# Cisco Catalyst 2960-S, 2960-X, and 2960-XR Stacking with FlexStack and FlexStack-Plus Technology: Description, Usage, and Best Practices 

## Introduction

Cisco ${ }^{\circledR}$ FlexStack and FlexStack-Plus stacking for Cisco Catalyst ${ }^{\circledR}$ 2960-S, 2960-X and 2960-XR Series Switches provides a true stacking solution, with all switches in a FlexStack or FlexStack-Plus stack acting as a single switching unit. FlexStack and FlexStack-Plus provide a unified data plane and single configuration for a group of stacked Cisco Catalyst 2960-S, 2960-X, and 2960-XR switches.

FlexStack and FlexStack-Plus lower the total cost of ownership with built-in 1: N redundancy, high availability, and preprovisioning, not available in standalone switches. High-availability features such as EtherChannel and FlexLinks will work across stack members, increasing uptime and network connectivity.

FlexStack-Plus is based on FlexStack. FlexStack was introduced on the 2960-S in 2010. FlexStack-Plus is being introduced in 2013 with the 2960-X and 2960-XR models of Cisco Catalyst Ethernet switches.

FlexStack-Plus is an improvement over FlexStack because it doubles the speed of the stack bandwidth and allows for more members to join the stack. The FlexStack-Plus bandwidth for a single stack member is 40Gbps compared to 20Gbps with FlexStack. With multiple stack members, the FlexStack-Plus bandwidth is 80 Gbps , and the FlexStack bandwidth is 40Gpbs. With FlexStack-Plus technology, eight 2960-X or XR members can be joined into a single stack. With FlexStack only four 2960-S members could be joined into a single stack.

FlexStack-Plus is backward compatible with 2960-S and FlexStack. Mixing 2960-X and 2960-S members in a stack is supported. When the stack membership is mixed with 2960-X and 2960-S models, FlexStack-Plus falls back to FlexStack capabilities of 40Gbps stack bandwidth and four members maximum per stack.

For the remainder of this document, the term 2960 will apply to all the 2960-S, 2960-X, and 2960-XR switches unless explicitly called out. Also, the term FlexStack will refer to FlexStack-Plus as well as FlexStack unless explicitly called out.

The differences between FlexStack and FlexStack-Plus are identified in Table 1.
For more information about the 2960-X and 2960-XR, see the Campus LAN Access Switches product pages on Cisco.com. "Campus LAN Access Switches on Cisco.com"

As part of the Borderless Network architecture, FlexStack and FlexStack-Plus allows scalable provisioning and management to deploy any number of switches quickly and easily.

## Why Stack Ethernet Switches?

Stacking Ethernet switches reduces the network administrator's total cost of ownership. The cost of maintaining the network is decreased because there are fewer devices to manage, and the network uptime is increased with builtin redundancy.

Figure 1 shows two ways to deploy Ethernet switches at the access layer. In each case four switches are connected to a Cisco Catalyst 6509 Switch. The deployment scenario on the left shows four switches that are colocated in the same wiring closet, but are not stacked. Each standalone switch on the left has two uplinks to the distribution layer switch that are combined in an EtherChannel group. The four switches on the right are stacked together with FlexStack (FlexStack links show in red). The stack has four uplinks to the distribution layer switch. Because FlexStack supports EtherChannel grouping of Ethernet interfaces across the stack members, the four uplinks are grouped together into a single EtherChannel group. All the benefits of EtherChannel grouping are retained when switches are placed into a stack.

Figure 1. Comparison of FlexStack and Nonstacked Configurations


Stacking Ethernet switches provides the network administrator with three major operational benefits:

- Single point of management: All switches in the stack are managed as one.
- Built-in redundancy and high availability: The high-speed FlexStack connections provide redundant communication for each stack member to every other member.
- Scalability to fit network needs: Installation of a new switch to the stack is easy. As the need for additional access ports grows, adding a new switch to an existing stack is easier and faster than adding a new standalone switch to the network.

The first operational benefit is fewer devices to manage. Multiple physical switches in a stack appear as a single logical switch. This eases management overhead because there are fewer devices in the network to manage. A single IP address is used to manage the logical switch. All manageable entities (for example, Ethernet interfaces and VLANs) on all physical switches can be configured and managed from the logical switch. The logical switch will appear as a single entity in the network. In a Layer 2 network, the logical switch will appear as a single spanningtree entity.

The second operational benefit that stacking provides is built-in redundancy and increased availability. Data path redundancy is built into the stacking architecture as there are two physical paths between any two stack members. Connecting stack members with the stacking cables provides data path redundancy for all stack members.
Stacking increases Ethernet switch availability by providing redundancy for both the physical switch and the uplink. Because different physical switches will connect to the upstream network, losing one switch or one uplink interface does not prevent connectivity to the network. Since the logical switch has multiple uplinks, the logical switch still has network connectivity because at least one uplink is still active.

In a logical switch, a single physical switch is acting as the stack master. The stack master manages all physical switches, including itself. If the master fails, another member automatically becomes the stack master following a well-documented election process (covered later). The configuration of the stack is preserved through a single switch failure or a reboot of all stack members. FlexStack provides $1: \mathrm{N}$ redundancy for the stack master, with the ability for any physical switch to back up the acting master. (See Figure 2.)

Figure 2. FlexStack with Redundant Links


Flexibility of FlexStack allows for modular stacking capability. Cisco Catalyst 2960 switches can be added to a stack at any time by adding the optional FlexStack module. As the demand for Ethernet access ports increases, Cisco Catalyst 2960 switches can be stacked together either to create a new stack or to add new Cisco Catalyst 2960 switches to an existing stack. This flexibility provides investment protection for the network administrator, who can add additional stack members as needed.

Adding members to a stack and replacing physical units are much easier when multiple switches are combined into a FlexStack group. Since the stack retains the configuration, there is no need to back up the configuration prior to removing the switch. Conversely, when a new member is inserted into the stack, the configuration of the stack is pushed to the new member. The network administrator need not explicitly recover the switch configuration since the stack took care of it.

## Stacking and Clustering

Stacking is not clustering. Clustering is a technology available on Cisco Catalyst fixed Ethernet switches that allows the network administrator to use a single public IP address to manage multiple physical switches. Clustering was created to assist network administrators by having a single point of management as well as in preserving valuable public IP addresses needed to manage the Ethernet switches. Stacking is much more than a single point of management and IP address preservation because it offers redundancy, availability, and ease of management. Cisco Catalyst 2960 supports clustering as well as stacking.

## What Are FlexStack and FlexStack-Plus?

FlexStack is the name of the stacking technology used by the Cisco Catalyst 2960-S Series of fixed Ethernet switches. FlexStack is specific to Cisco Catalyst 2960-S switches, and only Cisco Catalyst 2960-S Series Switches use FlexStack.

FlexStack-Plus is the name of the stacking technology used by the Cisco Catalyst 2960-X and 2960-XR Series of fixed Ethernet switches. FlexStack-Plus is specific to the 2960-X and 2960-XR Series of switches.

Both FlexStack and FlexStack-Plus require a specific external module in order to stack together. In addition to the FlexStack and FlexStack-Plus modules (Figure 3), a special FlexStack cable is also required to interconnect the modules. There are two different modules for FlexStack. The FlexStack module (PID: 2960S-STACK) is for 2960-S, and the FlexStack-Plus module (PID: 2960X-STACK) is for 2960-X and 2960-XR. These two modules are different sizes and cannot be inserted into the wrong 2960 model. Each FlexStack module supports two FlexStack ports. The FlexStack module is inserted into the rear of the Cisco Catalyst 2960 switch. Two FlexStack cables, inserted into the FlexStack module, provide data path redundancy for traffic flowing across the stack. Using FlexStack cables, the physical members of the stack form a ring, providing the built-in redundant data path for each member of the stack. FlexStack and FlexStack-Plus share the same stack cable types. You do not need different FlexStack cable types for FlexStack and FlexStack-Plus.

FlexStack and FlexStack-Plus are optional. In order for a 2960 switch to stack, a separate FlexStack or FlexStackPlus module (depending upon your exact 2960 model) must be inserted in the rear of the switch. Without this module, stacking cannot function. Figure 3 shows a FlexStack-Plus module to be inserted into the rear of a Cisco Catalyst 2960-X switch. All FlexStack module types are hot swappable.

Cisco Catalyst 2960 switches running FlexStack and FlexStack-Plus use a hop-by-hop method of transferring Ethernet packets across the stack. Packets traverse the stack by going from one member to another over the FlexStack links until the packet reaches its destination. This is a behavior similar to that of multiple standalone Ethernet switches forwarding packets from one switch to another.

Switches in a stack communicate with each other using the FlexStack protocol. The FlexStack protocol allows the FlexStack members to behave as one logical switch. The FlexStack protocol runs on each member. The FlexStack protocol is used by the master switch to manage the other members. The master tracks the presence of each member, the member's availability, the Cisco $I O S^{\circledR}$ Software image, and the status of each member's FlexStack connections. The FlexStack protocol is the same for FlexStack on the 2960-S models and the 2960-X and XR models. Because it is the same protocol for all models of 2960, the different 2960 model types can be stacked together.

Figure 3. FlexStack Module Inserted into Rear of 2960-X


A single FlexStack connection between two 2960 stack members is a full-duplex 10Gbps connection. A single FlexStack-Plus connection between two 2960-X members is full-duplex 20Gbps. Each Cisco Catalyst 2960-S and 2960-X and 2960-XR member supports two FlexStack connections, which doubles the line rate. Each FlexStack member can simultaneously send and receive Ethernet traffic over both stack links at line rate, effectively giving 40Gbps of stack bandwidth per 2960-X member and 20Gbs of stack bandwidth per 2960-S member. When two or more 2960-X members are stacked together, the effective stack bandwidth is 80Gpbs because each member is capable of sending and receiving 40Gbps simultaneously. Same for 2960-S, the effective stack bandwidth is 40Gbps because multiple members are sending and receiving 20Gbps simultaneously. (See Table 1.)

Table 1. Comparison of FlexStack-Plus, FlexStack and StackWise-480

| Stacking Item | C2960-X | C2960-XR | C2960-S | 3850 |
| :---: | :---: | :---: | :---: | :---: |
| Bidirectional stack bandwidth | 80Gbps | 80Gbps | 40Gbps | 480Gbps |
| Hot-swappable stack module | Yes | Yes | Yes | No (built-in) |
| Max stack members | 8 | 8 | 4 | 9 |
| Supports FlexStack | Yes | No | Yes | No |
| Supports FlexStack-Plus | Yes | Yes | No | No |

## Mixed Stacks: FlexStack and FlexStack-Plus

The 2960-S and 2960-X support a mixed stack of 2960 models. All these 2960 models run the FlexStack protocol, allowing them to be stacked together into a single stack. The exception is the 2960-XR does not stack with either the 2960-X or the 2960-S. The 2960-XR has the IP-Lite feature set, and this is not compatible with the LAN Base feature set on the 2960-X and 2960-XR. Table 2 shows all the allowed mixed stack combinations. What is crucial is the Cisco IOS Software feature set. All 2960 models with the LAN Base Cisco IOS Software feature set can stack together.

Table 2. Allowed Mixed Stack Combinations

| Mixed Stack Combination | 2960-XR IP Lite | 2960-X LAN Base | 2960-S LAN Base |
| :--- | :--- | :--- | :--- |
| 2960-XR IP Lite | Yes | No | No |
| 2960-X LAN Base | No | Yes | Yes |
| 2960-S LAN Base | No | Yes | Yes |

Table 3 shows FlexStack-Plus backward compatibility and performance of 2960-X and 2960-XR with the 2960-S switches.

Table 3. Backward Compatibility

| Scale of Mixed Stack Combination | Max Stack Bandwidth | Stack Limit | Cisco IOS Software Feature Set |
| :--- | :--- | :--- | :--- |
| 2960-XR IP Lite | 80 Gbps | 8 | IP Lite |
| 2960-X LAN Base | 80 Gbps | 8 | LAN Base |
| 2960-X, 2960-S LAN Base | 40 Gbps | 4 | LAN Base |

When the 2960-S and 2960-X members are stacked together, the entire stack (even the 2960-X members) fall back to FlexStack capabilities. Mixing 2960 members limits the max stack members to four, and 20Gbps stack bandwidth per member, and 40Gbps per stack.

## FlexStack Cabling

Special cables are used to connect Cisco Catalyst 2960 Series Switches together in a stack. These special FlexStack cables can only be used to stack Cisco Catalyst 2960 switches with other Cisco Catalyst 2960 switches that support FlexStack. Figure 4 shows a picture of two FlexStack cables connected to a fully inserted FlexStack module. Notice how the tabs on the cables are on opposite sides of each other. The FlexStack cables are keyed as well. Besides the tabs needing to be away from each other, the keying of the metal portion of the connector prevents the cable from being connected incorrectly.

Figure 4. FlexStack Module with Cables


When each member in the stack has two operational FlexStack links, then the stack is operating in a fully redundant mode. The dual FlexStack connections from each member to two other members are what provide the redundancy. Should any single FlexStack connection break or cease to operate, then the switches in the stack will use the remaining FlexStack connection that is provided.

## Deployment Topologies

Figure 5 shows a three-member stack with full bandwidth and with redundant FlexStack connections.
Figure 5. Fully Redundant Three-Member Stack


When one FlexStack link is not present, the stack still functions. Figure 6 shows a stack with incomplete FlexStack cabling. In Figure 6, all data traffic passes through the middle member. This stack is operating in a nonredundant mode. This stack provides only half of the possible bandwidth between members and does not have redundant connections. Only the middle member has full stack bandwidth. The top and bottom members are operating at half the fully redundant stack bandwidth. The fully redundant stack bandwidth for 2960-X and 2960-XR is 40Gbps per member. The fully redundant stack bandwidth for 2960-S is 20Gbps per member.

Figure 6. Nonredundant Three-Member Stack


Three different FlexStack-Plus cable lengths are available, as shown in Table 4.
Table 4. FlexStack Cable Lengths

| Product ID | Length | Description |
| :--- | :--- | :--- |
| CAB-STK-E-0.5M | 0.5 meters | This is the default cable that ships with the FlexStack or FlexStack-Plus module. |
| CAB-STK-E-1M | 1.0 meters |  |
| CAB-STK-E-3M | 3.0 meters |  |

The different lengths allow for deployment flexibility. Figure 7 shows how to stack four switches using only the 0.5meter cables. The stack member connections are interweaved using the stack cables. No cable extends more than two stack members. The interweaving of the stack links still provides redundant connections for the stack.

Figure 7. Four-Member Stack with 0.5 -Meter Cables


Figure 8 shows the conventional cabling deployment. The 3.0-meter cable is used to complete the redundant ring connection by connecting the top member with the bottom member. The other connections connect to directly adjacent members using the 0.5 -meter cables. The 3.0 -meter cable is not the default cable shipped for FlexStack and must be ordered. See the Cisco Catalyst 2960-S, 2960-X, or 2960-XR data sheet for ordering information.

Figure 8. Four-Member Stack with 3.0-Meter Cable


The 0.5-meter cable can be used to connect two Cisco Catalyst 2960 switches that are 4 rack units away from each other. In a stack of four, with all four members racked directly on top of each other, 0.5 -meter cables can be used to connect all stack members. When stack members are more than 4 rack units apart, then longer stack cables are required.

## FlexStack Protocol

FlexStack protocol allows every stack member to be in constant communication with the adjacent member and with the stack master. Each member is aware of the operational state of every stack port in the stack. FlexStack protocol is used to detect new member additions as well as member removal. The operational status for all stack members and their interfaces is communicated to each member through the FlexStack protocol.

The 2960-X with FlexStack-Plus has better reaction time to changes in the operation state of the FlexStack links than the 2960-S with FlexStack. The 2960-X has special hardware that is capable of detecting the change in the stack port operational state. The special hardware is capable of changing how packets are forwarded across the stack. This is the FlexStack link recovery time. Because it is being done in hardware, it is very fast. For packets traversing across the FlexStack-Plus stack, the traffic recovery time is 100 ms or less.

The 2960-S with FlexStack is not as fast. The recovery time for loss of a FlexStack link is 1 to 2 seconds. On the 2960-S with FlexStack, the change in stack port operational state is managed by the CPU in software. Because it is in software, the operational state change has to be processed, and the software has to reprogram the forwarding logic to forward traffic around the now nonoperational FlexStack link.

## Role of the Stack Master

The stack master controls the configuration and is the central point for management. All Layer 2 protocol traffic (for example, VLAN Trunking Protocol [VTP], Dynamic Trunking Protocol [DTP], Cisco Discovery Protocol, and Link Layer Discovery Protocol [LLDP]) is forwarded to the master regardless of where the protocol packet ingresses. The master will also transmit all Layer 2 protocol packets. If the egress interface is on another member, then the protocol packet is passed from the master along the stack interfaces to the destination member.

The same behavior is seen for management traffic. Simple Network Management Protocol (SNMP), HTTP, Telnet, and Secure Shell (SSH) Protocol type of management traffic is all forwarded to the master regardless of the ingress interface. The response from the master is sent through the stack to the destination interface.

On configuration changes, the stack master pushes a copy of the configuration to every member. This way all members have a copy of the saved configuration.

Stack upgrades occur on the master. The master pushes new Cisco IOS Software images to all members. Each member stores a copy of the Cisco IOS Software image on its local Flash.

## FlexStack LED Operation

The LEDs on the front of the Cisco Catalyst 2960-S, 2960-X, and 2960-XR can be used to view stack operation. When multiple Cisco Catalyst 2960 switches are stacked together, only the stack master will have the "MSTR" LED solid green. Other members of the same stack will have this LED dark.

To find the switch member number, press the mode button until the stack LED goes green. At this point the LEDs over the Ethernet ports are used. The port number matching the member will start to blink. For member 2, the LED for port 2 will blink. For member 3, port 3 will blink. Depending upon the number of members in the stack (up to four for 2960-S and up to 8 for 2960-X), the other ports will go solid green. This allows the network administrator to see the number of stack members just from the LEDs over the ports.

When a switch is operating standalone, the MSTR LED will be solid green.

## FlexStack Details

## Inserting the FlexStack Module

The FlexStack module will insert easily into the slot. Just before it is fully inserted, there will be "minor" resistance. This is normal. Push the module completely into the switch chassis. Tighten the screws no more than "finger" tight. The screws are meant to keep the module from slipping out, not to maintain stack link connectivity. Finger tight is sufficient to maintain a fully operational stack connection.

## FlexStack Redundant Power

Redundant power for each Cisco Catalyst 2960-S and 2960-X is provided by the Cisco Redundant Power System 2300 (Cisco RPS 2300). The 2960-XR has its own dual redundant power supplies and does not support the RPS2300. If the RPS2300 is being used to provide redundant power, then each stack member must have its own connection to the Cisco RPS 2300.

## FlexStack Operational Management

Managing FlexStack switches involves adding members to and removing members from an existing stack. This section will cover techniques for safely and efficiently adding members to and removing members from a stack of Cisco Catalyst 2960-S or 2960-X switches.

## FlexStack Operations

To see FlexStack operation and member status, use the Cisco IOS Software CLI command "show switch." This command gives the basic information, including number of members, which member is the master, and state of all members. The following is a sample output of the command for a stack of 2960-S members:


The preceding example shows that member 1 is the master. The master has stack priority of 14 . All other members have the default stack priority of 1 . The state of all members is ready.

Other possible stack member states are:

- "progressing": booting up.
- "mismatch": a hardware or software mismatch is preventing the member from fully joining.
- "provisioned": this member has been configured, but it is not part of the stack and has never been part of the stack.
- "removed": if there are stack members that are not currently powered on or were removed from the stack, these members would have the state "removed."

The MAC address of the master is used as the MAC address of the logical stack. This is the base MAC address from the master, and it will be used for all IP communications, as well as Cisco Discovery Protocol data exchange and Spanning Tree Protocol.

To view other operational aspects of the stack, these commands are useful:

- Show version: view the Cisco IOS Software version of all stack members. Useful if there is a mismatch.
- Show switch neighbors: view which members are connected to which stack ports.
- Show switch stack-ports: state of the stack ports on each member.

To determine stack port connectivity and operational status, use the commands "show switch stack-ports" and "show switch neighbors." Refer back to Figures 5 and 6 . One shows a three-member stack with a full redundant mode. The other shows a stack in nonredundant mode. Here is what the show commands will show for each stack, respectively.

Fully redundant three-member stack (Figure 5):


Nonredundant three-member stack (Figure 6):


The master of the stack will be the one management interface for the entire stack. All IP communications (that is, Telnet, SSH, SNMP, ping, HTTP, and so on) will be handled by the master regardless of the member the communication ingressed the stack. All console activity regardless of media (RJ45 or USB) is directed to the master regardless of the member to which the console port is connected.

The Cisco Catalyst 2960-S, 2960-X, and 2960-XR support the CISCO-STACKWISE-MIB for managing switch stacks. This is the same MIB that the Cisco Catalyst 3750, 3750-E, and 3750-X Series Switches support.

## Making a Specific Switch the Stack Master and Master Election Rules

## Stack Election Rules

When multiple physical members are present in the stack, one of the members will be elected as the stack master. There can only be one master and up to three backups 1:3.

The election rules used by FlexStack in choosing a master are applied in this order:

1. Switch that is the current master
2. Switch with higher priority
3. Switch that has configuration file
4. Switch with longest uptime
5. Switch with lowest MAC address block

When new switches join an existing stack, then rules 1 and 2 will be applied.
When two standalone switches are joined together, having no switch priority configured, and both having existing configuration files, then the switch with the longest uptime will be chosen.

To have a specific member elected as master, two methods can be used.

1. Set the priority of the desired member to be higher than all other members. This is rule 2.
a. Example of configuring member 1 to have priority 14 :

C2960-48 (config) \#switch 1 priority 14
b. A reboot of the current master will be required after this to change the active master. Just configuring the priority will not by itself change the active master. Save the configuration before rebooting.
2. When the switches boot, and after they have completed power-on self-test (POST), there is a delay of approximately 2 minutes while the switches wait for all switches that could be in the stack to boot. This 2minute delay is called the election window. Network administrators can take advantage of this window by first powering up only the desired master switch. Once the election window closes, this switch will become the master of a stack of one. Any subsequent switches that boot after the first switch will become members of the existing stack

The switch with the lowest MAC address is used to break the tie in those scenarios in which the other rules cannot be applied.

It is desirable for network administrators to make a specific switch in the stack the master. This could be the topmost switch or the bottommost switch in the stack. For switch identification purposes, knowing how the members are connected is an advantage. If a switch needs to be replaced, then it is very convenient to know where that switch is located in relation to the other stack members.

Note: All Cisco Catalyst 2960-S, 2960-X, and 2960-XR LAN Base switches are equally capable of becoming the stack master. Physical place in the stack does not affect switch behavior.

For brand-new switches, with no configuration, all powered at the same time, the switch with the lowest MAC address will be the master. It is not practical to orient and deploy switches in a stack based on MAC address. To create a new stack with the correct physical member as the master, the network administrator must use the "uptime" rule. To reliably make any new switch the master, it must be powered up first. Subsequent switches should not be powered up until the election window is complete and the first switch has become the master. This is about 2 minutes, or when the system LED and the MSTR LED turn solid green.

To change the master after the stack of Cisco Catalyst 2960 switches has been created and set up, use the global configuration command "switch X priority Y ," where X is the member number that is desired to be the master, and $Y$ is the priority value to use. To make member $X$ the master, you must assign it a priority value higher than that of the current master.

By default, all switches have priority 1. To view switch priority, use the "show switch" command:


To make sure that member 1 will become the master, the priority should be changed to any value greater than 1 .
Here is an example showing member 1 assigned priority 14:

```
C2960-48#config term
Enter configuration commands, one per line. End with CNTL/Z.
c2960-48(config)#switch 1 priority 14
```

If a switch other than the master is given higher priority, as shown in the preceding command, there will not be an immediate switchover of the master. The master will not change until the current master is rebooted, forcing a stack master election of remaining members.

Setting the priority is important when merging two stacks. Existing stacks are merged when the FlexStack cable from one stack is connected to another stack. It only takes one cable to merge two stacks. When two existing stacks are merged, a master election takes place. When two stacks merge, both stacks already have a stack master, so the first rule is not applied. The best way to determine which stack master will continue to be master of the new combined stack is with the priority value.

Note: A standalone switch is still a stack of one member. Adding a standalone switch that has been powered up to an existing stack is a still a stack merge. If the standalone switch has a higher priority than the master of the other stack, then the standalone switch will become the master of the combined stack. During a stack merge the new master is chosen during the election process. The election process is used during a stack merge and at stack power-on time.

As always during a stack merge, new members are rebooted after the election process has completed. Existing members are not booted.

## Configuring a Backup Master

It is a good practice to configure another switch to be the backup should the primary master fail. Configuring a backup master enables a predictable master election in case of primary master failure. To configure a secondary (or backup) master, assign a priority value less than that of the primary master but greater than that of the other members. In the example just given, the primary master has priority 14 , and all other members have priority 1 . Any member will be secondary master with a value greater than 1 and less than 14.

## Determining the Member Number

For brand-new switches the order in which switches are powered up initially is also the order the member numbers will be assigned. Brand-new switches do not have a member number and default to member 1 . As new switches join an existing stack, the new member is assigned the next available member number. When the second switch is powered on, it is assigned member number 2. When the third switch is powered on, it is assigned member number 3. And so on.

If new members are not powered on in any order, then member numbers can appear to be randomly assigned. Understanding FlexStack connections between members can potentially be confusing when the member numbers are not in an order that matches physical layout.

It is possible to change member numbers of switches in a stack. It is not required for stack operation. The stack will operate just fine when member numbers do not match the physical order of the stack members. In spite of this, it is desirable to have member numbers match the physical layout to ease managing the stack and its members.

Use the Cisco IOS Software CLI to change member numbers. The switches must be rebooted for the member number change to take effect. The following example will renumber member 4 to member 3:

```
Switch(config#) switch 4 renumber 3
```

Caution: Changing member numbers causes any configuration on the interfaces of the old member to be lost. All interface configurations are based on the member number. When the member is changed, the configuration for current members is NOT changed or copied to the new member number. In the preceding example, the configuration for all of the Ethernet interfaces on member 4 will be lost, and the configuration for member 3 is unknown. Take care when changing member numbers and be prepared to reconfigure the interfaces.

Each member retains its own member number through a reboot or a power cycle. When a member boots into a stack, the member announces its member number to the stack. If there is no conflict, the member number is retained. If there is a conflict with two different physical members having the same number, then the first member to join the stack retains its number, and the newest member must change its number. The next available member number is assigned to the switch in such a conflict situation.

Since members retain their member number through a reboot, when a member moves from one stack to another, the potential for member number conflict exists. The ability to modify the member number is critical to resolve member number conflicts.

## Adding Members to an Existing Stack

The relative ease with which new members can be dynamically added to an existing stack is a real strength of FlexStack. A brand-new switch added to an existing stack will join as a member. The new member will be given the next available member number. When a new member is added to a stack, it will be rebooted to clear any configuration it might have. When the new member completes the reboot, it will have the configuration of the stack.

It is recommended that FlexStack operate in a fully redundant mode. In order to add a member to an existing stack, at least one FlexStack link must be disconnected, causing the existing stack to operate in nonredundant mode. This disconnection will cause a minor traffic outage across the stack. Traffic will resume while the stack is in nonredundant mode. The network administrator does not need to connect the FlexStack link in a fully redundant mode for traffic to flow. When the FlexStack links return to a fully redundant mode, there will be another minor traffic outage across the stack. Network administrators should try to limit the number of FlexStack link operational changes to decrease the amount of traffic disruption.

Cisco Catalyst 2960-S, 2960-X, and 2960-XR switches do not need to be powered off to join a stack. However, when a new member joins an existing stack, the new member will be rebooted to clear its configuration anyway. The new member is automatically rebooted by the stack. A reboot is operationally equivalent to powering a switch off and then on again. Thus, it is a better practice to add new stack members to an existing FlexStack powered off. Connect the FlexStack cables, and then power on the new member.

The following are steps for adding a new member to an existing stack (Figure 9), and it is desired that the new member not become the master. In a rack deployment, new switches are typically added at the top or the bottom of an existing FlexStack. This example will show how to add switches to the bottom of an existing stack of switches.

1. Place the new switch in its final location. Do not insert the power cord. To minimize downtime, complete all cabling that would be required of any switch placed into production at this time.
2. Insert the FlexStack module in the back of the new member.
3. Disconnect the FlexStack cable from the existing member just above to the new member. This would be the cable that is connecting the bottom member with the top member. You will need this to connect the new member located at the bottom of the stack to the member at the top.
a. When the FlexStack cable is disconnected, this is when the stack enters into a nonredundant mode.
4. Connect the new member with a FlexStack cable to the member just above. The FlexStack port used to connect one member to another is not important. Stack members should be connected to each other any way that makes cable management easier for the network administrator.
5. Connect another FlexStack cable to the member at the top. Make sure that all cables are connected properly. A light tug on the FlexStack cable will make sure that it is properly seated.
6. Connect the AC power cord to the new member. It will power on. It will take a few minutes for the new member to complete the boot process.
7. The new member will join the existing stack as a member, not the master.
a. When the new member completes the boot process, this is when the stack enters into a fully redundant mode once again.
8. All the FlexStack links will be operational if all the connections were made properly.

Figure 9. Adding a New Member to an Existing Stack

1. Existing 3 member stack



New member
placed into final
position

2. Preparing to add new member

Disconnect stack cable from bottom member to make room for new member




5. Addition of new member is complete

| New member at |
| :--- |
| bottom with |
| FlexStack Cables |
| connected to top |
| member and |
| member just above |




member just above


If all went well, traffic on the exiting members was minimally interrupted. There were at least two interruptions. The first occurred when the FlexStack cable was removed in step 3. The other interruption occurred when the new switch joined the stack after being powered on (step 7). In between the two interruptions, traffic will continue to be switched over the existing members and the operational FlexStack links.

## Removing Members from an Existing Stack

FlexStack remembers all members that have been members. When a member is physically removed or just powered down, FlexStack still considers that switch a member, just not currently active. The reason that the stack remembers members that have been removed is because most of the time the member comes back. For example, when a member is powered off, then powered back on, the member will rejoin the existing stack, and all the member-specific configuration will be reapplied when it completes the boot process. This is the most common occurrence.

If the removal is permanent, then it is recommended that the stack configuration be "cleaned up." The member will be deleted from the stack configuration. Here is an example of a command to remove stack member 4 from the configuration after it has been physically removed from the stack:

```
C2960-48(config)#no switch 4 provision
```

If the network admin chooses not to "clean up" the configuration after a member has been removed, the configuration will remain the running-config and the startup-config for the stack. This extra configuration will make the configuration file "config.text" larger than it should be.

In a switch replacement situation, replacing a member switch with an identical model is very convenient with FlexStack stacking. FlexStack remembers the member model number, but not the MAC address. During a switch replacement, if the old member is replaced by the same Cisco Catalyst 2960-S, 2960-X, or 2960-XR model, when the new switch is inserted into the FlexStack stack, it will inherit the configuration of the previous member. If the new switch model does not exactly match the replaced switch, the old configuration is not used. If there is room in the stack for a new member, then the new switch is added to the stack. If the stack size is at the maximum, then the configuration for the replaced member is removed.

Replacing a master switch needs one additional item to be clarified. The switch priority (configured with the command "switch X priority Y ") is not stored in the running-config or the startup-config. When the replacement switch is placed into the stack, it will not inherit the switch priority. The new switch will have the default priority of 1. To make the new switch the master, the network administrator must configure the new switch to have a higher priority than any other member. The existing master must be rebooted (or the entire stack must then be rebooted) after this configuration to force a master change. Just configuring the priority to the new member is not enough to trigger a master change. If the entire stack was rebooted, the newly added switch with the highest priority will be the stack master. If only the master is rebooted, the existing members will have a master election. The switch with highest priority will become stack master. The old master will reboot and join the stack as a member.

## Upgrading the Cisco IOS Software Release on a Stack

FlexStack has the capability to upgrade all members with a single command. A single command eases the operational burden on the network administrator. The command to initiate an upgrade takes place on the master. The master makes sure that each member is properly upgraded.

The Cisco IOS Software CLI command to upgrade a FlexStack stack is:
"archive download-sw [options] <source-file>."
The default behavior is to upgrade all members.
The options for the 'archive download-sw' command allow for detailed control of the upgrade. Use the options to update only a single member instead of the entire stack or to reload the stack after the upgrade is successful instead of having the reboot be an extra step for the network admin.

The "source file" is a tar file downloaded from Cisco.com called an image file. The tar file can be copied over the network using Trivial File Transfer Protocol (TFTP) or FTP. The most common and recommended way to upgrade the stack of switches is using TFTP. The network admin should copy the tar file with the desired Cisco IOS Software image to a TFTP server in the network. In the 'archive download-sw' command, use the IP address of the TFTP server and the name of the image file as the source file. No destination file name is necessary.

Here is an example of using the 'archive download-sw' command to upgrade a stack of Cisco Catalyst 2960-S switches. Notice that the number of members in the stack is not relevant to this command. By default all members are upgraded unless the network admin uses the options to do otherwise. The TFTP IP address is 192.168.1.1, and the tar file name is c2960s-universalk9-tar.122.53.SE.tar.

[^0]Here is another example showing how to update a FlexStack stack that has a mix of 2960-S and 2960-X models. The same "archive download-sw" command is used. This time, two Cisco IOS Software images are identified on the command line. Since the 2960-S runs a different Cisco IOS Software image from the 2960-X, two different images are required. Even though two separate images are required, it is recommended that both images be on the same release. Keeping members on the same release is important. In a mixed stack of 2960-S and 2960-X members, if only one model type is updated, you run the risk at next stack reload of the stack not forming properly because of "mismatch".

```
C2960-48h# archive download-sw tftp://192.168.1.1/c2960s-universalk9-tar.150-
2.SE3 tftp://192.168.1.1/c2960x-universalk9-tar.150-2.SX
```

If a network connection is not possible, the stack can be upgraded by reading the image file(s) from a USB Flash drive. A USB Flash drive must be inserted in the USB A slot on the front of the switch. The tar file containing the Cisco Catalyst 2960-S or 2960-X Cisco IOS Software image must be on the Flash drive. The USB Flash drive is referenced by the member into which it was inserted. If the USB Flash drive was inserted into member 2 , then it will be referenced as "usbflash2:." Use the same download command used in the preceding example, except this time point to the USB Flash instead of the TFTP server:

```
C2960-48# archive download-sw usbflash2:c2960s-universalk9-tar.150-2.SE3
usbflash2:c2960x-universalk9-tar.150-2.SX
```

The Cisco IOS Software CLI command "archive download-sw ..." is very flexible and yet simple to use. It can retrieve source images from multiple locations (for example, FTP, TFTP, Flash). The command allows for detailed control of the upgrade process. For more information about the "archive download-sw" command, see the Cisco Catalyst 2960-X software configuration guide on Cisco.com.

## Preprovisioning Members

FlexStack allows for preprovisioning of future members. This allows the network administrator to configure Ethernet interfaces for members that have yet to join. In those instances where the network administrator knows how certain interfaces will be configured, the admin can do so in advance to ease the insertion of the new member into the stack and into the network. This is useful when all interfaces that will connect to end users will have the same configuration for every stack member. When new members are added to the stack, if they have been preprovisioned, then it is not necessary for a network administrator to return to the stack and configure the new interfaces.

As an example, the network administrator knows that interfaces 49 and 50 on a future member will be connecting to another Ethernet switch and will be used as uplinks. Interfaces 1-48 will be connecting to end-user desktops. The network administrator can make sure that VLANs will match the upstream switch by configuring the Ethernet interfaces 49 and 50 for the proper VLAN membership. The network administrator preprovisions the member number by matching the member number with a supported switch model. Ideally the admin would choose the next available member number, but can choose any available member number.

This example shows how to add a fourth switch to an existing stack of three. The next available member number is 4 based on the output from "show switch." After preprovisioning the member and matching the model, the admin can configure interfaces $\mathrm{Gi} 4 / 0 / 1-\mathrm{Gi} 4 / 0 / 48, \mathrm{Gi} 4 / 0 / 49$, and $\mathrm{Gi} 4 / 0 / 50$. The following are the CLI commands to implement this preprovisioning. After the switch has been preprovisioned, all interfaces on the future member can be configured.


```
C2960-48(config-if-range) # switchport mode trunk
c2960-48(config-if-range)# exit
C2960-48(config)# !!!!!
C2960-48(config)# !! now configure the access ports.
C2960-48(config)# !! quick example of configuring interfaces
C2960-48(config) # !!
C2960-48(config)# interface range Gi4/0/1 - 48
C2960-48(config-if-range)# switchport mode access
C2960-48(config-if-range)# switchport access vlan 20
C2960-48(config-if-range) # spanning-tree portfast
C2960-48(config-if-range)# exit
C2960-48(config) #exit
C2960-48#
```


## Recommended Deployment Strategies for Building FlexStack

These are the important items to consider when deploying FlexStack:

- Configure the stack so that the uplinks are present on at least two stack members.
- Use cross-stack EtherChannel to load balance the traffic on the uplinks.
- If there are only two physical links, and EtherChannel is not desired, then use the feature called backup interface.
- When using 10 Gigabit Ethernet as the uplinks, separate the 10 Gigabit Ethernet uplinks across stack members if possible.
- It is always recommended to have at least two uplinks from the stack to the aggregation device.
- This is standard best practice for any deployment regardless if stacking is present or not. It is recommended for physical link redundancy.
- It is recommended that the stack master not contain uplinks if it is possible to distribute the uplinks across other stack members.
- This recommendation should only be followed in stacks with three or more members.
- If the master fails and it has the uplinks, then those uplinks are lost as well. If the master fails, and it does not have any uplinks, this will cause a minimal amount of traffic loss.
- Configure one member to be the backup master. Use the Cisco IOS Software CLI configuration command "switch <member number> priority <1-15>" to configure a backup. In case of master failure, having a deterministic backup is useful.
- When using only two uplinks in a stack of four members, put the uplinks on members that are not directly connected with FlexStack links. Directly connected members are those connected together with FlexStack connections. This is only relevant on FlexStack groups that contain four members.
- This recommendation is purely for stack bandwidth efficiency. Placing the FlexStack uplinks on directly connected members is less efficient on stack bandwidth. In a four-member stack, with members connected 1-2, 2-3, 3-4, 4-1, put the uplinks on members 2 and 4 or 1 and 3 . (See Figure 10.)

Figure 10. Four-Member Stack with Two Uplinks


## FlexStack and FlexStack-Plus Comparison with StackWise Plus and StackWise-480

The FlexStack and FlexStack-Plus feature is based on StackWise ${ }^{\circledR}$ Plus. For those network administrators familiar with managing Cisco Catalyst 3750-E or 3750-X stacks running StackWise Plus, managing FlexStack is very similar. There are some differences between FlexStack, FlexStack-Plus, and StackWise Plus. This section will describe the similarities as well as the differences.

Operationally FlexStack is similar to StackWise Plus. The similarities are in the way the switch is managed. There is a single management point of contact for the stack. All switch configuration is done from the master and applies globally. The stack appears to have a single forwarding plane for all data traffic. FlexStack supports highavailability cross-stack features such as EtherChannel, FlexLinks, and Spanning Tree Protocol.

The operational differences are minor and are related more to Cisco IOS Software feature sets on the two different families of Ethernet switches, rather than on operational aspects of managing switch stacks. The Cisco Catalyst $3750-E, 3750-X$, and 3850 Series Switches use licensing to manage feature sets, whereas the Cisco Catalyst 2960-S and 2960-X switches do not have licenses for feature sets. The Cisco Catalyst 3750-E, 3750-X, and 3850 Series Switches support full Layer 3 functionality (for example, Layer 3 forwarding and IP routing), whereas the 2960-S and 2960-X support limited static routing with LAN Base. The 2960-XR supports full routing with the IP-Lite feature set.

Table 5 shows the differences (in gray) and the similarities (in orange) for FlexStack-Plus and StackWise Plus.
Table 5. Comparing FlexStack-Plus and StackWise Plus

| Stack Features | Cisco Catalyst 2960-S <br> FlexStack Stacking | Cisco Catalyst 2960-X <br> FlexStack-Plus Stacking | Cisco Catalyst 3750-X <br> StackWise Plus | Cisco Catalyst 3850 <br> StackWise-480 |
| :--- | :--- | :--- | :--- | :--- |
| Device limit | 4 units | 8 units | 9 units | 9 units |
| Stack bandwidth | 40Gbps | 80Gbps | 64Gbps | 480Gbps |
| Architecture | Hardware drop table | Hardware drop table | Ring (destination stripping) | Ring |
| Dynamic ring load | No | No | Yes | Yes |


| Stack Features | Cisco Catalyst 2960-S FlexStack Stacking | Cisco Catalyst 2960-X FlexStack-Plus Stacking | Cisco Catalyst 3750-X StackWise Plus | Cisco Catalyst 3850 StackWise- 480 |
| :---: | :---: | :---: | :---: | :---: |
| balancing |  |  |  |  |
| Stack convergence | 1-2 seconds | 100 milliseconds | Few milliseconds | Few milliseconds |
| Stack link failure detection | Software | Hardware | Hardware | Hardware |
| Stack quality of service (QoS) | Applied hop by hop | Applied hop by hop | Applied on ingress | Applied on ingress |
| Management | Single IP address, SNMP, syslog | Single IP address, SNMP, syslog | Single IP address, SNMP, syslog | Single IP address, SNMP, syslog |
| Configuration | Single config and CLI, auto image and config update | Single config and CLI, auto image and config update | Single config and CLI, auto image and config update | Single config and CLI, auto image and config update |
| Show and debug commands | Unified | Unified | Unified | Unified |
| Single forwarding and control plane | Synchronize Address Resolution Protocol (ARP), MAC address, Internet Group Management Protocol (IGMP), VLAN tables | Synchronize ARP, MAC address, IGMP, VLAN tables | Synchronize ARP, MAC address, IGMP, VLAN tables, routing tables | Synchronize ARP, MAC address, IGMP, VLAN tables, routing tables |
| Cross-stack features | Yes | Yes | Yes | Yes |
| Single bridge ID | Yes | Yes | Yes | Yes |
| Redundancy | Stack master 1:N redundancy | Stack master 1:N redundancy | Stack master 1:N redundancy | Stack active master with single standby predetermined |

The differences between FlexStack, FlexStack-Plus, and StackWise-480 are rooted in the different hardware architectures. StackWise-480 is the superior architecture. StackWise-480 is a dual redundant ring architecture with individual links operating at 240 Gbps bidirectionally. The StackWise-480 connections are not pure Ethernet links but rather a dual ring architecture, allowing packets to take either ring. The FlexStack and FlexStack-Plus links are full duplex 10Gbps or 20Gbps Ethernet links. The ring architecture gives StackWise-480 a significant advantage over FlexStack and FlexStack-Plus when it comes to high availability of the stack.

In Table 5 stack convergence is measured in milliseconds for StackWise-480, as opposed to seconds for FlexStack. This is because of the ring architecture. In StackWise-480 packets ingressing the stack can be forwarded on either ring. When one ring becomes inoperable, all ingressing packets are forced to a single ring. The speed of one ring is 240 Gbps and can easily accommodate an entire stack. The recovery time for a single link going inoperable is not dependent upon the number of members in the stack. It is consistent across all stack sizes. In FlexStack, when a single link becomes inoperable, there is a stackwide disruption in the forwarding of packets. The members must rediscover the new topology. This rediscovery time can take up to 1 to 2 seconds depending upon the number of members in the stack. Smaller stacks will reconverge more quickly than larger stacks.

Because FlexStack and FlexStack-Plus stack links are Ethernet links, forwarding packets from one stack member to another requires the packets to be forwarded by each member. This is referred to as hop by hop, meaning the packets must hop from the ingressing member to the egressing member across all members in between. Not all members see the packet. Only those members in the data path will forward the packet. This hop-by-hop behavior is the same as different Ethernet switches forwarding the packet among themselves. Hop by hop requires that each member must queue the packet. Queuing the packet means that QoS is applied at each hop. If a stack link between two members becomes congested, packets that successfully ingress on one member may be discarded by another member because of that congestion. There is no dynamic adjustment to the forwarding of packets when one link becomes congested.

With StackWise-480, once a packet is on the ring, it is guaranteed to be delivered to the destination interface. Add to this that StackWise-480 rings will dynamically load balance, and you can see that ring congestion is not an issue on StackWise-480.

## Troubleshooting Tips and Operational Details

## FlexStack Member Addition

When a new switch joins an existing stack (that is, a stack that has already elected a master), then the new switch will be rebooted after it joins. This is to remove any old configuration. When the new switch completes the reload, it will become a member of the existing stack. The new member will be assigned the next available unused member number if it was not preprovisioned. Preprovisioning of switch members is covered in the Cisco Catalyst 2960-S and 2960-X software configuration guide on Cisco.com. This behavior is covered in more detail in subsequent sections.

## Details on FlexStack Cables and Connections

Match the keys on the cable with the connector when connecting FlexStack cables. It should slide in easily. When a good connection is made, there should be a "click" as the cable connects correctly. If it does not slide in easily, try turning the cable 180 degrees to match the keys on the cable with the connector. The tabs on the cables face away from each other.

Figure 11 shows a stack with a single bad FlexStack cable or connection (indicated by link B). This stack provides only half bandwidth and does not have redundant connections. A broken FlexStack link is operationally the same as not having a FlexStack link.

To see if the stack has a broken or unconnected stack port, use the "show switch stack-ports" command. Stack port connections for the nonredundant stack shown in Figure 11:

Figure 11. Failed Stack Link


| C2960-NonRedundant\# show switch stack-ports |  |  |
| :---: | :---: | :---: |
| Switch \# | Port 1 | Port 2 |
| $--------~$ | ------ | ------ |
| 1 | Ok | Down |
| 2 | Ok | Ok |
| 3 | Ok | Down | | Only member 2 (in the middle) |
| :--- |
| has two active stack |
| connections. |



Figure 12 shows a stack without redundant connections. When the FlexStack link indicated by B breaks, the FlexStack will break into two logical switches. The top two members will form one logical stack with two members, and the other logical stack will contain a single member.

Figure 12. Broken Stack


Since the stack in the connection in Figure 12 has broken into two, only two members will have their stack port status shown. The third member has been removed from the stack.

Stack port connections for the broken stack shown in Figure 12:


## Conclusion

FlexStack and FlexStack-Plus technology provides network administrators the ease of use for managing multiple switches. Cisco Catalyst 2960-S, 2960-X, and 2960-XR switches are the fully featured Layer 2 switches that use FlexStack and FlexStack-Plus. FlexStack and FlexStack-Plus have built-in hardware redundancy along with built-in ease of use.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

## Asia Pacific Headquarters

 Cisco Systems (USA) Pte. Ltd. Singapore
## Europe Headquarters

Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.
Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)


[^0]:    C2960-48\# archive download-sw tftp://192.168.1.1/c2960s-universalk9-tar.150-2.SE3

