct attach cable but with a stacking connector at each end of the cable. It can join two switches together; therefore only one cable is required per switch.

- For switches running AlliedWare Plus software version 5.4.5 or later, any Allied Telesis TwinAx cable will work.

When creating a VCStack of units that are all located within **a meter** of each other, use local stacking components. [Figure 1](#) below shows three direct attach cables forming a VCStack of three switches.

Figure 1: Direct attach cables



Long-distance stacking using SFP+ modules

Requirements:


- x510 switches running AlliedWare Plus 5.4.3GA or later.
- On switches running AlliedWare Plus **prior** to 5.4.5, use AT-StackOP/0.3 (300m) or AT-StackOP/9.0 (9km) optical stacking modules (two per switch). A StackOP module is similar to an SFP+ module, therefore two modules are required per switch.
- On switches running AlliedWare Plus 5.4.5 or later, you can use any Allied Telesis SFP+ module.
- Fiber cable.

When creating a stack where switches are more than one meter apart, use long-distance stacking components.

If [Figure 1 on page 18](#) was configured using long-distance stacking components, then six SFP+ modules and three fiber cables would be required.

Note that a mix of local and long-distance stacking components can be used. This is useful if two groups of devices are separated by some distance. [Figure 2](#) below shows this scenario.

Figure 2: Stacking using SFP+ modules



Front-port stacking on CentreCOM XS900MX Series switches

For software version 5.4.6, the CentreCOM XS900MX Series switches enable you to stack two units together through ports S1 and S2 using SP10TW1-1 meter SFP+ direct attach cables. Connect S1 on one switch to S2 on the other.

The CentreCOM XS900MX Series consists of the following switches: XS916MXS and XS916MXT.

From software version **5.4.7** onwards, the CentreCOM XS900MX Series switches allow two units to be stacked using any of their ports. As the units house a mix of copper and fiber ports, this provides flexibility in which ports are used for stacking, and which remain available for network connectivity.

If stacking is enabled, all ports except ports 15 and 16 are network ports by default, and ports 15 and 16 are stacking ports by default. If stacking is disabled, all ports are network ports.


When changing a port to or from a stacking port, you must first configure the desired ports as stacking ports on each switch, using the **stackport** command. Then save the configuration file on each switch, restart both switches, and cable the switches together.

Stacking cables

Use one of the following cables to stack CentreCOM XS900MX Series switches. To connect:

- SFP+ ports, use SP10TW1-1 meter SFP+ direct attach cables.
- RJ-45 ports, Cat 6a or above cables

The stacking cables must form a ring, for example as shown in the diagram below, which shows stacking through RJ-45 ports 1 and 2:



Front-port stacking on CentreCOM GS900MX/MPX Series switches

The CentreCOM GS900MX/MPX Series switches allow 4 units to be stacked using front-port stacking, providing 40Gbps bandwidth. The GS900MX/MPX Series have four 10GbE SFP+ ports, two of which may be used for stacking instead of network connectivity.


The CentreCOM GS900MX/MPX Series consists of the following switches: GS924MX, GS924MPX, GS948MX, GS948MPX.

Stacking cables

Use the following cable to stack GS900MX/MPX Series switches:

- AT-SP10TW1-1 meter SFP+ direct attach cable

The stacking cables must form a ring, as shown in the diagram below:



Front-port stacking on CentreCOM FS980M Series switches

VCStacking is supported on the 28 and 52 port models (FS980M/28, FS980M/28PS, FS980M/52, FS980M/52PS).


From software release **5.4.7** onwards, the CentreCOM FS980M Series switches allow up to 4 units to be stacked using front-port stacking, providing 4Gbps bandwidth. The CentreCOM FS980M Series have four 1GbE SFP ports, two of which may be used for stacking instead of network connectivity.

Stacking cables

Use the following cable to stack CentreCOM FS980M Series switches:


- AT-SP10TW1-1 meter SFP direct attach cable

The stacking cables must form a ring. For example, in a stack of two FS980M/28, stacking port1.0.28 needs to be connected to port2.0.27 and port1.0.27 needs to be connected to port2.0.28, as shown in the diagram below:




For software version **5.4.7** onwards, up to four units may be used to form a VCStack. For example, in a stack of four FS980M/28, the cables are connected as follows:

- port1.0.28 <--> port2.0.27
- port2.0.28 <--> port3.0.27
- port3.0.28 <--> port4.0.27
- port4.0.28 <--> port1.0.27



Front-port stacking on the x310 switches

X310 switches can be connected in stacks of up to 4 units. The stacking ports are the S1 and S2 ports on the front of the unit. The stacking connections go from the S1 port of one switch to the S2 port of the next switch, and can be formed into a full ring.



The stacking ports S1 and S2 can be used as network ports, when stacking has been disabled by the **no stack enable** command, and the switch has been rebooted.

The port numbers allocated to the stacking ports when stacking has been disabled are the sequential numbers beyond the end of the rest of the port numbers, i.e. port1.0.27 and port1.0.28 on the 28-port switches, and port1.0.51 and port1.0.52 on the 52-port switches.

General Stacking Restrictions by Product

There are some restrictions to what products and connections you can combine in a single stack. In general:

- some switch families cannot be stacked together and you cannot combine different stacking methods or cables.

DC2552XS/L3

compatible stack partners	DC2552XS/L3 switches can only stack with other DC2552XS/L3 switches.
maximum stack size	2 switches.
cable type	QSFP+ Stack cables: <ul style="list-style-type: none"> ■ AT-QSFP1, 3 CU - QSFP+ copper cable 1m and 3m ■ AT-MTP12-1, 5 - MTP optical cable for AT-QSFPSR, 1m and 5m

SwitchBlade x8100 Series

compatible stack partners	A SwitchBlade x8100 switch can only stack with another SwitchBlade x8100 switch.
maximum stack size	2 switches.
cable type	Fibre: <ul style="list-style-type: none"> ■ Allied Telesis SFP+ module (two modules required per stacking link).

SwitchBlade x908

compatible stack partners	A SwitchBlade x908 switch can only stack with another SwitchBlade x908 switch.
maximum stack size	2 switches.
cable type	<ul style="list-style-type: none"> ■ For back-port stacking: AT-HS-STK-CBL650

x930 Series

compatible stack partners	x930 Series switches can only stack with other x930 switches. Within the x930 Series: You can create a VCStack of up to eight units: <ul style="list-style-type: none"> ■ Front port stacking (40G) - long distance with any SFP+ module. ■ Rear port stacking (160G) - StackQS module, 1 meter cables
maximum stack size	8 switches.
cable type	<ul style="list-style-type: none"> ■ AT-QSFP1CU (use with AT-StackQS module) - 1 meter QSFP+ direct attach stacking cable ■ AT-SP10TW1, 3, 7 - 1, 3, or 7 meter TwinAx direct attach cable ■ You can use any Allied Telesis SFP+ module for front port stacking, to get Long Distance Stacking (LD-VCStack)

x900 Series	compatible stack partners	x900 Series switches can only stack with other x900 switches. Within the x900 Series: <ul style="list-style-type: none"> ■ x900-12XT/S switches can only stack with x900-12XT/S switches ■ x900-24XS and x900-24XT can stack together
	maximum stack size	2 switches.
	cable type	<ul style="list-style-type: none"> ■ AT-XEM-STK-CBL350 (350mm length) or AT-XEM-STK-CBL2.0 (2 meter length).
x610 Series	compatible stack partners	x610 Series switches can stack with x600 switches, but the x600 stack restrictions will apply, so the whole stack is at the level of the x600. Within the x610 Series, you can stack any model with another model. This means that all the switches in the x610 Series can stack together without restriction.
	maximum stack size	8 switches.
	cable type	<ul style="list-style-type: none"> ■ AT-StackXG/0.5-00—0.5 meter cable for stacking ■ AT-StackXG/1-00—1 meter cable for stacking ■ AT-SP10TW1—1 meter SFP+ direct attach cable ■ AT-SP10TW3—3 meter SFP+ direct attach cable ■ AT-SP10TW7—7 meter SFP+ direct attach cable ■ Allied Telesis SFP+ module with AT-X6EM/XS2 for long distance stacking
x600 Series	compatible stack partners	x600 Series switches can stack with x610 switches, but the x600 stack restrictions will apply, so the whole stack is at the feature level of the x600. Within the x600 Series, you can stack any model with another model. This means that all the switches in the x600 Series can stack together without restriction.
	maximum stack size	4 switches.
	cable type	<ul style="list-style-type: none"> ■ AT-STACKXG/0.5 (0.5 meter length) or ■ AT-STACKXG/1 (1 meter length).

x510 Series **compatible stack partners** x510 Series switches can only stack with other x510 switches.

maximum stack size 4 switches.

cable type Standard Allied Telesis branded TwinAx direct attach cables or SFP+ modules can be used to stack x510 switches.

For **local** stacking:

- For switches running AlliedWare Plus **prior** to 5.4.5 - AT-StackXS/1.0 —1 meter stacking cable
- For switches running AlliedWare Plus 5.4.5 or later:
 - AT-SP10TW1—1 meter SFP+ direct attach cable
 - AT-SP10TW3—3 meter SFP+ direct attach cable
 - AT-SP10TW7—7 meter SFP+ direct attach cable

For **long distance** stacking:

- Fiber cable
 - For switches running AlliedWare Plus **prior** to 5.4.5 - AT-StackOP/0.3 (300m) or AT-StackOP/9.0 (9km) optical stacking modules (two per switch).
 - For switches running AlliedWare Plus 5.4.5 or **later**, you can use any Allied Telesis SFP+ module.
-

x310 Series **compatible stack partners** x310 Series switches can only stack with other x310 switches, (currently available in 10-100 switches only).

maximum stack size 4 switches.

cable type Standard Allied Telesis branded TwinAx direct attach cables can be used to stack x310 switches.

For local stacking:

- AT-StackXS/1.0 —1 meter stacking cable.
-

CentreCOM XS900MX Series **compatible stack partners** The CentreCOM XS900MX Series switches can only stack with other CentreCOM XS900MX Series switches.

maximum stack size 2 switches.

cable type Standard Allied Telesis branded TwinAx direct attach cables can be used to stack CentreCOM XS900MX Series switches.

For local stacking, use one of the following cables:

- AT-SP10TW1-1 meter SFP+ direct attach cable.
 - RJ-45 cable - copper cable Cat 6a or above
-

**CentreCOM
GS900MX/MPX
Series**

compatible stack partners	CentreCOM GS900MX/MPX Series switches can only stack with other CentreCOM GS900MX/MPX Series switches.
maximum stack size	4 switches.
cable type	Standard Allied Telesis branded TwinAx direct attach cables can be used to stack CentreCOM GS900MX/MPX Series switches. For local stacking: <ul style="list-style-type: none"> ■ AT-SP10TW1—1 meter SFP+ direct attach cable.

**CentreCOM
FS980M Series**

compatible stack partners	CentreCOM FS980M Series switches can only stack with other CentreCOM FS980M Series switches.
maximum stack size	4 switches.
cable type	Standard Allied Telesis branded TwinAx direct attach cables can be used to stack CentreCOM FS980M Series switches. For local stacking: <ul style="list-style-type: none"> ■ AT-SP10TW1—1 meter SFP+ direct attach cable.

How the Stack Communicates

The stack communicates through the stacking cables. You can also configure a resiliency link between stack members, which is used when a failover event occurs on the stack. See the section "[Failover and Recovery](#)" on [page 52](#) for more information.

One switch controls all switch management on a stack. Which stack member does this is discussed in "[The Roles of each Switch in a Stack](#)" on [page 28](#). The software version, startup configuration, and running configuration are exactly the same on each member of a stack. For information on how the stack synchronizes the files, see the section "[Software and Configuration File Synchronisation](#)" on [page 48](#).

The internal communication between stack members is carried out using IP packets sent over the stacking links. This stack management traffic is tagged with a specific VLAN ID and uses IP addresses in a specified subnet.

The default is:

- VLAN 4094
- Subnet 192.168.255.0/28

The management traffic is queued to egress queue 7 on the stack link. The section "[Configuring the Stack](#)" on [page 45](#) discusses the minor restrictions necessary when configuring VLANs, IP subnets, and QoS on the stack.

The Roles of each Switch in a Stack

Each switch in a stack acts in one role at any time. This is either as a **backup member** (also called **stack member**) or a **stack master** (normally as the **active master**). The stack members are controlled by the stack master. All configuration, including updating the switch software, is set on the stack members by the stack master.

The stack master performs a number of tasks that a stack member does **not** perform:

- It controls all switch management activity.
- It synchronises boot release and configuration files with stack members.
- All routing protocol packets are processed by the stack master. The stack master then transfers any requisite table updates to the stack members.

The stack master also handles communication on behalf of the stack.

- When you Telnet or SSH to the stack, you connect to a process running on the stack master.
- When you connect to the asyn port on a stack member, this automatically creates an SSH session to the master. This connection will behave as if you were connected to the asyn port on the master.
- The stack master handles SNMP interactions. It gathers MIB statistics from the stack members, and delivers these statistics in response to SNMP Get requests.

You can still execute some specific commands and manage files on an individual switch in the stack. See the section "[Executing Commands and Managing Files on Stack Members](#)" on page 57 for more information.

When a VCStack is operating normally, the stack master acts as the **active master**. However during some failover and resiliency situations, when a stack member cannot contact the **active master**, it may act as either a **disabled master**, or **fallback master**. See the section "[Failover and Recovery](#)" on page 52 for more information about the differences between these stack master roles.

Selecting the active master

The stack members negotiate among themselves as to which switch will become the active master. The election of the active master is based on two criteria:

- each stack member's priority setting
- each stack member's MAC address

For each of these criteria, a lower number indicates a higher priority. The active master is the switch with the **lowest** priority value. If multiple switches share the lowest priority value, then the switch with the lowest MAC address becomes the active master.

Manually changing the active master

You can choose a specific switch to be active master by changing its priority value. By default, a switch has a stack priority value of 128. The command to change a switch's priority value for stacking is:

```
awplus(config)#stack <switch-stack-ID> priority <0-255>
```

The stack ID is a unique identifier that is assigned to each switch in the stack. See the section "[Identifying each Switch with Stack IDs](#)" on page 29 for more information.

If a switch is already part of a stack, you can still set the priority value on each switch in the stack through the active master. However, if you set the priority value on a stack member lower than the active master's priority value, the active master does not immediately relinquish its role as master. The stack member with the lowest priority will take over as active master only if the current active master leaves the stack (this includes any reboot by the stack master).

If a switch has not yet joined a stack, you can still use the **stack priority** command to change the priority value before connecting it to the stack. However, even if the new stack member has a lower priority value than the active master, it will not take over as active master unless the current active master is removed or rebooted.

Identifying each Switch with Stack IDs

Each switch in a stack has an ID number, which can be an integer between 1 and 8. The default on each switch is a stack ID of 1.

The ID number of a stack member is an important identifier. All commands that are port or switch specific need to use the stack ID to identify which stack member the commands apply to.

The stack ID is also used to manage specific stack members. When you Telnet or SSH to the stack, your login session terminates on the active master. Similarly, if you connect to the console port of any stack member, your console login session is sent through to the active master. The active master switch is the automatic point of contact for any management session on the stack. If you want to examine something that physically resides on one of the other stack members, such as files in a stack member's Flash, or voltage levels on a stack member's power supply, then you can do this by specifying the stack ID of the stack member in commands.

The stack IDs are used frequently in the stack configuration scripts. For example, any command in the configuration script that refers to a physical port will include a stack ID in the port number. The section "[Port numbering](#)" on page 46 explains the port numbering scheme used in stacks.

For these reasons, the stack IDs on each switch within a stack are unique and a switch's stack ID normally does not change without user intervention.

The only exception to this is when a new switch is connected to a stack and the switch has the same stack ID number as another stack member—the new switch's ID will be automatically renumbered.

Note: There is no connection between stack ID and active master status. The active master switch does not have to be the switch with a Stack ID of 1.

Displaying the stack IDs

To see the stack ID on a switch before it is connected to a stack, use the command:

```
awplus#show stack
```

The output will show the current stack ID and any pending change to the stack ID.

```
x900#show stack
Virtual Chassis Stacking summary information

ID Pending ID MAC address Priority Status Role
1 - 0000.cd27.c4bf 128 Ready Active master

Operational Status Standalone unit
Stack MAC address 0000.cd27.c4bf
```

You can also use this command to see the stack numbers on a two-switch stack. To match the physical switch with the right stack ID number, look for the **active master** LED on the front panel. Then use the **show stack** command to show all members of the stack. You can match the LED status to the **show** command output.

```
x900#show stack
Virtual Chassis Stacking summary information

ID Pending ID MAC address Priority Status Role
1 - 0000.cd27.c4bf 128 Ready Active master
2 - 0000.cd28.0801 128 Syncing Backup Member

Operational Status Normal operation
Stack MAC address 0000.cd27.c4bf
```

The **show stack** command is available on all stacking switches running the AlliedWare Plus operating system.

On the front of the XEM-STK there is an indicator that displays the stack ID of the switch that the XEM is installed in.

On x600 and x610 Series switches, you can use the command:

```
awplus#show stack indicator <1-8>|all [time <1-500>]
```

This causes the **master** LED on the switch to flash in a sequence which indicates the stack ID number. The pattern will be a number of flashes in quick succession followed by a longer pause; where the number of flashes equals the stack member ID.

For example, on a four-switch stack with the stack IDs 1, 2, 3, and 4, the switch with stack ID:

- 1 will flash on and off without pausing * * * * *
- 2 will flash twice then pause * * * * *
- 3 will flash three times then pause * * * * *
- 4 will flash four times then pause * * * * *

You can extend the time that the **master** LEDs will flash by up to 500 seconds to give you more time to reach the stack's location, by using the optional **time** parameter with the command. By default the LEDs will flash for 5 seconds.

For details of how this is implemented on the SBx8100, under VCStack Plus, please see ["VCStack Plus – Stacking on the SBx8100" on page 63](#)

Assigning stack IDs

Stack IDs can be assigned in a number of ways. We recommend only assigning stack IDs when you set up the stack, as a change to stack ID numbers is not automatically reflected in configuration scripts.

You can assign stack IDs to switches before they are part of a stack:

- ["Manual assignment on a switch before stacking" on page 31](#)

To assign stack IDs once the stack is created, you can use the following methods:

- ["Automatic assignment as a switch joins the stack" on page 32](#)
- ["Manual renumbering of a switch after stacking" on page 32](#)
- ["Cascade renumbering of the stack" on page 33](#)
- ["Renumbering the whole stack using the XEM Select button" on page 34](#)

Manual assignment on a switch before stacking

You can manually assign the stack IDs on switches before you stack them together.

If your switch has never had a stack ID assigned to it, it will have the stack ID of 1. You can assign it a new stack ID with the command:

```
awplus(config)#stack 1 renumber <1-8>
```

If the switch has previously been in a stack and was assigned a stack ID, it keeps that stack ID with it, even if it is removed from the stack. The stack ID has become an inherent property of the switch. If you wish to renumber it you will need to specify the current stack ID.

For example to renumber the stack ID from 3 to 5, use the command:

```
awplus(config)#stack 3 renumber 5
```

The stack ID renumbering does not take effect until the switch is rebooted, so you will receive the following warning as shown in the next figure.

```
awplus(config)#stack 1 renumber 2
Warning: the new ID will not become effective until the stack-member
reboots.
Warning: the boot configuration may now be invalid.
```

Until the reboot occurs, the new stack ID value is noted as the 'pending' stack ID, as shown in the next figure.

```
x900#show stack
Virtual Chassis Stacking summary information

ID   Pending ID   MAC address           Priority   Status   Role
1    2             0000.cd27.c4bf       128      Ready   Active master

Operational Status           Standalone unit
Stack MAC address            0000.cd27.c4bf
```

Automatic assignment as a switch joins the stack

When a stack is established, the stack will automatically assign a new stack ID to a switch if it has the same stack ID as another member. This means that you can create a stack without previously changing the stack ID numbers from the default of 1. The stack master will be assigned stack ID 1, and the other switches will be automatically assigned other IDs.

Manual renumbering of a switch after stacking

If you want to manually renumber the switches when they are already part of a stack, use the command:

```
awplus(config)#stack <1-8> renumber <1-8>
```

You can issue this command from the active master to renumber any switch in the stack. To avoid duplicate IDs, a warning message appears if you assign to a stack member the same ID that is currently assigned to another stack member. However, you can continue to renumber the stack member IDs and fix potential ID duplications. Once you have removed any duplicate IDs, you can reboot the switches with ID changes to implement your changes.

If you do not remove the duplications, then when the stack reboots, the switch with the highest stack priority (the lowest priority value) is allocated this ID, and the competing switch is automatically assigned another ID.

Cascade renumbering of the stack

For larger stacks, it is useful to have all the switches numbered in sequence, based on the order of their stack connections. The command that can achieve this is:

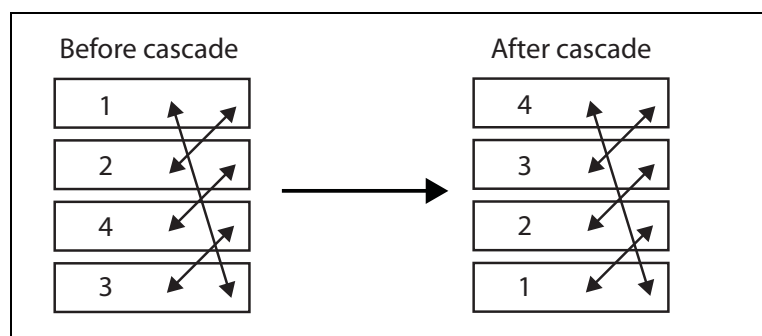
```
awplus(config)#stack <1-8> renumber cascade [<1-8>]
```

For example, you can enter the command:

```
awplus(config)#stack 3 renumber cascade 1
```

This starts the stack numbering on the switch that currently has Stack ID 3 and gives that switch Stack ID 1. The other switches in the stack are renumbered sequentially based on their connection to the switch that now has the stack ID 1.

The following diagram shows the change from this command:



If the second stack ID parameter is not supplied in the command, then the numbering begins from 1. When the numbering process hits the maximum existing stack ID value, it assigns the value 1 to the next switch in the stack.

We recommend using this command to ensure the switches are sequentially numbered, if you did not manually assign the switches with stack IDs before connecting them to the stack. Use the command straight after the stack is first connected.

Renumbering the whole stack using the XEM Select button

On switches connected with XEM-STKs, you can use the **Select** button on the front of the XEM-STK. This achieves the same effect as the **renumber cascade** command. When you press the **Select** button on a XEM-STK, the switch with that XEM-STK is assigned stack ID 1, and the other members of the stack are numbered sequentially from there.

Caution with stack ID renumbering

Stack IDs are critical in identifying each physical switch in the port numbering and configuration commands. However, the stack's configuration script is not altered when stack IDs are renumbered. This means that you may lose the connection between configuration commands and the physical switches, and commands could end up being applied to a different switch.

For example, if switch A in the stack currently has stack ID 2, then any commands in the configuration script that refer to stack ID 2 are applied to switch A. If the stack IDs are renumbered, so that stack ID 2 is now allocated to switch B, then the commands in the configuration script that refer to stack ID 2 will now be applied to switch B. If switch A is given a stack ID which is not in the configuration script, then it will not have any configuration applied. We recommend that you only renumber the stack when it is initially configured, or during a time of major stack reconfiguration.

Steps to Set up a VCStack

There are no set rules regarding the order in which stack configuration tasks need to be carried out. However, these steps provide a guideline to help ensure that the stack creation process goes smoothly.

Step 1: Prepare the switches

Before connecting any of the switches together:

- Ensure that all the switches have the same feature licences installed. If you have purchased feature licences to enable certain features to operate on the stack, then all stack members need to have the licences installed. If some stack members have feature licences installed for features that will not be used on the stack, and other switches do not have those licences installed, then remove those unnecessary licences.
- If you are using XEM stacking, you must install the XEMs into the switches before you can enter any stacking commands. An x900 Series switch will reject any stacking commands until a XEM-STK is installed.

```
x900(config)#stack 1 priority 0
% The unit has no XEM-STK installed, stacking commands are not allowed
```

- Decide which stack member will be the active master and set its priority to 0, to ensure it becomes the active master. The command is:

```
awplus(config)#stack 1 priority 0
```

(The switch, if it has not been stacked before, should have the default stack ID of 1.)

Step 2: Install and power the stack master

Install and power up the master switch. It will detect that there are no other members in the stack, so it will elect itself master. At bootup, it will output the following messages:

```
13:35:10 awplus-1 VCS[1164]: Member 1 (00-00-cd-28-07-c0) has become the
Active master
```

Then, after the switch has booted, the **show stack** command will show a stack with only one member:

```
x900#show stack
Virtual Chassis Stacking summary information

ID   Pending ID   MAC address           Priority   Status   Role
1    -             0000.cd27.c4bf       128      Ready   Active master

Operational Status           Standalone unit
Stack MAC address           0000.cd27.c4bf
```

Step 3: Install and power the backup member

Install the next switch, connecting the stacking cable from that switch to the master.

Note: Make sure the stacking cables are crossed over between the stack members. This means that stack port 1 on switch 1 should connect to stack port 2 on switch 2. If this is not done, the stack links will not come up and the stack will not form.

Power up the switch—it will detect that there is already an active master, and so will come up as a backup member. The active master will assign it the first available stack ID.

At bootup, the new stack member outputs the following messages:

```
Notice: Possible stack, please wait whilst searching for members.
12:26:57 awplus-2 VCS[1043]: Member 1 (0000.cd27.c4bf) has joined stack
12:26:57 awplus-2 VCS[1043]: Please configure 'stack virtual-mac' to
minimize network disruption from failovers
12:26:57 awplus-2 VCS[1043]: Member 1 (0000.cd27.c4bf) has become the
Active master
```

After bootup, the **show stack** command will show that there are two switches in the stack:

```
x900#show stack
Virtual Chassis Stacking summary information

ID  Pending ID  MAC address          Priority  Status  Role
1   -           0000.cd27.c4bf      128     Ready   Active master
2   -           0000.cd28.0801      128     Syncing Backup Member

Operational Status          Normal operation
Stack MAC address           0000.cd27.c4bf
```

The active master will check that the new stack member has the same software version as itself. If the software versions are different, the active master will use the software auto-synchronization mechanism to force the new stack member to run the same software version.

Step 4: Install and power the next backup member

Repeat [step 3](#) for each of the other switches in the stack, remembering to connect port 2 of each new switch to port 1 of its neighbor. For the last switch added to the stack, connect port 1 of this switch to port 2 of the first installed switch.

Step 5: Confirm that the stack is operating

Check that the stack links have all come up successfully. This can be done by:

- checking the LEDs on the switches or XEMs.
The port LEDs for all stack members should be green. Port LEDs that are off or flashing amber indicate that the stack is not operating correctly. The master or status LED will be green on the switch that is the stack master.
- using the **show stack detail** command.
This command provides a snapshot of the stack status.

See the section "[Checking stack status](#)" on [page 69](#) for more information about LEDs and the **show stack detail** command.

Step 6: Check stack numbering

If you are not satisfied with the stack IDs that the switches have been automatically assigned, then this is the right moment to remedy that. To change a stack ID, use the command:

```
awplus(config)#stack <1-8> renumber <1-8>
```

It is usually a good idea to give the stack ID 1 to the active master, as it is quite natural to associate the lowest ID with the active master switch.

To sequentially renumber all stack members, you can use the:

- **select** button on the XEM-STK

- **stack renumber cascade** command

To renumber the stack members and give stack ID 1 to the active master, use the command:

```
awplus(config)#stack <current ID on master switch> renumber cascade 1
```

Step 7: Reboot the switches

Reboot all the stack members, and check that they all come up with the stack IDs that you are expecting. You can use the command **show stack detail** to check the stack IDs and the status of the neighbor connections.

Step 8: Configure any priority changes

If there is another switch that you want to be the one that takes over as active master, if the active master goes down, then set its priority to a value lower than the other switches:

```
awplus(config)#stack <1-8> priority 10
```

Step 9: Configure the stack as one switch

You are now ready to configure the stack with port channels, VLANs, IP addresses, and so on. For example, to create a port channel that aggregates ports from two different stack members, you would configure as follows:

```
awplus(config)#configure terminal
awplus(config)#interface port1.0.1
awplus(config-if)#channel-group 1 mode active
awplus(config-if)#interface port2.0.1
awplus(config-if)#channel-group 1 mode active
```

Once you are happy with the stack configuration, make a **backup** copy of the configuration file.

Steps to Replace a Stack Member

If you need to replace a stack member, use the following steps to achieve a smooth transition:

Step 1: Configure the Stack ID on the replacement switch

Prepare the replacement switch by configuring it with the same stack ID as the switch that you are replacing.

```
awplus(config)#stack <current stack ID> renumber <1-8>
```

Step 2: Configure the feature licences

Ensure that the replacement switch is configured with the same set of feature licences as the existing stack members.

Step 3: (On an XS900MX switch) configure the stack ports

On the replacement XS900MX switch, if you are not using the default ports 15 and 16 for stacking, then configure the required ports as stack ports. For example, if ports 3 and 4 are the stack ports, use the following commands:

```
awplus(config)#interface port2.0.3-4
awplus(config-if)#stackport
```

Step 4: Remove the failed switch

Unplug the failed switch from the stack.

Step 5: Install the replacement switch

- Connect the stacking cables to the replacement switch. Power up the replacement switch. It will detect that there is already an active master, and so will come up as a stack member.
- The active master will check that the new stack member has the same software version as itself. If the software versions are different, the active master will use the software auto-synchronization mechanism to force the new stack member to run the same software version.
- The new switch will also receive the startup configuration from the active master. As the replacement switch has been configured with the same stack ID as the replaced switch, it will receive exactly the same configuration as the replaced switch, and will operate exactly as that switch had.

Provisioning

- Provisioning provides the ability to pre-configure ports that are not yet present in a switch or in a stack. If a switch can have XEMs hot-swapped into it, then provisioning allows the ports of those yet-to-be-inserted XEMs to be preconfigured ready for the arrival of the XEMs. Similarly, if you know that a unit is going to be added to a stack, you can pre-configure that new switch in anticipation of its addition to the stack.
- With provisioning you can configure stack members and their ports, even though they are not currently physically present, and configure them ready for future addition. This means that you can either pre-configure ports belonging to a bay or switch that has not yet been installed, or load a configuration that references these ports.
- Provisioning also automatically keeps track of the configuration that was present on XEMs that have been hot-swapped out of a switch, or on units that have been removed from a stack. Provisioning keeps a 'placeholder' for a XEM or switch which has been hot-swapped out.
- If you provision a switch or bay, then decide later to change the stack member ID or bay number before it has been installed, you must unprovision (**no switch <stack ID> bay/switch**) the switch or bay first.

Provisioning a bay

With the **show system** command we can see that the stack member 2— the x900-24XT switch— does not have a XEM in bay 2:

```
awplus#show sys
Stack System Status                                     Wed May 05 00:04:16 2015

Stack member 1:

Board          ID Bay  Board Name                      Rev  Serial number
-----
Base           271      x900-24XS                        B-0  P1HF7801H
Expansion      272 Bay1  XEM-1XP                          B-0  41AR67008
Expansion      285 Bay2  XEM-STK                          A-0  M1L18400R
PSU            212 PSU1  AT-PWR01-AC                      F-1  73173269
Fan module     214 PSU2  AT-FAN01                          F-1  73169578
-----
RAM: Total: 513372 kB Free: 396680 kB
Flash: 31.0MB Used: 15.9MB Available: 15.1MB
-----
Environment Status : Normal
Uptime              : 0 days 00:55:48
Bootloader version : 1.0.9

Stack member 2:

Board          ID Bay  Board Name                      Rev  Serial number
-----
Base           270      x900-24XT                        A-0  M1QH78003
Expansion      285 Bay1  XEM-STK                          A-0  M1L17400G
PSU            212 PSU2  AT-PWR01-AC                      B-1  61410709
-----
RAM: Total: 513372 kB Free: 410648 kB
Flash: 63.0MB Used: 30.9MB Available: 32.1MB
-----
Environment Status : Normal
Uptime              : 0 days 00:25:34
Bootloader version : 1.0.9

We can see that Stack member 1 is the master, and we are connected to
the console port on this switch:

awplus#show stack
Virtual Chassis Stacking summary information

ID Pending ID  MAC address          Priority  Status  Role
1   -           0000.cd27.c4bf       128     Ready   Active master
2   -           0000.cd28.0801       128     Ready   Backup Member

Operational Status          Normal operation
Stack MAC address           0000.cd27.c4bf
```


On the Stack master (stack member 1) we can provision a XEM-12 for Stack member 2 in bay 2 (which is currently empty):

Note that the switch automatically provisions all currently installed switches and XEMs as it boots up (it does not provision the actual stacking XEMs).

```
awplus(config)#switch 2 bay 2 provision xem-12

switch 1 provision x900-24
switch 1 bay 1 provision xem-1
switch 2 provision x900-24
switch 2 bay 2 provision xem-12
!
interface port2.0.1-2.0.24
 switchport
 switchport mode access
!
interface port2.2.1-2.2.12
 switchport
 switchport mode access
!
```

We can see above that we now have ports 2.2.1-2.2.12 available for configuration in the running-config, even though stack member 2 does not actually have a 12 port XEM (XEM-12) physically installed in bay 2 yet.

This means that these ports can now be configured, ready for when the XEM-12 is installed. Commands can refer to ports on that provisioned XEM as though it were already present:

```
awplus(config)#int port2.2.1
awplus(config-if)#switchport access vlan 2
```

Once a XEM is hot-swapped into bay 2 the switch 2 bay 2 provision XEM-12 will still show in the running configuration, along with the other installed switches and XEMs.

If the XEM is removed, the provisioning for it will remain along with the configuration for the ports associated with it.

What happens when a provisioned XEM gets hot-swapped out?

In the example below, stack member 1 has a XEM-1XP installed in bay 1 and its port (port1.1.1) has been configured as a trunk:

```
switch 1 provision x900-24
switch 1 bay 1 provision xem-1
switch 2 provision x900-24
!
interface port1.1.1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan all
  switchport trunk native vlan none
!
```

If the XEM-1XP is hot-swapped out of bay 1:

```
awplus#08:23:05 awplus HPI: HOTSWAP Pluggable 1.1.1 hotswapped out:
FTRX-1411-3
08:23:05 awplus HPI: HOTSWAP XEM 1 hotswapped out: XEM-1XP
08:23:05 awplus EXFX[1268]: Handle event: bay 1 hsState 4 bt 272 br 0
08:23:05 awplus NSM[1121]: Removal event on bay 1.1 has been completed
```

We can see that the configuration associated with this port is still in the running configuration:

```
interface port1.1.1
  switchport
  switchport mode trunk
  switchport trunk allowed vlan all
  switchport trunk native vlan none
!
```

What happens when the XEM is hot-swapped back in?

```
awplus#08:25:18 awplus HPI: HOTSWAP XEM 1 hotswapped in: XEM-1XP
08:25:18 awplus HPI: HOTSWAP Pluggable 1.1.1 hotswapped in: FTRX-1411-3
08:25:18 awplus EXFX[1268]: Handle event: bay 1 hsState 2 bt 272 br 1
08:25:22 awplus EXFX[1268]: Board XEM-1XP inserted into bay 1
08:25:22 awplus EXFX[1268]: Please wait until configuration update is
completed
08:25:22 awplus IMI[1123]: All users returned to config mode while
switch synchronization is in progress.
08:25:22 awplus VCS[1118]: XEM-1XP has been inserted into bay 1.1
08:25:22 awplus NSM[1121]: Insertion event on bay 1.1 has been completed
08:25:23 awplus IMI[1123]: Configuration update completed for port1.1.1
```

We can see above that port1.1.1 has had its configuration updated from the running configuration.

What happens if a different type of XEM is hot-swapped in?

If the XEM-1XP is hot-swapped out and a different type of XEM (in this case a XEM-12T) is hot-swapped into bay 1 instead:

```
awplus#08:28:48 awplus HPI: HOTSWAP Pluggable 1.1.1 hotswapped out:
FTRX-1411-3
08:28:48 awplus HPI: HOTSWAP XEM 1 hotswapped out: XEM-1XP
08:28:48 awplus EXFX[1268]: Handle event: bay 1 hsState 4 bt 272 br 0
08:28:48 awplus NSM[1121]: Removal event on bay 1.1 has been completed

awplus#08:29:05 awplus HPI: HOTSWAP XEM 1 hotswapped in: XEM-12T
08:29:05 awplus EXFX[1268]: Handle event: bay 1 hsState 2 bt 274 br 2
08:29:08 awplus EXFX[1268]: Board XEM-12T inserted into bay 1
08:29:08 awplus EXFX[1268]: Please wait until configuration update is
completed
08:29:08 awplus IMI[1123]: All users returned to config mode while
switch synchronization is in progress.
08:29:08 awplus VCS[1118]: XEM-12T has been inserted into bay 1.1
08:29:09 awplus NSM[1121]: Insertion event on bay 1.1 has been completed
08:29:11 awplus IMI[1123]: Configuration update completed for
port1.1.1-1.1.12
```

We can see that the provisioning has been modified to reflect the actual hardware installed, and the running configuration now has ports1.1.1-1.1.12, which are the 12 ports belonging to the XEM-12T in bay 1.

```
switch 1 provision x900-24
switch 1 bay 1 provision xem-12
switch 2 provision x900-24
!
interface port1.1.1-1.1.12
 switchport
 switchport mode access
!
```

Provisioning a switch

We have a standalone switch that we will be adding another switch to in the future to form a stack:

```
awplus#show sys
Switch System Status                               Wed May 05 14:34:12 2015

Board      ID  Bay  Board Name                               Rev  Serial number
-----
Base       287      x900-12XT/S                               A-0  M1NB7C023
Expansion  285  Bay1  XEM-STK                                   A-0  A1L18305D
-----

RAM:  Total: 513372 kB Free: 422964 kB
Flash: 63.0MB Used: 46.0MB Available: 17.0MB
```

The current switch has an ID (stack member) of 2.

```
awplus#show stack
Virtual Chassis Stacking summary information

ID Pending ID MAC address Priority Status Role
2 - 0000.cd28.bff7 128 Ready Active master

Operational Status Standalone unit
Stack MAC address 0000.cd28.bff7
```

We can provision stack member 1 so that we can configure the future stack member's ports before we actually have the second switch connected:

```
awplus(config)#switch 1 provision ?
x600-24 Provision an x610-24 switch
x600-48 Provision an x610-48 switch
x900-12 Provision an x900-12 switch
x900-24 Provision an x900-24 switch
x908 Provision an x908 switch
```

Select the model of switch that will be connected in the future (note that you can only stack, and therefore provision, switches of the same basic model).

If we try to provision an x900-24Ts/X switch for stack member 1, and the existing switch (stack member 2) is an x900-12XT/S, we get the following error message:

```
awplus(config)#switch 1 provision x900-24
% Board class x900-24 is incompatible with existing stack members
```

We can successfully provision an x900-12XT/S as follows:

```
awplus(config)#switch 1 provision x900-12
```

The running-config shows that we can now configure the ports (1.0.1-1.0.12) on provisioned stack member 1:

```
switch 1 provision x900-12
switch 2 provision x900-12
!
interface port1.0.1-1.0.12
 switchport
 switchport mode access
!
```

Note that the configuration applied to ports 1.0.1-1.0.12 is just the default port configuration. The port trunk configuration that had been provisioned for the XEM-1XP is completely discarded when the XEM-12S is hot-swapped in instead.

Re-provisioning

If you want to change the provisioning (for example, change a provisioned x610-24Ts to an x610-48Ts), you would first have to use the command **no switch x provision**, followed by **switch x provision x610-48**, as the command **switch x provision x610-48** will fail if there is existing provisioning. However this process means you will lose all of the configuration for portx.0.1-24.

Using **switch x reprovision x610-48** allows you to change the provisioning without losing any existing configuration (within the limits of the respective port counts of the two device types). In short, it allows you to change existing provisioning, provided no actual hardware is present.

We can also reprovision a XEM in a bay or an x900 or SBx908. Below we have provisioned a XEM-12 in bay 2 on switch member 2:

```
awplus(config)#switch 2 bay 2 provision xem-12
```

We could then configure port2.2.1 (the first port on the XEM-12) as follows:

```
awplus(config)#int port2.2.1
awplus(config-if)#swi access vlan 2
```

If we decided to use a XEM-1XP instead of the XEM-12, we can reprovision this change and keep the configuration for any ports that overlap - in this case only port2.1.1:

```
awplus(config)#switch 2 bay 2 reprovision xem-1
```

If we had used the **no provision** command followed by the **provision** command for the new XEM, the overlapping port (port2.2.1) would have been deleted and any configuration on it lost.

Configuring the Stack

The configuration script on a VCStack is shared between all stack members. When configuring the stack, you will need to be aware that the stack uses a specific port numbering scheme and that there are some minor restrictions to VLAN, IP subnet, and QoS configuration. There are also specific triggers available for stacking that you may wish to use. See the following sections for more detail:

- ["Port numbering" on page 46](#)
- ["VLAN and IP subnet restrictions" on page 46](#)
- ["Quality of Service \(QoS\) restriction" on page 47](#)
- ["Stacking triggers" on page 47](#)

For information about how the stack synchronizes the startup configuration and running configuration between stack members, see the section "[Software and Configuration File Synchronisation](#)" on page 48.

Port numbering

In the AlliedWare Plus operating system, switch port interfaces are always referenced as portx.y.z. The first number, or 'x', in the three number port identifier is the Stack ID.

For a stand-alone switch that has never been in a stack, the ports are always numbered 1.y.z. When configuring a stack, however, there will be stack members with other stack ID values. To configure ports on these switches, use the stack ID to refer to each physical switch, for example 2.0.1, 2.0.2, 2.0.3, for the first 3 ports on a switch with stack ID 2.

Please note that when a switch is removed from a stack, it maintains its stack ID number. If it is configured as a stand-alone switch, you will need to either change the stack ID back to 1, or use the current stack ID when configuring its ports.

VLAN and IP subnet restrictions

Internal communication between the stack members is carried out using IP packets sent over the stacking links. This stack management traffic is tagged with a specific ID and uses IP addresses in a specified subnet.

By default, the VLAN and subnet used are:

- VLAN 4094
- Subnet 192.168.255.0/28

You may need to change these values if they clash with a VLAN ID or subnet that is already in use in the network. Use the commands:

```
awplus(config)#stack management subnet <ip-address>
awplus(config)#stack management vlan <2-4094>
```

It is important that the settings for **management subnet** and **management vlan** are the same for all the switches in a stack. If a switch is added to a stack, and its setting for **management vlan** and/or **management subnet** differ from those on the other stack members, the new switch will not be joined to the stack and a log message similar to the following will be created:

```
06:51:41 awplus-vcs VCS[1066]: Member 2 (0009.41fd.dd0b) cannot join
stack - incompatible VCS management subnet
```

Quality of Service (QoS) restriction

The management communication that the stack members exchange over the stacking link is vital for the successful operation of the stack. Be careful to avoid interrupting that communication. This communication traffic is transmitted on egress queue 7 on the stacking ports.

We recommend that you avoid sending any user traffic into queue 7 on a VCStack. In any QoS configuration on a stack, prioritize traffic only into queues 0-6. Moreover, you should ensure that the CoS-to-Queue map does not automatically use queue 7 for the transmission of packets that are received with a CoS value of 7. To ensure this, change the cos-to-queue map to put packets with a CoS value of 7 into queue 6. Use the command:

```
awplus(config)#mls qos map cos-queue 7 to 6
```

Stacking triggers

There are five trigger types defined for particular stack events. Any scripts that you configure for triggers are copied from the active master to all stack members. The triggers are:

```
awplus(config-trigger)#type master-fail
awplus(config-trigger)#type stack member join
awplus(config-trigger)#type stack member leave
awplus(config-trigger)#type stack xem-stk up
awplus(config-trigger)#type stack xem-stk down
```

Licensing

On a stack, the **license** command will activate the license on all the current stack members. If other members are subsequently added to the stack, the license will not be automatically installed into them. They will connect to the stack successfully, even though they have a different license status to the other stack members. A log message warning you that a license mismatch now exists will be generated, but you must manually install the required license(s) into the new member(s). It is important to bring the new member(s) up to the same license level as the existing master. If this is not done, then if a license-lacking member becomes master at some point, then the whole stack will revert to the license level of that master.

To see the license status of stack members, use the command:

```
show license [<label>] member [<1-8>|all]
```

Just as the **license** command applies a license to all current stack members, so too does the **no license** command remove a license from all current stack members.

Software and Configuration File Synchronisation

A VCStack requires that the software version and the configuration files on all stack members are the same. If these files are not the same, then the stack cannot form or then operate correctly. To make stack formation easy, all these files are synchronized automatically on the stack by the stack master.

The following are synchronized:

- "Software release auto-synchronisation" on page 48
- "Shared running configuration" on page 49
- "Shared startup configuration" on page 50
- "Scripts" on page 50

Software release auto-synchronisation

The VCStack implementation supports a feature called **software auto-synchronization**. The effect of this feature is that when a new member joins a stack and has a software release that is different to the active master, then the active master's software release is copied onto the new member. The new member then reboots and comes up on that release.

The sequence of events that occurs at startup is shown in the output below:

```
Notice: Possible stack, please wait whilst searching for members.
12:37:30 awplus-2 VCS[1057]: Member 1 (0000.cd27.c4bf) has joined stack
12:37:30 awplus-2 VCS[1057]: Please configure 'stack virtual-mac' to minimize network
disruption from failovers
12:37:30 awplus-2 VCS[1057]: Member 1 (0000.cd27.c4bf) has become the Active master
12:37:34 awplus-2 VCS[1057]: Software incompatibility detected.
12:37:34 awplus-2 VCS[1057]: Auto synchronization starts.
Unit will reboot once that finishes. Please wait..
12:39:35 awplus-2 VCS: The new software x930 5.4.5-0.2.rel is set as preferred software
and used at the next reboot.
12:39:35 awplus-2 VCS: The current software x930 5.4.5-0.1.rel is set as backup
software.
12:39:35 awplus-2 VCS: VCS SW version auto synchronization successfully completed.

Restarting system
```

The software auto-synchronization feature is enabled on all switches by default. You can enable or disable it using the command:

```
awplus(config)#(no) stack <1-8> software-auto-synchronization
```

If you disable software auto-synchronization, you may disrupt a stack forming if the stack members have different software releases.

If a new member joins the stack, and is running a software version that is different to the active master, the active master will reject the new member from the stack if it cannot synchronism the software releases.

Note: This feature can result in the new stack member downgrading its software release if the active master is running an older software version.

Occasionally you will find that the software version running on the existing stack members and the software version on the new stack member are sufficiently incompatible that the stack will not be able to auto synchronism the software version on the new member.

On these occasions, the log on the stack master will contain a message:

```
user.alert awplus VCS[994]: Neighbour on link 1.0.2 cannot join stack-
incompatible stack S/W version
```

Also, the internal stacking log on the stack master will contain similar messages, that can be seen by using the command **show stack full-debug | include version**.

```
awplus#show stack full-debug|include version
2015 Jul 3 02:04:12 user.debug awplus-2 VCS[992]: STK TRACE: Topology
discovery version mismatch - received v23.4
2015 Jul 3 02:04:12 user.debug awplus-2 VCS[992]:Neighbour on link 2.0.1
cannot join stack - incompatible stack S/W version
```

When this occurs you will need to:

1. remove the new member from the stack
2. manually upgrade the software on that switch to the same version as is running on the existing stack members
3. re-connect the new member to the stack.

Shared running configuration

A single, unified running configuration is shared by all the switches in a stack.

This means that the running configuration on each stack member is exactly the same, and contains the configuration information for all stack members. For example, on a two-switch stack, switch A with stack ID 1 will show the configuration for switch B with stack ID 2, as well as its own configuration. If the running configuration on switch B with stack ID 2 is examined, the output will be exactly the same as that produced by switch A.

Shared startup configuration

Once a stack has formed, the startup configuration stored on each stack member is over written by the active master's startup configuration. This means all the switches in a stack have exactly the same startup configuration script.

Every time the startup configuration script on the active master is changed, for example when the running config is copied to the startup config, the updated startup script is copied to all the other stack members.

```
stack(config)#boot config-file interop.cfg
VCS synchronizing file across the stack, please wait..
File synchronization with stack member-2 successfully completed
[ DONE ]
```

It is important that you take this behavior into account when you first create a stack. If you have carefully crafted a startup configuration on the switch that you intend to be the active master, but some mistake results in a different switch becoming the active master when the stack forms, then the intended startup configuration will be over written by something else (assuming the same filename was used for both configurations).

We recommend that you:

- take care when first installing the VCStack to ensure that you configure the stack priority values correctly. If you do not, the wrong stack member may become the active master on first boot, and overwrite the stack with the wrong startup configuration.
- backup the startup configuration of the intended active master before connecting the switch to the stack. You can make another copy on the switch's Flash file system with the command:

```
awplus#copy startup_filename.cfg backup_filename.cfg
```

As well as the running configuration and startup configuration, the stack synchronizes the **boot backup configuration** file, and the **fallback configuration** file. Like the startup configuration file, these are synchronized from the active master's file system, so the same recommendations apply.

Scripts

Script files stored on the active master's file system are copied to the stack members.

Rolling Reboot

A major benefit of Virtual Chassis Stacking is that it provides unit resiliency—if one unit in the stack goes down, the other members of the stack will continue to forward data. It is highly desirable for this continuity of service to persist even when the stack is being rebooted. The purpose of the **reboot rolling** command is to reboot a stack in a manner that maintains continuity of service.

The **reboot rolling** command allows a stack to be rebooted in a rolling sequence so that no more than one unit of the stack is in reboot at any given time. First, the stack master is rebooted causing the remaining stack members to failover and elect a new master.

Here we have a stack of three x610 switches:

```
awplus#show stack
Virtual Chassis Stacking summary information

ID  Pending ID  MAC address      Priority  Status  Role
1   -           0015.77c2.4b7d   128      Ready   Backup Member
3   -           0015.77e8.a892   128      Ready   Backup Member
4   -           0015.77c9.73cb   128      Ready   Active master

Operational Status      Normal operation
Stack MAC address        0015.77c9.73cb
```

We can see that stack member **4** is the Active master.

Executing the **reboot rolling** command gives the following:

```
awplus#reboot rolling
The stack master will reboot immediately and boot up with the
configuration file settings. The remaining stack members will then
reboot once the master has finished re-configuring.

Continue the rolling reboot of the stack? (y/n):y

awplus#22:11:15 awplus VCS[995]: Automatically rebooting stack member-
4 (MAC: 00.15.77.c9.73.cb) due to Rolling reboot

URGENT: broadcast message:
System going down IMMEDIATELY!

... Rebooting at user request ...
```

As soon as the rebooted Active master has reloaded, it becomes the Active master again. While it is rebooting, another switch in the stack will assume the Active master role.

Immediately after the Active master has reloaded and assumed its role again, all of the other switches in the stack are rebooted at the same time.

```
Active master booting up:

Loading default configuration
.

done!
Received event network.configured

Rolling reboot, rebooting all other stack members, please wait for stack
to reform.
```

You can see in the Active master's log that the other stack members (1 and 3) have rebooted:

```
2015 May 10 22:12:11 user.crit awplus-4 VCS[995]: Member 4 (0015.77c9.73cb) has become
the Active master
2015 May 10 22:12:37 local6.notice awplus VCS[995]: Link down event on stack link 4.0.2
2015 May 10 22:12:37 local6.notice awplus VCS[995]: Link down event on stack link 4.0.1
2015 May 10 22:13:32 local6.notice awplus VCS[995]: Link up event on stack link 4.0.1
2015 May 10 22:13:32 local6.notice awplus VCS[995]: Link down event on stack link 4.0.1
2015 May 10 22:13:32 local6.notice awplus VCS[995]: Link up event on stack link 4.0.2
2015 May 10 22:13:33 local6.notice awplus VCS[995]: Link down event on stack link 4.0.2
2015 May 10 22:13:36 local6.notice awplus VCS[995]: Link up event on stack link 4.0.1
2015 May 10 22:13:36 user.crit awplus VCS[995]:Member 3 (0015.77e8.a892)has joined stack
2015 May 10 22:13:36 user.notice awplus VCS[995]: Link between members 4 and 3 is up
2015 May 10 22:13:37 local6.notice awplus VCS[995]: Link up event on stack link 4.0.2
2015 May 10 22:13:37 user.crit awplus VCS[995]:Member 1(0015.77c2.4b7d) has joined stack
2015 May 10 22:13:37 user.notice awplus VCS[995]: Link between members 4 and 1 is up
2015 May 10 22:13:37 user.notice awplus VCS[995]: Link between members 3 and 1 is up
```

Failover and Recovery

A VCStack behaves in a reliable and consistent manner if any component fails, whether the problem is with a stacking link or stack member. However, before looking at how the stack reacts to each of those scenarios, we first need to understand a generic mechanism that is available in Allied Telesis VCStack to help deal with failure scenarios—namely the resiliency link.


The resiliency link feature

If one or more stacking links fail, so that some stack members are no longer in contact with the active master switch, then the other stack members are left with a dilemma. Should they assume that the active master switch has gone down, and re-elect a new active master, or should they assume that the active master is still operating?

This problem is particularly important when a stack has multiple connections to edge switches in the network. In this network scenario, if a stack splits into two active sections

that are operating independently, then the edge switches will continue to treat their uplink ports as aggregated and share traffic across the links.

However, the broken stack link could mean that traffic arriving at some of the stack members cannot reach the intended destination.



It is necessary to have a backup mechanism through which stack members can check if the active master is still active in case the stack link communication to the active master is lost.

The mechanism provided in Allied Telesis VCStack is called the **resiliency link**, which may be either the eth0 port (only on SwitchBlade x908 or x900 Series switches) or a dedicated VLAN (resiliencylink VLAN) to which switch ports may become members. These resiliency ports are connected together, and the stack members all listen for periodic (one per second) Health Check messages from the master. As long as the active master sends Health Check messages, the other stack members know that the active master is still active.

This means that the stack members can know whether the active master is still operating if they lose contact with the active master through the stacking links.

The out-of-band Ethernet port is configured as a resiliency port with the command:

```
awplus(config)#stack resiliencylink eth0
```

Note that even if you configure the eth0 port as a resiliency port, you can still use it for out-of-band management. A VLAN, and switch port are configured for resiliency link connection with the commands:

```
awplus(config)#stack resiliencylink vlan1000
awplus(config)#interface port1.0.1
awplus(config-if)#switchport resiliencylink
```

Note also that this VLAN is dedicated to the resiliency link function and must not be the stack management VLAN or a customer data VLAN.

In the situation where a stack member loses contact through the stacking links, but continues to receive health check messages via the resiliency link, the stack member becomes a disabled master. It runs the same configuration as the active master, but it has its links shutdown.

There is a trigger that can be configured when a switch becomes a disabled master with the commands:

```
awplus(config)#trigger 82
awplus(config-trigger)#type stack disabled master
awplus(config-trigger)#script 1 flash:/disabled.scp
```

Although this command could activate any trigger script, the intention here is that the script will re-activate the links from their previously shutdown state, to enable the user to manage the switch.

Virtual MAC

To minimize data loss if a new stack member is required to become the master, enable a Virtual MAC. The VCStack then uses the virtual MAC address for communication with other devices. As this single virtual MAC address is used for the complete VCStack, there is no change of MAC address if a new stack member is required to become master. In conjunction with Fast Failover, this ensures maximum continuity of network service, as there is no need for other devices in the network to learn a new MAC address into their MAC or ARP tables.

You can configure the virtual MAC address manually by specifying a VCStack virtual chassis ID. Or the VCStack can use its default virtual chassis ID. The ID selected will determine which virtual MAC address the stack will use. The MAC address assigned to a stack must be unique within its network.

Enter the chassis ID in decimal format. The virtual chassis ID entered will form the last 12 bits of a pre-selected MAC prefix component; that is, 0000.cd370xxx.

For example:

```
awplus(config)#stack virtual-mac
awplus(config)#stack virtual-chassis-id 63
```

This will result in a virtual MAC address of: 0000.cd37.003f

Repairing a broken stack

When two stub stacks are reconnected, the stack will detect if there is more than one master (active, disabled, or fallback). You will see console messages and logs that report a **duplicate master** was detected. The stack carries out an election to choose the true active master, based on priorities and MAC addresses. Any losing master will reboot and become a backup member. The switches that have been running on their fallback configuration (most likely those in the same stub as the losing active master) will reboot and come up on the shared configuration pushed to them by the winning active master.

When a stack detects a duplicate master, a debug snapshot file is created and stored in Flash. The file is called debug-duplicate-master-*<time & date>*.tgz. The debug snapshot does not contain a .core file, as is produced when a process unexpectedly terminates on a switch. Instead, it contains a snapshot of various pieces of information within the software at the time when the stack detected the duplicate master.

Disabled master

A properly functioning VCStack contains an (active) master and one (backup) member. Under fault conditions it is possible for the stack member to lose connectivity with the stack master. In this situation the stack member without master connectivity will become a “disabled master”. Once in this state the disabled master will disable all of its own ports.

Apart from this, the operation and “look and feel” of a disabled master is very similar to an active master. The active master’s ports are unaffected by this change, and they will continue to forward traffic normally.

By disabling all the stub’s switchports, the disabled master avoids potential network connectivity problems that could result from by having two stack masters using the same configuration, or two switches in separated stubs trying to share the same “logical” communications paths such as a non functioning aggregated links. The active master’s ports are unaffected by this, and they will continue to forward traffic normally. Note that status information for members of the stack stub can be accessed by logging into the disabled master, in the same way as obtaining status information for a normal stack.

Disabled Master Monitoring (DMM)

The stack resiliency link and disabled master state offer a unique prevention of catastrophic network connectivity problems. However, when stack members become separated, the network is still left in a fragile state where the stack master no longer has the redundancy of the backup stack member. If the original stack master were to subsequently fail whilst the stack was separated, then all network connectivity would be lost if the disabled master’s switchports remained shut down.

The DMM feature avoids this situation by continuing to monitor the status of the original stack master (the active master) via the stack resiliency link. When the DMM feature is enabled, the disabled master can detect a failure of the original stack master within a few seconds.

If a failure is detected, the disabled master transitions to the active master state and automatically re-enables all its switchports. This allows traffic forwarding via the stack to continue.

STACK BEHAVIOR WITH DMM DISABLED	STACK BEHAVIOR WITH DMM ENABLED
<ul style="list-style-type: none"> ■ The VCStack breaks with the Stack Resiliency Link configured and enabled. 	<ul style="list-style-type: none"> ■ The VCStack breaks with the Stack Resiliency Link configured and enabled.
<ul style="list-style-type: none"> ■ The separated stack member becomes a disabled master. 	<ul style="list-style-type: none"> ■ The separated stack member becomes a disabled master.
<ul style="list-style-type: none"> ■ The disabled master does not monitor the active master. 	<ul style="list-style-type: none"> ■ The disabled master monitors the active master.
<ul style="list-style-type: none"> ■ If the active master fails then the disabled master does not become the active master (no state transition). 	<ul style="list-style-type: none"> ■ If the active master fails then the disabled master becomes the active master (disabled to active transition).
<ul style="list-style-type: none"> ■ No switchports are re-enabled on the disabled master. No traffic is forwarded. 	<ul style="list-style-type: none"> ■ Switchports on the disabled master are re-enabled. Traffic is still forwarded.

To enable the DMM feature, use the commands:

```
awplus#configure terminal
awplus(config)#stack disabled-master-monitoring
```

To disable the DMM feature, use the commands:

```
awplus#configure terminal
awplus(config)#no stack disabled-master-monitoring
```

To show the status of DMM on the VCStack, use the command:

```
awplus#show stack [detail]
```

To apply a trigger upon transition from active master state to disabled master state, use the command:

```
awplus#type stack disabled-master
```

To apply a trigger upon transition from disabled master state to active master state, use the command:

```
awplus#type stack master-fail
```

Note: A disabled master trigger allows you to specify a script to reconfigure the disabled master on the fly, should a catastrophic failure separate the stack. This is useful to configure an alternate IP address so you can still log in to the disabled master via an SSH or a Telnet connection. The trigger script should use the **no shutdown** command to re-enable any switchports needed for an SSH or a Telnet management connection.

Replacing a stack member

A stack member can be removed from a stack (hot-swapped out) with minimal impact on stack traffic. To do this, power-down the stack member, and disconnect its stacking ports. You can seamlessly swap a stack member switch into the stack to replace another with the same configuration. This provides a simple way to replace an out-of-service switch with minimal impact, and minimal administration requirement. Before inserting the replacement device into the stack, make sure that:

- The replacement device is running a compatible firmware version.
- You set the Stack ID on the replacement device to the same ID as the device being replaced.
- The replacement device is installed with the same licenses as the other stack members, this ensures consistent behavior across the stack.

Once these requirements are met, insert the new stack member, reconnect the stacking ports and power-up the new stack member.

Executing Commands and Managing Files on Stack Members

There are some limited actions that you can do on an individual switch member:

- **Executing commands**

You can use the **reload** command to reboot an individual switch, or use **show** commands to see the individual physical state of a switch and its file system. Show commands that relate to the ports, counters, or configuration are applicable for the stack only.

- **Managing files**

You can manage files on an **individual** switch. Note that some files are synchronised between switches. See the "[Software and Configuration File Synchronisation](#)" on [page 48](#) section for more information.

You can also delete files off **all stack members simultaneously**, see "[Deleting files from all stack members](#)" on [page 59](#).

Executing commands

If you want to run a command that displays information from a specific stack member, you can prefix the command with the syntax:

```
awplus#remote-command <stack-ID>
```

For example, to see the full state of all the file systems on the stack member with stack ID 3, use the command:

```
awplus#remote-command 3 show file systems
```

To see the processes running on the stack member with stack ID 2, use the command:

```
awplus#remote-command 2 show process
```

To use the **show system environment** command on the stack member with stack ID 2, use the command:

```
awplus#remote-command 2 show system environment
```

Output

```
x900#remote-command 2 show system environment
Stack Environment Monitoring Status
Stack member 2:
Overall Status: Normal

ID Sensor (Units)                Reading  Low Limit High Limit Status
1  Device Present                 Yes      -        -        Ok
2  PSU Overtemp                   No       -        -        Ok
3  PSU Fan Fail                   No       -        -        Ok
4  PSU Power Output               No       -        -        Ok

Resource ID: 2 Name: PSU bay 2 (AT-PWR01-AC)
ID Sensor (Units)                Reading  Low Limit High Limit Status
1  Device Present                 Yes      -        -        Ok
2  PSU Overtemp                   No       -        -        Ok
3  PSU Fan Fail                   No       -        -        Ok
4  PSU Power Output               Yes      -        -        Ok

Resource ID: 3 Name: x900-24XT
ID Sensor (Units)                Reading  Low Limit High Limit Status
1  Voltage: 2.5V (Volts)          2.578   2.344   2.865   Ok
2  Voltage: 1.65V (Volts)         1.629   1.488   1.816   Ok
3  Voltage: 3.3V (Volts)          3.369   2.973   3.627   Ok
4  Voltage: 1.8V (Volts)          1.797   1.615   1.979   Ok
5  Voltage: 12V (Volts)           11.938  10.813  13.188  Ok
6  Temp: Ambient (Degrees C)      23      -126    55      Ok
7  Temp: Mid Internal (Degrees C)  41      -126    85      Ok
8  Temp: Bk Internal (Degrees C)  31      -126    75      Ok

Resource ID: 4 Name: XEM-STK Bay: 2
ID Sensor (Units)                Reading  Low Limit High Limit Status
1  Voltage: 2.5V (Volts)          2.435   2.344   2.865   Ok
2  Voltage: 1.65V (Volts)         1.617   1.491   1.814   Ok
3  Voltage: 3.3V (Volts)          3.248   2.973   3.627   Ok
4  Voltage: 5V (Volts)            5.026   4.505   5.495   Ok
5  Voltage: 12V (Volts)           11.563  10.813  13.188  Ok
6  Voltage: 1.8V (Volts)          1.772   1.617   1.983   Ok
7  Temp: Internal (Degrees C)     40      78(Hyst) 80      Ok
```

Additionally, the **reload** command can take a stack ID as an extra parameter. To reboot just the stack member with stack ID 4, use the command:

```
awplus#reload stack-member 4
```

You will get the following question from the stack:

```
stack#reload stack-member 4
reboot stack member 4 system? (y/n):
```

Managing files

If you wish to perform an action on another stack member's file system, the syntax is:

```
awplus#<stack-member-name>/flash:[ / ]<filename>
```

The *stack-member-name* is not the stack ID, but is an extended hostname formed as:

```
awplus#<hostname>-<stack-ID>
```

If the hostname of the stack is **BlueCore**, then the *stack-member-name* for switch 2 in the stack is:

```
BlueCore-2
```

If you do not use the *stack-member-name* prefix, then the command refers to a file that resides on the stack master.

You can also use the *stack-member-name* syntax for the directory command. To view the contents of the Flash file system on a specific stack member, you can use the syntax:

```
awplus#dir <stack-member-name>/flash:/
```

Deleting files from all stack members

If you wish to remove file(s) from all stack members simultaneously, use the following command:

```
awplus#delete stack-wide force [recursive] <name>
```

where:

- **recursive**—deletes directories that match the name, including their contents
- **<name>**—the name of the files or directories to delete.

The filename can include the wildcard *

Use the wildcard with caution, because this command is non-interactive and does not ask for confirmation before deleting files. This is indicated by the mandatory **force** parameter.

You can use this command within an AMF working set.

For example, to delete a file that is located in Flash memory on all stack members, use the following command:

```
awplus#delete stack-wide force test.scp
```

To remove directories output1 and output2 from an external card on all stack members, use the command:

```
awplus#delete stack-wide force recursive card:output*
```

Remote login

Occasionally it can be desirable to login to a specific member of the stack (for example, to manage feature licences per individual unit). Remote login facilitates this by allowing a user on the master switch to log into the CLI of another stack member.

In most respects, once you are logged into the other stack member, the result of entering commands will be as if you were logged into the stack master, i.e. the command **show ip interface** will show all IP interfaces configured on all switches in the stack - not just those on the stack member that you have connected to with the **remote-login** command. Configuration commands are still broadcast to all stack members.

Some **show** commands that display physical attributes of the switch, and commands that access the file system, are executed locally. Also, commands related to feature licences are executed locally.

To login from the stack master (stack member 1 in this case) to stack member 2:

```
awplus#remote-login ?
  <1-8>  A specific stack member ID

awplus#remote-login 2
Type 'exit' to return to awplus.

AlliedWare Plus (TM) 5.3.4 05/04/10 11:59:17

awplus-2>en
awplus-2#
```

Notice that the prompt has changed to reflect the stack member (2) that we are currently connected to. A directory listing will now show the files on stack member 2 only:

```
awplus-2#dir *.cfg
  948 -rw- May  4 2015 20:59:48  flash:/default.cfg
  677 -rw- May  3 2015 18:39:04  flash:/zz.cfg
 2944 -rw- Mar 23 2015 12:55:40  flash:/ospfv3.cfg
```

We can delete a file from stack member 2 as if we were directly connected to it:

```
awplus-2#del zz.cfg
Delete flash:/zz.cfg? (y/n)[n]:y
Deleting..
Successful operation
awplus-2#
```

To return to the stack master, use the **exit** command.

```
awplus-2#exit
awplus#
```

Managing licenses on a stack member

Remote login makes managing feature licenses on the stack members easy. We can simply connect to the stack member:

```
awplus#remote-login 2
Type 'exit' to return to awplus.

AlliedWare Plus (TM) 5.3.4 05/04/10 11:59:17

awplus-2>en
awplus-2#
```

The **show license** command displays the current feature licenses on stack member 2:

```
awplus-2#show license
Software Feature Licenses
-----
Index                : 0
License name         : Base License
Customer name        : Base License
Quantity of licenses : 1
Type of license      : Full
License issue date   : 10-May-2015
License expiry date  : N/A
Features include     : VRRP OSPF-64 RADIUS-100 Virtual-MAC

Index                : 1
License name         : Base License
Customer name        : Base License
Quantity of licenses : 1
Type of license      : Full
License issue date   : 11-Aug-2014
License expiry date  : N/A
Features include     : BGP-64 PIM RIPNG VRRP OSPF-FULL VlanDT
OSPF-64 BGP-FULL IPv6Basic MLDSnoop BGP-5K RADIUS-100 RADIUS-FULL PIM-
100 ACCESS LAG-128 Virtual-MAC
```

Use the **license** command to apply the required feature license to the VCStack, as shown in the following output example. The **license** command will add a license to all stack members and the **no license** command will remove a license from all stack members.

```
awplus-2#license IPv6
Qd0NvZJ8DutyLAYbsM8pCpY1d8Ho9mzygweBp+paBqVu7By1bTZ+Jipo57

A restart of affected modules may be required.
Would you like to continue? (y/n): y
Stack member 1 installed 1 license

1 license installed.

awplus-2#show license

Software Feature Licenses
-----
Index                : 0
License name          : Base License
Customer name         : Base License
Quantity of licenses : 1
Type of license       : Full
License issue date    : 10-May-2015
License expiry date   : N/A
Features include      : VRRP OSPF-64 RADIUS-100 Virtual-MAC

Index                : 1
License name          : Base License
Customer name         : Base License
Quantity of licenses : 1
Type of license       : Full
License issue date    : 11-Aug-2014
License expiry date   : N/A
Features include      : BGP-64 PIM RIPNG VRRP OSPF-FULL VlanDT
OSPF-64 BGP-FULL IPv6Basic MLDSnoop BGP-5K RADIUS-100
RADIUS-FULL PIM-100 ACCESS LAG-128 Virtual-MAC

Index                : 1
License name          : Base License
Customer name         : Base License
Quantity of licenses : 1
Type of license       : Full
License issue date    : 09-May-2014
License expiry date   : N/A
Features include      : RADIUS-FULL
```

VCStack Plus – Stacking on the SBx8100

Introduction

To achieve excellent resiliency and simplified management, two SwitchBlade x8100s can be connected together to form a single stacked unit. This is only possible using the CFC960 controller cards, which have 4 x 10 GB SFP+ ports. Because the stacking connections are via fiber SFP+ ports, each SBx8100 switch can be located long distances apart, even kilometers apart if required. The stacking connections must use fiber links, and can be terminated by any Allied Telesis branded SFP+ modules.

Enabling stacking

Stacking is disabled by default on the SBx8100. When stacking is enabled, with the **stack enable** command, all of the CFC960 SFP+ ports will be reserved for stacking use - they will not be configurable as normal network switchports.

When stacking is enabled via the CLI, the stack Virtual-MAC feature will automatically be enabled as well.

Each stack member must have a unique stack ID. This is set with the command: **stack <existing stack-ID> renumber <new stack-ID>**. The current stack ID of each switch can be seen with the **show stack** command. The SBx8100 must be manually renumbered via the CLI before it is connected in a stack.

Once the stack has formed, the **show stack** command will show the Active CFC.

ID 1.5 refers to Stack member 1, the CFC in bay 5
 ID 1.6 refers to Stack member 1, the CFC in bay 6
 ID 2.5 refers to Stack member 2, the CFC in slot 5
 ID 2.6 refers to Stack member 2, the CFC in slot 6

```
awplus#show stack
Virtual Chassis Stacking summary information

ID    Pending ID  MAC address          Priority  Status  Role
1.5   -            eccd.6db3.586a      128     Ready   Active CFC
1.6   -            eccd.6db3.5919      128     Ready   Backup Member
2.5   -            eccd.6db3.5892      128     Ready   Backup Member
2.6   -            eccd.6db3.58f6      128     Ready   Backup Member

Operational Status          Normal operation
Stack MAC address           0000.cd37.0061 (Virtual MAC)
```

The **show stack detail** command displays a 'Last role change' date/time. Normally if all cards booted up at the same time, the date/time is the same for all cards.

If one card rebooted (or was inserted later), it will have a more recent date/time. If a master (Active CFC) failover occurred, then the stack master will have a more recent date/time than the other backup member CFCs.

```
awplus#show stack detail
Virtual Chassis Stacking detailed information

Stack Status:
-----
Operational Status           Normal operation
Management VLAN ID          4094
Management VLAN subnet address 192.168.255.0
Virtual Chassis ID           97 (0x61)
Virtual MAC address           0000.cd37.0061

Card 1.5:
-----
ID                             1.5
Pending ID                       -
MAC address                      eccd.6db3.586a
Last role change                 Thu Apr 10 23:38:47 2014
Product type                     AT-SBx81CFC960
AT-SBx81CFC960 Stacking Ports   Enabled
Role                             Active CFC
Status                           Ready
Priority                          128
Host name                        stk_a_1340_3001
Stack port1.5.1 status           Learnt neighbor 2.5
Stack port1.5.2 status           Learnt neighbor 2.6
Stack port1.5.3 status           Learnt neighbor 2.5
Stack port1.5.4 status           Learnt neighbor 2.6

Card 1.6:
-----
ID                             1.6
Pending ID                       -
MAC address                      eccd.6db3.5919
Last role change                 Thu Apr 10 23:38:48 2014
Product type                     AT-SBx81CFC960
AT-SBx81CFC960 Stacking Ports   Enabled
Role                             Backup Member
Status                           Ready
Priority                          128
Host name                        stk_a_1340_3001-1.6
Stack port1.6.1 status           Learnt neighbor 2.5
Stack port1.6.2 status           Learnt neighbor 2.6
Stack port1.6.3 status           Learnt neighbor 2.5
Stack port1.6.4 status           Learnt neighbor 2.6
```


Card 2.5:	

ID	2.5
Pending ID	-
MAC address	eccd.6db3.5892
Last role change	Thu Apr 10 23:38:47 2014
Product type	AT-SBx81CFC960
AT-SBx81CFC960 Stacking Ports	Enabled
Role	Backup Member
Status	Ready
Priority	128
Host name	stk_a_1340_3001-2.5
Stack port2.5.1 status	Learnt neighbor 1.5
Stack port2.5.2 status	Learnt neighbor 1.6
Stack port2.5.3 status	Learnt neighbor 1.5
Stack port2.5.4 status	Learnt neighbor 1.6
Card 2.6:	

ID	2.6
Pending ID	-
MAC address	eccd.6db3.58f6
Last role change	Thu Apr 10 23:38:47 2014
Product type	AT-SBx81CFC960
AT-SBx81CFC960 Stacking Ports	Enabled
Role	Backup Member
Status	Ready
Priority	128
Host name	stk_a_1340_3001-2.6
Stack port2.6.1 status	Learnt neighbor 1.6
Stack port2.6.2 status	Learnt neighbor 1.5
Stack port2.6.3 status	Learnt neighbor 1.5
Stack port2.6.4 status	Learnt neighbor 1.6

Software release and license requirements

- All CFCs must be running the same software release in order to form a stack. If a CFC is running a different software release, then that CFC will be automatically synchronized to the stack's software release by the VCStack software-auto-sync feature.
- In order to stack, both chassis also need to have their own VCStack-Plus license installed (specific to SBx8100 stacking).
- Both chassis also need to have the same feature licenses installed (this is the same as VCStack on other platforms).
- Under normal circumstances, CFC card 5 should end up the master of each chassis. One of these CFCs will be elected stack master.
- If CFC card 6 becomes master, rather than card 5, then it is likely that card 5 failed to boot or was significantly delayed in booting. There is a limit to how long the switch waits, and it will not change mastership if there is already a master present.

CFC operation

When two SBx8100s are stacked together, one CFC card will become the Active CFC for both chassis and the other three CFCs will be in Backup.

In the chassis that does not contain the Active CFC - i.e. just two backup CFCs - one of these backup CFCs becomes the **H/W master** for that chassis. The H/W master relates to how the SBx8100 hardware is designed - only one CFC on the chassis can access the H/W information for that chassis.

In a VCStack Plus stack, the H/W master doesn't have a lot of extra duties to do - the stack master (Active CFC) handles most things. In the backup chassis, the only things the H/W master would do differently to the H/W backup member would be:

- controlling the bootup of the LIFs (Line Interface Fabric/cards) i.e. serving out BOOTP IP and TFTP release.
- environmental monitoring, i.e. temperature/fault sensors in fan tray.
- PoE power allocation (as each chassis has its own power to distribute).

The **Active** CFC will do all of the CPU processing for the stack.

You can infer which CFC is the H/W master in the backup chassis from the highest **uptime** displayed in the **show card** command output.

LIF operation

All cards communicate with all other cards in the stack, regardless of what chassis they are in, depending on the actual communication.

For example:

- **stackd** communicates with all other cards.
- **ospfd** only communicates with the other CFCs.
- the **BOOTP** client on the LIFs only communicates with the local H/W master.

The insertion or removal of LIF cards in either chassis is controlled by the Active CFC.

The maximum number of supported ports in the SBx8112 stack is **400** switchports in a fully-populated SBx8112 stack with 24 cards in total.

File synchronization

- **VCStack SW auto-sync** makes sure the same software release is used across the stack, and auto-upgrades the software if needed.
- **VCStack file replication** makes sure necessary files on the master get synchronized with the backup member.

Stack failover

A single CFC960 will be elected as the stack master or 'Active CFC' for the entire SBx8100 VCStack. When the Active CFC is removed or rebooted, the backup CFC on the same chassis will take over as master. However, if the backup CFC on the same chassis is not in the equivalent of the VCStack 'Ready' state, then a 'Ready' CFC on the peer chassis will take over as master.

When the entire SBx8100 containing the Active CFC is failed over (e.g. power-cycled), the CFC960 that has H/W mastership on the peer SBx8100 will take over as master. In other words, when an entire SBx8100 chassis is failed, the CFC with the longest system up-time will become the Active CFC.

Failing over any CFC960 in the stack will temporarily halve the available stacking backplane bandwidth. This may cause network disruption if traffic across the stacking backplane exceeds 40Gbps.

CFC failover

When a single Active CFC failover occurs (i.e. the new master CFC is still in the **same** chassis), the interruption to traffic will not be significantly different from CFC failover on a standalone SBx8100.

Chassis failover

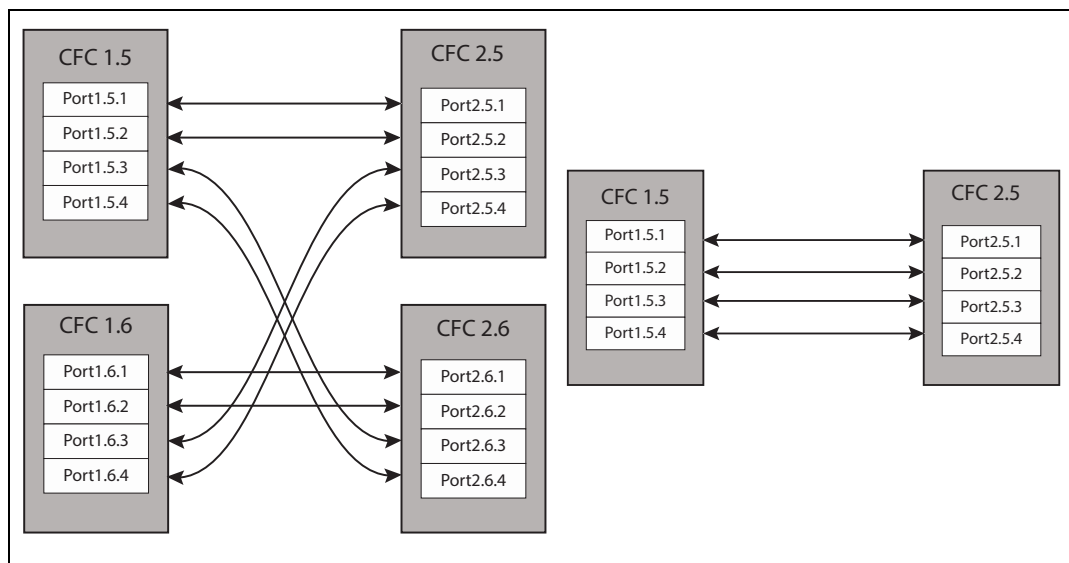
When a entire SBx8100 failover occurs (i.e. the new master CFC is in a **different** chassis), the interruption to static Layer 2 and Layer 3 unicast traffic forwarding will be < 3 seconds. Interruption to other protocol traffic, such as multicast, will be best effort, however disruption may be significantly higher. Failing over both CFCs in the same chassis at the same time will result in all LIFs in that chassis rebooting as well.

Troubleshooting

The command **show counter stack | grep Master** (note the capital 'M') will show you any failovers or conflicts (duplicate masters) that have occurred. Check **show reboot history** and look for any unexpected reboots (note that this will include any times the chassis has been power-cycled). Use the **show stack detail** command to check the 'Last role change' date/time.

Cabling

- If both SBx8100s have two CFC960 controller cards installed, then the cabling is as shown on the left.
- If both SBx8100s only have a single CFC960 controller card installed, then the cabling is as shown on the right:



Note: Any number of cables can be used to interconnect the CFCs, but reducing the number of cables will mean: reduced bandwidth between VCStack Plus members and reduction in resiliency capabilities.

Monitoring and Troubleshooting

The physical connection and the signaling communications between stack members are vital to stacking. Because of this, the **show** commands and debugging facilities for VCStack are oriented around stack port status and stack signaling communications.

Checking stack status

You can check that the stack links have come up successfully by:

- checking the LEDs on the switch or XEM
- using the **show stack detail** command

The tables following describe the LED state and functions for the XEM-STK, SwitchBlade x908, SwitchBlade x8100 CFC960 control card, x600, x510, and x610 Series switches.

LEDs

The LEDs on the **XEM-STK** show the following:

LED	STATE	MEANING
Port 1 and Port 2	Green	A stacking link is established.
	Amber (flashing slowly)	The link has transmission fault.
	Off	The stacking link is down.
Status	Green	The switch is the stack master.
	Amber	The switch is a backup member.
	Green (flashing)	The stack is selecting a stack master.
	Off	The switch is not a stack member.
Numeric ID	1 to 8	The stack member ID.
	Off	The switch is not a stack member.

The front panel of the **SwitchBlade x8100 CFC960 control card** indicates the control card's status within the stack.

LED	STATE	MEANING
CFC	Green	Active master
	Flashing green	Startup
	Amber	Backup member (standby)
	Flashing red	Disabled

The front panel of the **SwitchBlade x908** has LEDs for monitoring back-port stacking:

LED	STATE	MEANING
Port 1 and Port 2	Green	A stacking link is established.
	Amber (flashing slowly)	The link has transmission fault.
	Off	The stacking link is down.
master	Green	The switch is the stack master.
	Amber	The switch is a backup member.
	Green (flashing)	The stack is selecting a stack master.
	Off	The switch is not a stack member.

The front panel of the **x600 Series** switch has the following LEDs for monitoring stacking:

LED	STATE	MEANING
MSTR	Green	The switch is the stack master.
	Off	The switch is a backup member.
L/A 1 and L/A 2	Green	A stacking link is established on that link.
	Green (flashing)	The link is transmitting or receiving data.
	Off	The stacking link is down.
PRES	On	An AT-STACKXG module is correctly installed in the switch.
	Off	There is no AT-STACKXG installed in the switch, or the module is installed incorrectly.

The front panel of the **x610 Series** switch has the following LEDs for monitoring stacking:

LED	STATE	MEANING
MSTR	Off	The switch is not part of a stack or is a member unit of the stack.
	Solid Green	The switch is the master unit of the stack.
L/A 1	Off	Stack port 1 has not established a link to a stacking port on another VCStack stacking module.
	Solid Green	Stack Port 1 has established a link to a stacking port on another VCStack stacking module.
	Flashing Green	Stack Port 1 has established a link to a stacking port on another VCStack stacking module and is sending or receiving packet traffic.
L/A 2	Off	Stack port 2 has not established a link to a stacking port on another VCStack stacking module.
	Solid Green	Stack Port 2 has established a link to a stacking port on another VCStack stacking module.
	Flashing Green	Stack Port 2 has established a link to a stacking port on another VCStack stacking module and is sending or receiving packet traffic.
PRES	Off	Expansion slot for VCStack stacking module is empty.
	Solid Green	A VCStack stacking module is installed in the switch.

The stacking LEDs on the **x510** are quite different to the other products:

- there is no master LED.
- there is a 7-segment LED that displays the actual stack ID number.

Each switch must be assigned an ID number. The range is 1 to 4 and the default is 1. The ID numbers are displayed on the ID LEDs on the front panels of the units.

In addition, there is an LED associated with each SFP socket that is used for stack connections.

The table below shows the LED states when the socket contains a stacking transceiver.

LED	STATE	MEANING
Link activity	Off	The slot is empty, the stacking transceiver has not established a link to a network device, or the LEDs are turned off. To turn on the LEDs, use the eco-friendly button.
	Solid green	The stacking transceiver has established a link at 10 Gbps to another switch in the stack.
	Flashing green	The stacking transceiver is receiving or transmitting packets.

The stacking LEDs on the **x310** are the same as the x510:

- there is no master LED.
- there is a 7-segment LED that displays the actual stack ID number.

Each switch must be assigned an ID number. The range is 1 to 4 and the default is 1. The ID numbers are displayed on the ID LEDs on the front panels of the units. In addition, there is an LED associated with each SFP socket that is used for stack connections.

The table below shows the LED states when the socket contains a stacking transceiver.

LED	STATE	MEANING
Link activity	Off	The slot is empty, the stacking transceiver has not established a link to a network device, or the LEDs are turned off. To turn on the LEDs, use the eco-friendly button.
	Solid green	The stacking transceiver has established a link at 10 Gbps to another switch in the stack.
	Flashing green	The stacking transceiver is receiving or transmitting packets.

The stacking LEDs on the **CentreCOM GS900MX/MPX** and **CentreCOM XS900MX Series** switches are the same as the x310 and x510:

- there is no master LED.
- there is a 7-segment LED that displays the actual stack ID number.

Each switch must be assigned an ID number. The range is 1 to 4 (GS900MX/MPX) and 1 to 2 (XS900MX) and the default is 1. The ID numbers are displayed on the ID LEDs on the front panels of the units. In addition, there is an LED associated with each SFP socket that is used for stack connections.

The table below shows the LED states when the socket contains a stacking transceiver.

LED	STATE	MEANING
Link activity	Off	The slot is empty, the stacking transceiver has not established a link to a network device, or the LEDs are turned off. To turn on the LEDs, use the eco-friendly button.
	Solid green	The stacking transceiver has established a link at 10 Gbps to another switch in the stack.
	Flashing green	The stacking transceiver is receiving or transmitting packets.

Show stack detail command

The **show stack detail** command provides a snapshot of the stack status. It includes a full summary of the status of all the stack members and the status of their connections to the other member. This allows you to check that all the stack members have active connections to each other and have recognized correctly which neighboring switch is connected to each of their stacking ports.

```
x900#show stack detail
Virtual Chassis Stacking detailed information
Stack Status:
-----
Operational Status                Normal operation
Management VLAN ID                4094
Management VLAN subnet address    192.168.255.0
Virtual Chassis ID                 2528 (0x9e0)
Virtual MAC address                Disabled
Disabled Master Monitoring         Enabled

Stack member 1:
-----
ID                                 1
Pending ID                         -
MAC address                        0000.cd29.95f7
Last role change                   Thu Mar 16 08:31:49 2017
Product type                       SwitchBlade x908
Role                               Active Master
Status                             Ready
Priority                            128
Host name                          awplus
S/W version auto synchronization   On
Resiliency link status             Not configured
Stack port1.0.1 status             Learnt neighbor 2
Stack port1.0.2 status             Learnt neighbor 2

Stack member 2:
-----
ID                                 2
Pending ID                         -
MAC address                        0000.cd28.5377
Last role change                   Thu Mar 16 08:33:22 2017
Product type                       SwitchBlade x908
Role                               Backup Member
Status                             Ready
Priority                            128
Host name                          awplus-2
S/W version auto synchronization   On
Resiliency link status             Not configured
Stack port2.0.1 status             Learnt neighbor 1
Stack port2.0.2 status             Learnt neighbor 1
awplus#
-----
```

Stack debug output

The clearest way to receive stack debug output is to enable console logging of VCStack messages. Using this method, you receive just the stack debug messages without the other messages that are seen when you enable terminal monitor.

To enable console logging of VCStack messages, use the command:

```
awplus(config)#log console program VCS
```

You must enter the parameter **VCS** in uppercase letters when you enter this command. This command enables you to receive a record of stack link and communication events.

For example if a stack port goes down, and then comes up again, the series of messages output is:

```
awplus#conf t
Enter configuration commands, one per line. End with CNTL/Z.
awplus(config)#conf tlog console program VCS
awplus(config)#exit
awplus#
awplus#08:37:01 awplus VCS[670]: Link down event on stack link port1.0.2
08:37:01 awplus VCS[670]: STK TRACE: Xbar update event is queued.
08:37:01 awplus VCS[670]: STK TRACE: Link port1.0.1: --> 2 (0000.cd28.5377)
08:37:01 awplus VCS[670]: STK TRACE: Link port1.0.2: N/A
08:37:01 awplus VCS[670]: STK TRACE: Stack topology has changed - updating stack H/W routes for L2 connectivity
08:37:02 awplus VCS[670]: STK TRACE: stack H/W route update complete

08:37:45 awplus VCS[670]: Link up event on stack link port1.0.2
08:37:45 awplus VCS[670]: Beginning neighbor discovery on link port1.0.2
08:37:45 awplus-2 VCS[669]: Link up event on stack link port2.0.1
08:37:46 awplus VCS[670]: STK TRACE: > Txd Hello msg on port1.0.2 (nbr ?):
08:37:46 awplus VCS[670]: STK TRACE: < Rxd Hello msg on port1.0.2 (nbr ?):
08:37:46 awplus VCS[670]: Already have one link between Member 1 and Member 2
08:37:46 awplus VCS[670]: STK TRACE: Xbar update event is queued.
08:37:46 awplus VCS[670]: STK TRACE: Link port1.0.1: --> 2 (0000.cd28.5377)
08:37:46 awplus VCS[670]: STK TRACE: Link port1.0.2: --> 2 (0000.cd28.5377)
08:37:46 awplus VCS[670]: STK TRACE: Ring topology is now complete
08:37:46 awplus VCS[670]: STK TRACE: > Txd Hello msg on port1.0.2 (nbr 2):
08:37:46 awplus VCS[670]: STK TRACE: > Txd Topology DB msg on L2: sender 1
08:37:46 awplus VCS[670]: Neighbor discovery on link port1.0.2 has successfully completed
08:37:46 awplus VCS[670]: STK TRACE: Stack topology has changed - updating stack H/W routes for L2 connectivity
08:37:46 awplus VCS[670]: STK TRACE: < Rxd Topology DB msg on L2: sender 2
08:37:46 awplus VCS[670]: STK TRACE: Received topology DB message from unconnected neighbor Member 2
08:37:47 awplus VCS[670]: STK TRACE: stack H/W route update complete

awplus#
```

If stack-link communication to a stack member is completely lost, the series of messages output is:

```
BlueCore#
BlueCore#01:04:55 BlueCore VCS[994]: STK TRACE: < Rxd Link-down msg on L2: 4 (00
15.77c9.
73cb) <--> 2 (0015.77e8.a87d)
01:04:55 BlueCore VCS[994]: Link between members 4 and 2 is down
01:04:55 BlueCore VCS[994]: STK TRACE: Link 1.0.1: --> 4 (0015.77c9.73cb)
01:04:55 BlueCore VCS[994]: STK TRACE: Link 1.0.2: --> 3 (0015.77e8.a892) --> 2
(0015.77e8.a87d)
01:04:55 BlueCore VCS[994]: STK TRACE: Stack topology has changed - updating st
ack H/W routes for L2 connectivity
01:04:55 BlueCore-3 VCS[996]: Link down event on stack link 3.0.2
01:04:56 BlueCore VCS[994]: STK TRACE: stack H/W route update complete
01:04:55 BlueCore-4 VCS[995]: Link down event on stack link 4.0.1
01:04:56 BlueCore VCS[994]: STK TRACE: < Rxd Link-down msg on L2: 3 (0015.77e8.
a892) <--> 2 (0015.77e8.a87d)
01:04:56 BlueCore VCS[994]: Link between members 3 and 2 is down
01:04:56 BlueCore VCS[994]: STK TRACE: Link 1.0.1: --> 4 (0015.77c9.73cb)
01:04:56 BlueCore VCS[994]: STK TRACE: Link 1.0.2: --> 3 (0015.77e8.a892)
01:04:56 BlueCore VCS[994]: Member 2 (0015.77e8.a87d) is leaving the stack (unr
eachable via stack links)
01:04:56 BlueCore VCS[994]: STK TRACE: Shutting down access to member 2's file
system
01:04:56 BlueCore VCS[994]: STK TRACE: Member 2 (0015.77e8.a87d) state Backup M
ember --> Leaving
01:04:56 BlueCore VCS[994]: Member 2 (0015.77e8.a87d) has left stack
01:04:56 BlueCore VCS[994]: STK TRACE: Stack topology has changed - updating st
ack H/W routes for L2 connectivity
01:04:56 BlueCore VCS[994]: STK TRACE: stack H/W route update complete
01:04:56 BlueCore VCS: Updating stack file system access...
01:04:57 BlueCore VCS: Shutting down access to member-2's file system
01:04:59 BlueCore VCS[994]: HA monitoring detected member-1 no change
01:04:59 BlueCore VCS[994]: HA monitoring detected member-3 no change
01:04:59 BlueCore VCS[994]: HA monitoring detected member-4 no change
01:04:59 BlueCore VCS[994]: HA monitoring detected member-2 left stack
01:04:59 BlueCore NSM[997]: Removal event on unit 2.0 has been completed
```

To see a very detailed log of stacking related events after they have occurred, and other VCStack debug information, use the command:

```
awplus#show stack full-debug [<1-8>]
```

Even if you had not been capturing stack log output at the moment an event occurred, you can still retrospectively obtain the logging information by using this command. If you do not specify a stack ID, then each stack member's output is displayed, one after the other.

This command produces a large amount of output, as shown in the following figure. The ["Stack debug output" on page 74](#) section contains the detailed log information.

BlueCore#show stack full-debug

Detailed debugging information for stack member 1

VCS host configuration

```
-----
Unit stackable, stack H/W present
VCS mgmt VLAN 4094, subnet 192.168.255.0
Topology: Ring
Neighbor port-1: 4
Neighbor port-2: 3
Stack Virtual MAC: feature disabled
```

ID	Mac Address	PP	LJ	Cfg	IP	Status	Role
> 1	0015.77c2.4b7d	00	N	Y	192.168.255.1	Ready	Active master
2	0015.77e8.a87d	00	Y	Y	192.168.255.2	Syncing	Backup Member
3	0015.77e8.a892	00	N	Y	192.168.255.3	Ready	Backup Member
4	0015.77c9.73cb	00	N	Y	192.168.255.4	Ready	Backup Member

VCS Inter-process connectivity configuration

```
-----
Type          Lower      Upper      Port Identity      Publication
-----
```

0	16781313	16781313	<1.1.1:14753793>	14753794	zone
	16781314	16781314	<1.1.2:3494871041>	3494871042	
	16781315	16781315	<1.1.3:3173122049>	3173122050	
	16781316	16781316	<1.1.4:1199587329>	1199587330	
1	1	1	<1.1.1:14761987>	14761988	node
9500	1	1	<1.1.1:14745638>	14745639	cluster
	2	2	<1.1.2:3494862887>	3494862888	
	3	3	<1.1.3:3173113895>	3173113896	
	4	4	<1.1.4:1199579174>	1199579175	
	100	100	<1.1.1:14753794>	14753795	node
	200	200	<1.1.1:14786564>	14786565	node
			<1.1.1:14745679>	14745680	node
			<1.1.1:14753852>	14753853	node
			<1.1.1:14745674>	14745675	node
			<1.1.1:14745672>	14745673	node
			<1.1.1:14745670>	14745671	node
			<1.1.1:14745663>	14745664	node
			<1.1.1:14876727>	14876728	node
			<1.1.1:14819381>	14819382	node
			<1.1.1:14745651>	14745652	node
			<1.1.1:15073328>	15073329	node
			<1.1.1:14991406>	14991407	node
			<1.1.1:14745630>	14745631	node
			<1.1.1:14745629>	14745630	node
			<1.1.1:14745625>	14745626	node
			<1.1.1:14753816>	14753817	node
			<1.1.1:14745623>	14745624	node
			<1.1.1:14770198>	14770199	node
			<1.1.1:14753813>	14753814	node
			<1.1.1:14745620>	14745621	node
			<1.1.1:14762003>	14762004	node
			<1.1.1:14762002>	14762003	node
			<1.1.1:14745617>	14745618	node
			<1.1.1:14762000>	14762001	node
			<1.1.1:14753807>	14753808	node
			<1.1.1:14753805>	14753806	node
			<1.1.1:14761996>	14761997	node
			<1.1.1:14761995>	14761996	node
			<1.1.1:14745609>	14745610	node

[Continued on next page]

			<1.1.1:14753802>	14753803	node
			<1.1.1:14770184>	14770185	node
			<1.1.1:14761991>	14761992	node
			<1.1.1:14761990>	14761991	node
			<1.1.1:14958597>	14958598	node
9501	1	1	<1.1.1:14770221>	14770222	cluster
	2	2	<1.1.2:3494862892>	3494862893	
	3	3	<1.1.3:3173163049>	3173163050	
	4	4	<1.1.4:1199579177>	1199579178	
9506	1	1	<1.1.1:14852155>	14852156	cluster
	2	2	<1.1.2:3494862922>	3494862923	
	3	3	<1.1.3:3173220410>	3173220411	
	4	4	<1.1.4:1199579196>	1199579197	
9513	1	1	<1.1.1:14753806>	14753807	cluster
	2	2	<1.1.2:3494879245>	3494879246	
	3	3	<1.1.3:3173130253>	3173130254	
	4	4	<1.1.4:1199587342>	1199587343	
9550	9550	9550	<1.1.4:1199710258>	1199710259	
			<1.1.2:3494862906>	3494862907	
			<1.1.1:14934066>	14934067	cluster
			<1.1.3:3173187633>	3173187634	
9551	9551	9551	<1.1.1:14745676>	14745677	cluster
			<1.1.2:3494862919>	3494862920	
			<1.1.4:1199579208>	1199579209	
			<1.1.3:3173113917>	3173113918	
9600	1	1	<1.1.1:14794792>	14794793	cluster
	2	2	<1.1.2:3494862893>	3494862894	
	3	3	<1.1.3:3173113898>	3173113899	
	4	4	<1.1.4:1199628328>	1199628329	
9601	1	1	<1.1.1:14745641>	14745642	cluster
	2	2	<1.1.2:3494862894>	3494862895	
	3	3	<1.1.3:3173113899>	3173113900	
	4	4	<1.1.4:1199595562>	1199595563	
9602	1	1	<1.1.1:14745643>	14745644	cluster
	2	2	<1.1.2:3494862896>	3494862897	
	3	3	<1.1.3:3173113901>	3173113902	
	4	4	<1.1.4:1199579180>	1199579181	
9603	1	1	<1.1.1:14745642>	14745643	cluster
	2	2	<1.1.2:3494862895>	3494862896	
	3	3	<1.1.3:3173122092>	3173122093	
	4	4	<1.1.4:1199579179>	1199579180	
9604	1	1	<1.1.1:14745644>	14745645	cluster
	2	2	<1.1.2:3494862897>	3494862898	
	3	3	<1.1.3:3173138478>	3173138479	
	4	4	<1.1.4:1199579181>	1199579182	
9034	1	1	<1.1.1:15351834>	15351835	cluster
	2	2	<1.1.2:3495411739>	3495411740	
	3	3	<1.1.3:3173662747>	3173662748	
	4	4	<1.1.4:1200439322>	1200439323	
9040	2	2	<1.1.2:3494871042>	3494871043	
	3	3	<1.1.3:3173122050>	3173122051	
9041	1	1	<1.1.1:14745657>	14745658	cluster
	2	2	<1.1.2:3494862915>	3494862916	
	3	3	<1.1.3:3173113935>	3173113936	
	4	4	<1.1.4:1199587385>	1199587386	

[Continued on next page]

```

Link <multicast-link>
  Window:20 packets
  RX packets:102 fragments:0/0 bundles:0/0
  TX packets:42244 fragments:0/0 bundles:0/0
  RX naks:118 defs:2 dups:115
  TX naks:4 acks:5 dups:231
  Congestion bearer:0 link:0  Send queue max:18 avg:5

Link <1.1.1:vlan4094-1.1.2:vlan4094>
  ACTIVE MTU:1500 Priority:10 Tolerance:3000 ms Window:50 packets
  RX packets:261 fragments:0/0 bundles:0/0
  TX packets:236 fragments:0/0 bundles:0/0
  TX profile sample:28 packets average:42 octets
  0-64:93% -256:7% -1024:0% -4096:0% -16354:0% -32768:0% -66000:0%
  RX states:115 probes:29 naks:0 defs:0 dups:0 tos:0
  TX states:130 probes:41 naks:0 acks:1 dups:0
  Congestion bearer:0 link:0  Send queue max:1 avg:0

Link <1.1.1:vlan4094-1.1.3:vlan4094>
  ACTIVE MTU:1500 Priority:10 Tolerance:3000 ms Window:50 packets
  RX packets:1401 fragments:0/0 bundles:0/0
  TX packets:1368 fragments:0/0 bundles:0/0
  TX profile sample:86 packets average:43 octets
  0-64:98% -256:2% -1024:0% -4096:0% -16354:0% -32768:0% -66000:0%
  RX states:7745 probes:1234 naks:0 defs:0 dups:0 tos:0
  TX states:10440 probes:3858 naks:0 acks:1 dups:0
  Congestion bearer:0 link:0  Send queue max:5 avg:0

Link <1.1.1:vlan4094-1.1.4:vlan4094>
  ACTIVE MTU:1500 Priority:10 Tolerance:3000 ms Window:50 packets
  RX packets:187 fragments:0/0 bundles:0/0
  TX packets:179 fragments:0/0 bundles:0/0
  TX profile sample:48 packets average:42 octets
  0-64:98% -256:2% -1024:0% -4096:0% -16354:0% -32768:0% -66000:0%
  RX states:7994 probes:1333 naks:0 defs:0 dups:0 tos:0
  TX states:10532 probes:4009 naks:0 acks:0 dups:0
  Congestion bearer:0 link:0  Send queue max:5 avg:0

VCS management traffic
-----
Mon May 17 01:09:05 UTC 2015

vlan4094 Link encap:Ethernet HWaddr 00:15:77:C2:4B:7D
  inet addr:192.168.255.1 Bcast:192.168.255.15 Mask:255.255.255.240
  inet6 addr: fe80::215:77ff:fec2:4b7d/64 Scope:Link
  UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
  RX packets:138575 errors:0 dropped:0 overruns:0 frame:0
  TX packets:205773 errors:0 dropped:0 overruns:0 carrier:0
  collisions:0 txqueuelen:0
  RX bytes:13436002 (12.8 MiB) TX bytes:23187142 (22.1 MiB)

          Pkts          Bytes
Non-VCS Q7: 0              0
Rx AIS:    21173         2559494
Rx mcast:  1912

VCS packet Replication
Type      Tx      Rx      Drops
Total     0       0       0
Bytes     0       0
STP       0       0       0
[Continued on next page]

```

```

EPSR          0          0          0
LACP          0          0          0
Mcast        0          0          0
sflow        0          0          0
BcExc        0          0          0
Other        0          0          0
PortAuth     0          0          0

```

TxBuff min: 952 86268ms ago

TxBuff cur: 1008

Last Rx: 11ms ago

Last Tx: 14ms ago

```

          CPU0
16:      4135  IPIC  Level Enabled  0    serial
17:     186739 IPIC  Level Enabled  0    linux-kernel-bde
22:         0  IPIC  Level Enabled  0    LM81 1
23:         0  IPIC  Level Enabled  0    LM81 2
24:        287 IPIC  Level Enabled  0    mpc83xx_spi
25:        749 IPIC  Level Enabled  0    i2c-mpc
26:         0  IPIC  Level Enabled  0    i2c-mpc
74:         0  IPIC  Edge  Enabled  0    GPIO
75:         0  IPIC  Edge  Enabled  0    GPIO

```

BAD: 0

VCS debug

```

-----
2015 May 16 23:39:23 kern.info awplus kernel: TIPC: Activated (version 1.6.4 com
piled May 4 2010 11:45:08)
2015 May 16 23:39:23 kern.info awplus kernel: TIPC: Started in single node mode
2015 May 16 23:39:25 user.info awplus VCS[960]: Parsing 'stack' commands from co
nfig file /flash/default.cfg
2015 May 16 23:39:25 kern.info awplus kernel: TIPC: Started in network mode
2015 May 16 23:39:25 kern.info awplus kernel: TIPC: Own node address <1.1.1>, ne
twork identity 4711
2015 May 16 23:39:31 user.debug awplus VCS[994]: SFL: Base feature license alloc
ated
2015 May 16 23:39:31 user.info awplus VCS[994]: SFL: [stackd] LicenceCheck: Virt
ual-MAC is active
2015 May 16 23:39:31 user.info awplus VCS[994]: SFL: [stackd] LicenceCheck: retu
rns Success.
2015 May 16 23:39:31 user.debug awplus VCS[994]: 2. init app stk-topo-event sess
Pt=0x100696e8, addr=0x4800e000
2015 May 16 23:39:31 user.debug awplus VCS[994]: 2. init app stk-port-1 sessPt=0
x10067480, addr=0x4800f000
2015 May 16 23:39:31 user.debug awplus VCS[994]: 2. init app stk-port-2 sessPt=0
x10069c90, addr=0x48010000
2015 May 16 23:39:31 user.debug awplus VCS[994]: 2. init app stk-topo-msg sessPt
=0x10069b40, addr=0x48011000
2015 May 16 23:39:31 user.debug awplus VCS[994]: 2. init app stk-topo-error sess
Pt=0x10069b58, addr=0x48012000
2015 May 16 23:39:31 user.debug awplus VCS[994]: 2. init app stk-resiliency-link
sessPt=0x10066418, addr=0x48013000
2015 May 16 23:39:31 user.debug awplus VCS[994]: STK DEV : Stacking Resiliency
Link counters register is successful.
2015 May 16 23:39:45 user.debug awplus VCS[994]: STK DISC : Member-1 XEMs presen
t: Bay 0: XEM-STK,
2015 May 16 23:39:45 user.info awplus VCS[994]: Parsing 'stack resiliencylink vl
an' commands from config file /flash/default.cfg
2015 May 16 23:39:45 user.info awplus VCS[994]: Stacking Ports were discovered o
n the mainboard of member 1

```

Counters

You can obtain detailed counters relating to stack events and signaling packets with the **show counter stack** command. You can use these for tracking down whether signaling packets are being lost, by checking if there are discrepancies between the number sent from one switch and the number received by its neighbor. The event counters make it possible to see if unexpected events have been occurring on the stack. The following figure is an output example:

```
BlueCore#show counter stack
Virtual Chassis Stacking counters

Stack member 1:

Topology Event counters
Units joined           ..... 4
Units left             ..... 1
Links up               ..... 8
Links down             ..... 4
ID conflict            ..... 0
master conflict        ..... 0
master failover        ..... 0
master elected          ..... 1
master discovered      ..... 0
SW autoupgrades        ..... 0

Stack Port 1 Topology Event counters
Link up                ..... 2
Link down              ..... 1
Nbr re-init            ..... 0
Nbr incompatible       ..... 0
Nbr 2way comms         ..... 1
Nbr full comms         ..... 2

Stack Port 2 Topology Event counters
Link up                ..... 1
Link down              ..... 0
Nbr re-init            ..... 0
Nbr incompatible       ..... 0
Nbr 2way comms         ..... 1
Nbr full comms         ..... 0

Topology Message counters
Tx Total                ..... 37
Tx Hellos                ..... 3
Tx Topo DB               ..... 2
Tx Topo update           ..... 2
Tx Link event            ..... 0
Tx Reinitialise          ..... 0
Tx Port 1                ..... 3
Tx Port 2                ..... 2
Tx 1-hop transport       ..... 5
Tx Layer-2 transport     ..... 32
Rx Total                 ..... 87
Rx Hellos                ..... 4
Rx Topo DB               ..... 2
Rx Topo update           ..... 10
Rx Link event            ..... 6
[Continued on next page]
```



```

Rx Reinitialise          ..... 0
Rx Port 1                ..... 2
Rx Port 2                ..... 2
Rx 1-hop transport      ..... 4
Rx Layer-2 transport    ..... 14

Topology Error counters
Version unsupported     ..... 0
Product unsupported     ..... 0
XEM unsupported         ..... 0
Too many units         ..... 0
Invalid messages       ..... 0

Resiliency Link counters
Health status good     ..... 0
Health status bad      ..... 0
Tx                     ..... 0
Tx Error               ..... 0
Rx                     ..... 0
Rx Error               ..... 0
Stack member 3:

Topology Event counters
Units joined           ..... 4
Units left             ..... 1
Links up               ..... 8
Links down             ..... 4
ID conflict            ..... 0
master conflict        ..... 0
master failover        ..... 0
master elected          ..... 0
master discovered     ..... 1
SW autoupgrades        ..... 0

Stack Port 1 Topology Event counters
Link up                ..... 1
Link down              ..... 0
Nbr re-init            ..... 0
Nbr incompatible      ..... 0
Nbr 2way comms        ..... 1
Nbr full comms        ..... 3

Stack Port 2 Topology Event counters
Link up                ..... 4
Link down              ..... 3
Nbr re-init            ..... 0
Nbr incompatible      ..... 0
Nbr 2way comms        ..... 2
Nbr full comms        ..... 0
[Continued on next page]

```

```

Topology Message counters
Tx Total                ..... 37
Tx Hellos                ..... 4
Tx Topo DB              ..... 3
Tx Topo update          ..... 3
Tx Link event           ..... 1
Tx Reinitialise         ..... 0
Tx Port 1               ..... 2
Tx Port 2               ..... 5
Tx 1-hop transport      ..... 7
Tx Layer-2 transport    ..... 30
Rx Total                ..... 87
Rx Hellos                ..... 6
Rx Topo DB              ..... 3
Rx Topo update          ..... 9
Rx Link event           ..... 5
Rx Reinitialise         ..... 0
Rx Port 1               ..... 2
Rx Port 2               ..... 5
Rx 1-hop transport      ..... 7
Rx Layer-2 transport    ..... 80

Topology Error counters
Version unsupported     ..... 0
Product unsupported    ..... 0
XEM unsupported         ..... 0
Too many units         ..... 0
Invalid messages       ..... 0

Resiliency Link counters
Health status good     ..... 0
Health status bad      ..... 0
Tx                     ..... 0
Tx Error               ..... 0
Rx                     ..... 0
Rx Error               ..... 0

Stack member 4:

Topology Event counters
Units joined           ..... 4
Units left             ..... 1
Links up               ..... 8
Links down             ..... 4
ID conflict            ..... 0
master conflict        ..... 0
master failover        ..... 0
master elected          ..... 0
master discovered      ..... 1
SW autoupgrades        ..... 0

Stack Port 1 Topology Event counters
Link up                ..... 5
Link down              ..... 4
Nbr re-init            ..... 0
Nbr incompatible       ..... 0
Nbr 2way comms         ..... 4
Nbr full comms         ..... 5
[Continued on next page]

```

```

Stack Port 2 Topology Event counters
Link up                ..... 1
Link down              ..... 0
Nbr re-init           ..... 0
Nbr incompatible      ..... 0
Nbr 2way comms        ..... 1
Nbr full comms        ..... 0

Topology Message counters
Tx Total               ..... 52
Tx Hellos              ..... 9
Tx Topo DB             ..... 5
Tx Topo update         ..... 4
Tx Link event          ..... 3
Tx Reinitialise        ..... 0
Tx Port 1              ..... 9
Tx Port 2              ..... 3
Tx 1-hop transport    ..... 12
Tx Layer-2 transport   ..... 40
Rx Total               ..... 76
Rx Hellos              ..... 9
Rx Topo DB             ..... 5
Rx Topo update         ..... 7
Rx Link event          ..... 3
Rx Reinitialise        ..... 0
Rx Port 1              ..... 10
Rx Port 2              ..... 3
Rx 1-hop transport     ..... 13
Rx Layer-2 transport   ..... 63

Topology Error counters
Version unsupported    ..... 0
Product unsupported    ..... 0
XEM unsupported        ..... 0
Too many units         ..... 0
Invalid messages       ..... 0

Resiliency Link counters
Health status good     ..... 0
Health status bad      ..... 0
Tx                     ..... 0
Tx Error               ..... 0
Rx                     ..... 0
Rx Error               ..... 0

```

Converting Stacking Ports to Network Ports

x610 Series: Reconfiguring AT-x6EM/XS2 stacking module ports as network ports

The AT-Stack module AT-x6EM/XS2 contains two 10 GbE SFP+ ports. By default, these ports are configured for stacking. However, you can reconfigure them as network switch ports. To do this, reconfigure the switch as a non-stacked switch, using the command **no stack <stack-ID> enable**, then reboot to apply this configuration.

In the non-stacked mode these ports will appear as configurable switch ports, even without the SFP+ AT-Stack module being inserted within the switch. When configuring these ports, you can identify them by their value 1 in the middle address component of the port number triplet. For example the port number 1.1.2 has the components: **1** (stack member identifier; always a 1 in non-stack mode) **.1** (the board identifier; always a 1 for the AT-Stack module) **.2** (the port number on the AT-Stack module: can have the value 1 or 2).

x930 Series: Reconfiguring AT-StackQS stacking module ports as network ports

On x930 Series switches, you can configure each port on the AT-StackQS card as:

- a stacking port, or
- one 40Gbps port, or
- four 10Gbps ports (28-port models only).

The ports are configured as stacking ports by default. When converted to network switch ports, they operate as 40Gbps ports by default.

To convert the ports to network switch ports, you need to disable VCStack on the ports. There are two options for doing this:

- make the switch into a standalone switch, by running the command:
no stack <stack-id> enable, or
- use the 10Gbps front-panel SFP+ ports for stacking, by running the command:
stack enable builtin-ports

Operating the ports as four 10Gbps ports

To use an AT-StackQS port as four 10Gbps ports, you need a AT-QSFP-4SFP10G-3CU or AT-QSFP-4SFP10G-5CU breakout cable. Then use the command **platform portmode interface** to change the port or ports to 10Gbps.

For example, to change both the stacking ports into 10Gbps ports, use the commands:

```
awplus#configure terminal
awplus(config)#no stack 1 enable
awplus(config)#platform portmode int port1.1.1,port1.1.5 10gx4
```

In 10Gbps mode, the ports are numbered as follows:

SLOT NUMBER	PORT NUMBER	BECOMES
1	1.1.1	1.1.1, 1.1.2, 1.1.3, 1.1.4
2	1.1.5	1.1.5, 1.1.6, 1.1.7, 1.1.8

When changing the portmode setting, you must also remove any interface and channel-group configuration from the specified ports, and then reboot the switch. Note that the AT-StackQS ports can only operate as four 10Gbps network switch ports on 28-port switch models, not on 52-port switch models.

DC2552XS/L3 Series: Reconfiguring the front QSFP+ ports as network ports

On DC2552XS/L3 series switches, the front 4 QSFP+ ports can be configured as stacking ports or network ports. By **default**, the ports are configured as **stacking** ports. Note that when stacking is enabled, all four QSFP+ ports become stacking ports, even if you only use one or two ports for stacking.

To convert stacking ports to **network** ports, VCStack needs to be disabled using the following command:

```
awplus(config)#no stack <stack-id> enable
```

When operating as network ports the four QSFP+ ports are numbered as follows:

```
Stack 1
Port 1.0.49, 1.0.53, 1.0.57, 1.0.61
```

Operating the QSFP+ ports as four 10Gbps ports

To use a QSFP+ port as four 10Gbps ports:

- a AT-QSFP-4SFP10G-3CU or AT-QSFP-4SFP10G-5CU breakout cable is required.
- the QSFP+ port must be configured for 10Gbps mode using the following command:

```
awplus(config)#platform portmode interface <port> 10gx4
```

In 10Gbps mode, the ports are numbered as follows:

PORT	BECOMES
1.0.49	1.0.49, 1.0.50, 1.0.51, 1.0.52
1.0.53	1.0.53, 1.0.54, 1.0.55, 1.0.56
1.0.57	1.0.57, 1.0.58, 1.0.59, 1.0.60
1.0.61	1.0.61, 1.0.62, 1.0.63, 1.0.64