Omni Switch 6450/ 6250 / 6350

Release 6.7.1.137.R04

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

Important Notice: For a copy of software release not posted on the Web or if you have any question or concern please contact Alcatel's Technical Support Department.

Problems Fixed Between Builds 76 and 108

PR	**221471**	Build:	6.7.1.C.21
Summary:	OS6250 does not accept redirect-URL which has more than 128 characters.		
Explanation:	Changes are done to accept 252 characters as the redirect URL length.		

PR	**222426**	Build:	6.7.1.93.R04
Summary:	Switch rebooting with pmd file after entering no snmp command station		
Explanation:	Code changes done to update the community map only if username is present.		

PR	**221826**	Build:	6.7.1.C.21
Summary:	OS6450 - LBD auto-recovery timer beyond 300 secs does not work		
Explanation:	LBD auto recovery timer beyond 300 secs does not work		

PR	**220749**	Build:	6.7.1.C.19
Summary:	unable to ssh/console to the switch and mempart alloc warning messages seen in logs		
Explanation:	Memory leak prevented in IPMS module.		

PR	**222067**	Build:	6.7.1.89.R04
Summary:	Primary switch in the 2 unit stack crashed and generated the PMD file..		
Explanation:	IPMS CMM does a sync of the same flow information repeatedly. When the OMEM buffers are consistently being used and if the sync takes place, more and more memory is allocated which leads to memory depletion. Code changes done to avoid un-necessary flow information sync. .		

PR	**215398**	Build:	6.7.1.C.20
Summary:	OS6450 DHCP packets looping on linkagg between 6900-VC and 6450 stack. Inconsistent DHCP issue during lease refresh with Wyse thin client.		
Explanation:	To bring the logic of identifying if a packet has already been routed in the hardware in 66x like that of 64x.		

PR	**222017**	Build:	6.7.1.97.R04
Summary:	OS6450 multicast source timeout is not proper with slow source packet rate.		
Explanation:	Fix is to timeout the source flows with respect to "ip multicast source timeout" command with more accuracy.		

PR	**220353**	Build:	6.7.1.98.R04
Summary:	Issue with authentication for supplicant 802.1x devices caused by 'reason 38 authentication timeout.		
Explanation:	Removed old authentication context, to enable smooth authentication of the user in fresh context.		

PR	**222666**	Build:	6.7.1.101.R04
Summary:	801x display issue is notice after the takeover.		
Explanation:	Update the secondary cmm database regarding the flushing of entries when the state changes from connecting to disconnected		

PR	**220724**	Build:	6.7.1.86.R04
Summary:	Memory leak notice due to SSH session.		
Explanation:	Memory leak prevented with SW modification		

PR	**223300**	Build:	6.7.1.108.R04
Summary:	Even after the polling interval, Radius server operation status remains down after re-enabling polling with radius-health-check.		
Explanation:	Display correct status of Radius Server when Radius Health Check is enabled		

Problems Fixed Between Builds 109 and 137

PR	**223529**	Build:	6.7.1.111.R04
Summary:	NI out of resource msg handled properly		
Explanation:	NI out of resource handled properly		
PR	**223528**	Build:	6.7.1.111.R04
Summary:	system is busy message when "show configuration snapshot" command is ran even without any active session doing a "show configuration snapshot"		
Explanation:	Code changes done to avoid memory leak in Configuration manager while executing multiple "show configuration snapshot".		
PR	**223544**	Build:	6.7.1.111.R04
Summary:	system is busy message when "show configuration snapshot" command is ran even without any active session doing a "show configuration snapshot"		
Explanation:	Code changes done to avoid memory leak in Configuration manager while executing multiple "show configuration snapshot".		
PR	**223658**	Build:	6.7.1.112.R04
Summary:	EAP failure packet received when disconnect/reconnect supplicant IP Phone		
Explanation:	Auth-server down re-auth time period re-authentication should not be triggered when the Client is classified with Auth-server down Voice policy		
PR	**223530**	Build:	6.7.1.112.R04
Summary:	Client authentication issue even when RADIUS server is reachable		
Explanation:	Last probed time was set to 0 for first time when radius-health check enabled		
PR	**223838**	Build:	6.7.1.112.R04
Summary:	No understandable logs generated from switch when a loop is detected		
Explanation:	Changed level of BPDU debug to lower level		
PR	**223711**	Build:	6.7.1.113.R04
Summary:	high cpu due to webview task during captive portal re-direction in BYOD		
Explanation:	Mac present in BYOD context is removed and added again during re-authentication		
PR	**224320**	Build:	6.7.1.115.R04
Summary:	high cpu due to tWirlpool task		
Explanation:	Avoid high CPU due to twhirlpool task		
PR	**224518**	Build:	6.7.1.116.R04
Summary:	Three EAP failures sent by the switch to supplicant IP phone		
Explanation:	Remove previous mac context before adding new authenticated mac		
PR	**213255**	Build:	6.7.1.118.R04
Summary:	No SNMP trap sent when primary link of linkagg is disconnected		
Explanation:	Delay added in sending trap , when primary port of linkagg is down		
PR	**223577**	Build:	6.7.1.119.R04
Summary:	health-check doesn't work during 4-5 min after a takeover		
Explanation:	Sending a dummy request to radius server every 10 sec from takeover in the newly formed primary as the ip stack is already running and the task has spawned.		

PR	**224568**	Build:	6.7.1.119.R04
Summary:	2xOS6900 - packet loss in adjacent switch (OS6450) connected to slave unit when the slave unit is powered OFF electrically.		
Explanation:	Code changes done to prevent intermittent packet loss in ERP ring when ring goes to protection state.		

PR	**223679**	Build:	6.7.1.119.R04
Summary:	port going down due to STP violation even through no bpdu is received on the port		
Explanation:	Code changes done to arrest port shutdown due to invalid frames		

PR	**225012**	Build:	6.7.1.120.R04
Summary:	onex debugs redirected to swlog		
Explanation:	Onex Debugs redirected to swlog		

PR	**223328**	Build:	6.7.1.121.R04
Summary:	POE firmware (1.90) upgrade		
Explanation:	POE Firmware upgrade to 1.90		

PR	**225273**	Build:	6.7.1.122.R04
Summary:	DDM traps implemented		
Explanation:	DDM traps implemented		

PR	**226138**	Build:	6.7.1.123.R04
Summary:	AOS switch does not generate a logs message when the violation occured due to VRRP/OSPF packet.		
Explanation:	Swlog added for OSPF/VRRP violation		

PR	**225547**	Build:	6.7.1.125.R04
Summary:	show aaa-device all-users unp <user network profile name >, does not list the users associated with a specific user network profile, when UNP mobile rules are applied using an Access Guardian Group Mobility device classification policy		
Explanation:	Code change done to display the list of users associated with the given user network profile given under group mobility rule by filtering based on the group mobility policy type also.		

PR	**226041**	Build:	6.7.1.126.R04
Summary:	After installing KB3212646 in Windows 2012 Radius server, fragmented EAP-TLS header are stripped by the switch to the client.		
Explanation:	Code changes has been done that the Radius packets are processed if the value of EAP fragment is of 1 byte. The data in the packet will be forwarded to the client correctly and hence the authentication will succeed.		

PR	**226615**	Build:	6.7.1.127.R04
Summary:	show running directory" command shows that it is synchronized even though AAA configuration were changed.		
Explanation:	show running directory command shows stack is not synchronised when onex configuration are done.		

PR	**224913**	Build:	6.7.1.127.R04
Summary:	OS6450-U24SXM: User Port is showing up even if it is admin down (100MB SFP).		
Explanation:	Workaround done to power down the port when admin status is down with 100M SFP		

PR	**224136**	Build:	6.7.1.130.R04
Summary:	OV2500 unable to read serial number. of the SEC/SLAVE units power supply.		
Explanation:	Changes done to show serial number of the power supplies of SEC/SLAVE units other than 900W power supplies		

PR	**224051**	Build:	6.7.1.130.R04
Summary:	LBD enhancement		
Explanation:	Loopback detection Enhancement		
PR	**224085**	Build:	6.7.1.131.R04
Summary:	Linkagg timeout with reason " linkAggNi main info(5) lacp_rxm_expired 1/1/9(8)"		
Explanation:	Code changes done to update timeout in linkagg ports during run time		
PR	**223191**	Build:	6.7.1.131.R04
Summary:	Dying Gasp trap not seen randomly after the cold reboot		
Explanation:	Increased the priority of the dying gasp packet which is send to SNMP station.		
PR	**221808**	Build:	6.7.1.131.R04
Summary:	6450 client MAC learnt on 802.1x and MAC table (connected by hub) though the device is disconnected.		
Explanation:	Check added to delete a Mac-address in OnexCmm		
PR	**225608**	Build:	6.7.1.131.R04
Summary:	OS-6450-P10 hanged and rebooted		
Explanation:	Changes done to handle memLeak in taUdldni		
PR	**227285**	Build:	6.7.1.131.R04
Summary:	Systrace Error "taRadiusst [CTRACE] Task 7e2e530 call circ_trace_put3 of task AAA(7e43e20)" clarification		
Explanation:	systrace message in "taRadiusStats " should not access the circular buffer owned by AAA . . Message should log only in SYSTRACE.		
PR	**227348**	Build:	6.7.1.131.R04
Summary:	AOS 6x sending LLC packet with 64 bytes data for AMAP packets -seen as malformed in Wireshark		
Explanation:	Code change done to calculate and send the correct length value in the amap packet.		
PR	**226054**	Build:	6.7.1.134.R04
Summary:	LFP Enhancement		
Explanation:	LFP feature check in		
PR	**227720**	Build:	6.7.1.135.R04
Summary:	Webview is allowing to configure multiple source port, stacking ports, but same is not reflected in CLI and Webview		
Explanation:	Corrected LFP webview port display		
PR	**227498**	Build:	6.7.1.135.R04
Summary:	show ip interface dhcp-client shows admin down while it should show admin up while no response from DHCP Server		
Explanation:	Display the admin state to be enabled even if the ip of dhcp client is 0.		
PR	**227830**	Build:	6.7.1.136.R04
Summary:	Synchronization statistics not updated when configuring SSP		
Explanation:	update the synchronisation statistics while configuring SSP		
PR	**227833**	Build:	6.7.1.136.R04
Summary:	CPU Spike Occurred when LFP source port went to down during SSP		
Explanation:	Code changes done to prevent the loop causing high cpu in LFP		
PR	**227777**	Build:	6.7.1.137.R04
Summary:	show vlan table alignment is not proper		
Explanation:	Code changes has been done that the alignment of "show vlan" is uniform.		

PR	**226404**	Build:	6.7.1.137.R04
Summary:	672R01-MR: After Hard rebooting NI2 CPU spike occurred in NI1, device able to recovered after reload all command		
Explanation:	Changes done to avoid CPU hogging by vstk cmm		

Known Issues:

PR	**222688**	Build:	
Summary:	Traffic is getting dropped while testing open flow with actions drop/out port/set with vlan.		
Explanation:	Expected behaviour: Apply actions are supported only in software, the rate of packets handled in software will depend on the CPU load at that point of time.		

PR	**222243**	Build:	
Summary:	Set source mac bucket action fails to work (for write actions)		
Explanation:	Set source mac action cannot be supported in hardware for both write actions and apply action.		

PR	**222276**	Build:	
Summary:	Untagged frames gets processed with given set of group actions for 'with VLAN tag' flow condition		
Explanation:	Match criteria set as untagged packets in API mode allows tagged traffic as well.		

PR	**222521**	Build:	
Summary:	Functionality of IDLE_TIMEOUT is similar to HARD_TIMEOUT. Flow entry gets removed after the IDLE_TIMEOUT.		
Explanation:	With idle timeout configured, flow is removed from the switch when ideal timeout expires even with continuous traffic flow i.e. Ideal timeout behaves as hard time out.		

PR	**220338**	Build:	
Summary:	High CPU due to the task "taSLNEvent"		
Explanation:	Whenever MAC movement (from one port to another) happens, this is processed in software resulting in higher CPU utilization. This is expected SW behavior.		

PR	**222922**	Build:	
Summary:	In API mode when TOS is set as match criteria , input packet Flows are not matched		
Explanation:	In API mode TOS value cannot be used as a match criteria along with a match criteria requiring a specific Ethernet type.		

PR	**227844**	Build:	
Summary:	Interfaces retaining down state even though physical LED is up after the SSP state comes from Protection to Active		
Explanation:	With Stack Split Protection configured and when the stack rejoins post the split, few connected ports Link/Oper status is shown as DOWN in CLI (ex: show interfaces <slot/port> port), even while the front panel LED status is UP & traffic is flowing. Below workaround can solve this issue:		

1. Admin down / up the interface manually.
2. Physical cable pull and resection.
3. Reloading the particular NI.

PR	**217634**	Build:
Summary:	SFP ports do not come up after a reboot or disconnection of SFP due to uplink ports auto-neg issue.	
Explanation:	There is an interop issue for auto-negotiation working, when 6450 and 6350 are connected using UPLINK ports at both ends. Workaround in this case, is to disable auto-negotiation and force set say, 1000 Mbps full duplex on both ends. This issue is not seen when connection between 6450 & 6350 is done using uplink port on one end and network/user port at another end.	

New Features

1. Source Learning Enable/Disable for Non-Metro units

Platforms Supported: OmniSwitch 6450, OmniSwitch 6250 and OmniSwitch 6350

Hosted AOS SW Release: 671.108.R04

Hash Chain Length enhancement feature is to modify the depth of the hash chain length of FBD table (bucket size) used while writing MAC addresses in ASIC in the AOS platform.

Collisions while writing MAC addresses in ASIC occurs mainly due to poor hashing algorithm being used. Hashing mode XOR has more chances of collision, i.e. same hash-index being assigned to different set of source/destination address. Changing the hash-mode to CRC is a more efficient technique than XOR mode. Since the mac-collision is hardware and algorithm dependent, we may further reduce the probability of mac-collision by using the provision in hardware to increase the bucket size from current set value of 4 to a higher value of 8. This setting can be controlled from our software.

This feature requirement is to change the Hash Length of the FDB table from DEFAULT (bucket size is 4) to EXTEND (bucket size is 8) through a configuration command with mode options as DEFAULT or EXTEND. By default, the Hash Length value will be DEFAULT i.e. bucket size as 4.

The modification of hash length done at runtime will be effective only during boot-up and hence when this command is executed by the user, a warning message is displayed on the console that the change will be effective only after the next reload / reboot of the switch / stack.

Usage:

hash-control chain-length {DEFAULT | EXTEND}

This command configures the hash chain length in the HW. Depending upon this configuration, the hashing bucket size for the hardware table will be decided.

The allowed CLI combinations are as follows:

hash-control chain-length default

hash-control chain-length extend

Syntax Definitions:

DEFAULT	If configured, Hash Chain Length will be set to 4.
EXTEND	If configured, Hash Chain Length will be set to 8.

Defaults:
DEFAULT is the default value for this CLI

Usage Guidelines:

1. Use this command to change the hash chain length from EXTEND to DEFAULT mode and vice-versa.
2. After using this command, save the configurations using “write memory” command and reload the switch to reflect the hash length changes in the switch.

Display Commands
show hash-control chain-length

This CLI displays the configured value for the depth of the hashing bucket

Usage Guidelines
The * symbol displayed in the show output (FDB Hash Chain Length = EXTEND*) indicates that the configured hash chain length will be applied only after reloading the switch.

Example
/* sample output */
➔ show hash-control chain-length
FDB Hash Chain Length = DEFAULT

➔ show hash-control chain-length
(*=new hash chain length config will be applied after reboot)
FDB Hash Chain Length = EXTEND*

Configuration Snapshot
➔ Show configuration snapshot chassis
! Chassis :
hash-control chain-length EXTEND

Limitations:

If any modification in hash chain length is made by the user, it is important to reload/reboot of the switch/stack. Only the configured value will be displayed for SNMP and Webview. Any configuration change made with respect to hash chain length in CLI/SNMP/Webview requires switch reboot to get the configuration changes applied in the switch.

Without doing reboot (after change in hash length), actions like inserting a new NI or doing takeover should not be done. Therefore a warning message is reflected in console to indicate that the user must reboot/reload the switch after the change in hash chain length

2. Source Learning Enable / disable for Non-Metro units

Platforms Supported: OmniSwitch 6450, OmniSwitch 6250 and OmniSwitch 6350

Hosted AOS SW Release: 671.108.R04

Currently we have a limitation in AOS, that enabling / disabling of source-learning can be done only in Metro OmniSwitch units. This feature implementation now allows source-learning enable / disable for non-metro units also.

This provides the user an option to enable or disable source MAC learning on a specified port or linkagg. Any port, except those already activated with 'software' learning, can be set to source learning enabled / disabled by users. This feature is restricted to maximum of 48 ports (including Linkagg ports) across system.

No new CLI is added. Same CLI earlier used for metro units can be now used for non-metro units as well.

Limitations:
None

3. Support for packet modification actions for group type ALL in OpenFlow

Platforms Supported: OmniSwitch 6450 and OmniSwitch 6250

Hosted AOS SW Release: 671.108.R04

Current design of QOS can execute a set of actions for a packet when there is a match for one condition. In order to support a flow which can execute more than one set of actions for a particular condition, we use the concept of Groups and Buckets in OpenFlow. The limitation earlier in Groups and Buckets is that, packet modification actions were not supported. This feature implementation enables the support of packet modification action for group type ALL.

Usage:

debug OpenFlow flow-id <value>

Output will change to reflect the new flow information in the form of buckets

Example;

debug OpenFlow flow-id 5

Flow ID: 5
Logical Switch: LS_1
Priority: 0
Idle Timeout: 0
Hard Timeout: 0
Flow Type: Wildcard
Match: Ingress Port: 0/5, Src MAC: 00:00:00:00:00:01/ff:ff:ff:ff:ff:ff, Dst MAC: 00:00:0d:00:00:01/ff:ff:ff:ff:ff:ff, VLAN Priority: 3, Ether-type: 800
Apply Action(s): VLAN: 142, VLAN Priority: 0x300, Dst MAC: 00:00:00:00:0d:01, Out: 1/1;,
Write Action(s): Group: 1
Bucket: 1
Action: VLAN: 143, VLAN Priority: 0x4, Drop
Bucket: 2
Action: VLAN: 144, VLAN Priority: 0x5, Out: 1/1;
Bucket: 3
Action: VLAN: 145, Out: 1/3; 1/5;

For the single flow match criteria different bucket actions can be applied using Write actions.

Use case:

OpenFlow controller attempts to push the below flow to OpenFlow enabled 6450.

- *in_port=1, actions=output:2,push_vlan:0x8100,set_field:5->vlan_vid,output:3.*

This flow means that: The untagged frame from port 1, should be sent out of port 2 without any modification to the packet. But as a tagged frame out of port 3. For the frame sent out of port 3 having tag id 5 (VLAN 5).

Earlier, in the switch side, validation of the flow fails because packet modification action is not allowed for group type ALL (SET_VLAN of VLAN 5) and the below errors are seen in the switch,

```
>>>>>>
        OFCMM_LOG_GROUP   ofcmm_group_validate:222   OFLS ID: 1 gid: 1 type: 0 bkts count: 96
        OFCMM_LOG_GROUP   ofcmm_bucket_validate:174   OFPGMFC_BAD_BUCKET return
        OFCMM_LOG_PROTO   ofcmm_proto_send_error:108   OFP error msg. Type: 6 code: 12
        OFCMM_LOG_PROTO   ofcmm_proto_send:48   send 4 1 108 0x4
        OFCMM_LOG_PROTO   ofcmm_protov4_rcv_group_mod:1327
        OFPETV4_GROUP_MOD_FAILED. cmd: 0 xid: 4 len: 96
<<<<<<
```

As per Openflow 1.3.1 specification, ALL is a required group type. With this new implementation above use case would be accepted and traffic would work functionally as expected in the use-case.

Limitations:

1. The support for different types of actions can be done only in software so, packets which exceeds the predefined rate limit for OpenFlow (1024 pps) will be dropped.
2. The concept of group type ALL does not work when we try to send the packets to non-primary Nis.

4. Change in the order of actions in reply messages

Platforms Supported: OmniSwitch 6450 and OmniSwitch 6250

Hosted AOS SW Release: 671.108.R04

Multipart flow statistics messages and Multipart group descriptor messages show the list of actions associated with the flow and group respectively. The current implementation fills the actions in the multipart flow statistics reply messages and multipart group descriptor messages in the same order the actions were received in the 'flow mod' message from the OpenFlow controller.

Usage:
The output of "*debug OpenFlow flow-id all*" is expected to display the actions in the order it was received in the 'flow mod' message from the OpenFlow controller.

Example*:*
Debug OpenFlow flow-id 7.
Flow ID: 7
Logical Switch: LS_1
Priority: 0

Idle Timeout: 0
Hard Timeout: 0
Flow Type: Wildcard
Match: Ingress Port: 1/27, Src MAC: 00:00:0a:00:00:18/ff:ff:ff:ff:ff:ff, Dst MAC: 00:00:00:00:00:18/ff:ff:ff:ff:ff:ff, VLAN: 140, VLAN Priority: 5,
Ether-type: 800, Src IP Address: 172.16.0.0/16;
Apply Action(s): Strip VLAN, Push VLAN, VLAN: 145, VLAN Priority: 0x500, Dst MAC: 00:00:00:00:0b:53, IP TOS: 0x28, Out: 1/37;

Actions mentioned under Apply actions need to follow the order mentioned in the controller and the Reply part messages should possess the same order of actions

Limitations:
Openflow version 1.0 does not support concept of Multipart reply for groups. But for individual flows statistics message (OFPST_FLOW) we send out the list of actions associated with the flow in the actual order it was received

5. Support for PUSH_VLAN

Platforms Supported: OmniSwitch 6450 and OmniSwitch 6250

Hosted AOS SW Release: 671.108.R04

PUSH VLAN action inserts new VLAN tags to the incoming packets. It is similar to the Q-in-Q service. Newly pushed tags should always be inserted as the outermost tag in the outermost valid location for that tag i.e. when a new VLAN tag is pushed, it should be the outermost tag inserted, immediately after the Ethernet header and before other tags.

Related functions done as part of this feature:

1. OmniSwitch has to count the number of VLAN tags, number of SET_VLAN and VLAN_PCP actions that the flow is trying to push. When the controller tries to insert a flow which has more than 1 tag, the actions get replaced if it is already present.
2. In a case where there are PUSH_VLAN actions set without SET_VLAN, the switch should respond with an error OFPBACV4_BAD_SET_ARGUMENT.
3. PUSH VLAN cannot be supported in hardware. So, VLAN headers needs to be pushed by the software only.
4. A priority of zero and the tag of zero are used for the new tag inserted using PUSH VLAN.
5. Modifies the output of ‘debug openflow flow-id all” if needed.

Usage:
The output of “debug OpenFlow flow-id all” is expected to display the PUSH VLAN actions received from the OpenFlow controller.

debug OpenFlow flow-id 7.
Flow ID: 7
Logical Switch: LS_1
Priority: 0
Idle Timeout: 0
Hard Timeout: 0
Flow Type: Wildcard
Match: Ingress Port: 1/27, Src MAC: 00:00:0a:00:00:18/ff:ff:ff:ff:ff:ff, Dst MAC: 00:00:00:00:00:18/ff:ff:ff:ff:ff:ff, VLAN: 140, VLAN Priority: 5, Ether-type: 800, Src IP Address: 172.16.0.0/16;

Apply Action(s): Strip VLAN, Push VLAN, VLAN: 145, VLAN Priority: 0x500, Dst MAC: 00:00:00:00:0b:53, IP TOS: 0x28, Out: 1/37;

Limitations:

1. Ether types other than 0x8100 is not supported
2. PUSH_VLAN only adds VLAN tags. The other fields in VLAN tag cannot be modified with PUSH_VLAN.
3. For every push VLAN action there has to be a corresponding set or modify VLAN action field.
4. A maximum of 8 VLAN tags can be inserted to a packet. If we have a case where packets coming in with VLAN tags, and if we try to add more tags which leads to the number of tags in the packet becoming greater than 8, we stop adding more tags to the packets

6. Support for Apply Actions

Platforms Supported: OmniSwitch 6450 and OmniSwitch 6250

Hosted AOS SW Release: 671.108.R04

Earlier, AOS fully supported only the Write Actions instruction type in OpenFlow. Instructions of type *Apply-Actions* were also considered as *Write-Actions* when a flow is saved in the database and action structure used Write Actions. Now with this feature implementation, AOS look at these OpenFlow actions as two separate entities and supports the actions separately for both these types.

Related functions done as part of this feature:

1. Output of “debug OpenFlow flow-id all” should reflect the Apply-Actions also.
2. Multipart reply packets for Apply-Actions maintains the order in which the flow was received.
3. Incoming packets to OF-NI matching a flow which has Apply-Actions applies the actions to the incoming packet in the order the actions were received in the first place.
4. If we have a flow with both Apply-Actions and Write-Actions, we start with Apply-Actions first and then the Write-Actions.
5. If there are more than 2 instructions of the same type, Switch responds with error to controller.

Limitations:

1. Support a maximum of 32 actions for Apply-Actions type in every flow. If the number of actions are greater than 32, switch should respond with an error.
2. Support for Apply-Actions can be done only in software so, packets which exceeds the predefined rate limit for OpenFlow will be dropped.

7. Link Fault Propagation

Platforms Supported: OmniSwitch 6450, OmniSwitch 6250 and OmniSwitch 6350

Hosted AOS SW Release: 671.137.R04

Link Fault Propagation (LFP) feature provides a mechanism to propagate a local interface failure into another local interface. In many scenarios, a set of ports provide connectivity to the network and if all these ports go down, the connectivity to the network is lost. However, the remote end users remain unaware of this loss of connectivity and therefore keeps sending traffic which is unable to reach the network. LFP provides a solution to this problem by shutting down the user ports that connect the user to the network if all the ports providing connectivity to the network are down. LFP monitors a set of interfaces. Another set of interfaces (destination ports) that connect to the network are brought down once all the monitored interfaces (source ports) are down.

Usage:

a. *[no] link-fault-propagation group {<num> | <num-num>}*

This command is used to creates/deletes a Link Fault Propagation Group.

Syntax Definitions:

num	:	Indicates the unique group id. The allowed range is 1-8.
num-num	:	Range of unique group ids.

Usage Guidelines:

1. Range option will only be available for “no” CLI.

b. *[no] link-fault-propagation group <num> {source | destination} {port <slot/port [-port2]> <slot/port [-port2]> | linkagg <aggid [-aggid2]>}*

This command is used to configures the source port(s) and(or) destination port(s) for a group.

Syntax Definitions:

num : Indicates the unique group id.

aggid[-aggid2] : Link Aggregate Identifier. Aggid2 refers to the last aggid in the range of aggregates

slot/por>[-port2] : The slot number for the module and the physical port number on that module (e.g. 2/1 specifies port 1 on slot 2). Port2 refers to the last port in the range of ports.

Usage Guidelines:

1. A group must exist before adding source/destination port(s) to the group.
2. A group can have maximum 48 source ports and 48 destination ports.
3. A maximum of 48 link aggregates will be supported irrespective of the number of ports in the each aggregate in a group.
4. A port/linkagg added as a source port in a group cannot be added as a destination port for this group or any other group.
5. A Port/linkagg added as a destination port in a group cannot be added as a source port for this group or any other group.
6. If port is recovered due to interface recovery timer, then the port will go back to the shutdown state if the error persists.

c. *link-fault-propagation group <num> wait-to-shutdown <value>*

This command is used to configures the value of wait to shutdown timer. The destination ports are shutdown after the expiry of this timer if all the source ports are down.

Syntax Definitions:

num : Indicates the unique group id.

Value : Indicates the value of timer. Allowed range is 0-300 in multiples of 5.

Usage Guidelines:

1. By default, the value is 0.
2. The value of 0 implies that the Wait To Shutdown timer is disabled

d. *link-fault-propagation group <num> [admin-status {enable | disable}]*

This command is used to allows to administratively enable/disable link-fault-propagation on a group(s).

Syntax Definitions:

num : Indicates the unique group id. The allowed range is 1-8..

Usage Guidelines:

1. By default, the admin status of a group is disabled.
2. While group creation admin status option will be available.

e. *show link-fault-propagation group [<num>]*

This command is used to displays a Link Fault Propagation Group.

Syntax Definitions:

num : Indicates the unique group id.

Usage Guidelines:

1. none.

Example

```
show link-fault-propagation group 2
Group Id : 2
  Source Port(s)         : 0/1-2 1/1-5 1/7,
  Destination Port(s)    : 0/3 1/10-13,
  Group-Src-Ports Status  : up,
  Admin Status           : enable,
  Wait To Shutdown       : 10

show link-fault-propagation group 6
Group Id : 6
  Source Port(s)         : 1/2 1/6 1/9,
  Destination Port(s)    : 1/10-11 1/13,
  Group-Src-Ports Status  : down,
  Admin Status           : enable,
  Wait To Shutdown       : 5

show link-fault-propagation group
Group Id : 2
  Source Port(s)         : 0/1-2 1/1-5 1/7,
  Destination Port(s)    : 0/3 1/10-13,
  Group-Src-Ports Status  : up,
  Admin Status           : enable,
```

```
 Wait To Shutdown       : 10

Group Id : 6
 Source Port(s)         : 1/2 1/6 1/9,
 Destination Port(s)    : 1/10-11 1/13,
 Group-Src-Ports Status  : down,
 Admin Status           : disable,
 Wait To Shutdown       : 5

Group Id : 7
 Source Port(s)         : 1/1 1/3,
 Destination Port(s)    : 0/3 1/5 1/7 1/11 1/13 1/15 1/17 1/19 1/21 1/23,
 Group-Src-Ports Status  : up,
 Admin Status           : enable,
 Wait To Shutdown       : 100
```

f. *show interfaces {<slot/port [-port2]> | <slot>} port*

This command is used to displays the admin status and link status for the specified port(s) along with the reason for violation in case link status of port is down.

Syntax Definitions:

slot/port[-port2] : The slot number for the module and the physical port number(s) on that module (e.g., 3/1 specifies port 1 on slot 3). Port2 refers to the last port in the range of ports

slot : Slot number

Usage Guidelines:

2. none.

Example

```
show interfaces 1 port
Legends: WTR - Wait To Restore
       #  - WTR Timer is Running & Port is in wait-to-restore state
       *  - Permanent Shutdown

Slot/   Admin    Link    Violations  Recovery   Recovery     WTR          Alias
Port    Status   Status              Time       Max        (sec)
------+---------+--------+----------+----------+----------+----------+-----------------------------------
 1/1   enable    down    LFP         300        10         0 ""
 1/2   enable    down    none        300        10         0 ""
 1/3   enable    down    LFP         300        10         0 ""
 1/4   enable    down    none        300        10         0 ""
       .
       .
       .
 1/48  enable    up      none        300        10         0 ""
 1/49  enable    down    LFP         300        10         0 ""
 1/50  enable    up      none        300        10         0 ""
 1/51  enable    up      none        0          0          0 ""
 1/52  enable    up      none        0          0          0 ""
```

```
show interfaces 1/1-3 port
Legends: WTR - Wait To Restore
       #  - WTR Timer is Running & Port is in wait-to-restore state
       *  - Permanent Shutdown

Slot/   Admin    Link    Violations  Recovery   Recovery     WTR          Alias
Port    Status   Status               Time      Max        (sec)
------+----------+---------+----------+----------+----------+----------+-----------------------------------
 1/1    enable     down     LFP          300        10         0 ""
 1/2    enable     down     none         300        10         0 ""
 1/3    enable     down     LFP          300        10         0 ""

show interfaces 1/1 port
Legends: WTR - Wait To Restore
       #  - WTR Timer is Running & Port is in wait-to-restore state
       *  - Permanent Shutdown

Slot/   Admin    Link    Violations  Recovery   Recovery     WTR          Alias
Port    Status   Status               Time      Max        (sec)
------+----------+---------+----------+----------+----------+----------+-----------------------------------
 1/1    enable     down     LFP          300        10         0 ""
```

g. *show configuration snapshot link-fault-propagation*

This command is used to displays all the link fault configurations in the CLI.

Syntax Definitions:

1. None

Usage Guidelines:

2. none.

Example

```
PTDUT1-STACK-OF-3-> show configuration snapshot link-fault-propagation
! Link-fault-propagation :
link-fault-propagation group 1
link-fault-propagation group 1 wait-to-shutdown 100
link-fault-propagation group 2 admin-status enable
link-fault-propagation group 7
link-fault-propagation group 8 admin-status enable
```

Limitations:

1. Since vlan stacking is not supported in OS6350, the vlan stacking related configuration will not be supported for Link-fault-propagation whereas it is supported in OS6450.
2. Since ERP is not supported in OS6350, the ERP related configuration will not be supported for Link-fault-propagation whereas it is supported in OS6450.