

CN129E&CN93240YC-FX2 Release Note 9.3.15

已修复 bug 清单

Bug ID	Description
CSCvg06451	<p>Headline: Remove CLI to configure Fill Pattern as only IDLE is supported.</p> <p>Symptoms: The switchport fill-pattern is not a supported configuration on ICN 9000 switches, as only IDLE patterns are supported. Hence, configuration CLI is removed.</p> <p>Workarounds: None.</p>
CSCvx24733	<p>Headline: snmp-server enable traps ospf 1 is removed from show run ospf after reloading the device.</p> <p>Symptoms: After the reload, the snmp-server enable traps ospf command does not show up in the output of the show run ospf command but is present in the show running-config command. Additionally, the command cannot be removed from the running-configuration.</p> <p>Workarounds: Write erase and reload the box.</p>
CSCvy60055	<p>Headline: When object does not exist, no track force is not showing any error message.</p> <p>Symptoms: INOS-CN does not generate error message when removing a non-existent track object from the configuration. No error message is generated even when the object was removed by previous command as there is one more track object in the configuration. Removing track 2 and then removing the track object again from configuration results in INOS-CN generating error message as there are no track objects left in the configuration.</p> <p>Workarounds: This is a cosmetic issue; however, configuration is removed properly.</p>
CSCwc72568	<p>Headline: Generation of non-sysmgr_core while collecting pss dump of a corrupted db.</p> <p>Symptoms: When any application crashes, sysmgr triggerscore collection for the application. The pss db dump for the application is also collected as part of core collection. During the process of core collection, if the pss db of the application is already corrupted (due to some memory corruption), then the collection utility (pss2dump utility) crashes. This leads to generation of a nonsysmgr core along with the application core.</p> <p>Workarounds: None</p>
CSCwd89402	<p>Headline: BFD auth configuration is not displayed in the show run command output after upgrade to ICN INOS-CN Release 9.3(10), without any functional impact.</p> <p>Symptoms: BFD authentication configuration gets lost from running-config after upgrading to ICN 9000 to 9.3(10). However, it is present in the startup-config, and sessions remain UP using authentication.</p> <p>Workarounds: Configure no BFD echo before the upgrade. This works as both preventive and reactive workaround.</p>

CSCwd92065	<p>Headline: The start time and end time exported in the NetFlow are showing incorrect values.</p> <p>Symptoms: Wrong timestamp is found in the NetFlow exported data, where the time is ahead of the actual time or the system uptime.</p> <p>Workarounds: None.</p>																																																																						
CSCwe09300	<p>Headline: Internal BGP routes are getting installed as External routes with an AD of 20 in the Routing Table.</p> <p>Symptoms: installs ECMP route with iBGP and eBGP next-hops, eventhough the maximum pathseibgp command is not configured.</p> <p>Workarounds: Workaround is to re-initiate route installation in RIB. Examples of possible workarounds:</p> <ol style="list-style-type: none">1. Re-announce affected network2. Remove multipath command3. Change bestpath selection algorithm "best path med non-determenestic"4. Restart BGP process																																																																						
CSCwe25343	<p>Headline: ICN 9000 VXLAN virtual peer-link tunnel recovery failure.</p> <p>Symptoms: After a specific failure of the virtual peer-link the tunnel used to traverse CFS traffic may fail to recover after the virtual peer-link is brought backup. The virtual peer-link shows UP but no VLANs forwarding over the peer-link and CFS traffic does not make it through the tunnel.</p> <p>Workarounds: Flapping the uplinks that are used to build the virtual peer-link forces the virtual peerlink to re-initialize and should recover from this state. If flapping the uplinks does not recover the peer-link, a reload is required.</p>																																																																						
CSCwe31550	<p>Headline: LPM entry gets stuck after multiple bgp withdraws.</p> <p>Symptoms: In an ECMP environment where the BGP prefixes have constant withdraws, it is possible for the "single entries" (no ECMP) to get stuck in LPM table. As a consequence, the ECMP entry ends up installed at the end and deferred, because we honor the first one (older entry).</p> <p>show hardware internal tah l3 v4lpm prefix 20.20.48.0/24 table 1" **EPE label *Flags: CC=Copy To CPU, SR=SA Sup Redirect, DR=DA Sup Redirect, TD=Bypass TTL Dec, DC=SA Direct Connect,DE=Route Default Entry, LI=Route Learn Info,HR=Host as Route</p> <table><tr><th>Idx</th><th>vrf</th><th>ip/len</th><th>mpath</th><th>nump</th><th>base/l2ptr</th><th>cc</th><th>sr</th><th>dr</th><th>td</th><th>dc</th><th>de</th><th>li</th><th>hr</th></tr><tr><td>13/711</td><td>1</td><td>x.x.x./24</td><td>0</td><td>0</td><td>0x50005</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr></table> <p>not cleared <<<<<<< "No ECMP" path entry</p> <table><tr><td>7/799</td><td>1</td><td>x.x.x.x/24</td><td>1</td><td>2</td><td>0x400004e1</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr></table> <p>ECMP entry <<<<< new</p> <p>Workarounds: The workaround is as follows:</p> <ol style="list-style-type: none">1. Clearing the route works as a temporary workaround, as this forces INOS-CN to put the ECMP entry at the top. However, if another withdraw occurs, the condition can repeat. clear ip route x.x.x.x/x2. Clearing the LPM entry works as a better workaround. However, if the constant withdraws continue, the condition can repeat. Attach module 1 ; clear hardware internal tahl3 ipv4 tcam-bank ztcam-index y <table><tr><th>Idx</th><th>vrf</th><th>ip/len</th><th>mpath</th><th>nump</th><th>base/l2ptr</th><th>cc</th><th>sr</th><th>dr</th><th>td</th><th>dc</th><th>de</th><th>li</th><th>hr</th></tr><tr><td>Z/Y</td><td>1</td><td>x.x.x.x/x</td><td>1</td><td>2</td><td>0x-----</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr></table>	Idx	vrf	ip/len	mpath	nump	base/l2ptr	cc	sr	dr	td	dc	de	li	hr	13/711	1	x.x.x./24	0	0	0x50005									7/799	1	x.x.x.x/24	1	2	0x400004e1									Idx	vrf	ip/len	mpath	nump	base/l2ptr	cc	sr	dr	td	dc	de	li	hr	Z/Y	1	x.x.x.x/x	1	2	0x-----								
Idx	vrf	ip/len	mpath	nump	base/l2ptr	cc	sr	dr	td	dc	de	li	hr																																																										
13/711	1	x.x.x./24	0	0	0x50005																																																																		
7/799	1	x.x.x.x/24	1	2	0x400004e1																																																																		
Idx	vrf	ip/len	mpath	nump	base/l2ptr	cc	sr	dr	td	dc	de	li	hr																																																										
Z/Y	1	x.x.x.x/x	1	2	0x-----																																																																		

CSCwe35981	<p>Headline: Unicast route create fails and ICN 9000 tahusd crashes.</p> <p>Symptoms: ICN 9000 switch running with "LPM-heavy mode" has Unicast route create failed logs. After 1 second, Service "tahusd" crashes.</p> <p>Workarounds: None.</p>
CSCwe51271	<p>Headline: The show process memory command should be available for network-operator role.</p> <p>Symptoms: The show process memory command is not available for the network-operator role and the Permission denied for the role error is seen.</p> <p>Workarounds: None.</p>
CSCwe55237	<p>Headline: Configuring overlapping IPv6 address with different mask on L3 intf does not throw error.</p> <p>Symptoms: If two IPv6 addresses which would be part of the same logical subnet based on bits used on two different L3 interfaces on the same box using different subnet mask length are configured (an overlap of IPv6 address), no warning message is seen. However, this is not the case with IPv4.</p> <p>Workarounds: Use a subnet calculator to ensure that the addresses do not overlap.</p>
CSCwe65091	<p>Headline: The gnmic reply is missing key-value pairs due to PropertyName=<> is_set=0 and is_default_defined=0.</p> <p>Symptoms: The gnmic client call receives a reply but the payload is missing some key:value pairs that are reported by the show system internal dme running-config all command.</p> <p>Workarounds: None.</p>
CSCwe67953	<p>Headline: Error or incorrect result when computing multiple file hashes simultaneously.</p> <p>Symptoms: The switch throws an error when multiple file hashes are being computed at the sametime by multiple SSH sessions: ck.out Empty / cksum or md5sum or sha256sum or sha512 calculation Error.</p> <p>Alternately, the wrong hash can be reported for some of the files computed.</p> <p>Workarounds: Wait for File A to hash before starting the hash for File B.</p>
CSCwe67996	<p>Headline: Routes with BGP backup route do not get advertised to a BGP peer.</p> <p>Symptoms: BGP backup route does not get advertised to Peer.</p> <p>Workarounds: Clear ip route <>.</p>
CSCwf03457	<p>Headline: Auto-complete for VRF name can cause unexpected config changes.</p> <p>Symptoms: When the first letter of a VRF is typed and the enter key is pressed, if there is only one VRF starting with that letter, the switch may not try to auto-complete the VRF name depending on the command that is used. For commands such as "show run vrf" or "vrf member" the auto-complete feature does not work but for the command "novrf context" the auto-complete feature works as expected. This means that the switch will delete the VRF starting with that letter instead of showing a message that the VRF doesn't exist, this behavior can lead customer to delete a VRF by mistake causing network disruptions.</p> <p>Workarounds: The workaround is as follows;</p> <ol style="list-style-type: none"> 1. Write all VRF names while deleting a VRF. 2. Use TAB key to auto-complete, however, sometimes this is not possible, for example, when using scripts for automation.

CSCwf19968	<p>Headline: vsh.bin fails after setting a SPAN capture with thousands of source VLANs in a single session.</p> <p>Symptoms: When configuring thousands of source VLANs the device becomes unresponsive, the switch hangs and closes the ssh session, then the vsh.bin process reloads unexpectedly.</p> <p>Workarounds: Avoid setting a higher amount of source VLANs as supported. The SPAN limit for the total source VLANs supported in a single session is 32 VLANs.</p>
CSCwf21554	<p>Headline: In ICN INOS-CN, special character ">" causing issues with clear-text key-string in key chain.</p> <p>Symptoms: Using a key-string with ">" character may cause the parser to not capture the string, resulting in an empty string and missing configuration.</p> <p>Workarounds: Avoid using the ">" character in a clear-text (pre encrypted) key-string configuration.</p>
CSCwf27236	<p>Headline: The show file INOS-CN.9.3.10.bin md5sum NXAPI through curl causes further request to fail server busy.</p> <p>Symptoms: When show file INOS-CN.9.3.10.bin md5sum is issued to the ICN switches through a curl statement and if any further statement is sent, then the request fails with an error message.</p> <p>Workarounds: Access the ICN switch using the feature bash as follows: config(switch)# feature bash config(switch)# run bash sudo su bash-4.3# echo 3 > /proc/sys/vm/drop_caches bash-4.3# echo 3 > /proc/sys/vm/drop_caches bash-4.3#</p>
CSCwf32297	<p>Headline: md5sum verification error with Value too large for defined data type.</p> <p>Symptoms: When checking md5sum for any file > 2 GB while running an impacted 32-bit INOS-CN image, the md5sum validation fails with the following error: # show file INOS-CN64-cs.10.3.2.F.bin md5sum Error: /bootflash/INOS-CN64-cs.10.3.2.F.bin: Value too large for defined data type</p> <p>Workarounds: None</p>
CSCwf33807	<p>Headline: Kernel logs are saved in the tmp_logs directory.</p> <p>Symptoms: After a crash in the FEX and after a reload, the kernel traces are not saved in the platform.</p> <p>Workarounds: None.</p>
CSCwf34708	<p>Headline: ICN INOS-CN installation with no-reload option can cause BFD down.</p> <p>Symptoms: BFD cannot go up after upgrading OS as follows:</p> <ol style="list-style-type: none"> 1. Run the install all INOS-CN [os image] no-reload command on all switches. 2. Reload one of the switches. 3. BFD neighbor is down and cannot come up until reload of the opposite switch. Besides, OSPF neighbor with BFD is still in full state even though BFD is down. <p>Workarounds: Reload other switches.</p>

CSCwf37914

Headline: Reduce severity of syslog when the mc-drop command is enabled.

Symptoms: After enabling hardware qos pf mc-drop, a lot of sev2 messages, that are strictly informational, flood the logs.

Workarounds: Change global logging level (not recommended) Logging discriminator.