Release Notes for PICOS 4.1

These notes summarizes PICOS 4.1 new features, new hardware, known bugs, and bug fixes. Best practices recommend that you read all the content before upgrading to this release. For more detailed feature information, refer to the configuration guides.

New Features

Layer 2 and Layer 3

Bug ID	Release	Description
13358	4.1.0	**DHCPv6 Relay** DHCPv6 relay is the IPv6 equivalent of the DHCP relay for IPv4. The DHCPv6 relay agent passes DHCPv6 solicitation or request to the DHCPv6 server on a different subnet and get the returned DHCPv6 advertise or reply from the DHCPv6 server back to the DHCPv6 client.
13440	4.1.0	**Support OSPFv3 for IPv6** OSPFv3 (Open Shortest Path First version 3) is added to support IPv6. Please note that VRF over OSPFv3 is not available in this version (4.1.0).
13487	4.1.0	**Multiple OSPF Instances** Eight OSPF processes are started if multi-instance is enabled. Each OSPF instance, identified by an unique number (1 - 8), can be configured separately. Please note that multiple OSPF instances are supported only in the default VRF with OSPFv2.
13087	4.1.0	**Support VRRPv3 for IPv6** VRRPv3, with IPv6 introduced, can bundle 2 physical PICOS routers as one single virtual router. The ND (Neighbor Discovery) protocol is enabled on the virtual IPv6 VRRP interfaces by default. In particular, the RA (Router Advertisement) packets will be sent out from the VRRP virtual interfaces periodically.
13386	4.1.0	**Update ping and traceroute for IPv6** With upgrade to Debian 10.9, CLI commands, ping and traceroute, are updated to support IPv6. And ping6 and traceroute6, are removed from CLI under operational mode.
13195	4.1.0	**Multi-line Banner** Added the support of multi-line banners with the newly-added CLI command "set system login multiline-banner ...".
13584	4.1.1	**VRF Route Leaking** Routes can be leaked from one VRF to a different VRF. In version 4.1.1, we support VRF route leaking in case of static routing.
13631	4.1.1	**IPv6 Support in ACL** Support IPv6 in ACL rules including firewall filter and NAC downloadable/dynamic ACL rules. IPv4 and IPv4 ACL rules can be configured in parallel. If a packet matches both a IPv4 rule and a IPv6 rule, the IPv6 rule has higher priority.
13618	4.1.1	**Specify Source Interface for SNMP Trap Configuration** Support to source interface for SNMP Traps. The source interface could be loopback or l3-interface. The specified source interface is used to derive the source IP address for the SNMP traps sent, so that traps received from each switch will always have a single consistent source IP address.

13557	4.1.1	**Loopback Inband Interface** IP addresses configured on loopback interface in default VRF namely lo can be used for inband management.
13195	4.1.1	**Multi-line Announcement** The content of login annoucement with multiple lines can be configured by new added CLI command "set system login multiline-announcement ...".
13186	4.1.2	**LACP Fallback** LACP is a mechanism to keep a link available on a LACP LAG port even though the LACP is not ready on the peer device such as the case of PXE (Preboot eXecution Environment) client-server. In particular, LACP fallback can be enabled on MLAG port.
13675	4.1.2	**MSDP Support and Anycast RP** MSDP, Multicast Source Discovery Protocol, is used for sharing active multicasting sources between different RPs. Anycast RP is a specific application scenario of MSDP, which can provide load sharing and redundancy in PIM-SM networks. In case of anycast RP, multiple RPs are configured with the same IP address on loopback interfaces. A source or receiver will select the closest RP based on the routing distance.
13685	4.1.2	**Timestamp on NAC Session** For a NAC session which is authorized for a specific port (802.1X) or client (MAB), when execute command "run show dot1x interface gigabit-ethernet xxxx", display the timestamp indicating that when the session is ready.
13731	4.1.2.2	**Add a new VXLAN decapsulate-mode** Add a new decapsulate-mode "service-vlan-per-port". The decapsulated packet can be tagged or untagged dynamically based on the setting on the output port.

Interface Management

Bug ID	Release	Description
13470	4.1.1	**Enable FEC on 25G Ports** Support FEC (Forward Error Correction) on 25G ports under both OpenFlow and L2/L3 mode. Please note that CL74 (Base-R) instead of RS-FEC is supported for 25G ports on Tomahawk/Tomahawk+ platforms such as AS7312 and AG5648.

Hardware Support

Bug ID	Release	Description
13643	4.1.2	**AS6812_32X Support** AS6812_32X is available in 2.11.x. Support AS6812_32X in 4.1.2.

Crossflow

Bug ID	Release	Description
13582	4.1.1	**Multi-Actions under Crossflow Mode** If no flow hit by the incoming packet, it will be trapped to CPU and sent to the Controller. With multi-actions flow such as "in_port=5,actions=normal,controller", the packets received on port 5 will be sent to controller as well as be forwarded out normally.

OEM

Bug ID	Release	Description
13571	4.1.1	**Migrate NPB to 4.1.1** Migrate NPB application to 4.1.1 from 3.2.x.

Fixed Issues

Layer 2 and Layer 3

Bug ID	Release	Description
13145	4.1.0	**DNS name-server and DNS Search List** If users configure DNS name-server or DNS search list from PICOS CLI, the associate configuration included in /etc/resolv.conf should not be updated from DHCP server.
13436	4.1.0	**Cannot Access to SNMP Agent** If users enable management VRF and Inband connection as following CLI commands, SNMP agent stopped responding. set system inband enable true set system management-vrf enable true
13419	4.1.0	**The Size of NETCONF.events Increases without Limit** If users configure NETCONF on the switch, the size of NETCONF log file /tmp/stream/NETCONF.events will increase without limit. Add a mechanism of rotation to fix this issue. Keep the size of this file within 2M.
13457	4.1.0	**Crash if Users Configure port-mode via NETCONF** The NETCONF process might crash if configure port-mode to "trunk" and then back to "access" repeatedly via NETCONF.
13579	4.1.0.1	**Multicast Routes in VRF not Applied to ASIC** The multicast routes specific to a VRF cannot be applied to the hardware ASIC. Therefore, the multicast traffic within this VRF is pumped up to CPU and routed out by the multicast routes in the kernel. This issue was fixed in 4.1.0.1. Additionally, as a known issue of FRR, if configure join-prune-interval or keep-alive-timer or register-suppress-time, they will be duplicated in the configuration both under default VRF and other VRFs even if PIM is used there.
13585	4.1.0.1	**Multicast Packet Drop Caused by Aging-out of Multicast Routes in the Kernel** Generally, the multicast traffic hits the multicast routes in ASIC and go through the data plane. The corresponding multicast routes in the kernel will be aged out if no traffic hit within a specific timeout. The multicast routes in the kernel should be updated based on the status of the multicast routes in ASIC.
13556	4.1.1	**Upgrade2 from PICOS 3.x.x to PICOS 4.x.x** The configuration in 3.x.x will not be brought into 4.x.x potentially. Instead the configuration /pica/config-4.x/pica_startup.boot will be loaded when update to 4.x.x. In case of upgrade2, the old version will be kept in the switch in order for returning to it latter. Additionally, the 3.x.x configuration will be copied to /udata/pica/config-3.x/picos_startup.boot. The platforms with SquashFS including N3000 and N3100 don't support this feature.
13601	4.1.1	**NETCONF Update** NETCONF is updated to support 4.x.x changed functionalities including static routing, VRRP VXLAN and L3 interface. The Yang model support for BGP and OSPFv2 & OSPFv3 are not added in 4.1.1.
13570	4.1.1	**Encrypt Passwords Configured in BGP & OSPF** Encrypt the passwords appearing in BGP & OSPF set commands.
13575	4.1.1	**Include FRR Config in the tech_support** By including FRR config in the tech_support which would be sent back from customer for trouble shooting, it is much helpful to identify configuration consistency issues between PICOS CLI and FRR runtime configuration.
13597	4.1.1	**SNMPwalk Fails** It's possible that SNMPwalk fails when switch has multiple VLANs enabled inband because the returned SNMP reply message can only be sent back via the L3 interface on which the associate SNMP query message is received. It is fixed in 4.1.1 by lookup the L3 routing tables to send the SNMP reply messages to SNMPwalk client.
13376	4.1.1	**Image File is Not Removed after Upgrade** In case of upgrade2, if place the image file of new version in /udata instead of /cftmp, this image file will not be removed when update to new version.
13595	4.1.1	**TACACS+ User in PICOS CLI Prompt** If a TACACS+ user login to the switch, the TACACS+ user name should be showed in PICOS CLI prompt in stead of the associate local user names such as admin or operator.

13617	4.1.1	**Keep DNS Servers in Configuration Order** The nameservers stored in resolv.conf should be in the order we configure "system dns-server-ip" in the switch.
13611	4.1.1	**Suppress VRRP rsyslog Messages** It does not make sense to print the same rsyslog messages "Cannot find IF vlanXXXX" repeately if there are 2 VRRP groups configured on the same subnet.
13593	4.1.1	**The VXLAN Traffic cannot be Terminated** The VXLAN packet with a VRRP virtual IP address as the underlay external destination IP address cannot be terminated. This issue is fixed in 4.1.1.
13377	4.1.1	**Suppress False FRR Error Messages** When configure OSPF, it may print below false rsyslog message. Cannot stop staticd: pid file not found Cannot stop zebra: pid file not found Cannot stop ospfd: pid file not found Cannot stop bgpd: pid file not found
13654	4.1.2	**802.1X Authorization Failure on VXLAN Access Port** When reset VXLAN configuration from AmpCon SDN controller, 802.1X authorization on the VXLAN access port may be failed. This issue is fixed in 4.1.2.
13586	4.1.2	**Different MAC Addresses of VRRP IP Returned** In case of active-active VRRP over MLAG, when a client requests the MAC address binding to VRRP virtual IP address via ARP /NS, different MAC address may be returned. This issue is fixed in 4.1.2.
13739	4.1.2.2	**Fixed IPv6 Routing Errors** For specific route entry in hardware, if the next hop is changed to an immediately connected host or re-directed to a l3 interface, this route entry may be removed mistakenly. This issue was fixed in 4.1.2.2. Additionally, IPv6 NA (Neighbor Advertisement) packets should not be routed out of the switch. If a NA is pumped to CPU, will check its source IP address against the ingress L3 interface and make sure the associate IPv6 neighbor is added to the right IPv6 subnet.
13727	4.1.2.2	**Remove MTU Restriction on VXLAN Traffic** The restriction that packet with size greater than 1422 bytes cannot go through VXLAN tunnel is removed in 4.1.2.2.

OVS Features

Bug ID	Release	Description
13338	4.1.0	**Multiple Output Tunnel Ports in a VXLAN Flow** Multiple tunnel ports can be added as the output ports in a VXLAN flow. It requires that the VXLAN ports must have the same VNI (Virtual Network Identifier) on the ingress side.