

## 17.3R3-S8 - List of Known issues

PR Number	Synopsis	Category: ESWD
<a href="#">1192520</a>	GARPs being sent from the switch once in 10 minutes Product-Group=junos	GARPs were being sent whenever there is a mac (fdb) operation (add or delete). This is now updated to send GARP when interface is UP & I3 interface attached to the vlan.
PR Number	Synopsis	Category: Kernel
<a href="#">1376362</a>	The kernel crash when GRES configuration is enabled and committed Product-Group=junos	Mid Scale Licenses (ML) is installed and GRES is not enabled in the chassis. This means, forwarding information base (FIB) scale would be 512k in primary and 256k in backup. When the primary RE is operating at scale beyond 256K (e.g 300K), Under this state, once GRES configuration is enabled and committed, a ksyncd core and a vmcore can be seen on backup RE.
PR Number	Synopsis	Category: EX4300 PFE
<a href="#">1369461</a>	ex4300 : firewall filter applied on L3 interface is not counting eBGP packets Product-Group=junos	firewall filter applied on L3 interface cannot count eBGP packets eBGP packets comes with ttl=1 so they cannot hit the interface filters. We can catch ttl=1 packets only on loopback filter. Ttl=1 packets are not treated as routable packets. This is an expected behavior in case of ttl=1.
PR Number	Synopsis	Category: QFX Multichassis Link Aggregate
<a href="#">1454764</a>	ARP reply unicast packets might be flooded to all interfaces in the vlan Product-Group=junos	Flooding of ARP reply unicast packets is seen as a result of an ARP request sent for the device's VRRP MAC address. The ARP reply, which is flooded in the VLAN by the device, has the correct DMAC of the originator of the ARP request. In other words, the ARP reply is flooded but with the correct unicast DMAC. The ARP reply is not broadcasted. It is independent of MCLAG and VRRP scenario.
PR Number	Synopsis	Category: QFX PFE L2
<a href="#">1339348</a>	On QFX platforms the Q-in-Q encapsulated IFL interface shows all statistics as zero Product-Group=junos	Statistics at the logical level (IFL) are not supported in QFX platforms for layer 2 interfaces. This limitation is not specific to service-provider or enterprise style configuration. It is valid for all layer 2 interfaces, trunk or access.

<a href="#">1447333</a>	Packet might be dropped on QFX5100/QFX5200/EX4600 Product-Group=junos	On QFX5100/QFX5200/EX4600, in Junos 18.1 release, if there are multiple layer 3 sub interfaces on multiple interfaces, packet might be dropped when multiple layer 3 subinterface config added and rollbacks happens.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: QFX L3 data-plane/forwarding</b>
<a href="#">1407175</a>	JDI-RCT : QFX-5100 VC/VCF : Observed Error BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(),128:l3 nh 6594 unintsall failed in h/w with Mini-PDT base configurations Product-Group=junos	On QFX5100 Virtual Chassis or Virtual Chassis Fabric, when the mini-PDT-base configuration is issued, the following error message is seen in the hardware: BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(), 128:l3 nh 6594 unintsall failed. There is no functionality impact because of this error message.
<a href="#">1443507</a>	IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present Product-Group=junos	On all Junos platforms which are enabled with MultiChassis Link Aggregation Group (MC-LAG), if there are multiple Integrated Routing and Bridging (IRB) interfaces present, and the Inter Chassis Link (ICL) is also connected over an IRB interface, when both MC-LAG peers have not learnt link-local addresses and IPv6 ping is firstly initiated from the remote peer, the Neighbor Solicitation (NS) packet might take ICL path and couldn't get answered properly.
<a href="#">1444845</a>	CRC errors might be seen on QFX5100-VC Product-Group=junos	In QFX5100 Virtual Chassis(VC) scenario, if the VC connections are disconnected for any reason, like rebooting the switch or pulling out the optical module, the CRC errors and packets loss might be seen when the VC connections resume working again. Due to the VCP ports are not getting initialized properly.
<a href="#">1459329</a>	OSPF stuck at Exchange state for lag interfaces in a QFX5100-EX4300 mixed VC setup whose child members belong to EX4300 after Rebooting VC. Product-Group=junos	OSPF stuck at Exchange state for lag interfaces in a QFX5100-EX4300 mixed VC setup whose child members belong to EX4300 after Rebooting VC.
<a href="#">1475825</a>	IPv6 NA (ff02::1 - All node multicast) is not forwarded through QFX5110-VC with IRB. Product-Group=junos	IPv6 NA packets are being dropped on FPC2 due to incorrect filter (icmpv6-nd) is hit which has drop action. IPv6 NA packets must hit -ipv6_linklocal dynamic filter. -ipv6_linklocal filter is not hit because InPorts qualifier is incorrect.
<a href="#">1485612</a>	FPC may go to "NotPrsnt" state after upgrading with non-tvp image in VC/VCF setup Product-Group=junos	On EX4600/QFX5100 platform, there are two types of PIC (Physical Interface Card). The first one is PIC with the integrated PHY capability (called PHYLESS). The second one is PIC with an external PHY capability (called PHY). If VCPs (Virtual Chassis Port) are configured on external

		PHY capability PIC(s), the FPC(s) might go to "NotPrsnt" state after upgrading with non-tvp image in VC/VCF setup. The affected FPC(s) cannot be used to forwarding traffic.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: JUNOS kernel/ukernel changes for ACX</b>
<a href="#">1076943</a>	SNMP MIB walk/get/set on jnxDomCurrentTable and jnxDomNotifications might fail on ACX platforms Product-Group=junos	SNMP MIB walk/get/set on jnxDomCurrentTable and jnxDomNotifications might fail on ACX platforms while relevant CLI shows proper output.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: ACX MPLS</b>
<a href="#">1488614</a>	Local switching in VPLS in ACX does not carry the inner vlan with input-vlan-map pop and output-vlan-map push operation Product-Group=junos	In VPLS local switching scenario, the packet can never be double tagged even if we have a push operation in output vlan map. As it is a Product limitation.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: "agentd" software daemon</b>
<a href="#">1363199</a>	Streaming telemetry data might not be received by one client when two clients subscribe to the same path with same frequency at the same time. Product-Group=junos	Streaming telemetry data might not be received by one client when two clients subscribe to the same path with same frequency at the same time. This is because there could be a potential race condition for sensor ID allocation.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: MPC Fusion SW</b>
<a href="#">1508794</a>	A Regression issue introduced by PR1463859 causing WAN-PHY interface continuously flaps with default hold-time down of 0 Product-Group=junos	PR1463859 introduces a software defect that causes a 10GE interface to flap continuously when configuring with the WAN-PHY framing with the default "hold-down" timer (0). Once upgrading a router to an affected software release, the interface may flap continuously. This is not applicable to an interface with the default framing - LAN-PHY.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: MX Layer 2 Forwarding Module</b>

<a href="#">1424635</a>	RE kernel crashes may be seen in EVPN scenario when proxy arp is enabled Product-Group=junos	In EVPN scenario when proxy arp is enabled and the lowest 3 bytes of irb logical interface's address are zero, RE kernel crashes may be seen if IPv6 address is configured on irb logical interface and then got removed.
<a href="#">1451959</a>	Junos OS: MX Series: Crafted packets traversing a Broadband Network Gateway (BNG) configured with IPv6 NDP proxy could lead to Denial of Service (CVE-2020-1633) Product-Group=junos	Due to a new NDP proxy feature for EVPN leaf nodes introduced in Junos OS 17.4, crafted NDPv6 packets could transit a Junos device configured as a Broadband Network Gateway (BNG) and reach the EVPN leaf node, causing a stale MAC address entry. This could cause legitimate traffic to be discarded, leading to a Denial of Service (DoS) condition. Refer to <a href="https://kb.juniper.net/JSA11012">https://kb.juniper.net/JSA11012</a> for more information.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: BBE interface related issues</b>
<a href="#">1480154</a>	commit check is needed for vlan-overlap configuration Product-Group=junos	commit check is needed for vlan-overlap config
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: MIBs related to BBE</b>
<a href="#">1476596</a>	The bbe-mibd might be crashed on MX platform in subscriber enviroment Product-Group=junos	On all MX platform with subscrirber enviroment, when about 600 subscribers on line and Simple Network Management Protocol (SNMP) query OID from IPv6IfTable only constantly, memory leaking with leaking size about 1KB each time might be happened. The bbe-mibd will be crashed as memory leaking size keep increasing. Broadband Edge (BBE) related Management Information Base (MIBS) would be affected during bbe-mibd self-restart.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Bi Directional Forwarding Detection (BFD)</b>
<a href="#">1354409</a>	AE interface and BFD session remain down after interface disable/enable Product-Group=junos	With Bidirectional Forwarding Detection (BFD) configured on an aggregated Ethernet interface, if you disable/enable the aggregated Ethernet interface, then that interface and the BFD session might not come up.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Border Gateway Protocol</b>

<a href="#">1265504</a>	BGP MIBv2 enterprise MIB objects for InetAddress types not properly generating OIDs Product-Group=junos	When generating SNMP traps or notifications for BGP events from the jnxBgpM2 MIB, Junos OS does not emit objects of type InetAddress with the expected length field. This will cause compliant SNMP tools to be able to parse the contents of those objects properly. In particular, the length field for the InetAddress OBJECT-TYPE is omitted. Using the set protocols bgp snmp-options emit-inet-address-length-in-oid command causes these objects to emit in a compliant fashion. Given the length of time that this error has been in place, it was decided to leave the existing non-compliant behavior in place to avoid breaking tools that had accommodated the existing behavior as the default.
<a href="#">1396344</a>	Processing a large scale as-path regex will cause the flap of the route protocols Product-Group=junos	When 'as-path-group' is configured under BGP, if a configuration with a large scale as-path regex is committed, the route protocols flap might be seen.
<a href="#">1432100</a>	The "dead" next-hop might stay in the forwarding table in a BGP-LU scenario after the primary interface recovers Product-Group=junos	In BGP Labeled Unicast (BGP-LU) scenario, if the device works as penultimate hop and receives BGP-LU routes with indirect next-hop from an egress router, after the operational next-hop interface corresponding to those labeled routes flaps, a "dead" next-hop type (discard action is performed for this type) may be set for the related clone routes (s=0) and still there even the next-hop interface is operational again.
<a href="#">1481589</a>	The rpd process might crash with BGP multipath and route withdraw occasionally Product-Group=junos	In BGP multipath scenario on all Junos platforms, if original active/best BGP multipath routes got flapped occasionally, some routes, which are changed from best routes to alternative now, are removed from the BGP multipath route list, but its multipath flag are not cleared, due to some temporarily incorrect route reorder in BGP. And when the currently active route got withdrawn for some reason, and the original best routes again become the best routes then rpd might hit an assert and crash.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: BBE Remote Access Server</b>
<a href="#">1402012</a>	The authd crash might be seen due to a memory corruption issue Product-Group=junos	In subscriber scenario, the authd might crash multi-times due to a memory corruption issue.
<a href="#">1402653</a>	The subscribers might need to take login retry in the scenario with high usage of the address pool Product-Group=junos	On all Junos platforms running in subscriber scenario with address pool configured, if the address pool has high usage so that only few addresses are free to allocate, when there are subscribers logout and login, the released address by one subscriber might get re-used and

		allocated to the other subscriber very soon. Due to this issue, syslog error messages might be seen, and the affected subscriber might need to take retry for login.
<a href="#">1449064</a>	Subscribers login fails when PCRF server is unreachable Product-Group=junos	In Gx-Plus for Provisioning Subscribers scenario, when the PCRF (Policy and Charging Rules Function) server is unreachable or the diameter protocol is down, the subscriber login might fail to successfully establish a session or the subscribers might fail to bind a service policy by Gx-Plus after the PCRF Server connectivity is restored.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Cassis pfe microcode software</b>
<a href="#">1298161</a>	FPC ppe_img_ucose_redistribute Failed to evict needed instr to GUMEM error reported randomly Product-Group=junos	In some MX Series deployments running Junos OS, random syslog messages are observed for FPC cards: "fpcx ppe_img_ucose_redistribute Failed to evict needed instr to GUMEM - xxx left". These messages are not an issue and might not have a service impact. These messages will be addressed as INFO level messages. On a Packet Forwarding Engine, there are dedicated UMEM and shared GUMEM memory blocks. This informational message indicates some evicting events between UMEM and GUMEM and can be safely ignored.
<a href="#">1303489</a>	MPC Major alarm, with logs: XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) Product-Group=junos	In some scenarios with MPC, Major alarm and following messages are generated. this Major error is triggered due to parity error, and the impacted queue might drop packets, This might impact the forwarding, to recover MPC card need to be rebooted messages log: fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) alarmd[3158]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC x Major Errors Major alarm set, FPC x Major Errors fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) cli> show chassis alarms 1 alarms currently active Alarm time Class Description 2019-01-25 15:18:03 UTC Major FPC x Major Errors
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: MX Platform SW - UI management</b>
<a href="#">1498538</a>	SNMP polling does not show correct PSM jnxOperatingState when one of the PSM Inputs failed Product-Group=junos	SNMP polling does not show correct PSM jnxOperatingState when one of the PSM Inputs failed or not exists
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Class of Service</b>

<a href="#">1359767</a>	Configuring host-outbound-traffic under class of service might cause certain devices to stop. Product-Group=junos	When host-outbound-traffic is configured under class-of-service on an affected platform, a corruption of the TTP packets related to class-of-service marking on the PFE can cause the device to repeatedly crash.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Firewall Filter</b>
<a href="#">1473093</a>	Traffic might not be forwarded into the right queue but the default queue when VPLS traffic has three or more VLAN tags with VLAN priority 5 Product-Group=junos	On the MX platform with MPC line card (except DPC line card) used, if an input firewall filter is configured at the ingress VPLS interface, the packet with a VLAN priority of 5 with three or more VLAN tags might be forwarded into the wrong queue. When this occurs, it might cause traffic loss due to congestion as all traffic is forwarded into the default queue.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Covers Application classification workflows apart from custo</b>
<a href="#">1455465</a>	The traffic loss might occur when application service is configured Product-Group=junos	On vSRX3.0 platform, traffic loss might occur when application service is configured.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: EVPN control plane issues</b>
<a href="#">1362222</a>	EVPN type-5 route might be lost if knob "chained-composite-next-hop ingress no-evpn" is configured Product-Group=junos	In the scenario of EVPN type 5 route with MPLS encapsulation for EVPN-MPLS on MX Series platforms, if the "chained-composite-next-hop ingress no-evpn" statement is configured, the EVPN type 5 route might be lost in the EVPN routing table.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Express PFE FW Features</b>
<a href="#">1312874</a>	Discrepancy in 'show pfe filter hw summary" output in 17.2R1 release Product-Group=junos	The "show pfe filter hw summary" command was displaying proper results in 17.2R1. This PR fixed the issue and the output will look like: regress@velit-r0> show pfe filter hw summary Slot 0 Chip Instance: 0 HW Resource Capacity Used Available ----- ----- Filters 8191 2 8189 Terms 65536 81 65455 {master:0}
<a href="#">1372944</a>	Junos OS: PTX1000 and PTX10000 Series, QFX10000 Series using non-AFT architectures: A specific genuine packet inspected by sFlow will cause	An improper initialization of memory in the packet forwarding architecture in Juniper Networks Junos OS non-AFI/AFT platforms which may lead to a Denial of Service (DoS) vulnerability being exploited when a genuine packet is received and inspected by non-AFT/AFI sFlow and when the

	a reboot. (CVE-2020-1617) Product-Group=junos	device is also configured with firewall policers. Refer to <a href="https://kb.juniper.net/JSA11000">https://kb.juniper.net/JSA11000</a> for more information.
<a href="#">1380917</a>	FPC might crash on PTX or QFX10000 after lo0 filter change Product-Group=junos	On PTX or QFX10000 platform with vmhost based FPC (i.e. PTX-FPC3, PTX1000), FPC crash might be seen after changing filter under lo0.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: SRX1500 platform software</b>
<a href="#">1403727</a>	Throughput or latency performance of all traffic drops when TCP traffic is passing through the device. Product-Group=junos	On vSRX, SRX1500, SRX4100, SRX4200 and SRX4600 platforms, when TCP Traffic is passing through the device for a certain period, throughput performance of all traffic is dropped about two thirds and latency performance of all traffic is increased up to around 20 ms.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: PTX Express ASIC interface</b>
<a href="#">1428307</a>	Interface does not come up after interface flapping and FPC reboot Product-Group=junos	In PTX with FPC3-PTX and QSFP28 PIC, or MX platforms with EA/ZT-chip based line cards, one of the interfaces on them might not come up after an interface of peer device flapping in short intervals and then restart the local FPC. Due to the BCM8238x chip of Broadcom with a wrong re-timer leading to the local interface remain in "down" state.
<a href="#">1445473</a>	The PTX BCM firmware needs to be upgraded to DE2E Product-Group=junos	On the PTX5000/PTX3000 router, the t6e-pic (15x100G PIC, 96x10G PIC) Broadcom firmware needs to be upgraded to DE2E through Junos software upgrade. The firmware upgrade process will take up to 5 minutes before the firmware upgrade complete interface will stay in downstate.
<a href="#">1453217</a>	The 100G interface might not come up after flapping on PTX5000 Product-Group=junos	On PTX5000 Router, the 100-Gbps interface might not come up after flapping due to optic reliability issues.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: PTX Express ASIC platform</b>
<a href="#">1384435</a>	An enhancement of optimizing the report to the single bit error check Product-Group=junos	Due to transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary hardware replacements. This change applies to all Platforms using Hybrid Memory Controller (HMC).

PR Number	Synopsis	Category: Kernel software for AE/AS/Container
<a href="#">1370015</a>	Kernel crash might be seen after committing demux related config Product-Group=junos	In subscriber management scenario, if an AE interface is associated as the underlying-interface of a demux0 unit and both demux0 unit and AE unit (corresponding to the above AE interface) are configured with a duplicated VLAN id, kernel may crash after committing the config. Same core-dump may appear even in a non subscriber-management environment, while adding/modifying an IFL (Any logical link) which has the duplicate VLAN configuration or a bulk configuration.
<a href="#">1390367</a>	Traffic destined to VRRP VIP gets dropped as filter is not updated to related logical interface Product-Group=junos	On MX platform with enhanced-ip and VRRP configured, if remove/add a child link from AE bundles, traffic destined to VRRP VIP might be dropped.
<a href="#">1440033</a>	Traffic loss is observed on newly added interfaces in AE on "MX and EX platforms" Product-Group=junos	On MX and EX series devices that support the "enhanced-ip" feature and when the new interface is added to aggregated ethernet (ae), output traffic is observed as 0 on the existing member links. This is being observed due to software issues as wrong weights are being set for existing child interfaces in the AE bundle.
PR Number	Synopsis	Category: Optical Transport Interface
<a href="#">1297164</a>	100-Gigabit Ethernet interfaces might not come up when "otn-options laser-enable" is configured on PTX Series platforms. Product-Group=junos	On PTX3000/PTX5000 platforms with P1-PTX-2-100G-WDM Physical Interface Card (PIC), if "otn-options laser-enable" is configured, the 100G interfaces might not come up.
<a href="#">1475777</a>	The interface on MIC3-100G-DWDM might be going down after performing an interface flap Product-Group=junos	On MX with MIC3-100G-DWDM installed, after performing an interface flap, the interface on 100G DWDM MIC might be going down.
PR Number	Synopsis	Category: Integrated Routing & Bridging (IRB) module
<a href="#">1461677</a>	In EVPN scenario memory Leak might be observed when proxy-macip-advertisement is configured Product-Group=junos	In EVPN scenario if "proxy-macip-advertisement" knob is configured, it might cause memory leak. Traffic would be impacted in case the memory leak is not stopped.

PR Number	Synopsis	Category: ISIS routing protocol
<a href="#">1482983</a>	The output of CLI command of 'show isis interface detail' might be incorrect if 'wide-metrics-only' is enabled for IS-IS and the ASCII representation of the metric in decimal is more than 6 characters long Product-Group=junos	If 'wide-metrics-only' is enabled for any IS-IS level and a metric configured on the IS-IS enabled interface for that level has ASCII representation in decimal more than 6 characters long, this interface's metric for that level will be merged with 'priority' field value in the output of 'show isis interface detail'.
PR Number	Synopsis	Category: track re issu control procedure bugs
<a href="#">1256113</a>	Traffic disruption seen when secondary node FPCs come online Product-Group=junos	On SRX Chassis cluster, Traffic disruption seen on primary node when secondary node FPCs come online
PR Number	Synopsis	Category: Flow Module
<a href="#">1465944</a>	Junos OS: High-End SRX Series: Multicast traffic might cause all FPCs to reset. (CVE-2020-1634) Product-Group=junos	On High-End SRX Series devices, in specific configurations and when specific networking events or operator actions occur, an SPC receiving genuine multicast traffic may core. Refer to <a href="https://kb.juniper.net/JSA11014">https://kb.juniper.net/JSA11014</a> for more information.
PR Number	Synopsis	Category: JSR Infrastructure
<a href="#">1445791</a>	The show security flow session command fails with error messages when SRX4100 or SRX4200 has around 1 million routing entries in FIB. Product-Group=junosvae	On SRX4100/SRX4200 platforms, once 1 million Routing Information Base (RIB)/Forwarding Information Base (FIB) routes entries are present on the device, an error might be returned after issuing "show security flow session" or other Command-Line Interface (CLI) which requires the information from Packet Forwarding Engine (PFE).
PR Number	Synopsis	Category: PFE infra to support jvision
<a href="#">1456275</a>	Queue data might be missing from path '/interfaces/interface/state' Product-Group=junos	On subscribing to '/interfaces/interface/state/' it might get subscribed to '/junos/system/linecard/interface/traffic/' internally, instead of '/junos/system/linecard/interface/queue/' which might impact traffic.

PR Number	Synopsis	Category: Layer 2 Control Module
<a href="#">1505710</a>	The l2cpd crash might be seen if adding/deleting ERP config and then restart l2cpd Product-Group=junos	Core files are generated if you add or delete ERP configuration multiple times and restarted l2cpd or rebooted the box. If this issue is happened, l2cpd does not recover again and generates core file continuously.
PR Number	Synopsis	Category: Layer2 forwarding on EX/NTF/PTX/QFX
<a href="#">1492746</a>	L2ALD memory leak observed in EVPN/VXLAN deployments Product-Group=junos	Frequent changes in EVPN/VXLAN mac learning process may trigger L2ALD memory leak
PR Number	Synopsis	Category: AgentSmith MPC Platform
<a href="#">1433948</a>	Interfaces on MPC-3D-16XGE-SFPP may go down due to CBO clock failure Product-Group=junos	On MX Series platforms, if the clock frequency is slowly changing on CBO (slow drift), the clock source for MPC-3D-16XGE-SFPP may not be changed to CB1, which will cause interfaces on it to go down and remain in the downstate.
PR Number	Synopsis	Category: Multiprotocol Label Switching
<a href="#">1181407</a>	Using CCC for remote-interface-switch or LSP-switch requires self-ping to be disabled Product-Group=junos	When configuring CCC remote-interface switch or LSP switch, self-ping should be disabled on the LSPs, referred to in the CCC configuration, by configuring the following: [edit protocols mpls label-switched-path lsp1] + no-self-ping. If this configuration is not set, LSPs will not complete the make-before-break (MBB) process.
<a href="#">1406400</a>	The rpd might crash when RSVP bypass path flaps Product-Group=junos	In the MPLS-TE with RSVP scenario, in rare cases, when bypass path of RSVP goes down or comes up, CSPF re-calculation is triggered by this update event of state changing and rpd needs to try to remove/add this path from the TED (Traffic Engineering Database), but at that moment, the TED database corruption happens and that may cause the rpd crash.
<a href="#">1460283</a>	Pervious configured credibility preference it is not considered by CSPF despite the configuration is	After configuring the credibility, the new credibility preference value will be stored internally and its not cleared or consider by the CSPF module, incase if the perviously configuration of

	deleted or changed to prefer another protocol in TED Product-Group=junos	"traffic-engineering credibility-protocol-preference" was deleted or if you configure "traffic-engineering credibility-protocol-preference" under another protocol (for example ISIS)
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Multi Protocol Label Switch OAM</b>
<a href="#">1399484</a>	The rpd process might crash when executing "traceroute mpls bgp" Product-Group=junos	When traceroute is performed to a remote host for an MPLS LSP using the command "traceroute mpls bgp", in very rare cases, it is possible that mplsoam daemon is holding the stale BGP instance handle in the query to the rpd process to get the information for the Forwarding Equivalence Class (FEC). Hence rpd crash might occur because of the invalid instance and cause traffic impact till rpd comes back up.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Multicast for L3VPNs</b>
<a href="#">1442054</a>	Memory leak might happen if PIM messages received over an MDT (mt- interface) in Draft-Rosen MVPN scenario Product-Group=junos	On all Junos platforms configured in the Draft-Rosen Multicast Virtual Private Network (MVPN) scenario, if Protocol Independent Multicast (PIM) messages are received over an Multicast Distribution Tree (MDT) tunnel logical interface (mt- interface), there might be memory leak which will lead to the rpd process crash.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: OS IPv4/ARP/ICMPv4</b>
<a href="#">1442815</a>	ARP resolution might fail after ARP HOLD NHs are added and deleted continuously. Product-Group=junos	ARP (Address Resolution Protocol) address resolution might fail after ARP HOLD NHs (next-hop) are getting added and deleted from ARP entries continuously.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: FreeBSD Kernel Infrastructure</b>
<a href="#">1198395</a>	The /var/run is in storage file system but it should be in memory file system Product-Group=junos	The /var/run is in the storage file system but it should be in the memory file system.
<a href="#">1360444</a>	Error messages might be seen when the system boots up Product-Group=junos	nfsd not running causing the core to fail to be written to the memory ms22 /etc/mount-re: Mounting 128.0.0.1:/var/tmp/pics on /var/re: mount_nfs failed/timed out Capture below to validate if you see the above logs. root@ms22% mount root@ms22% df -k

<a href="#">1367477</a>	The FPC might go down on some vmhost based PTX/QFX platforms Product-Group=junos	On PTX1000/PTX10001-20C/PTX10002-60C/QFX10002-60C, the Flexible PIC Concentrator (FPC) might reboot which might result in the FPC not coming up or the system becoming unresponsive.
<a href="#">1410542</a>	The chassisd process might crash due to a thread locking defect Product-Group=junos	The chassisd crash with core dump file might be seen if some error happens in chassisd syslog functionality. Traffic is impacted as FPCs are restarted.
<a href="#">1463802</a>	The scheduled tasks might not be executed if "cron" daemon goes down without restart automatically Product-Group=junos	The service utility "cron" runs in the background and regularly checks /etc/crontab for tasks to execute and searches /var/cron/tabs for custom crontab files. These files are used to schedule tasks which "cron" runs at the specified times. "cron" daemon is started during boot. If for some reason, the "cron" process exits, the scheduled tasks will not be executed. "cron" was not restarted automatically and had to be started manually.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: "ifstate" infrastructure</b>
<a href="#">1439906</a>	FPC might reboot if jlock hog occurs on all Junos VM based platforms Product-Group=junos	On a JUNOS VM using TSC clocking from the host system, "jlock hog" messages may be seen. This may lead to FPCs reboot.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: OSPF routing protocol</b>
<a href="#">1489637</a>	The rpd crashes when reset OSPF neighbours Product-Group=junos	On all Junos platforms, if more than one million routes are installed into Open Shortest Path First (OSPF) routing table, also, either Nonstop Active Routing (NSR) or Graceful Routing Engine Switchover (GRES) with warm standby is enabled, rpd might crash on backup routing engine (RE) when OSPF neighbors are reset. This is a scaled issue. This issue could be recovered after rpd on backup RE self-restart.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: PTX5KBroadway based PFE IPv4, IPv6 software</b>
<a href="#">1254415</a>	PTX FPC is reporting TQCHIP : Fatal error pqt_min_free_cnt is zero Product-Group=junos	PTX Series platforms with the FPC-PTX-P1-A or FPC2-PTX-P1A line card might encounter a single event upset (SEU) event that can cause a linked-list corruption of the TQCHIP. The following syslog message is reported: Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9

		08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Junos OS Chassis Management Error handling does detect such condition, and raises an alarm and performs the disable-pfe action for the affected PFE entity. To recover this PFE entity, a restart of the FPC is needed. Soft errors are transient or non-recurring. FPCs experiencing such SEU events do not have any permanent damage. Contact your Juniper support representative if the issue is seen after a FPC restart.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: vMX Platform Infrastructure related issue tracking</b>
<a href="#">1343170</a>	The vFPC might get absent resulting in the total loss of traffic Product-Group=junos	On the MX150 platform, the vFPC (virtual Flexible PIC Concentrator) might get absent and the l2cpd and chassisd might crash. Hence all the traffic is discarded.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: PTP related issues.</b>
<a href="#">1461031</a>	The PTP function may hog kernel CPU for a long time Product-Group=junos	If the PTP function is configured, the process pcmd may be starved due to the defect that PTP function hog kernel CPU for a long time (>100ms), which can cause the "keep-alive" of the corresponding functions timeout such as pcmd based functions - LFM, BFD.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Chassis mgmt for all QFX systems - chassis MIB, alarms, CLI</b>
<a href="#">1357971</a>	Some harmless logs might be seen Product-Group=junos	Due to a bug in Broadcom code for older SDK version 6.5.8 there is a possibility one could see errors related to "soc_ser_counter_info_get" and "ser mem info is not initialized" on "vty fpc0" logging. These messages are harmless and does not impact any functionality in Broadcom SDK.
<a href="#">1475851</a>	[Tencent] ULC-30Q28 FPC major error after system boot up or fpc restart Product-Group=junos	FPC is reporting Major Error because of SHUTDOWN ERI failure during BIST with repair on HMC having FW version >= 0x9c. BIST with repair is incorporated during boot up, so this Major Error will be reported during boot up having FW>=0x9c on HMC. During debugging, its found that this ERI is getting issued even before the HMC was brought up which is resulting into this Error. So added bringup steps for HMC before the BIST procedure starts as done for vale-ptx.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: QFX Platform related (SYSLOG/ALARMS/miscellaneous)</b>

<a href="#">1409448</a>	The 10G fiber interfaces might flap frequently when they are connected to other vendor's switch Product-Group=junos	When QFX5100 and other vendor's switch are connected via 10G fiber link, flaps may happen sometimes.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: QFX PFE Class of Services</b>
<a href="#">1445960</a>	CoS classifier might not work as expected Product-Group=junos	On QFX5000 Series platforms(except for the QFX5100) in the VxLAN scenario, the traffic is not classified properly on the UNI interface which has multiple VLANs configured.
<a href="#">1480522</a>	QFX5100 VC : CoS ETS/Scheduler not getting applied on VCP Ports : CoS not getting programmed at PFE Product-Group=junos	QFX5100 VC : CoS ETS/Scheduler config may not work on VCP Ports. Switch won't through a commit error on applying the CoS config on but VCP CoS won't get programmed at PFE. Issue resolved in 18.1R3-S9 & 17.3R3-S8.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: QFX L2 PFE</b>
<a href="#">1414213</a>	QFX5K: EVPN / VxLAN: Mutlicast NH limit is 4K Product-Group=junos	In QFX5K, multicast next hop limit is 4K. so based on the configuration, if there are more than 4K multicast NHs getting created, some of NHs won't be installed and you may see traffic drop for those groups.
<a href="#">1500825</a>	On QFX5100, ERPS may not work correctly Product-Group=junos	On QFX5100, ERPS may not work correctly on branch which as 1473610 fix, due to stp instance programming failure in hardware.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: QFX VC Infrastructure</b>
<a href="#">1394060</a>	The unknown log frequently show up on VC console Product-Group=junos	This issue is due to the DMA access. This issue was with the old BSD kernel and It is fixed on the latest BSD kernel .
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: RPD Next-hop issues including indirect, CNH, and MCNH</b>
<a href="#">1439317</a>	Packet drop might be seen on ACX Series platform when chained composite next hop is enabled for L3VPN. Product-Group=junos	On ACX platform, packet drop might be noticed at Packet forwarding Engine (PFE) if chain composite for L3VPN is enabled in PE-PE directly connected scenario for those destinations that points to chain composite. The issue is seen when RPD ends up creating route pointing to chain composite instead of indirect nexthop for PE-PE directly connected.

PR Number	Synopsis	Category: Resource Reservation Protocol
<a href="#">1368177</a>	RPD might restart after an MPLS LSP flap if "no-cspf" and "fast-reroute" are configured in an LSR ingress router. Product-Group=junos	RPD may restart unexpectedly after an MPLS LSP flap when "no-cspf" and "fast-reroute" are configured in LSR ingress router
<a href="#">1401152</a>	A single-hop bypass LSP might not be used for traffic when both transit chaining mode and sensor-based-stats are used. Product-Group=junos	On JunOS platforms with transit chaining mode enabled, if Resource Reservation Protocol (RSVP) link/node protection is enabled and sensor-based-stats is used, a single-hop bypass label-switched path (LSP) nexthop might not be installed in forwarding information base (FIB) even it is in routing information base (RIB). Hence the single-hop bypass LSP will fail to forward traffic when needed.
PR Number	Synopsis	Category: SW PRs for SCBE3 fabric
<a href="#">1491968</a>	FPCs might stay down or restart when swapping MPC7/8/9 with MPC10/11 or vice versa in the same slot Product-Group=junos	In MX240/MX480/MX960 routers with SCB3E or MX2010/2020 with SFB3 scenario, if MPC7E/8E/9E is swapped with MPC10E/11E each other or vice versa in the same slot, the different encoding mode between two MPCs might cause SCB3E/SFB3 to not change the mode gracefully according to the new MPC type inserted. This causes fabric destination errors which can trigger fabric healing mechanisms and cause system-wide impact due to fabric planes and FPCs getting reset. [TSB17748]
PR Number	Synopsis	Category: IPSEC functionality on M/MX/T ser
<a href="#">1354757</a>	Newly provisioned IPsec tunnel could not forward traffic Product-Group=junos	A newly provisioned IPsec tunnel may not forward traffic. This issue gets triggered in certain specific conditions wherein the RE daemon(kmd) processes the Outside MS-IFL UP event followed by Inside MS-IFL UP event. Normally, kmd receives Inside MS-IFL UP event followed by Outside MS-IFL UP event and this is one of the main reason that this issue is very hard to reproduce.
<a href="#">1444183</a>	The kmd process might crash and restart with a kmd core file created if IP of NAT mapping address for IPsec-VPN remote peer is changed Product-Group=junos	The kmd (Key Manager Daemon) process is mainly responsible for IPsec key negotiation. When IPsec-VPN peers enable Network Address Translation-Traversal (NAT-T) and established IKE SA (IPsec security associations) with Dynamic Endpoint (DEP) tunnel through the intermediate NAT device, the kmd might crash when IP of NAT mapping for IPsec-VPN remote peer is changed. The kmd crash may result in IPsec traffic loss. When kmd crashes, the established IPsec tunnel

		will not be affected, unless the IPsec SA re-negotiate happens to take place during the kmd restarting. For the new establishing IPsec tunnel, it cannot be established until kmd comes back up automatically. In rare cases, the kmd will restart, but it may crash again.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Generic platform and infra issues for MS-MIC and MS-MPC(XLP)</b>
<a href="#">1102367</a>	MS-MIC, MS-MPC might generate coredump upon receiving fragmented traffic Product-Group=junos	On MX Series routers where MS-MIC or MS-MPC is inserted, certain combinations of fragmented packets might lead to an MS-MIC or MS-MPC coredump.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: SFW, CGNAT on MS-MIC/MS-MPC (XLP)</b>
<a href="#">1366259</a>	MS-MPC/MS-PIC might be crash in NAT scenario. Product-Group=junos	MS-MPC/MS-PIC might be crash if two or more service sets configured with the same prefix lists and SIP ALG is configured in NAT scenario.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: MPC7, MPC8, MPC9, SFB2, MRATE &amp; 8x100 MICs</b>
<a href="#">1354070</a>	The log of "SMART ATA Error Log Structure error: invalid SMART checksum." might be seen on FPC with WINTEC mSata SSD Product-Group=junos	The log of "SMART ATA Error Log Structure error: invalid SMART checksum." might be seen on FPC with WINTEC mSata SSD
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: MPC7/8/9 Interface Issues</b>
<a href="#">1441816</a>	Egress stream flush failure and traffic blackhole might occur Product-Group=junos	Egress stream flush failure and silent dropping of traffic could occur in a rare occasion for a repeatedly flapping link on MPC7E, MPC8E, MPC9E cards, MX204 and MX10003.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: MX10003/MX204 MPC defects tracking</b>
<a href="#">1336575</a>	Random JUNOS MPC7 8 9E cards failed to start after software upgrade with PFE syslog message "BIST has	Random JUNOS MPC7 8 9E card BIST has detected error, err code id and fail to initialize after software upgrade In software fix, we enhance HMC for NVM write.

	detected error, err code id ". Internal BIST failed Product-Group=junos	
<a href="#">1474231</a>	QSA adapter Lane 0 port might be also brought down when disabling one of the other lanes Product-Group=junos	When QSA adapter is installed, the Lane 0 port might be also in down state when disabling one of the other lanes (1, 2 or 3) due to the chan number not entertained. It is not expected behaviour and it might affect service.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Trio pfe qos software</b>
<a href="#">1382288</a>	One single port with Dual stack subscribers pppoe/dhcpv6 drop all the connections and no subscribers seen now. Product-Group=junos	One single port with dual stack subscribers pppoe/dhcpv6 drop all the connections and no subscribers are seen.
<a href="#">1418602</a>	FPC log messages: "Q index(xxxxx) is not allocated" Product-Group=junos	The cause of the messages is a race condition. For each IFL, IFLSET or IFD object, HALP statistics module reads stats values from hardware for the individual queues. HALP stats module is implemented as a separate stats thread and while stats thread is waiting, PFEMAN thread may assign different queues to the IFL/IFLSET/IFD object. After stats thread gets CPU back, before reading HW for the next queue, stats thread needs to validate that queue is still applicable, which fails if we see the message. The messages are harmless and can be ignored.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Trio pfe bridging, learning, stp, oam, irb software</b>
<a href="#">1412457</a>	Ethernet MAC addresses may not be learnt after performing the "clear bridge mac table" Product-Group=junos	An LU-base MPC may not learn Ethernet MAC address after the "clear bridge mac table" command is issued. Examples of LU-based MPC are the MPC2/3/4 and MPC-3D-16XGE
<a href="#">1420626</a>	The unicast traffic to destination reachable over IRB and LSI with two next-hops might be dropped due to PFE mis programming Product-Group=junos	On MX Series routers with Trio chip set based MPCs, unicast traffic might get dropped when the destination is reachable over an integrated routing and bridging (IRB) interface and a label-switched interface (LSI) with two next hops.
<a href="#">1472222</a>	JDI-RCT:M/Mx: Linecard Errors found at HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invlaid IRB topo/ IRB ifl zero in I2 nh	Linecard Errors found at HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invlaid IRB topo/ IRB ifl zero in I2 nh 40495 add IRB

	40495 add IRB Product-Group=junos	
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Trio pfe I3 forwarding issues</b>
<a href="#">1381527</a>	Constant memory leak might lead to FPC memory exhaustion Product-Group=junos	On MX, EX9200, SRX4600 or SRX5000 platforms, constant memory leak might occur on a Flexible PIC Concentrator (FPC). This might finally lead to memory exhaustion and the FPC might crash and generate a core file. [TSB17775]
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Trio pfe multicast software</b>
<a href="#">1457166</a>	JDI-RCT:M/Mx: Expected pim joins are not learnt after performing GRES Product-Group=junos	JDI-RCT:M/Mx: Expected pim joins are not learnt after performing GRES
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Issues related to port-mirroring functionality on JUNOS</b>
<a href="#">1486749</a>	VPLS port-mirroring is sending packet without ethernet header on the outbound direction Product-Group=junos	Egress monitored traffic is not mirrored to destination for Analyzers on MX router, where handling of traffic was not supported in egress mirroring case for L3 as input interface.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: UI Infrastructure - mgd, DAX API, DDL/ODL</b>
<a href="#">991081</a>	The configuration change on Backup RE cannot be synchronized while deleting inactivated interfaces/routing-instances from primary Product-Group=junos	In the dual Routing Engine scenario, the backup Routing Engine does not sync up the configuration change while deleting an inactivated interface/routing-instances from the primary. So after the operation, the inactivated interface/routing-instances still exists on the backup Routing Engine.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: PTX/QFX10002/8/16 specific software components</b>

<a href="#">1412318</a>	hostname does not update at FPC shell after system configuration change on CLI Product-Group=junos	On PTX platform, hostname does not update at FPC shell after host name change unless FPC reboot.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: VMHOST platforms software</b>
<a href="#">1349373</a>	FPCs may reboot continuously until the Routing Engine reboot Product-Group=junos	On a next-generation Routing Engine (NG-RE), a failure of the Hardware Random Number Generator (HWRNG) will leave the system in a state where there is not enough entropy available to operate. Please refer to TSB17734( <a href="http://kb.juniper.net/InfoCenter/index?page=content&amp;id=TSB17734">http://kb.juniper.net/InfoCenter/index?page=content&amp;id=TSB17734</a> ) for more details.
<b>PR Number</b>	<b>Synopsis</b>	<b>Category: Virtual Private LAN Services</b>
<a href="#">1295664</a>	LSI interface might not be created causing remote MACs not being learned with an error log of "RPD_KRT_Q_RETRIES: ifl iff add: Device busy" Product-Group=junos	With VPLS being configured, after upgrade to 15.1/16.1/17.x releases, in some circumstances VPLS LSI interface is not correctly created, causing remote MACs not being learned and L2 VPLS outage. The issue is not reproduced and the code change is not a fix but adds a instrumentation using a hidden command 'show vpls ipc-history', which should be captured right away when the issue is seen on latest releases. show vpls ipc-history <<<<< show vpls connections show krt queue show route forwarding-table extensive /var/log/messages