

Release Notes - Maintenance

OmniSwitch 6860/6860E

Release 8.2.1.335.R01

The following is a list of issues that have been identified and corrected in this AOS software release. This document is intended to be used as a pre-upgrade guide and does not replace the GA Release Notes which are created for every GA release of software.

Contents

Contents 2

Fixed Problem Reports Between Builds 305 and 335 3

Fixed Problem Reports Between Builds 298 and 304 8

Fixed Problem Reports Between Builds 289 and 297 9

Fixed Problem Reports Between Builds 279 and 288 11

Fixed Problem Reports Between Builds 277 and 278 11

Fixed Problem Reports Between Builds 270 and 276 11

Fixed Problem Reports Between Builds 259 and 269 13

Fixed Problem Reports Between Builds 255 and 258 14

Open Problem Reports and Known Issues 14

New Features Introduced in 8.2.1.335.R01 15

New Features Introduced in 8.2.1.297.R01 21

Technical Support 26

Appendix A: General Upgrade Requirements and Best Practices 27

Appendix B: Standard Upgrade - Standalone/Virtual Chassis 31

Appendix C: ISSU - OmniSwitch Virtual Chassis 33

Fixed Problem Reports Between Builds 305 and 335

The following issues were fixed between AOS releases 8.2.1.305.R01 and 8.2.1.335.R01.

PR	Description
198939	<p>Summary: Unable to display correct return attributes which configured on NPS server.</p> <p>Explanation: To display correct return attributes which are configured on NPS server.</p>
202110	<p>Summary: Security vulnerability: Port scanning test provides the information regarding the open "non-well known" port.</p> <p>Explanation: Open port vulnerability addressed for application saaCmm and slbCmm.</p>
210937*	<p>Summary: OS6900 issue with ARP resolution on UNP port.</p> <p>Explanation: Issue occurs when dynamic SAPs ageout and are later recreated. Fixed use-after-free problem that manifests when multiple dynamic non-xlate SAPs are configured on the same interface. Also fixed the mac-learning aging-time issue to allow the max age time configuration of 1000000 sec.</p>
212427	<p>Summary: ICMP traffic is block due to DHCP snooping feature enabled.</p> <p>Explanation: Enhanced the DHCP snooping handling to include support for some clients that send DHCP INFORM packets with "your IP address (yiaddr)" set to zero.</p>
213533*	<p>Summary: portMgrNi main error(2) : [pmAppSendPortStateUpdate:2443] Invalid SetType in App Registration Entry.</p> <p>Explanation: Corrected a logic in port manager to avoid unnecessary error message.</p>
214569*	<p>Summary: OS6900 isis VC crash.</p> <p>Explanation: Properly delete adjacency from database.</p>
214476	<p>Summary: OS6860 Unable to ssh to the switchn using "ssh -l" command.</p> <p>Explanation: Allow SSH to remote host using "ssh -l admin" command.</p>
214757	<p>Summary: OS6860: bcmd sdk info(5) Unit: 0 CDC RX FIFO entry 7 double-bit ECC error.</p> <p>Explanation: Allow the switch to function without ECC error in the presence of packets that have nine preamble/StartOfFrameDelimiter bytes.</p>
214763*	<p>Summary: OS6900: BGP does not advertise Loopback0.</p> <p>Explanation: Allow Loopback0 to be advertised using BGP.</p>
214780*	<p>Summary: 'show interface status' command in Master Split-Topology Switch displays that the interfaces are admin enabled.</p> <p>Explanation: Set the interfaces status to disable if chassis is in SPLIT topology.</p>

215699*	<p>Summary: SPB- If VLAN 1 has an IP interface, Service Port configuration is not loaded after reboot.</p> <p>Explanation: Preserve the original default vlan when configure or unconfigure an access port.</p>
215806*	<p>Summary: Query regarding memory status in a 2XOS6900-X20 VC.</p> <p>Explanation: Update the memory benchmark values after System Ready to avoid out of sync memory stats between CMM and NI.</p>
216256*	<p>Summary: show mac-learning summary timeout due to slNi task in loop</p> <p>Explanation: Workaround to breakout of slNi loop during 'show mac-learning' processing.</p>
216461	<p>Summary: OS6860 code 8.2.1.255.R01 "modify running directory" command clarification.</p> <p>Explanation: The user can change the running directory through "modify running directory" command only if the current running directory is CERTIFIED. If the current running directory is not CERTIFIED then the "modify running directory" will fail with an error.</p>
216492	<p>Summary: Need to check CVE-2016-2108 CVE-2016-2107 CVE-2016-2105 CVE-2016-2106 CVE-2016-2109 CVE-2016-2176.</p> <p>Explanation: Updated openssl package to fix the following vulnerabilities: CVE-2016-2108 CVE-2016-2107 CVE-2016-2105 CVE-2016-2106 CVE-2016-2109 CVE-2016-2176.</p>
216514*	<p>Summary: Switch-1 chassis- slave was not operational in 2XOS6900 VC; however, chassis-2 master remained operational.</p> <p>Explanation: Corrected an invalid string copy which causes a crash in pmmcmd task.</p>
216710	<p>Summary: OSPF route flap issue.</p> <p>Explanation: Update LSA age correctly when we are flooding them to our neighbors.</p>
216721*	<p>Summary: Radius CLI task suspension multiple times in the PRI units of 6860.</p> <p>Explanation: Crash in the RadCli task is fixed, and debug logs for Radius/AAA/LPS is enhanced.</p>
216724*	<p>Summary: SPB service mapped to wrong port after ISSU from 734204.R02 to 734.248.R02.</p> <p>Explanation: Make sure the current shortest path to the next hop is updated, when the reloaded chassis (through which the data path should flow) rejoins the VC during ISSU, to prevent packet loss.</p>
217053*	<p>Summary: OS6860-24: No rebooting reason logged in switch logs.</p> <p>Explanation: Fix a rare occurrence where kernel crash event fails to be logged in the</p>

	swlog.
217319	<p>Summary: Storm control: port is shutdown immediately once there is ONE broadcast packet.</p> <p>Explanation: For the rate set at PPS 244, there was statistics calculation issue. Hence, the port was going to violation even for single packet. This calculation has been taken care in the fix checked in.</p>
217445	<p>Summary: OS10K - Write memory flash-synchro via OmniVista SSH does not show any progress.</p> <p>Explanation: Print the progress of flash synchro with continuous dots.</p>
217605*	<p>Summary: 'aaa test-radius-server' command selects incorrect source and NAS-IP-address IP address.</p> <p>Explanation: Correct NAS IP address and Destination IP address in Radius packets.</p>
217606*	<p>Summary: Switch picks incorrect NAS-IP-address during client authentication.</p> <p>Explanation: Make sure NAS IP address always has the correct format.</p>
217760*	<p>Summary: Unit 2 and Unit 3 in a VC of 3 OS6900 crashed due to Spin lock issue</p> <p>Explanation: Avoid spinlock by always releasing lock before exiting.</p>
217804	<p>Summary: 802.1x authentication failed in OS6860.</p> <p>Explanation: RADIUS client before this fix used a fixed UDP port (998) for communicating with RADIUS server. The usage of this fixed port was due to some restriction in IPNI in past. In this PR the problem was because this UDP port 998 was not allowed in the firewall. As a fix now RADIUS client instead of using fixed UDP port it uses ephemeral port for communicating with RADIUS server.</p>
217903	<p>Summary: SES AAA error(2) Error 3: Operation canceled [in catchAllAndLog()].</p> <p>Explanation: AAA Refresh message not applicable to authentication via RADIUS service.</p>
218046*	<p>Summary: ERROR: Service (11) does not exist.</p> <p>Explanation: Fix overflow display for SPB service.</p>
218055*	<p>Summary: OS10K: Unexpected BGP crash.</p> <p>Explanation: Avoid parsing attributes that are not present.</p>
218147	<p>Summary: Do not forward IPv6 Network Discovery packets with hop limit not equal to 255.</p> <p>Explanation: Drop packets if hop limit does not equal 255.</p>
218298	<p>Summary: OS 6860 UNP user status shows Active with no MAC learnt.</p> <p>Explanation: The Learning Status for the user on UNP+LPS port is fixed to reflect the</p>

	correct status.
218395	<p>Summary: OS10K is not using managed interface for TACACS request after upgrade to 734.248R02 from 732.689R01.</p> <p>Explanation: The TACACS client has no MIP interface and all the configurations to it are governed by AAA module. AAA server (TACACS) configuration takes VRF name as one of the input. During boot up when this server configuration is executed the AAA fails to get VRF name to id mapping and hence the failure. The failure happens because the cslib shared database is not ready or populated.</p>
218434*	<p>Summary: Unable to create static SAP when the dynamic rule is active on the switch.</p> <p>Explanation: Allow creating a static SAP with a dynamic service - no snapshot support yet.</p>
218556*	<p>Summary: Different OID's for AlcatelIND1Base.mib</p> <p>Explanation: New set of MIB's were produced with different MIB file naming convention. ALCATEL-IND1-BASE.mib is not called ALCATEL-ENT1-BASE.mib and similar naming convention for the rest of the Alcatel MIBs. The ALCATEL-ENT1-BASE.mib was further modified to change the name of few objects to have different names than what is in 6.x MIB and to differentiate them from 6.x MIB objects. This was done to allow both (6.x and 7.x) sets of MIB's to co-exists in the customer NMS software. The modified MIB's are not part of the release. They are separately posted on the customer support site.</p>
219200*	<p>Summary: OS6900 STP topology age display issue.</p> <p>Explanation: The Max limit for Topology age as per existing calculation was approximately 268 days. The logic for Topology age calculation has been modified to accomodate higher values.</p>
219392	<p>Summary: When enable accounting , for example 'aaa accounting session FR' all accounting packet type send out from Omniswitch is 'Accounting-On'</p> <p>Explanation: On enabling AAA accounting for ASA users the switch is supposed to send "Accounting On" only once, and subsequently send Accounting Start and Accounting stop on session start and end respectively. This PR was result of a wrong merge which is corrected now.</p>
219631*	<p>Summary: OS6900 switch is not forwarding DHCP packets to server.</p> <p>Explanation: There was an issue handling multiple DHCP servers in ip-helper, when one of them is not reachable, which is fixed.</p>
219642*	<p>Summary: ipmsNi ipms warning error on OS10K in a VC and Crash of the XNI-U32S NI upon forming VC.</p> <p>Explanation: Update multicast flow(index) limit in a VC, depending on the capability of lowest capable NI.</p>
219676	<p>Summary: 6860: No Communication via SPB network.</p>

	Explanation: Reset port group flags on non SPB access port.
219725*	Summary: Unexpected reload / takeover in an OS-6860 VC. Explanation: Fix memory leak in SNMP set operation.
219774*	Summary: 2XOS6900- VM Mac-address are not learned in SPB network; multicast stale entries on takeover during. Explanation: As a result of a takeover or ISSU, the real forwarding state was being modified during the audit between SvcMgr and IPMS-CMM. Ensure that the shadow copy of the forwarding state is used during the entire audit.
219789*	Summary: OS6900 swlog error "svcCmm mMIP error" RCA. Explanation: When customer did SNMP WALK and there is no remote RFP-SPB information OBJECT exist then RFP-SPB application throws errors instead of NO SUCH INSTANCE. This results continuously error logs onto the console. Following are the Problem scenario : 1. When RFP-SPB configuration doesnt not exist, but still user/customer try to do SNMP WALK , we will see this issue. 2. When customer configured RFP-SPB but there is no remote information learnt then we will see this issue. Solution : When there is no remote RFP-SPB information OBJECT exist we should return SNMP NO SUCH OBJECT instead of error.
219932*	Summary: 2xOS10K VC the SFP "JDSU" is not working after the upgrade of VC to 7.3.4.273.R02. Explanation: Fixing VFL error frames by setting proper frame size.
220019	Summary: OS6860 - handling of LDAP QoS policies with keyword "trap". Explanation: Reject rules from LDAP that have trap keyword in them.
220040	Summary: LPS static mac-address binding behavior in the OS6860. Explanation: This fix is required for" 1. Configuring static MAC if the MAC has been learned as Filtering 2. Warning message at the time of configuration of duplicate static MAC. 3. Legends were added in cli "show port-security"
220137*	Summary: Clarification on 2XOS10K VC error messages. Explanation: Remove unwanted swlog errors for linkagg/vfl member ports.

220187	<p>Summary: 4xOS6900VC NTP synchronization issue in slave unit.</p> <p>Explanation: Corrected NTP synchronization between master and slave chassis if NTP was started in a non-default VRF.</p>
220209	<p>Summary: Need to Check Vulnerability for the CVE-2016-5696 for 7x and 8x switches.</p> <p>Explanation: Increased the default value of tcp_challenge_ack_limit to a large number to prevent the attacker from inferring any additional data about the client-server connection.</p>
220341	<p>Summary: OS6860 -UNP : Failure-policy mac-authentication is not working as expected.</p> <p>Explanation: Due to corruption of Radius accounting message, mac-authentication was not working as expected under this scenario.</p>
220345*	<p>Summary: OS6900: Client not receiving IP from DHCP server over GRE.</p> <p>Explanation: OS6900: Client not receiving IP from DHCP server over GRE.</p>
220525*	<p>Summary: BGP cluster id list is not displayed in the 7x and 8x.</p> <p>Explanation: Update BGP cluster ID properly from te update packet.</p>

Fixed Problem Reports Between Builds 298 and 304

The following issues were fixed between AOS releases 8.2.1.298.R01 and 8.2.1.304.R01. This release addresses the “ALE Support Advisory (SA-N0034) - AOS ISSU Postponement” announced on May 20, 2016 and allows for the resumption of ISSU.

PR	Description
215517	<p>Summary: SSH session syslog missing the host name.</p> <p>Explanation: Code modified to pass hostname in the syslog entry for SSH.</p>
216065	<p>Summary: When a Master VC lost power and rejoined a VC of 8, it rebooted 2 times before joining the VC successfully.</p> <p>Explanation: Enabled TCP keep alive on the system to ensure proper socket disconnection and added defensive mechanism on linkagg code to properly handle improper socket disconnection.</p>

Fixed Problem Reports Between Builds 289 and 297

The following issues were fixed between AOS releases 8.2.1.289.R01 and 8.2.1.297.R01.

PR	Description
211328*	<p>Summary: 2XOS6860 LACP flapping issue.</p> <p>Explanation: Corrected endian issue when configuring LACP long timeout.</p>
212121	<p>Summary: OS6860 NTP issue in 8.2.1 code.</p> <p>Explanation: Enabled changing of server parameters.</p>
213084	<p>Summary: OS6860: Able to webview to switch even if webview accesss is disabled.</p> <p>Explanation: Prevent webview access to the switch if webview access is disabled.</p>
212311*	<p>Summary: OS6860 swlogd: lpCmm LanCmmMip info(5) lpTrapPethPsePortOnOff 167: chassisId 8 slot 1 port 4 message flooding switch logs.</p> <p>Explanation: Implemented different logic to suppress informational lpTrapPethPsePortOnOff traps for non-PoE ports.</p>
213223	<p>Summary: CVE-2016-0778 and CVE-2016-0777 has been fixed in the 8.2.1.</p> <p>Explanation: Fixed issues raised by CVE-2016-0778 and CVE-2016-0777.</p>
213380*	<p>Summary: 2xOS6860: swlog ?stpCmm library(plApi) error(2).</p> <p>Explanation: Changes made to avoid port library errors in STP Cmm.</p>
213662	<p>Summary: -2015-7547 and CVE-2015-0235 in AOS 7 & 8 in OS6860 & OS6900.</p> <p>Explanation: Fixed issues raised by CVE-2015-7547 and CVE-2015-0235.</p>
214060*	<p>Summary: On 6860E : Issue with error message when we applied the QoS setup.</p> <p>Explanation: Do not configure QoS rules on a slot which is not part of QoS condition.</p>
214061	<p>Summary: OS686E-Power used "display 0" with web view</p> <p>Explanation: Removed the Power Used column in the Power Supplies page for OS6860 since it is unavailable due to hardware limitations; and updated corresponding help page content.</p>
214073*	<p>Summary: After upgrade from 8.1.1.497 to 8.2.1.255.R01 using RCL successfully, switch still shows running from 8.1.1.</p> <p>Explanation: Error handling is done in the curl script for SFTP during RCL for 6860.</p>
214103	<p>Summary: OS6860 link-fault-propagation error after reboot on 8.2.1 code.</p> <p>Explanation: Linkagg member port of LFP is now added if the configuration is applied</p>

	before system ready.
214326	<p>Summary: 2xOS6860: stpCmm library(plApi) error(2) plGetGportFromIfIndex_f@2751: Get port info (ifIndex -10924)</p> <p>Explanation: Made changes to avoid error messages during port conversion in STP.</p>
214368	<p>Summary: OS6860 switch port going shutdown state when LLDP packet is received.</p> <p>Explanation: When the port was being shutdown due to a violation the output of the command 'show violation' was showing invalid source and invalid reason. This has been corrected.</p>
214382*	<p>Summary: OS6860 Need commands to get the OSPF LSA details</p> <p>Explanation: Corrected byte ordering issue when handling CLI command 'show ip ospf lsdb'.</p>
215065*	<p>Summary: Multicast traffic take long time to recover on ERP ring when it is restore.</p> <p>Explanation: Software was modified to have ERP promptly inform VLAN Manager for any ERP port state change.</p>
215219*	<p>Summary: OS6860 with code 8.2.1.258.R01 filling with "etherCmm library(portmgrlibcmm) error".</p> <p>Explanation: Made changes to avoid port library error messages in interfaces.</p>
215230*	<p>Summary: OS6860: Ni-2 reloaded out of VC-8 and had issues with voice VLAN.</p> <p>Explanation: The fix contains preventive array-out-of-bounds check in message-handler from Master to Slave AgCmm to update UNP users information on slave.</p>
215275	<p>Summary: BW oper value(qos qsi ouput) is not shown correctly for linkagg in OS6860.</p> <p>Explanation: Operational Bandwidth is now shown as a percentage.</p>
215317	<p>Summary: Switch crashed due trapmgr stack while removing snmp configuration from switch.</p> <p>Explanation: SNMP station is now properly deleted when using the 'no snmp station' CLI.</p>
215388*	<p>Summary: OS6860: Incorrect spelling for violation messages by LBD seen on swlogs.</p> <p>Explanation: Correcting spelling error of 'violation'.</p>
215492	<p>Summary: 'unp port 1/1/1 vlan 10' command is accepting in UNP port with edge template. However, we are unable to find in the switch configuration.</p> <p>Explanation: An error message is now displayed when trying to configure unp-port-level static vlan on an UNP port when the port is already attached with an edge-template.</p>

215717*	<p>Summary: NTP source interface not used even after configuration.</p> <p>Explanation: Software updated to use the proper NTP source interface.</p>
215923*	<p>Summary: OS6860 unable to create port monitoring in tag port.</p> <p>Explanation: Added port monitoring check to not allow the App-monitoring port as source port.</p>

Fixed Problem Reports Between Builds 279 and 288

The following issues were fixed between AOS releases 8.2.1.279.R01 and 8.2.1.288.R01.

PR	Description
214005 214421	<p>Summary: Chassis 2-4 in a VC of 7 rebooted. VC crash, reboot and split into 1 Master and 1 failure-shutdown state.</p> <p>Explanation: Fixed pktdrv buffer leak when sending to an invalid port.</p>

Fixed Problem Reports Between Builds 277 and 278

The following issues were fixed between AOS releases 8.2.1.277.R01 and 8.2.1.278.R01.

PR	Description
212889* 213167* 213480*	<p>Summary: Chassis in a VC may drop out or are unable to be synchronized.</p> <p>Explanation: Modified the ISIS-VC LSP expiry handler to improve handling of missed LSP packets.</p>

Fixed Problem Reports Between Builds 270 and 276

The following issues were fixed between AOS releases 8.2.1.270.R01 and 8.2.1.276.R01.

PR	Description
211946	<p>Summary: App-mon production kit is not working after reload and 2nd takeover.</p> <p>Explanation: Application signatures of production kit was not getting detected by app-mon after reload and 2nd takeover though production kit was installed properly. Fix was provided to correct this behavior.</p>
212040*	<p>Summary: 410 Siemens phone not negotiating with 6860.</p>

	Explanation: Clean up stale software context which causes authentication issue.
212071*	Summary: OS6860: Ping issue after applying QoS. Explanation: Match ARP packets correctly to the QoS policies configured.
212325	Summary: High memory due to lpCmm task. Explanation: Incorrect PoE disconnection from peers due to incorrect information being received from a library call has been fixed. Additionally, an unbounded retransmission queue used to communicate with peers during congestion is now limited.
212343*	Summary: High memory due to slbcmmmd on OS6860. Explanation: Disconnected sockets are now properly handled.
212311*	Summary: OS6860 swlogd: lpCmm LanCmmMip info(5) lpTrapPethPsePortOnOff 167: chassisId 8 slot 1 port 4 message flooding switch logs. Explanation: The logged messages are informational and are created due to an SNMP trap being generated when PoE ports are powered/unpowered. This is the correct operation since PDs are not expected to change state while under normal operation. Non-PoE terminals and unterminated PoE ports should not be powered. The lanpower port <chassis>/<slot>/<port> admin-state disable command may be used to disable the 802.3af/at power for ports that are not connected to PDs. Additionally, a range of ports may be given as in the following example: lanpower port 1/1/1-13 admin-state disable .
212477	Summary: Unable to ping IP interface when L2 connection is moved from the primary VC to Slave VC. Explanation: AOS updated to not create a dhcp-client interface in the IP(ni) module when the dhcp-client doesn't have an IP address. This was causing a drop of all the IP packets.
211133*	Summary: kernel: [689541.680000] error writing 94 to 13, read back fffffff5/-11 ret -11 count 5 Explanation: Changed kernel log text to avoid being misinterpreted as error log.
212554*	Summary: OpenSSL and vulnerabilities: CVE-2015-1794, CVE-2015-3193, CVE-2015-3194, CVE-2015-3195, CVE-2015-3196. Explanation: Code updated to OpenSSL 1.0.2e version to fix the listed vulnerabilities.
212712*	Summary: Chassis 1 at 60% CPU due to stpni task in a VC of 8 6860Es. Explanation: Fixed internal software loop in stpNi task which was causing high CPU utilization.
212715	Summary: Flow count for monitor app-list is not incremented if the same flow exists in the monitor flow table.

	Explanation: Fix provided to increment the gross flow count when the same flow has been previously detected and exists in the monitor flow table.
--	--

Fixed Problem Reports Between Builds 259 and 269

The following issues were fixed between AOS releases 8.2.1.259.R01 and 8.2.1.269.R01.

PR	Description
210473*	Summary: Parity Errors caused VC malfunction (chassis 2 not reachable). Explanation: Implemented Broadcom patch to clear the parity error.
210492*	Summary: T6860-P48 issue - Device not able to connect - Parity error BD. Explanation: Implemented Broadcom patch to properly clear the L2_ENTRY parity error.
211220*	Summary: OS6860: VC of 5 and no interfaces seen other than on units 1 and 5. Explanation: Various VC Improvements implemented: a) CPU queueing for VC protocol packets; b) additional logs for VC topology change; c) fix bug of false chassis deletion
211459*	Summary: OS6860: lpNi LanNi error(2) lpNiPollTimer 2227: Bad SendlpNi LanXtr error(2) lp69xGetPowerSupplyParameter 2130: No buffer for send lanpower errors. Explanation: Software updated to recover and reallocate buffer pool memory.
211650	Summary: In a VC setup, using NTP source-ip loopback address, changed the system time, only the master is able to synch up the time with NTP server. The slave unit is not able to synch up. Explanation: Software updated to not use NTP source IP when configuring the NTP clients running on VC slave chassis when connecting to the NTP master on the master chassis.
211687	Summary: After running couple days, 100M Full Duplex stopped sending/receiving traffic, toggle the auto-neg fixed the problem. Explanation: Fixed issue with 100Mbps port not passing traffic.
211884	Summary: Cosmetic Issue - OS6860 LED ports color issue. Explanation: Fixed issue with port's LED color change.
212122	Summary: The byte/packet counts in DPI csv file for long running flows are accumulated in each successive record. They should be incremented. Explanation: Software updated to provide packet/byte increment for interval specific updates in csv file.

Fixed Problem Reports Between Builds 255 and 258

The following issues were fixed between AOS releases 8.2.1.255.R01(GA) and 8.2.1.258.R01.

PR	Description
207292	<p>Summary: Occasionally at boot up the system may display Buffer I/O errors similar to the following. This has not resulted in any functional failures:</p> <pre>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]</pre> <p>Explanation: Update made to the page allocation memory flags.</p>

Please Note: PRs identified with an asterisk have been addressed and are considered to be fixed in AOS. However, the status of the PR may still be in 'Verify'. This is likely due to the issue only being seen in very specific configurations or the issue is seen intermittently making the exact issue difficult to reproduce in a lab environment.

Open Problem Reports and Known Issues

The following issues are identified in AOS Release 8.2.1.R01.

IP over SPB Loopback

On an OS6860, when a packet with a destination address of a router MAC or a MAC associated with a VRRP interface, is received on the SAP side of the loopback cable the VLAN tag in the frame is removed to allow it to be processed properly by the switch. The tag needs to be restored so that the frame can be properly forwarded on the egress SAP by configuring the following:

- VLAN translation must be enabled on the port and SAP level on one side of the loopback cable
- All VLAN tags need to be explicitly configured on the SAP side of the loopback cable

New Features Introduced in 8.2.1.335.R01

Allow UNP-Profile Change Through Re-classification of UNP User Upon IP Address Change

When there are IP-Based UNP classification rule(s) configured in AOS, such that an administrator expects certain users to be learned via one of those UNP classification rules, it is important that the first unknown packet from the user carry the IP address information, matching the configured UNP IP-Based classification rule(s).

In most practical scenarios, such classification rules can't be utilized as desired for learning the UNP users which obtain their IP address dynamically, mainly because the first packet from the user may be either :

1. A non-IP packet or L2 frame (LLDP, EAP etc).
2. Invalid IP-Pkt (Invalid ARP/GARP request/reply or other pkts - one with sender IP: 0.0.0.0 or 169.254.0.0/16).
3. An IP packet (DHCP Discover with srcIP=0, NetBIOS pkts etc) with source IP address other than final IP address that the user is expected to obtain via DHCP or through static IP assignment .
4. Once the user is learned in a UNP profile on the UNP port via any other means (i.e. by using a default-unp-profile on the port and not via the configured IP-based classification rule), the user isn't allowed to change its UNP profile anymore. Meaning it is not allowed to undergo re-classification even if it gets a final IP address leased from a DHCP server (or, when an administrator statically changes the assigned final IP address for the user).

Existing Solution

Issues (1) and (2) above were solved by introducing “force-l3-learning” mode using the command **unp force-l3-learning** in an earlier in 821.R01 release. This mode drops any L2-frames and invalid IP packets from users for learning on UNP Ports. **Note:** This mode was enforced only if there existed at least one IP-based UNP classification rule configured in the system. In this mode only a valid IP packet could be used for learning the user.

Limitations with existing solution:

1. The solution worked only for users with statically assigned IP addresses which sends out a valid IP packet having the final configured IP address which could be used by the switch for learning on UNP ports using configured IP-based classification rule.
2. When “force-l3-learning” mode was enabled globally, supplicants were prevented from being learned in the system, as the EAP/EAPOL being L2-frames were dropped from learning in this mode. Thus, this mode didn't support learning of 802.1x users on UNP ports.

New Implementation:

In order support use case in as in point (3) above, for dynamically assigned IP users, an extension to the earlier solution using “force-l3-learning” mode is required. The new implementation introduces the following enhancements:

- Upon enabling the “force-l3-learning” mode, EAP frames are exempted from being dropped, which allows the learning of supplicants through 802.1x. **Note:** Supplicants should not undergo classification, as 802.1x has a higher precedence over classification on UNP port.
- With “force-l3-learning” mode enabled, once an user is learned through a temporary IP packet in a temporary UNP profile (i.e. default-unp-profile), if there is any change in the IP address of the user (i.e. once the user gets an IP address leased in the learned VLAN), UNP will re-trigger the classification engine. This provides a second chance for the user to get classified using an IP-based classification rule and ultimately classifying the user into a final UNP profile obtained through IP-based classification.

- Upon re-classification of the user there is a choice given to the administrator to either do a port-bounce or simply re-learn the user in the final unprofiled VLAN. This is achieved through a new configuration parameter **port-bounce**, introduced under **force-l3-learning** as an extension to the earlier supported command.

Recommendations:

- It is advisable to have port-bounce enabled, if the final unprofiled VLAN (obtained after re-classification of the user) and the initial unprofiled VLAN (e.g: default-unprofiled vlan) are different.
- Upon re-classification of the user in a new profile and VLAN, the old context of the user will be deleted and newly learned user context updated.
- If port-bounce is enabled, upon re-classification of the user in another profile with different VLAN, it will result in deleting the re-classified context for the user. The current implementation expects the user traffic to be again seen with updated IP address to finally create the context for the user.
- If port-bounce isn't enabled, the solution might work only for the users getting IP addresses leased in the same subnet whether it is learned in the initial unprofiled VLAN-X or in the final unprofiled VLAN-Y. This can happen if the DHCP-Server is configured to lease out IP address based on MAC address rather than subnet.
- Port-bounce isn't supported for UNP linkagg.

There is a provision to enable/disable “force-l3-learning” mode and its associated “port-bounce” configuration at the UNP port and UNP edge-template levels.

Note: If an UNP port has this mode configured through port-level or edge-template level, it will have higher precedence over the global configuration for “force-l3-learning mode”.

CLI Commands**Global Level Syntax**

```
-> unprofiled force-l3-learning {enable|disable} port-bounce {enable|disable}
```

Port-level Syntax:

```
-> unprofiled port <c/s/p> force-l3-learning {enable|disable } port-bounce  
{enable|disable}
```

Edge-template-level Syntax:

```
-> unprofiled edge-template <name> force-l3-learning {enable|disable} port-bounce  
{enable|disable}
```

Usage Guidelines:

- By default “force-l3-learning” is **Disabled**, and “port-bounce” is **Enabled** at Global, UNP Port and UNP Edge-Template levels on the switch.

Note: The default configuration values for “force-l3-learning” and “port-bounce” won't appear in “show config snapshot”

- Executing the global level command for “force-l3-learning” and “port-bounce”, will override the existing port-level config values on all the existing UNP ports/linkaggs which don’t have any UNP edge-template attached.
- Executing the global level command for “force-l3-learning” and “port-bounce”, won’t affect/change the existing config values in any of the UNP edge-templates.
- Executing the global level command for “force-l3-learning” and “port-bounce”, won’t affect/change the existing config values on any UNP Port/Linkagg which has an edge-template applied. These config values on such UNP port/linkagg would only be dictated by changing values in the attached UNP edge-template.
- Any new UNP port/linkagg created in the system would derive the “force-l3-learning” and “port-bounce” config values from globally configured values.
- “force-l3-learning” and “port-bounce” config values on a UNP port/linkagg (without any edge-template attached) can be changed individually by using port-level configuration.
- “force-l3-learning” and “port-bounce” config values on a UNP port/linkagg (with an edge-template attached) can be changed individually by using edge-template-level configuration on the associated edge-template.
- On any UNP port/linkagg, which has configuration for “force-l3-learning” and associated “port-bounce” derived using any of the above mentioned means, the “force-l3-learning” feature/mode would be enforced only when there exists at least one IP-based UNP classification rule configured in the system (either of these IP-Rule, IP + Port Rule, IP + Group-ID Rule, IP + Port + Group-ID Rule, IP + MAC + Port Rule, IP + MAC + Group-ID Rule, Extended Rule using IP condition)

Show Commands:

The configured “force-l3-learning” mode and its associated “port-bounce” status at global, port and edge-template levels are available through following commands:-

Global Level Show

```
-> show unp global configuration
```

Example:

```
-> show unp global configuration
```

```
Mode : Bridge
  Dynamic Vlan Configuration      = Disabled,
  Dynamic Profile Configuration  = Disabled,
  Auth Server Down UNP           = -,
  Auth Server Down Timeout       = 60,
  Auth Server Down VXLAN UNP     = -,
  Auth Server Down VXLAN Timeout = 60,
  Force L3-Learning              = Disabled
  Force L3-Learning Port Bounce = Enabled
```

```
Mode : Edge
  Auth Server Down UNP           = -,
  Auth Server Down Timeout       = 60,
  Redirect Port Bounce           = Enabled,
  Redirect Pause Timer           = -
  Redirect http proxy-port       = 8080
  Redirect Server IP             = -
  Allowed IP                     = -
  Force L3-Learning              = Disabled
  Force L3-Learning Port Bounce = Enabled
```

Port-level Show:

```
-> show unsp port <c/s/p> config
```

Example:

```
-> show unsp port 1/1/13 config
Port 1/1/13
  Port-Type = Bridge,
  802.1x authentication = Disabled,
  802.1x Pass Alternate Profile = -,
  802.1x Bypass = Disabled,
  802.1x failure-policy = default,
  Mac-auth allow-eap = none,
  Mac authentication = Disabled,
  Mac Pass Alternate Profile = -,
  Classification = Disabled,
  Default Profile = -,
  Customer Domain ID = 0,
  Port Control Direction = Both,
  Egress Flooding = Not Allowed,
  Trust-tag Status = Disabled,
  Force L3-Learning = Disabled,
  Force L3-Learning Port Bounce = Enabled,
  802.1x Parameters:
    Tx-Period = 30,
    Supp-Timeout = 30,
    Max-req = 2
```

Edge-Template-level Show:

```
-> show unsp edge-template <name> config
```

Example:

```
-> show unsp edge-template config
Edge Template: ET
  802.1x Authentication = Disabled,
  802.1x Pass Alternate Profile = -,
  Mac Authentication = Disabled,
  Mac-Auth Pass Alternate Profile = -,
  Classification = Enabled,
  Default Edge Profile = -,
  Group-ID = 0,
  AAA Profile = -,
  Redirect Port Bounce = Enabled,
  Port Control Direction = Both,
  802.1x Tx-Period = 30,
  802.1x Supp-Timeout = 30,
  802.1x Max-Req = 2
  802.1x Bypass = Disabled
  802.1x failure-policy = default
  Mac-auth allow-eap = none
  Trust-Tag = Disabled
  Force L3-Learning = Disabled
  Force L3-Learning Port Bounce = Enabled
```

Show Config Snapshot:

-> show configuration snapshot da-unp

Example1:

```
-> show configuration snapshot da-unp
! DA-UNP:
unp edge-profile abc
unp vlan-mapping edge-profile abc vlan 10
unp edge-template ET
unp edge-template ET classification enable
unp linkagg 20 port-type edge
unp port 1/1/13 port-type bridge
unp port 1/1/14 port-type edge
unp port 1/1/14 default-edge-profile "abc"
unp port 1/1/14 force-l3-learning enable
unp port 1/1/15 port-type edge
unp port 1/1/15 classification enable
```

Example2:

```
-> show configuration snapshot da-unp
! DA-UNP:
unp edge-profile abc
unp vlan-mapping edge-profile abc vlan 10
unp force-l3-learning enable
```

Use Cases Supported

The following use cases are supported by this solution:

- Case 1: All UNP ports of the system have only one client connected to them and are not 802.1x supplicants and require the force-l3-learning feature.
 - Enable force-l3-learning using global configuration.
 - Port bounce can be enabled or disabled for force-l3-learning at global level.
- Case 2: All UNP ports of the system have only one client connected to them and are 802.1x supplicants connected to them.
 - Force-l3-learning is not applicable in this case.
 - Port bounce is not applicable in this case.
- Case 3: All UNP ports of the system have a supplicant and a non-supplicant client connected to them. The non-supplicant requires force-l3-learning support with port-bounce
 - Enable Force-l3-learning using global configuration.
 - Port bounce can be enabled or disabled for force-l3-learning at global level.
 - Force-l3-learning will be applicable only for the non-supplicants and clients which went through classification.
 - Since port bounce was enabled and force-l3 learning resulted in vlan change due to UNP profile change, then all clients on the port are affected.
- Case 4: All UNP ports of the system have a supplicant and non-supplicant client connected to them. The non-supplicant requires force-l3-learning support without port-bounce.
 - Enable Force-l3-learning using global configuration.
 - Port bounce can be enabled or disabled for force-l3-learning at global level.
 - Force-l3-learning will be applicable only for the non-supplicants and clients which went through classification.

- Since port bounce was not enabled and if force-l3-learning results in vlan change due to UNP profile change, it may affect clients which needs to get an IP address in the new vlan.
- Case 5: Some UNP ports of the system have only supplicants. Some ports have only non-supplicants, and some UNP ports have both supplicant & non-supplicant clients connected to them.
 - Enable Force-l3-learning using global configuration. This will get applied to all the UNP Ports in the system.
 - Port bounce can be enabled or disabled for force-l3-learning at global level.
 - Disable Force-l3-learning on UNP ports where only supplicants are connected. Supplicants won't go for re-classification.
 - Force-l3-learning will be applicable only for the non-supplicants and clients which went through classification.

Note : Case 1 to 5 can also be configured using per port or edge-template based commands.

New Features Introduced in 8.2.1.297.R01

Transparent Bridging

The transparent bridging enhancement associates NNI ports with all VLANs (1 - 4094) even if they are not created in the switch. Currently AOS can support this by creating all possible VLANs (1 - 4094) and associating them to NNI ports. The transparent bridging enhancement has following advantages over the conventional configuration approach:

- Reduces the administrative effort of configuring VLANs from 1 to 4094 and associated VPAs.

Transparent bridging associates all VLANs from 1 to 4094 to the specified NNI port and spanning tree group 1. This feature is typically limited to a “ring” topology where there are only 2 NNI ports/LAGs per switch.

Related CLI

Global enable and disable of transparent bridging:

```
-> ethernet-service transparent-bridging {enable/disable}
```

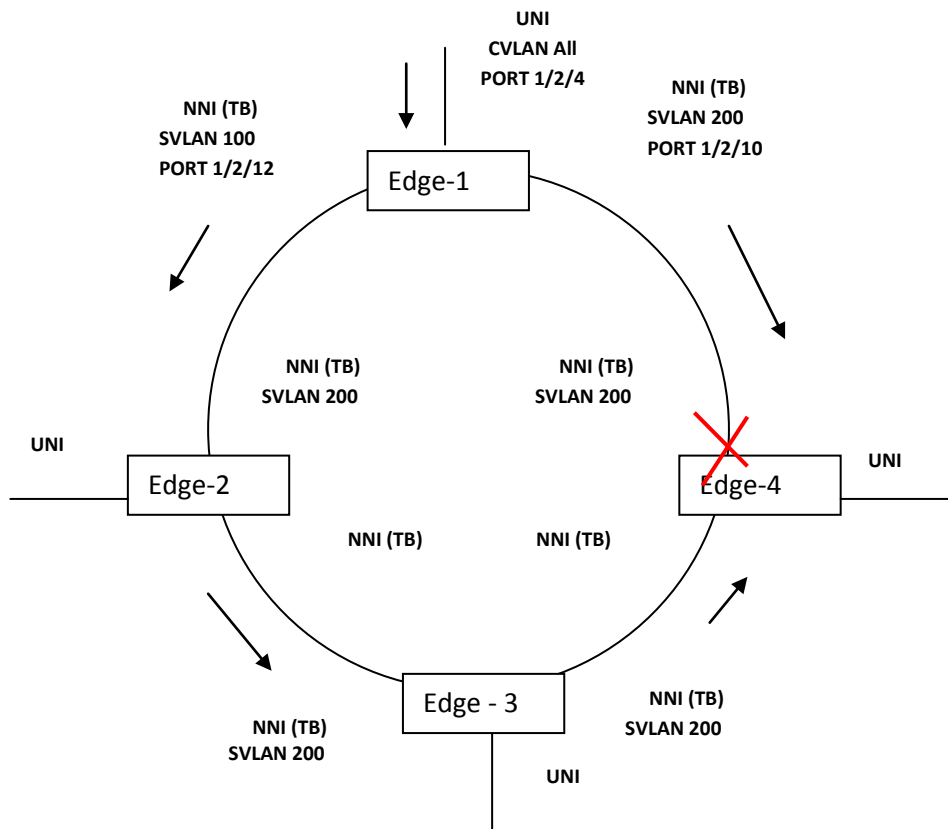
Enable transparent bridging per port:

```
-> ethernet-service transparent-bridging {enable/disable}
-> ethernet-service nni port 1/1/5 transparent-bridging {enable/disable}
-> ethernet-service nni linkagg 5 transparent-bridging {enable/disable}
-> show ethernet-service nni
```

Port	TPID	Legacy BPDU		Transparent Bridging
		stp	mvrp	
2/10	0x8100	Disable	Disable	Enable
2/11	0x8100	Disable	Disable	Enable

```
-> show ethernet-service transparent-bridging
Global Transparent Bridging : disabled,
```

Transparent Bridging - Use Case 1



Transparent Bridging - Use Case 1 Diagram

In the above topology, VLANs 100 and 200 are configured on Edge-1 NNI ports. Only VLAN 200 is configured on all other edge switches NNI ports. On Edge-1 CVLAN 10 is mapped to SVLAN 100. Since transparent bridging is enabled on all the NNI ports of all the edge switches of this topology, though VLAN 100 is not configured on NNI ports of Edges 2, 3 and 4 explicitly, the traffic with VLAN 100 flows through the Edge-2, Edge-3 and Edge-4. Since transparent bridging is enabled only when STP mode is 'FLAT', one of the links in the ring goes to blocking state preventing loops.

Edge -1 Configurations

```
! VLAN :
ethernet-service svlan 100 name "VLAN 100"
ethernet-service svlan 200 name "VLAN 200"
! VLAN STACKING:
ethernet-service svlan 100 nni 1/2/12
ethernet-service svlan 200 nni 1/2/10
ethernet-service service-name "cust1" svlan 100
ethernet-service sap 1 service-name "cust1"
ethernet-service sap 1 uni 1/2/4
ethernet-service sap 1 cvlan 10
ethernet-service transparent-bridging enable
ethernet-service nni port 1/2/10 transparent-bridging enable
ethernet-service nni port 1/2/12 transparent-bridging enable
```

Edge-2, Edge-3 and Edge-4 Configurations

```
! VLAN :
ethernet-service svlan 200 name "VLAN 200"
! VLAN STACKING:
ethernet-service svlan 200 nni <PORT>
ethernet-service svlan 200 nni <PORT>
ethernet-service service-name "cust1" svlan 200
ethernet-service transparent-bridging enable
ethernet-service nni port <PORT> transparent-bridging enable
ethernet-service nni port <PORT> transparent-bridging enable
```

Guidelines:

- Transparent bridging supports both global and port level enable/disable commands.
- If transparent bridging is globally disabled and then re-enabled all existing port level configuration will be automatically re-applied.
- Port level transparent bridging is allowed only when there is at least one SVLAN configured on NNI port.
- Transparent bridging is only supported on NNI ports.
- Transparent bridging can only be configured when STP is configured in flat mode.
- Transparent bridging cannot be configured when STP protocol mode is set to MSTP.
- DHL and transparent bridging are not supported on the same NNI port.

Layer 2 Tunneling Protocol

The new L2TP behavior is as follows:

Protocol Destination MAC: 01:00:0c:cc:cc:cc						
UDLD Global Disable	UNI UN-TAG	UNI TAG	NNI UN-TAG	NNI TAG	LEGACY UN-TAG	LEGACY TAG
Existing Behavior	FWD	FWD	DROP	DROP	FWD	FWD
New Behavior	FWD	FWD	DROP	FWD	FWD	FWD
UDLD Global Enable						
Existing Behavior	DROP	DROP	DROP	DROP	DROP	DROP
New Behavior	FWD	FWD	DROP	FWD	DROP	DROP
UDLD Enabled on Port						
Existing Behavior	TRAP	TRAP	TRAP	TRAP	TRAP	TRAP
New Behavior	TRAP	TRAP	TRAP	FWD	TRAP	TRAP

Existing/New behavior of UDLD Destination MAC

Destination MAC 01:80:c2:00:00:08 (PVSTP)

This MAC is used as the Destination MAC for Provider STP BPDU.

PROTOCOL DEST MAC: 01:80:C2:00:00:08	UNI UN-TAG	UNI TAG	NNI UN-TAG	NNI TAG
Existing Behavior	FWD	FWD	TRAP	TRAP
New Behavior	FWD	FWD	TRAP	FWD

Existing/New behavior of Provider STP Destination MAC

Destination MAC 01:80:c2:00:00:0d (PVGVRP)

This MAC is used as the Destination MAC for Provider GVRP BPDU.

PROTOCOL DEST MAC: 01:80:C2:00:00:0d	UNI UN-TAG	UNI TAG	NNI UN-TAG	NNI TAG
Existing Behavior	FWD	FWD	TRAP	TRAP
New Behavior	FWD	FWD	TRAP	FWD

Existing/New behavior of Provider GVRP Destination MAC

Protocols in UNI Profile.

There is no change in the protocols in the UNI profile.

PROTOCOL DEST MAC	UNI UN-TAG	UNI TAG	NNI UN-TAG	NNI TAG
STP: 01:80:c2:00:00:00	Act as per UNI-Profile (tunnel, drop)	Act as per UNI-Profile (tunnel, drop)	TRAP	FWD
802.1x: 01:80:c2:00:00:03	Act as per UNI-Profile (tunnel,peer,drop)	Act as per UNI-Profile (tunnel,peer,drop)	TRAP	FWD
802.3AD: 01:80:c2:00:00:02	Act as per UNI-Profile (tunnel,peer,drop)	Act as per UNI-Profile (tunnel,peer,drop)	TRAP	FWD
802.1AB: 01:80:c2:00:00:0e	Act as per UNI-Profile (tunnel,peer,drop)	Act as per UNI-Profile (tunnel,peer,drop)	TRAP	FWD
AMAP: 00:20:da:00:70:04	Act as per UNI-Profile (tunnel, drop)	Act as per UNI-Profile (tunnel, drop)	TRAP	FWD
MVRP: 01:80:c2:00:00:21	Act as per UNI-Profile (tunnel, drop)	Act as per UNI-Profile (tunnel, drop)	TRAP	FWD

Existing/New behavior of UNI Profile Protocols

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
European 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 chassis-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.

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	Upgrading from 8.2.1
OS6860 - VC	ISSU - Supported (script file required) Standard Upgrade - Supported	ISSU - Supported Standard Upgrade - Supported
OS6860 - Standalone	ISSU - Not Supported Standard Upgrade - Supported	ISSU - Not Supported Standard Upgrade - Supported
Notes:	<ul style="list-style-type: none"> If upgrading from an 8.1.1 release the additional step of running a script file is required prior to performing an ISSU upgrade. Please see step 9 in Appendix C. If upgrading from 8.1.1.663.R01 maintenance release please contact Service & Support prior to upgrade. ISSU from 8.2.1 forward will not require the use of scripts. Please refer to the Switch Management Guide for additional information on ISSU and managing system files. 	

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 OS6860-48 8.2.1.258.R01 Service Release, November 18, 2015.,
```

```
Object ID: 1.3.6.1.4.1.6486.801.1.1.2.1.11.1.3,
```

```
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: THU NOV 19 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 synchronized, 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 to the 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** command.

```
6860-> show microcode
 /flash/working
Package      Release      Size  Description
-----+-----+-----+-----
Uos.img      8.2.1.335.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 to the 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 (only required when upgrading from 8.1.1)

Note: The following examples use `issu_dir` as an example ISSU directory name. However, any directory name may be used. Additionally, if an ISSU upgrade was previously performed using a directory named `issu_dir`, it may now be the *Running Configuration*, in which case a different ISSU directory name should be used.

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
```

Chas	MAC-Address	Address Local IP	Address 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. (only required when upgrading from 8.1.1) 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** command.

```
6860-> show microcode
 /flash/working
Package      Release      Size  Description
-----+-----+-----+-----
Uos.img      8.2.1.335.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
```