Release Notes

OmniSwitch 6860/6860E

Release 8.2.1.R01

These release notes accompany release 8.2.1.R01 software which is supported on the OmniSwitch 6860/6860E platforms. These release notes provide important information on individual software features and hardware modules. Since much of the information in these release notes is not included in the hardware and software user manuals, it is important that you read all sections of this document before installing new hardware or loading new software.

Contents

Related Documentation

These release notes should be used in conjunction with OmniSwitch AOS Release 8 User Guides. The following are the titles and descriptions of the user manuals that apply to this release. User manuals can be downloaded at: http://enterprise.alcatel-lucent.com/UserGuides

OmniSwitch 6860/6860E Hardware User Guide
Describes the hardware and software procedures for getting an OmniSwitch up and running as well as complete technical specifications and procedures for all OmniSwitch chassis, power supplies, and PoE.

OmniSwitch AOS Release 8 CLI Reference Guide
Complete reference to all CLI commands supported on the OmniSwitch. Includes syntax definitions, default values, examples, usage guidelines, and CLI-to-MIB variable mappings.

OmniSwitch AOS Release 8 Switch Management Guide
Includes procedures for readying an individual switch for integration into a network. Topics include the software directory architecture, software rollback protections, authenticated switch access, managing switch files, system configuration, using SNMP, and using web management software (WebView).

OmniSwitch AOS Release 8 Network Configuration Guide
Includes network configuration procedures and descriptive information on all the major software features and protocols included in the base software package. Chapters cover Layer 2 information, Layer 3 information, security options, and Quality of Service.

OmniSwitch AOS Release 8 Advanced Routing Configuration Guide
Includes network configuration procedures and descriptive information on all the software features and protocols included in the advanced routing software package. Chapters cover multicast routing (DVMRP and PIM), BGP, IS-IS, OSPF, and OSPFv3.

OmniSwitch AOS Release 8 Transceivers Guide
Includes transceiver specifications and product compatibility information.

Technical Tips, Field Notices
Contracted customers can visit our customer service website at: service.esd.alcatel-lucent.com.

System Requirements

Memory Requirements

OmniSwitch 6860/6860E Series Release 8.2.1.R01 requires 2GB of RAM and 2GB flash memory. This is the standard configuration shipped.

Configuration files and the compressed software images—including web management software (WebView) images—are stored in the flash memory.

UBoot and FPGA Requirements

OmniSwitch 6860/6860E models will be factory shipped with the correct Uboot/FPGA. They do not need to be upgraded and should not be downgraded.

OmniSwitch 6860/6860E (All models) - AOS Release 8.2.1.255.R01(GA)

Model	Uboot	FPGA
OS6860/OS6860E (except U28)	8.1.1.70.R01	Version 0.9
OS6860E-U28	8.1.1.70.R01	Version 0.14

New Hardware Support in 8.2.1.R01

SFP-10G-ZR

10-Gigabit SFP+ optical transceiver with DDM support. Supports transmission at 1550nm up to 80km with single mode fiber using an LC connector.

QSFP-40G-AOC20M

40-Gigabit QSFP+, 20m direct-attached active optical transceiver with DDM support.

SFP-10G-GIG-LR

Dual-speed SFP+ optical transceiver with DDM support. Supports transmission at 850nm up to 10km with single mode fiber using an LC connector.

New Features Summary

The following software features are being introduced with the 8.2.1.R01 release, subject to the feature exceptions and problem reports described later in these release notes:

Features listed as ‘Base’ are included as part of the base software and do not require any license installation. Features listed as ‘Advanced’ require the installation of that license.

AOS 8.2.1.R01 Feature Summary Table

Feature	Platform	License
Hardware / Virtual Chassis Feature Support		
lldp PoE	OS6860/6860E	Base
Remote Stacking	OS6860/6860E	Base
Manageability Feature Support		
Telnet/SSH clients per VRF - Management VRF	OS6860/6860E	Base
Additional SWLOG message for link up/down	OS6860/6860E	Base
Zero Touch Provisioning - Opex OV Cloud integration	OS6860/6860E	Base
OPenFlow VC mode	OS6860/6860E	Base
Openflow API mode in VC	OS6860/6860E	Base
default user profile for Admin users	OS6860/6860E	Base
IP Managed Services	OS6860/6860E	Base
Layer 2 Feature Support		
Loopback Detection (LBD)	OS6860/6860E	Base
STP Loopguard	OS6860/6860E	Base
DHL	OS6860/6860E	Base
Layer 3 Feature Support		
Internal DHCP Server IPv4/v6	OS6860/6860E	Base
IPv4 over SPB / Export Loopback0 to SPB IP	OS6860/6860E	Advanced
IPv6 Feature Support		
ISIS IPv6	OS6860/6860E	Advanced

Feature	Platform	License
M-ISIS	OS6860/6860E	Advanced
QoS Feature Support		
QoS ingress/egress bandwidth via RADIUS	OS6860/6860E	Base
QoS per port rate limiting	OS6860/6860E	Base
Tri color marking(SrTCM/TrTCM policy action and 802.1ad DEI bit mapping/marking)	OS6860/6860E	Base
Multicast Feature Support		
MAC address boundary out of range	OS6860/6860E	Base
Initial MC Packet	OS6860/6860E	Base
PIM-BFD Multicast subsecond convergence	OS6860/6860E	Advanced
Monitoring/Troubleshooting Feature Support		
Interface violation recovery	OS6860/6860E	Base
MIB Bit/s second	OS6860/6860E	Base
Additional storm control	OS6860/6860E	Base
Security/Access Guardian		
Application Monitoring and Enforcement	OS6860/6860E	Base
Legacy UNP/VNP from 7.X	OS6860/6860E	Base
LPS Sticky	OS6860/6860E	Base
Port Bounce on VLAN Change	OS6860/6860E	Base
UPNP/DLNA Relay	OS6860/6860E	Base
Access Guardian Enhancements	OS6860/6860E	Base
show 802.1x enhancement	OS6860/6860E	Base

New Features Descriptions

Hardware/Virtual Chassis Features

lldp PoE

With power-via-mdi configured the power for the powered device is negotiated using the optional power via MDI TLV in the LLDPDU. The powered device can request additional power using the power via MDI TLV. The switch will check the current PoE budget and if power is available the switch will provide the requested power to the powered device. If power is unavailable, the switch will respond with the existing maximum power information.

- Power negotiation is supported for Class 4 powered devices.
- The maximum power a powered device can request cannot exceed the maximum power allowed for the PoE class in which the powered device is detected.
- If the port is manually configured with a maximum power value, the powered device cannot receive more power than the maximum configured value.

Remote Stacking / Virtual Chassis

This feature expands the Virtual Chassis (stacking) capability of the OS6860/OS6860E to allow a VC to be configured over long distances. In addition to the dedicated VFL ports, the SFP+ ports can now be configured as VFL ports using various transceivers to increase the supported VFL distances.

AOS release 8.2.1 uses the automatic VFL feature to create a VFL between chassis. Once a port is configured as an auto VFL port the chasssis will attempt to automatically create a VFL between chassis. This auto VFL capability is automatically enabled on the dedicated 20G VFL ports.

Manageability Feature Support

Telnet/SSH clients per VRF - Management VRF

This feature allows management services to be enabled or disabled in a VRF other than the default VRF. This allows for the creation of a single management VRF instance, a VRF per management service, or multiple VRFs for a service. Depending on the type of service there are different levels of management allowed as described below.

Level 0 - The management service may only appear in the Default VRF.

Level 1 - User may specify a single VRF that all management services can be configured in. For example, both RADIUS and LDAP can use vrf-1.

Level 2 - Each management service or multiple management services can be configured for a different VRF. For example, RADIUS in vrf-1, LDAP in vrf-2, SNMP in vrf-3.

Level 3 - A management service may appear in multiple VRFs. For example, SSH and Telnet in vrf-1 and vrf-2.

Level	Description	Telnet/SSH/SFTP/SCP	Radius/SNMP/HTTP/HTTPS/ NTP/LDAP/TACACS+/Syslog
0	Default VRF only	Yes	Yes
1	Single VRF for all services	Yes	Yes
2	Single VRF per service, each service can be on a different VRF	Yes	Yes

3	Multiple VRFs per service, any service on any VRF	Yes	No

Additional SWLOG message for link up/down

This enhancement will generate a syslog message as well as a trap when trap generation is enabled for a port link up / link down event.

Zero Touch Provisioning - Opex OV Cloud integration

This feature modifies the Automatic Remote Configuration DHCP client process on an OmniSwitch to give priority to a DHCP response from OmniVista. If a DHCP response is received on VLAN 1 from a DHCP server other than OmniVista, the response will be temporarily stored. The DHCP client will continue to wait for a 30 second window to see if a DHCP response is received from the higher priority OmniVista DHCP server.

- Priority 1 – OmniVista DHCP server (VSI = alcatel.nms.ov2500)
- Priority 2 – OXO DHCP Server – (VSI = alcatel.a.a4400.0)
- Priority 3 – All other DHCP servers

If no DHCP response is received from the OmniVista DHCP server within the 30 second window, the stored response will be applied. If a DHCP response is received from the OmniVista DHCP server it will be immediately applied.

Additionally, the OmniSwitch will send the vendor class and switch type in the DHCP Discover/Request packets. The vendor class and switch type are sent in option-60 as OmniSwitch-*moduleType*. For example, OmniSwitch-OS6860E-P48.

OpenFlow in VC mode

OpenFlow is supported on both standalone and Virtual Chassis.

Default user profile for Admin users

Currently, a user can configure personal settings comprising the prompt, the more settings and the aliases. The settings are only valid during the life of the session and are lost once the user logs out. To save the personal settings, the “user profile save” command is used. The command is enhanced to add a default profile for ALL administrative users connecting to the switch.

The “user profile save” command has been modified to specify that the current settings must be in a global profile file or in the user specific profile file -> user profile save [global-profile]

The settings are then saved in a text file and there is one file per user. When a user logs in, the personal settings from the profile file are loaded. During the session, the user can remove the personal settings with the “user profile reset” command and go back to the factory default settings (i.e. “prompt” command, no more, no alias). This does not delete the user profile file.

The personal settings are configurable and can be saved by any users, regardless of their privileges. For instance, a read only user can configure the prompt, more and aliases and can save the settings in the file.

IP Managed Services

By default, most applications that run over IP use the egress IP interface address as the source IP while communicating with a peer/server. However, it may be desirable to have some applications use a specific source IP for the packets that are sent out. This feature provides ('ip service source-ip' command) the ability to configure a permanent source IP interface to send packets. The source IP interface can be the Loopback0 address or user defined IP interface.

Layer 2 Feature Descriptions

Loopback Detection (LBD)

LBD can detect and prevent L2 forwarding loops on port either in the absence of other loop-detection mechanisms like STP/RSTP/MSTP or when the mechanism can't detect it. Sometimes the STP/RSTP/MSTP based loop detection can't be used due to the following:

- There is a client's equipment that drops or cuts the BPDUs.
- The STP protocol is restricted on edge Network

The LBD feature detects that a port has been looped back or looped. If a loop-back/loop is detected, the port is disabled (forced down) and the appropriate Error Log is issued.

The remote-origin LBD capability allows processing the LBD frames received from a remote system. The MAC address of the remote system from which the LBD frames are received is recorded and the origin port is shut down. The remote-origin LBD can be configured globally and on per port.

Ethernet switch periodically sends out L2 Ethernet frame (LBD frame) from all loop-back detection enabled ports. The LBD frame is not a BPDU frame. In normal state of the access line this frame is removed from the network segment by the subscriber equipment. In case of failure (cable fault, NIC incorrect work, etc) switch receives back the control frame on the port. After receiving the frame, switch should force the access port down and issues a SNMP trap. In addition the port also can be re-enabled by user through cli command.

STP Loopguard

This feature is intended to prevent loops in a spanning tree bridged network when a device is unable to receive BPDUs on a non-designated port in a timely manner. Loop formation can occur when a bridge hosting a blocking port transitions that port to forwarding erroneously. This can lead to a temporary or even a permanent loop. This feature can be enabled either on a port or link aggregate and can be configured for any spanning tree mode (flat, 1x1, STP, RSTP, MST, PVST). Loopguard effectively protects each STP instance when configured on a port that supports multiple spanning tree instances.

DHL

Dual-Home Link (DHL) Active-Active is a high availability feature that provides fast failover between core and edge switches without using Spanning Tree. To provide this functionality, DHL Active-Active splits a number of VLANs between two active links. The forwarding status of each VLAN is modified by DHL to prevent network loops and maintain connectivity to the core when one of the links fails. This implementation of DHL is Active-Active. The DHL Active-Active feature is configurable on regular switch ports and on logical link aggregate ports (linkagg ID) as well as LACP aggregated ports.

Layer 3 Feature Support

Internal DHCP Server IPv4/v6

The OmniSwitch now supports an internal DHCP Server compliant with RFC 2131 and RFC 3315 based on Vital QIP 8.0 release. This feature can be used to provide IP addresses for small offices, management network, or local phone services. The following files are used to configure the internal DHCP server setting on the OmniSwitch:

- IPv4 Policy Files- dhcpd.conf, dhcpd.pcy
- IPv6 Configuration Files – dhcpd6.conf, dhcpd6.pcy

DHCP Policy files - The dhcpd(v6).pcy files initialize the global attributes for the DHCP server.

DHCP Configuration files - The dhcpd(v6).conf files are used to configure specific DHCP server settings on the switch such as the following:

- MAC pool allowed (for DHCPv4)
- MAC pool excluded (for DHCPv4)
- Subnet pools
- Dynamic scopes
- Static scopes
- IP range, mask, DNS, Default router, NetBIOS configurations for DHCPv4.
- User class specific configs.
- Vendor class specific configs.
- DUID Pool (for DHCPv6 only).
- Excluded DUID Pool (for DHCPv6 only)
- Manual DUID mapping (for DHCPv6 only).
- Other options that need to be sent to the client.

IP over SPB

The previous implementation of Shortest Path Bridging MAC (SPBM) provides L2 VPN capability that bridges L2 customer LAN segments. Customer edge (CE) devices form peers and exchange routing information, as well as perform the necessary IP forwarding. Then the SPBM Backbone Edge Bridges (BEBs) bridge the already routed IP traffic across the SPBM backbone.

This release now provides IPv4 over SPBM capability that consolidates the routing functionality of CE devices into BEB devices. The VRF instances on different BEBs are tied together via backbone service instance identifiers (I-SIDs) across the same SPBM backbone that is used to support Layer 2 VPNs.

The OmniSwitch IP over SPBM solution supports two methods for combining Layer 3 routing and SPBM in the same chassis: VPN-Lite and L3 VPN.

VPN-Lite

The VPN-Lite method provides a ‘gateway’ between a regular SPBM service and a router within the same OmniSwitch chassis. This solution provides a specific advantage in that it allows a single box to represent two ‘tiers’ in a typical ‘fat tree’ network, which is popular in datacenter deployments.

In addition, a VPN-Lite configuration can act purely as a L3 VPN when configured correctly. In this mode, existing routing protocols can form adjacencies across the SPBM Provider Backbone Bridge (PBB) network. To keep it purely a L3 VPN, the administrator makes sure that no SPBM Service Access Points (SAPs) that can inject bridged flows are allowed to attach to the VPN's I-SID.

The VPN-Lite approach uses the SPBM network in the same way a VLAN is used for transporting L3 frames. Each BEB or host can inject frames into the I-SID as needed, and BEBs can decide to bridge our route those frames based on their inner and outer destination MAC address.

L3 VPN

When the L3 VPN method is implemented, the OmniSwitch acts as an access or edge router to multiple VRFs and connects these VRFs across an SPBM managed PBB network. This solution is based upon the IETF drafts "IP/IPVPN services with IEEE 802.1aq SPB(B) networks" and uses the proposed IS-IS TLVs to exchange routes between the BEBs that host the same VPN services.

When the L3 VPN approach is implemented, each VPN is identified by a VRF locally on each BEB and globally in the backbone by an I-SID in the PBB header. SPB IS-IS will import/export routes from the local routing protocols running inside their respective VRFs. In essence, SPB IS-IS is creating tunnels between BEBs through which routed frames are sent to reach their target networks.

The L3 VPN solution gives an administrator the ability to build VPNs and extend them over a SPBM core without having to define routes and VRFs across that core by hand. The core boxes need only run SPBM.

Export Loopback0

The IPv4 network address of the Loopback0 interface is included when IPv4 routes are route-leaked into I-SIDs as part of the IP over SPBM feature.

IPV6 Feature Support

ISIS IPv6

Intermediate System-Intermediate System (IS-IS) is a shortest path first (SPF) or link-state protocol. IS-IS is an interior gateway protocol (IGP) that distributes routing information between routers in a single autonomous system (AS) for IP (IPv4 and IPv6). This feature allows a single routing protocol to support pure IP and dual environments. Integrated IS-IS is also deployed extensively in an IP-only environment.

M-ISIS

Multi-topology (M-ISIS) support is necessary in IS-IS to support network domains in which non-dual stack IS-IS routers exist. The default protocol behavior of IS-IS is to construct shortest paths through the network using the routers' MAC addresses with no regard to the different IP address families supported. This behavior may result in black-holed routing when there are some IPv4-only or IPv6-only routers in an IS-IS routing domain, instead of all dual-stack routers. M-ISIS mechanism runs multiple, independent IP topologies within a single IS-IS network domain, using separate topology-specific SPF computation and multiple Routing Information Bases (RIBs). M-ISIS is advised in networks containing ISIS enabled routers with different topologies of IPv4 and IPv6 capable routers.

QoS Feature Support

QoS ingress/egress bandwidth via RADIUS

This feature applies maximum ingress and egress bandwidth limiting on a port on the basis of UNP classification. When a user is successfully authenticated under a UNP policy either through RADIUS returned UNP attribute or through a local UNP policy, bandwidth limitations are applied on the port. If multiple user devices are classified into different profiles but learned on the same UNP port, the bandwidth parameter values obtained for the last user device learned are applied on the port. Parameter values applied through previously learned users are overwritten.

- “Per user" bandwidth profiling is not supported.
- User device can be a supplicant, non-supplicant, or a Captive Portal client.
- Supported only on UNP Edge and Bridge profiles.

QoS per port rate limiting

Per-port rate limiting allows configuring a policy rule that specifies a rate limiter for a group of ports or for each individual port in a group. In other words, rate limiting actions specify values that are shared by all the ports in a specific source port group, or the rate limiting actions are applied individually to each port that is a member of the source port group.

How the rate limiting actions are applied to group ports is determined by the type of source port group specified by the policy rule condition. There are two types of configurable source port group policy conditions: group and split group. Both types of group conditions represent an aggregate of member ports. The difference between the two types is how rate limiting actions are applied to the member ports.

- A regular source port group. When this type of group is defined as a condition for a policy rule, any rate limiter actions in that rule are applied to the group as a whole. All ports in that group share the rate limiting values. For example, if a maximum bandwidth value of 10M is applied, the member ports of that group share the 10M maximum bandwidth value.

- A source port split group. When this type of group is defined as a condition for a policy rule, any rate limiter actions in that rule are applied to each member port in that group. The rate limiting values are not shared between all the members of the group. For example, if a maximum bandwidth value of 10M is applied, each member port of the split group is given a maximum bandwidth value of 10M.

Using a split group condition in a rate limiting policy rule avoids having to create a separate policy rule for each individual port. This reduces the amount of configuration and switch resources required to apply rate limiting actions to individual ports.

A source port group is the only QoS group that supports the split group option; other QoS groups (for example, a destination port group or network group) do not support the split group option.

Tri color marking(SrTCM/TrTCM policy action and 802.1ad DEI bit mapping/marking)

Tri-Color Marking (TCM) provides a mechanism for policing network traffic by limiting the rate at which traffic is sent or received on a switch interface. The TCM policer meters traffic based on user-configured packet rates and burst sizes and then marks the metered packets as green, yellow, or red based on the metering results. TCM policer meters each packet and passes the metering result along with the packet to the Marker. Depending upon the result sent by the Meter, the packet is then marked with either the green, yellow, or red color. The marked packet stream is then transmitted on the egress based on the color-coded priority assigned. The TCM Meter operates in Color-Blind mode (the Color-Aware mode is not supported). In the Color-Blind mode, the Meter assumes that the incoming packet stream is uncolored. However incoming packets with the CFI/DEI bit set are automatically given an internal lower priority.

There are two types of TCM marking supported:

Single-Rate TCM (srTCM) according to RFC 2697—Packets are marked based on a Committed Information Rate (CIR) and two associated burst size values: Committed Burst Size (CBS) and Peak Burst Size (PBS).

Two-Rate TCM (trTCM) according to RFC 2698—Packets are marked based on a CIR value and a Peak Information Rate (PIR) value and two associated burst size values: CBS and PBS.

Both srTCM and trTCM handle the burst in the same manner. The main difference between the two types is that srTCM uses one rate limiting value (CIR) and trTCM uses two rate limiting values (CIR and PIR) to determine packet marking

Multicast Feature Support

MAC address boundary out of range

Multicast boundaries confine multicast addresses to a particular domain. Confining multicast addresses helps to ensure that multicast data traffic passed within a multicast domain does not conflict with multicast users outside the domain.

Multicast addresses 239.0.0.0 through 239.255.255.255 have been reserved by the IANA as administratively scoped addresses for use in private multicast domains. These addresses cannot be used for any other protocol or network function. Because they are regulated by the IANA, these addresses can theoretically be used by network administrators without conflicting with networks outside of their multicast domains. However, to ensure that the addresses used in a private multicast domain do not conflict with other domains (for example, within the company network or out on the Internet), multicast address boundaries can be configured.

AOS supports configuration of multicast route boundaries for the entire multicast group including scoped multicast addresses (224.0.0.0 through 239.255.255.255).

By default, route boundary configuration is supported for the scoped addresses (239.0.0.0 to 239.255.255.255). Use "ip mroute-boundary extended" command to allow multicast route boundary configuration on the complete multicast group range (224.0.0.0 to 239.255.255.255).

Initial MC Packet

Multicast is often used for audio\video streaming applications where the first packet may be dropped as it is used for learning the new flow. However, some multicast applications require the initial packets sent by the multicast source to be received. The packet buffering functionality can be enabled to prevent those first multicast packets from being dropped.

PIM-BFD Multicast subsecond convergence

This feature is to minimize the delay at the time of failure in the primary path forwarding multicast data packets by deploying BFD and MoFRR in Multicast Routing Protocols – in both PIM DM and PIM SM. On intimation from BFD about the primary link (neighbor) failure, sub second convergence could be achieved by a redundant path to carry forward the source traffic immediately. And also to minimize the delay in resuming the data packet flow in the alternate path by deploying the redundant path functionality.

With this feature BFD and MoFRR can be enabled for PIM neighbors (this is similar to the existing feature available for OPSF). BFD needs to be enabled on the PIM interfaces and MoFRR on all routes. Also a global control to enable BFD and MoFRR on PIM SM/DM is also available. This feature will reduce the time delay in reduction of PIM neighbor timeout detection.

Monitoring and Troubleshooting Feature Support

Interface violation recovery

The OmniSwitch allows features to shutdown an interface when a violation occurs on that interface. To support this functionality, the following interfaces violation recovery mechanisms are provided:

- Manual recovery of a downed interface using a CLI command.
- An automatic recovery timer that indicates how much time a port remains shut down before the switch automatically brings the port back up
- A maximum number of recovery attempts setting that specifies how many recoveries can occur before a port is permanently shutdown
- A wait-to-restore timer that indicates the amount of time the switch waits to notify features that the port is back up
- An SNMP trap that is generated each time an interface is shutdown by a feature. This can occur even when the interface is already shutdown by another feature. The trap also indicates the reason for the violation.
- An SNMP trap that is generated when a port is recovered. The trap also includes information about how the port was recovered.

MIB Bit/s second

In order to align the show output for the interfaces counter with the SNMP the following new SNMP objects are added to view the information:

- inBitsPerSec Counter64,: The average number of Bits Received per second
- outBitsPerSec Counter64,: The average number of Bits Transmitted per second
- ifInPauseFrames Counter64,: The average number of Pause Frames Received per second
- ifOutPauseFrames Counter64,: The average number of Pause Frames Transmitted per second
- ifInPktsPerSec Counter64,: The average number of Packet Received per second
- ifOutPktsPerSec Counter64,: The average number of Packets Transmitted per second

The "show interfaces counters" MIB objects are part of ifTable which is defined in the standard MIB File IF-MIB.mib. Since the standard MIB file cannot be edited for new variables or tables, it is done by using the AUGMENTS functionality.

The new MIB variables are defined in the interfaceStatsTable which is an expansion of ifEntry.

Additional storm control

When the incoming traffic flow of a port exceeds the configured high threshold value, the storm has to be controlled. This can be done by either rate limiting the traffic or blocking the traffic on that port. The traffic storm control continues to monitor the incoming traffic level even for the blocked port. When the traffic on the violated port reaches the configured low threshold value, the port state is reset to normal state. If the low threshold is not configured, incoming traffic level is not monitored. You can configure the violation mode to Shutdown, Trap, or Default when the ingress traffic exceeds the configured threshold value.

When traffic (broadcast, multicast or unknown unitcast) flows on a port for 5 seconds at an average speed above the configured upper threshold value, then the port is considered to be in storm state and actions would be taken as any one of those below:

- Default – The traffic gets rate limited to the upper threshold value and user will not get any indication. This is pre-existing behavior
- Trap – The traffic gets rate limited to the upper threshold value and also user will be notified by a trap message.

- Shutdown – The corresponding port will go down and also a trap will be generated to alert the user.
- The storm state of the port can be recovered by both manually and automatically. The below procedure is to recover the port manually.
    - Interfaces slot/port admin down/up
    - Port plug out/ plug in
    - Interface clear all violations (Only applicable for Shutdown action)

Also the port can be automatically recovered form storm state if the port is configured with lower threshold value and if the traffic on the port, where storm occurs, reaches below that lower threshold value.

Security Feature Support

Application Monitoring and Enforcement

The OmniSwitch Application Monitoring and Enforcement (AppMon) feature addresses the key challenges of real time classification of flows at the application level by providing differential QoS treatment in the form of higher priority marking and security policies at the application level. The AppMon feature improves the quality of the user experience through application aware network optimization and control.

The AppMon Enforcement feature allows the switch to differentiate between different traffic flows and assign the proper QoS and security policies. The feature also provides appropriate QoS marking to application flows as per the application-aware user configured QoS policies. An administrator can also associate QoS policies with UNP profiles and provide user level policy treatment.

Application Monitoring feature collects and reports the application specific flow information over a period. Based on this data, application specific enforcement policies can be designed.

Flow identification is done based on the following 5-tuple match of the IP packets:

- Source IP Address (IPv4 or IPv6)
- Destination IP Address (IPv4 or IPv6)
- Source L4 port
- Destination L4 port
- IP Protocol (TCP or UDP)

The following are the key components for AppMon:

- Application Signature Kit File - the file containing application signatures.
- Application Pool - pool of supported applications.
- Application List - list of applications supported for monitoring or enforcement. Applications can be added to this list using the application name or application group.
- Application Group - a group of applications. The group can be added to the application list.
- QoS policy: QoS policy configuration at application level or application group level.

Note: AppMon is supported in a virtual chassis of OmniSwitch 6860 and OmniSwitch 6860E platforms where at least one OmniSwitch 6860E is mandatory for the feature to work.

The following applications are supported by default on the OmniSwitch. However, when used in conjunction with OmniVista's signature upgrade capability, thousands of signatures can be supported. Please refer to the OmniVista documentation for more information when available.

- amazon, facebook, sip, skype, tftp, twitter, viber, webex, whatsapp, youtube

Legacy UNP/VNP from 7.X

The UNP Bridge (VLAN) and UNP Shortest Path Bridging (SPB) access modes are supported along with the UNP Edge mode. The following functionality is available to support the Bridge and VLAN modes:

- VLAN and SPB profile configuration.
- Classification rules for VLAN and SPB profiles (MAC address, MAC address range, IP address, VLAN tag).
- Dynamic VLAN and Profile configuration options for Bridge mode.
- Bridge and SPB access port type options for UNP ports.
- UNP port configuration options for VLAN and SPB profiles, authentication, and classification.

LPS Sticky

LPS sticky mode provides the following enhancements in the learning window:

Automatic conversion of MAC addresses to static MAC addresses:

When the "learn-as-static" option is enabled, MAC addresses are automatically learned as static MAC addresses during the learning window, even if the "convert-to-static" option is disabled. There is no need to manually enable the "convert-to-static" option on a per-port or global basis or wait until the learning window closes for the MAC addresses to convert to static addresses. The "learn- as-static" option can be used only when the "no-aging" option is also enabled.

MAC movement for the pseudo static MAC:

When the "mac-move" option is enabled, a pseudo-static MAC address learned on one port can move to another port in the same VLAN without getting dropped. The "mac-move" option can be used only when the "no-aging" option is enabled.

Infinite learning window:

Setting the LPS learning window time to zero configures an infinite source learning time period for all LPS ports. The learning of MAC addresses on LPS ports never times out. When an infinite learning window is set, all learning window options except the "convert-to-static" option are still valid.

Port Bounce on VLAN Change

Internal Captive Portal authentication is a configurable option for a UNP Edge profile that is applied after a user is assigned to the profile (after the initial 802.1X or MAC authentication or classification process). This type of authentication may result in a change to the initial Edge profile assignment for the user device, thus changing the VLAN assignment for the device. A new VLAN assignment is functional only after a port bounce or pause timer operation is completed. Existing BYOD global commands are leveraged to configure and apply the port bounce or pause timer settings to ensure a clean re-authentication process for non-supplicant devices when there is a VLAN change for the device.

UPNP/DLNA Relay

The Digital Living Network Alliance (DLNA) is a standards organization that defines the guidelines for multimedia devices. It also certifies communication between devices allowing them to discover/recognize each other and share digital content. DLNA uses Universal Plug and Play (UPnP) for media management, discovery and control. DLNA/UPNP uses the Simple Service Discovery Protocol (SSDP) to discover services, similar to Bonjour using mDNS for the same. In the ARUBA AirGroup solution, the ARUBA WLAN Controllers act as Bonjour and DLNA gateways allowing L2 discovery protocols, such as mDNS and SSDP, to extend across L3 boundaries through the gateway.

Along the lines of zero network configuration already supported by the OmniSwitch with mDNS, support for SSDP Relay for DLNA/UPnP enables the OmniSwitch to allow non-Apple devices to also discover the services with minimal configuration by the administrator. DLNA/ UPnP uses SSDP for dynamic discovery of services. The ARUBA controller Airgroup feature has support for DLNA and acts as a DLNA controller in addition to the support for mDNS. Similar to the OmniSwitch implementation of mDNS, the OmniSwitch relays SSDP packets to the ARUBA controller via an L2 GRE tunnel.

All the SSDP packets coming in on an OmniSwitch are intercepted and tunneled through GRE tunnel to the WLAN controller (acting as a gateway). The GRE tunnel is setup between the switch and the WLAN controller to tunnel both mDNS and SSDP frames. Similarly, traffic towards the SSDP clients/servers are sent back from the WLAN controller to the switch through the GRE tunnel. The reverse traffic is also intrcepted and then unicast or multicast from the switch to the respective ports.

Access Guardian Enhancements

The following Access Guardian enhancements are available to support similar functionality provided in other releases:

- Configuring an untagged VLAN-port association between a UNP port and a VLAN ID. This allows a UNP port to join one or more VLANs in order to facilitate broadcast traffic to the port even though no traffic was classified into those VLANs on the port. Supported on UNP Edge and Bridge ports.
- VLAN tag classification rule and VLAN tag as an option to existing classification rules.
- Trust tag option for UNP ports.

show 802.1x enhancement

The following fields are added in the “show 802.1x user” output:

- Auth Failure Reason: Reason for authentication failure as “SERVER UNREACHABLE” or “AUTHENTICATION FAILURE”. In the case of successful authentication “- “ is returned
- Auth Retry Count: Number of times the switch re-transmits a request for authentication information to the 802.1x user
- Last Successful Auth Time: Latest successful authentication time. If port was not authenticated before then “-“ is returned.

Unsupported Software Features

The following CLI commands and Web Management options may be available in the switch software for the following features. These features are not supported:

Feature	Platform
SAA	All
Traffic Anomaly Detection (Network Security)	All
VXLAN and VXLAN snooping	All
BGP 4-Octet ASN BGP AS Path Filtering for IPv6 BGP Password Support for IPv6 BGP Route Reflector for IPv6	All
Distributed ARP	All
IP Routed Port	All
Embedded Pyton Scripting / Event Manager Support	All

Unsupported CLI Commands

The following CLI commands may be available in the switch software for the following features. These commands are not supported:

Software Feature	Unsupported CLI Commands
SAA	All 'saa' commands
Network Security	All 'netsec' commands
BGP 4-Octet ASN	ip bgp autonomous-system
BGP AS Path Filtering for IPv6	ip bgp policy prefix6-list
BGP Password Support for IPv6	ipv6 bgp neighbor md5 key
BGP Route Reflector for IPv6	ipv6 bgp neighbor route-reflector-client
Distributed ARP	ip distributed-arp admin-state enable
IP Routed Port	ip interface rtr-port
Embedded Pyton Scripting / Event Manager Support	event-action trap
Per command authorization for TACACS	aaa tacacs command-authorization enable

Open Problem Reports and Feature Exceptions

The problems listed here include problems known at the time of the product's release. Any problems not discussed in this section should be brought to the attention of the Alcatel-Lucent Technical Support organization as soon as possible. Please contact customer support for updates on problem reports (PRs) where no known workaround was available at the time of release.

DHCP

PR	Description	Workaround
205933	The dhcp-server lease count for IPv4 may not show the proper count when both IPv4 and IPv6 clients are configured.	This is a display issue issue, there is no functional impact. Restarting the DHCP server fixes the issue.

ISSU

PR	Description	Workaround
209388	Some error messages may be displayed during an ISSU upgrade such as: **Message 1** - Thu Aug 28 19:31:33 : appMonCmm GENERAL error message: +++ ERROR:(1409254293.176)msgHandler[615]invalid msgId receieved 88 **Message 2** - Mon Sep 7 17:11:59 : slb rs info message: +++ Unknown RS message 00000000 **Message 3** - Fri Mar 21 20:26:27 : slCmm EXCEPT alert message: +++ EXCEPT:(18126.17)msgHandler[1876]invalid msgId receieved a000d **Message 4** - Oct 2 20:48:39 : bcmd rpcs alert message: +++ slnHwlrnCbkHandler:662 no buffer ALERT!! **Message 5** -Thu Sep 3 17:56:29 : iprm_0 ip6ni error message: +++ iprmNi6MsgHandler: Unexpected msg (0x7d0205) from c2 ni6=1 **Message 6** - Mon Aug 31 12:06:21 : iprm_4 ip6ni error message: +++ iprmNi6MsgHandler: Unexpect	These are display issues only during the transition between two releases. Once the upgrade is complete, the messages will no longer be displayed.

PoE

PR	Description	Workaround
192875	When two 600W or two 920W power supplies are connected and the power drawn is more than 450W for 600W power supply and more 780W for a 920W power supply, lanpower for one of the ports (lowest priority port) resets when the AC power chord is removed and re-inserted on any one of the power supplies.	If the power consumption is within the supported 450W for a 600W power supply and within 780W for a 920W power supply, this issue is not seen.

Port Mirroring

PR	Description	Workaround
210135	DSCP field of the port mirroring/monitoring traffic (IP packets only) is getting overwritten with the local chassis-id value when app-mon feature is enabled. Original traffic isn't affected, the DSCP field is modified only in the mirrored copy of the packets. When app-mon feature is not enabled, DSCP field is not modified for mirrored traffic.	There is no workaround at this time if the app-mon feature is enabled. If the app-mon feature is not enabled, the problem is not seen.

System

PR	Description	Workaround
207292	Occasionally at boot up the system may display Buffer I/O errors similar to the following. This has not resulted in any functional failures: Starting 6860 Boot Process [31.030000] Result: hostbyte=0x07 driverbyte=0x00 [31.060000] cdb[0]=0x28: 28 00 00 34 40 3e 00 00 08 00 [31.090000] end_request: I/O error, dev sda, sector 3424318 *[output truncated]*	There is no known workaround at this time.

Transceivers

PR	Description	Workaround
204991	DDM information for 1G and 10G may display the same value on some SFP-10G-GIG-LR/SR transceivers but should be different based on speed.	There is no known workaround at this time.

VRF

PR	Description	Workaround
210404	Additional VRF's cannot be configured if memory usage is greater than 80%.	There is no known workaround at this time.

Webview

PR	Description	Workaround
210599	In Webview the ingress depth is set to wrong value for a UNP edge profile.	Use the CLI to set the ingress depth.

Hot Swap Guidelines

Hot Swap Feature Guidelines

- Hot swap of like power supplies is supported.
- Hot swap of unlike power supplies is not supported.
- Hot insertion, the insertion of a power supply into a previously empty slot, is supported.
- Mixing of different wattage power supplies in the same chassis is not supported.

Hot Swap Procedure

The following steps must be followed when hot-swapping power supplies.

1. Disconnect the power supply cord from the power supply.
2. Extract the power supply.
3. Insert replacement power supply of same type.
4. Connect the power supply cord to the new power supply.

Technical Support

Alcatel-Lucent technical support is committed to resolving our customer's technical issues in a timely manner. Customers with inquiries should contact us at:

Region	Phone Number
North America	800-995-2696
Latin America	877-919-9526
Europe Union	+800 00200100 (Toll Free) or +1(650)385-2193
Asia Pacific	+65 6240 8484

Email: esd.support@alcatel-lucent.com

Internet: Customers with Alcatel-Lucent service agreements may open cases 24 hours a day via Alcatel-Lucent's support web page at: service.esd.alcatel-lucent.com.

Upon opening a case, customers will receive a case number and may review, update, or escalate support cases on-line. Please specify the severity level of the issue per the definitions below. For fastest resolution, please have telnet or dial-in access, hardware configuration—module type and revision by slot, software revision, and configuration file available for each switch.

Severity 1 Production network is down resulting in critical impact on business—no workaround available.

Severity 2 Segment or Ring is down or intermittent loss of connectivity across network.

Severity 3 Network performance is slow or impaired—no loss of connectivity or data.

Severity 4 Information or assistance on product feature, functionality, configuration, or installation.

Third Party Licenses and Notices

Legal Notices applicable to any software distributed alone or in connection with the product to which this document pertains, are contained in files within the software itself located at: **/flash/foss**.

Appendix A: General Upgrade Requirements and Best Practices

This section is to assist with upgrading an OmniSwitch. The goal is to provide a clear understanding of the steps required and to answer any questions about the upgrade process prior to upgrading. Depending upon the AOS version, model, and configuration of the OmniSwitch various upgrade procedures are supported.

Standard Upgrade - The standard upgrade of a standalone chassis or virtual chassis (VC) is nearly identical. All that's required is to upload the new image files to the *Running* directory and reload the switch. In the case of a VC, prior to rebooting the Master will copy the new image files to the Slave and once the VC is back up the entire VC will be synchronized and running with the upgraded code.

ISSU - The In Service Software Upgrade (ISSU) is used to upgrade the software on a VC or modular chassis with minimal network disruption. Each element of the VC is upgraded individually allowing hosts and switches which are dual-homed to the VC to maintain connectivity to the network. The actual downtime experienced by a host on the network should be sub-second in most cases but can vary depending upon the overall network design and VC configuration. Having a redundant configuration is suggested and will help to minimize recovery times.

Virtual Chassis - The VC will first verify that it is in a state that will allow a successful ISSU upgrade. It will then copy the image and configuration files of the ISSU specified directory to all of the Slave chassis and reload each Slave chassis from the ISSU directory in order from lowest to highest chassis-id. For example, assuming chassid-id 1 is the Master, the Slave with chassis-id 2 will reload with the new image files. When Slave chassis-id 2 has rebooted and rejoined the VC, the Slave with chassis -id 3 will reboot and rejoin the VC. Once the Slaves are complete they are now using the new image files. The Master chassis is now rebooted which causes the Slave chassis to become the new Master chassis. When the original Master chassis reloads it comes back as a Slave chassis. To restore the role of Master to the original Master chassis the current Master can be rebooted and the original Master will takeover, re-assuming the Master role.

Modular Chassis - The chassis will first verify that it is in a state that will allow a successful ISSU upgrade. It will then copy the image and configuration files of the ISSU specified directory to the secondary CMM and reload the secondary CMM which becomes the new primary CMM. The old primary CMM becomes the secondary CMM and reloads using the upgraded code. As a result of this process both CMMs are now running with the upgraded code and the primary and secondary CMMs will have changed roles (i.e., primary will act as secondary and the secondary as primary). The individual NIs can be reset either manually or automatically.

Guidelines - Depending on the topology, the following configuration guidelines can be used to help improve ISSU convergence times and connectivity during ISSU:

- Dual-homed hosts and switches can maintain connectivity during the VC upgrade process.
- Redundant L2 and L3 connections are suggested to help maintain connectivity and reduce recovery times.
- Graceful restart support enabled for OSPF.
- OSPF sub-second flag set: "debug ip ospf set subsecond 1"
- SFP Timer configured: delay=1, hold=2

Supported Upgrade Paths and Procedures

	Upgrading From 8.1.1
OS6860 – VC	ISSU - Supported Standard Upgrade - Supported
OS6860 – Standalone	ISSU – Not Supported Standard Upgrade - Supported
Notes:	If upgrading from 8.1.1.663.R01 maintenance release please contact Service & Support prior to upgrade.

If upgrading a standalone chassis or VC using a standard upgrade procedure please refer to Appendix B for specific steps to follow.

If upgrading a VC using ISSU please refer to Appendix C for specific steps to follow.

Prerequisites

This instruction sheet requires that the following conditions exist, or are performed, before upgrading. The person performing the upgrade must:

- Be the responsible party for maintaining the switch's configuration.
- Be aware of any issues that may arise from a network outage caused by improperly loading this code.
- Understand that the switch must be rebooted and network access may be affected by following this procedure.
- Have a working knowledge of the switch to configure it to accept an FTP connection through the EMP or Network Interface (NI) Ethernet port.
- Read the GA Release Notes prior to performing any upgrade for information specific to this release.
- Ensure there is a current certified configuration on the switch so that the upgrade can be rolled-back if required.
- Verify the current versions of Uboot and FPGA. If they meet the minimum requirements, (i.e. they were already upgraded during a previous AOS upgrade) then only an upgrade of the AOS images is required.
- Depending on whether a standalone chassis or VC is being upgraded, upgrading can take from 5 to 20 minutes. Additional time will be needed for the network to re-converge.

The examples below use various models and directories to demonstrate the upgrade procedure. However any user-defined directory can be used for the upgrade.

If possible, have EMP or serial console access to all chassis during the upgrade. This will allow you to access and monitor the VC during the ISSU process and before the virtual chassis has been re-established.

- Knowledge of various aspects of AOS directory structure, operation and CLI commands can be found in the Alcatel-Lucent OmniSwitch User Guides. Recommended reading includes:
    - Release Notes - for the version of software you're planning to upgrade to.
    - The AOS Switch Management Guide
        - Chapter - Logging Into the Switch
        - Chapter - Managing System Files
        - Chapter - Managing CMM Directory Content
        - Chapter - Using the CLI
        - Chapter - Working With Configuration Files
        - Chapter - Configuring Virtual Chassis

Do not proceed until all the above prerequisites have been met. Any deviation from these upgrade procedures could result in the malfunctioning of the switch. All steps in these procedures should be reviewed before beginning.

Switch Maintenance

It's recommended to perform switch maintenance prior to performing any upgrade. This can help with preparing for the upgrade and removing unnecessary files. The following steps can be performed at any time prior to a software upgrade. These procedures can be done using Telnet and FTP, however using SSH and SFTP/SCP are recommended as a security best-practice since Telnet and FTP are not secure.

1. Use the command '**show system**' to verify current date, time, AOS and model of the switch.

```
6860-> show system
System:
  Description:  Alcatel-Lucent OS6860E-P48 8.2.1.108.R01 Development, July 25, 2015.,
  Object ID:    1.3.6.1.4.1.6486.801.1.1.2.1.11.1.8,
  Up Time:      3 days 21 hours 23 minutes and 2 seconds,
  Contact:      Alcatel-Lucent, http://enterprise.alcatel-lucent.com,
  Name:         OS6860,
  Location:     Unknown,
  Services:     78,
  Date & Time:  MON AUG 03 2015 11:53:38 (UTC)
Flash Space:
    Primary CMM:
      Available (bytes):  847790080,
      Comments          :  None
```

2. Remove any old tech_support.log files, tech_support_eng.tar files:

```
6860-> rm *.log
6860-> rm *.tar
```

3. Verify that the **/flash/pmd** and **/flash/pmd/work** directories are empty. If they have files in them check the date on the files. If they are recently created files (<10 days), contact Alcatel-Lucent Service & Support. If not, they can be deleted.

4. Use the '**show running-directory**' command to determine what directory the switch is running from and that the configuration is certified and synchronized:

```
6860-> show running-directory

CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : WORKING,
  Certify/Restore Status   : CERTIFIED
SYNCHRONIZATION STATUS
  Flash Between CMMs       : SYNCHRONIZED,
  Running Configuration    : NOT SYNCHRONIZED

```

If the configuration is not certified and syncronized, issue the command **'write memory flash-synchro**':

```
6860-> write memory flash-synchro
```

6. If you do not already have established baselines to determine the health of the switch you are upgrading, now would be a good time to collect them. Using the show tech-support series of commands is an excellent way to collect data on the state of the switch. The show tech support commands automatically create log files of useful show commands in the **/flash** directory. You can create the tech-support log files with the following commands:

```
6860-> show tech-support
6860-> show tech-support layer2
6860-> show tech-support layer3
```

It is a good idea to offload these files and review them to determine what additional data you might want to collect to establish meaningful baselines for a successful upgrade.

Appendix B: Standard Upgrade - Standalone/Virtual Chassis

These instructions document how to upgrade an OS6860 standalone or virtual chassis using the standard upgrade procedure. Upgrading using the standard upgrade procedure consists of the following steps. The steps should be performed in order:

1. Download the Upgrade Files

Go the to Alcatel-Lucent Service and Support website and download and unzip the upgrade files for the appropriate model. The archives contain the following:

- OS6860 Image Files **-** Uos.img

2. FTP the Upgrade Files to the Switch

FTP the image files to the *Running* directory of the switch you are upgrading. The image files and directory will differ depending on your switch and configuration.

3. Upgrade the image file

Follow the steps below to upgrade the image files by reloading the switch from the *Running* directory.

```
6860-> reload from working no rollback-timeout
Confirm Activate (Y/N) : y
This operation will verify and copy images before reloading.
It may take several minutes to complete....
```

If upgrading a VC the new image file will be copied to all the Slave chassis and the entire VC will reboot. After approximately 5-20 minutes the VC will become operational.

4. Verify the Software Upgrade

Log in to the switch to confirm it is running on the new software. This can be determined from the login banner or the **show microcode** commmand**.**

```
6860-> show microcode
   /flash/working
   Package          Release                Size     Description
-----------------+-------------------------+--------+-----------------------------------
Uos.img          8.2.1.255.R01            210697424 Alcatel-Lucent OS
```

```
-> show running-directory

CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : WORKING,
  Certify/Restore Status   : CERTIFY NEEDED
SYNCHRONIZATION STATUS
  Running Configuration    : SYNCHRONIZED
```

Note: If there are any issues after upgrading the switch can be rolled back to the previous certified version by issuing the **reload from certified no rollback-timeout** command.

5. Certify the Software Upgrade

After verifying the software and that the network is stable, use the following commands to certify the new software by copying the *Running* directory to the Certified directory.

```
6860-> copy running certified
Please wait……………………………………….

-> show running-directory
CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : WORKING,
  Certify/Restore Status   : CERTIFIED
SYNCHRONIZATION STATUS
  Running Configuration    : SYNCHRONIZED
```

Appendix C: ISSU - OmniSwitch Virtual Chassis

These instructions document how to upgrade an OS6860 virtual chassis using ISSU. Upgrading a VC consists of the following steps. The steps should be performed in order:

1. Download the Upgrade Files

Go the to Alcatel-Lucent Service and Support Website and download and unzip the ISSU upgrade files. The archive contains the following:

- OS6860 Image Files **-** Uos.img
- ISSU Version File – issu_version
- Upgrade Script – OS6860_upgrade

2. Create the new directory on the Master for the ISSU upgrade:

```
6860-> mkdir /flash/issu_dir
```

3. Clean up existing ISSU directories

It is important to connect to the Slave chassis and verify that there is no existing directory with the path **/flash/issu_dir** on the Slave chassis. ISSU relies upon the switch to handle all of the file copying and directory creation on the Slave chassis. For this reason, having a pre-existing directory with the same name on the Slave chassis can have an adverse affect on the process. To verify that the Slave chassis does not have an existing directory of the same name as the ISSU directory on your Master chassis, use the internal VF-link IP address to connect to the Slave. In a multi-chassis VC, the internal IP addresses on the Virtual Fabric Link (VFL) always use the same IP addresses: 127.10.1.65 for Chassis 1,127.10.2.65 for Chassis 2, etc. These addresses can be found by issuing the debug command '**debug show virtual-chassis connection**' as shown below:

```
6860-> debug show virtual-chassis connection
                             Address          Address
 Chas  MAC-Address          Local IP         Remote IP         Status
-----+------------------+--------------------+-------------------+-------------
 1     e8:e7:32:b9:19:0b  127.10.2.65      127.10.1.65      Connected
```

4. SSH to the Slave chassis via the internal virtual-chassis IP address using the password 'switch':

```
6860-> ssh 127.10.2.65
Password: switch
```

5. Use the **ls** command to look for the directory name being used for the ISSU upgrade. In this example, we're using **/flash/issu_dir** so if that directory exists on the Slave chassis it should be deleted as shown below. Repeat this step for all Slave chassis:

```
6860-> rm –r /flash/issu_dir
6860-> rm vc811Issu
```

6. Log out of the Slave chassis:

```
6860-> exit
```

```
logout

Connection to 127.10.2.65 closed.
```

7. On the Master chassis copy the current *Running* configuration files to the ISSU directory:

```
6860-> cp /flash/working/*.cfg /flash/issu_dir
```

8. FTP the new image files and the “issu_version” file to the ISSU directory. Once complete verify that the ISSU directory contains only the required files for the upgrade:

```
6860-> ls /flash/issu_dir

Uos.img       issu_version  vcboot.cfg    vcsetup.cfg
```

9. FTP the “OS6860_upgrade” file to the /flash directory and execute the script. These commands create a file named “vc811Issu” on the /flash directory of all the slaves chassis which indicates ISSU will be performed from 8.1.1.R01 to 8.2.1.R01.

```
6860-> chmod a+x /flash/OS6860_upgrade
6860-> /flash/OS6860_upgrade create
6860-> Please enter password for user admin:

..... Creating vc811Issu in slave chassis id 2
```

10. Upgrade the image files using ISSU:

```
6860-> issu from issu_dir
Are you sure you want an In Service System Upgrade? (Y/N) : y
```

During ISSU ‘**show issu status**’ gives the respective status(pending,complete,etc)

```
6860-> show issu status
Issu pending
```

This indicates that the ISSU is completed

```
6860-> show issu status
Issu not active
```

Allow the upgrade to complete. DO NOT modify the configuration files during the software upgrade. It normally takes between 5 and 20 minutes to complete the ISSU upgrade.

11. Verify the Software Upgrade

Log in to the switch to confirm it is running on the new software. This can be determined from the login banner or the **show microcode** commmand**.**

```
6860-> show microcode

  /flash/working

  Package          Release                Size     Description
```

```
-----------------+------------------------+--------+-----------------------------------
Uos.img          8.2.1.255.R01            210697424 Alcatel-Lucent OS


6860-> copy running certified
Please wait…………………………………….

-> show running-directory
```

```
CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : issu_dir,
  Certify/Restore Status   : CERTIFY NEEDED
SYNCHRONIZATION STATUS
  Flash Between CMMs       : SYNCHRONIZED
  Running Configuration    : SYNCHRONIZED
```

12. Certify the Software Upgrade

After verifying the software and that the network is stable, use the following commands to certify the new software by copying the *Running* directory to the Certified directory:

```
6860-> copy running certified
Please wait…………………………………….

-> show running-directory


CONFIGURATION STATUS
  Running CMM              : MASTER-PRIMARY,
  CMM Mode                 : VIRTUAL-CHASSIS MONO CMM,
  Current CMM Slot         : CHASSIS-1 A,
  Running configuration    : issu_dir,
  Certify/Restore Status   : CERTIFIED
SYNCHRONIZATION STATUS
  Flash Between CMMs       : SYNCHRONIZED
  Running Configuration    : SYNCHRONIZED
```