BROADCOM®

EFOS

Release Notes

Release Notes
Software Release 3.4.3

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Table of Contents

Chapter 1: General Information

This document contains the release notes for the Ethernet Fabric Operating System (EFOS) software, including:

- Support of the Linux operating system on the Broadcom® reference platforms listed in Section 1.1, Environments.
- Platform-specific functionality of features such as Data Center, Switching, Routing, IPv6 Routing, BGP, IP Multicast, and Management.

NOTE: The suite of features EFOS supports is not available on all the platforms to which EFOS software has been ported.

1.1 Environments

1.1.1 Linux

EFOS software has been tested with the Linux kernel. The build and runtime environment is based on network platforms with the following:

- Ubuntu 16.04 LTS/CentOS 7.4 with kernel version 4.4.117

1.1.2 SDK

EFOS uses the Broadcom SDK, version 6.5.13 on all platforms.

1.1.3 Toolchain Support

EFOS supports the following toolchains:

- x86_64 cross-toolchain with multilib support, gcc-4.9.4, glibc-2.27 built using Crosstools-NG

1.1.4 Broadcom StrataXGS V

EFOS software has been tested with the following Broadcom StrataXGS® V reference platforms:

- BCM56870 platform with Intel ATOM processor and 4GB memory(ports: 48 × 25GbE, 8 × 100GbE)

Chapter 2: Operational Characteristics and Known Issues

2.1 Linux

Defect numbers are included for known issues (not for operational characteristics).

- Adding a default route to the routing function may disrupt connectivity to the service port in a device running Linux. The Linux networking stack will use the first entry in its routing table that matches a particular address. Multiple default routes will result in behavior that is dependent on the order in which those routes appear in the routing table.
- When using U-Boot firmware (or its predecessor PPCBoot), ensure that the “ethaddr” environment variable is set to match the MAC address that the EFOS software assigns to the service port (for example, base MAC address + 1). Similar precautions must also be taken if using other boot firmware.
- Devices with Ethernet service ports are generally supported with network drivers included as part of the standard Linux kernel sources. Some of these drivers choose to log status changes to the system console. Such console messages (generally referring to “eth0”) can be safely ignored.
- FP-8122. Several automated test suites report failures or warnings based on the behavior of networking functions of the Linux kernel used for this release. Such issues do not present problems for the operation of EFOS or any other component running on the system.
- FP-35136. It is possible to obtain neighbor table related messages on the console if the kernel IPv6 neighbor table is exceeded.
- FP-45979. On startup, EFOS Linux packages ignore the saved serial speed configuration. Instead, for the serial console on standard input/output, they use whatever serial speed is already set by the boot loader and/or kernel. This ensures that the serial speed will be consistent throughout the whole boot process. If U-Boot is used as a boot loader, EFOS has a built-in integration so that when a configuration is saved, U-Boot updates its saved serial speed. Define `BAUDRATE_USE_BOOTENV` for your platform to enable this. (EFOS reference packages before 5.0.0 do not define this). Otherwise, alter the function `osapi_nvstore_save()` to get your kernel and boot loader to use the saved serial speed. This behavior can be removed by altering the function `main()` to remove the call to `simSerialBaudRateOverride()`.
- FP-79682. File download using 1K XMODEM protocol is not supported using HyperTerminal. Other 1K XMODEM implementations function normally.
- FP-80650. Update bootcode is not available on 8548-based platforms (for example, GTO). Updating the GTO's bootloader (CFE) will overwrite the ptable entries configured in CFE. In order to avoid having to reconfigure CFE or maintain a private version of CFE to support EFOS, this command is not supported on the GTO platforms. If a customer platform requires this feature, and uses this BSP as a reference, this feature can be supported by the following:
  - Creating the directory bsp/cpu/$L7_CPU/link/target and adding the following files to the directory.
  - UPDATE_BOOTCODE—a bash script that is target specific (basically identifies which mtd partition contains the bootrom) that erases the partition and copies the contents of u-boot.bin to the boot partition.
  - u-boot.bin—a binary image of the bootloader. (It can be a renamed cfe.bin.)
- FP-167946. On XLP208B0XMC-based platforms, when EFOS firmware is upgraded from the EFOS Startup Utility Menu, firmware cannot be read by uboot which causes the device card to stop at the uboot prompt. A workaround is to manually erase the FLASH and then load the image from the uboot.
- FP-215848. Bonding interface configuration via Linux /etc/network/interfaces with bond-slaves, bond-master and related parameters is not supported. Bond enslaving configuration is possible using `interface up` commands for member interfaces in order to execute bonding commands.
- FP-214129. Use of the `ifconfig` command to add or delete CIDR IPv4 addresses can cause additional routes to get added/deleted. It is recommended that only the `ip addr` command is used for address management.
- FP-237427. On BCM56160 based platforms that use NOR flash to store image and configuration, it takes a longer time to boot to the operational state. This is due to the use of a large size JFFS2 partition and speed of flash part.

- FP-241666 When using ethtool -S <interface-name> command to view front panel interface details from Linux shell, 50 Gb/s port speed is reported as Unknown. Also, in the same command, for a 100 Gb/s capable interface, which can also be configured to support 40 Gb/s, supported speeds do not include 40 Gb/s.
- FP-246222: A 6to4 tunnel has to be unconfigured and configured again if the reachability to the tunnel endpoint is disrupted in the presence of an active 6to4 traffic flow(s).
- FP-247356: DAD failed addresses show up as Kernel routes in the CLI. There is no functionality impact; it is just a display issue.
- FP-259982: After upgrading to new EFOS code version for BCM958525XMC CPU card-based switches, users may observe slowness when accessing the management interface via serviceport. This can be resolved by upgrading the u-boot using the **update bootcode** command.
- FP-260746: Spectre variant 2 vulnerability has not been resolved on EFOS switches. Software fix for this issue is not complete without fixes to the rest of the user space Linux applications. For the EFOS 3.4.3 release, such an upgrade to Linux distribution has not been performed.

2.2 Broadcom StrataXGS V

Defect numbers are included for known issues (not for operational characteristics).

2.2.1 XGS Base

- FP-267800. On BCM56870-based platforms, ports in the same SERDES core must be configured to the same base frequency. On the IX8/B, every group of four ports must be configured as 10G or 25G as these ports are mapped sequentially to the SERDES cores.
- FP-268376. On BCM56870-based platforms, auto DoS control with tcpoffset does not function properly. As a workaround, administrators should create ACLs to trap these packets.
- FP-268646. On BCM56870-based platforms, DoS control for IP packets with TCP header size greater than configured value with offset 1 does not function properly. As a workaround, administrators should create ACLs to trap these packets
- FP-268647. On BCM56870-based platforms, DoS control for icmpv4 echo packets are not dropped when payload is greater than the configured value. As a workaround, administrators should create ACLs to trap these packets
- FP-268653. On BCM56870-based platforms, removing an RSPAN tag on egress does not function properly. This issue will be addressed in a future release. As a workaround, administrators should configure the RSPAN VLAN on a VLAN aware destination switch.

2.2.2 XGS Switching

- MAC address entries may not be inserted in or removed from the CPU’s copy of the forwarding database after being learned on the underlying hardware. This is most likely to occur under heavy periods of learning or aging. This does not affect the normal operation of the switch.
- With storm control enabled on an interface, the threshold (number of packets per second) is calculated on the basis of the speed of the port and the configured threshold level. The threshold is based on an average packet size of 512 bytes.
- The maximum number of MAC addresses supported by the hardware is offered but not guaranteed. The total number of MAC addresses that can be learned is dependent upon how the VLAN/MAC keys are hashed. The hardware table is broken up into buckets. Once a bucket's entries are used, any other VLAN/MAC keys that hash to that same bucket are not learned. Additionally, L2 multicast and L3 interfaces consume MAC addresses.
- If you enable port security on a port that is connected to another bridge, you must add the MAC address of the connected interface to the list of entries allowed. Create the static MAC address entry by using the `port-security mac-address mac addr vlanid` command. This will ensure that control frames like BPDUs and LAC PDUs from the partner device will not be dropped if the dynamic limit for PML is reached.
- For port mirroring, in case of unknown DAs and broadcasts egressing the multiple mirrored ports, only one copy of the packet will be seen on the probe, even though the same packet goes out the multiple mirrored ports.
- FP-12425. The 1519–1522 frame counter is incremented only for VLAN-tagged frames. Untagged frames with that size increment the >1522 counter.
- FP-14166. When creating static port MAC locking entries on a port-based routing interface, these entries will not have the reserved VLAN ID that the Broadcom driver reserves for the interface. None of these entries will come into effect as long as the port-based routing interface is active, and these frames will be discarded. The maximum dynamic entries can be set to something other than zero to allow for traffic to get through. VLAN-based routing interfaces do not have this issue.
- FP-25830. StrataXGS IV has the ability to limit the maximum bandwidth to the CPU on a per-queue basis in hardware. EFOS utilizes this mechanism to limit the low priority queues to values deemed to be appropriate for an 8245-based processor complex. Only the low priority queues (queues 0–5) are limited. The high priority queues (queues 6–7) are not limited because these are used for control and protocol traffic. The low priority queue limits will affect throughput tests for ARP packets, ping packets and any other traffic that utilizes queues 0–5. Source code customers may limit

additional queues, remove limits on these queues, or change the bandwidth limits by updating the cpu_kbps_rate_limit_per_cos[] and cpu_pps_rate_limit_per_cos[] array and recompiling. Use caution when rate limiting the high priority queues.

- FP-25971. For Minimum TCP header size DoS prevention, although EFOS UI specifies to enter the minimum TCP header size, you must enter the minimum TCP payload size. Also, the TCP Fragment Mode DoS feature must be enabled for the Minimum TCP header DoS prevention feature to work.
- FP-27237. MAC addresses can take longer than the age time to age from the hardware MAC table. If the age time is not being changed, it can take up to twice the age time to delete the MAC address. If the age time is changed, it can take up to three times the age time to delete the MAC address.
- FP-38460. Flow control may not be functional when multiple ingress ports are over-subscribing the egress port. EFOS uses default MMU settings from the SDK that divide the MMU packet buffer evenly among CoS queues/ports for fairness. This is a good general setting and suitable for most applications. Source code customers can change the MMU settings to match their application needs.
- FP-45367. For short cables, the length reported using cable diagnostics may not be accurate.
- FP-123821. When a packet is marked for dropping by the switching logic and an ICAP entry selects the redirect action, the ICAP entry will override the switching logic and the packet will be redirected.
- FP-137728. When two QSFP ports are connected where one is configured in 40G mode and the other is configured in 4x10G port, the 40G port correctly shows link down, while the 4x10G ports incorrectly show link-up.
- FP-149918. Untagged multicast traffic with reserved MAC addresses 01:80:c2:00:00:14 to 01:80:c2:00:00:15 as destination addresses are not forwarded.
- FP-151777. UDLD frames are not dropped by “deny any” MAC ACL applied in an outbound direction.
- FP-158985. Unicast Packets Received counter value is not incrementing properly upon receiving unicast traffic with random frame size.
- FP-175112. Ping does not function properly when SNAP encapsulation is enabled on an interface.
- FP-260599. When remote VLAN is configured on BCM56870, clear configuration command may lead to error log “Could not do rspanMirrorConfig for gport” and port-mirroring related data does not get cleared from hardware.

2.2.3 XGS Routing

- FP-192827. The record route IP option is not supported in this release for the BCM56850.
- FP-202278. There is only one set of IP statistics for the whole router. Per-VRF IP statistics are not supported in this release. Some IP statistics counters will show identical values for all VRFs.
- FP-263702. Few kernel ARP entries may get added in the non-default VRF upon ARP entry timeout. This happens if dynamic renew option is disabled and those entries ARP entries are not actively used for forwarding in the hardware. Such Kernel ARP entries need to be removed from the Linux shell explicitly by the user.
- FP-265798. EFOS switch platforms having less than 16GB RAMs and using ALPM SDM routing templates can exhibit slowness in processing certain system calls like system(), popen(). This can cause some fast adjacency detection protocols like BFD to timeout. It is recommended that switch platforms with 16GB memory be used or BFD timers can be increased to 200ms or greater values.

2.2.4 XGS QoS

- ACL operational characteristics:
  - ACLs are supported in the inbound direction only (except on the BCM56820 which also supports the outbound direction).
  - A single ACL is limited to the number of user-configurable rules supported by the platform. Multiple L2 and L3 ACLs can be applied to the same interface, as long as the total number of combined rules fit within the platform-specific limit.
- CoS Queuing operational characteristics:
  - Egress CoS minimum bandwidth guarantees on a LAG assume equal traffic distribution across all members.

  - Egress port rate shaping on LAGs assume equal traffic distribution across all members.
  - The actual guaranteed minimum bandwidth allocation might differ slightly from the configured value. The deviation varies depending on the average frame size.
  - The upper limit on the guaranteed minimum bandwidth for a single queue is 68%.
  - FP-27031. On some platforms, it is not possible (due to resource limitations) to have a unique 802.1p priority mapping on every port.
  - FP-27778. IP Precedence to priority mapping is not supported.
  - FP-27783. In Trust IP-DSCP mode, the DSCP table maps to 802.1p priority directly. This means that tagged IP packets mapped via the DSCP table will be marked with a new 802.1p priority upon egress, as well as being assigned the specific CoS queue.
- DiffServ operational characteristics:
  - ACLs and DiffServ policies cannot coexist on the same interface. However, it is possible for DiffServ to emulate ACL functions using the policy attribute *drop*.
  - Policing policies can be color-aware based upon IP Precedence or IP DSCP only.
  - For policing statistics, all non-conforming traffic is counted as *in-discarded-packets* even if the non-conforming action is not set to discard.
  - DSCP and IP precedence can be marked for both policing conforming traffic and non-conforming traffic.
  - FP-29181. Policing rate is configurable in increments of 64 Kbps. If the specified value is not an integral multiple of 64, it is rounded down to the immediately lower multiple of 64.
- FP-17550. IP traffic received by a routing port that results in an L3 destination lookup failure will be sent to the CPU, regardless of any DiffServ policy or ACL rules on that port that would otherwise cause those packets to be dropped.
- FP-126438. When configuring Traffic Class Groups (TCGs), the user must assign queues to TCGs in order from lowest TCG to highest TCG. When only TCG0 and TCG2 are used, and TCG1 is added later, all queues must first be assigned to TCG0, then all three TCGs may be configured in order. If queues are moved from one TCG to another TCG, the TCG is considered unused and the ordering requirement applies.
- FP-150886. Not able to apply an ACL with maximum rate-limit rules. A simple meter in DiffServ or ACL rate limiter occupies two meter entries in the meter H/W table - A pair of even/odd meters. To support the stats for in-profile and out-of-profile packets, EFOS uses two counter entries as well. Taking the example of Triumph2, the ICAP (IFP) has 16 slices of depth 512 = 8K. The meter table depth is also 8K. The counter table depth is also 8K. You can consider that there are 512 meters for every slice and 512 counters for every slice.

  The hardware has a limitation that the ICAP slice can only address counters in its corresponding counter hardware slice. For example, slice 0 can address only counters 0-511. Slice 1 can address counters 512- 1023 only, but meters are a global resource. For example, slice 0 can address meters anywhere in 8K meter table on StrataXGS IV devices.

NOTE: EFP on StrataXGS IV devices has counters as a global resource. Any slice can use any counter in the Counter Hardware table

If the slice width is single-wide, you can install 512 policies. Assume that there are no policies installed in this slice. Now If you try to install 257 policies with action as meter for each, each policy will require 2 meters and 2 counters each. Therefore, 256 policies will use the available 512 meters from anywhere in the hardware table. 256 policies will use 512 counters from the same slice (assuming none of the counters are already used from this slice). Now the 257th policy will use 2 meters from the global meter table. But the 257th policy would not find a counter to use because all counters from this slice are exhausted. Therefore, the policy will not be installed.

- FP-206662. After issuing the **clear counters all** command, congestion drops, but the counter does not get cleared. Users need to compare the previous reported values with the current values to determine the delta.

2.2.5 XGS Data Center

- FP-174098. When Priority Flow Control is configured manually at the same time as automatic mode, invalid configurations are improperly allowed. Once allowed, the user cannot remove the automatic configuration manually. The user must remove the configuration when auto mode is enabled.

- FP-204709. In a setup using VxLAN/NVGRE, if a PVID is used on an access port, only untagged packets can be sent to that port. Any tagged packets sent to the port, will also be received on the tenant network.
- FP-207077. When encapsulated packets in a VxLAN/NVGRE setup are mirrored, they will show source and destination MAC addresses as 0.
- FP-225389. The user is not being prompted to save the configuration after an “application install” config operation by the switch automatically. This issue is only seen when users try to install their own applications on the switch. The workaround for this issue is to save the configuration manually after the 'application install' config operation.

2.3 Software Features

Defect numbers are included for known issues (not for operational characteristics).

2.3.1 EFOS Base

- FP-104194. If an SNMP set is used to set the download image name to Image1, and then a get upload image name is performed, the name returned is the download image name. This is due to using a single variable in the application. No upload/download functionality is impacted; it is only a display issue.

2.3.2 EFOS Switching

- In each release, the values of porting constants for the reference platforms (for example, the sizes of the routing table or forwarding database) may be modified to be different than in the previous release. This should be considered when porting a new release.
- On some platforms, mismatches may occur between the CPU's copy of the forwarding database and the copy in the underlying NPU or ASIC hardware. This is most likely to occur under heavy periods of learning or aging, or if addresses are being learned or aged during a topology change (such as links being added or removed to a LAN, or deleting a VLAN). This does not affect the normal operation of the switch.
- On some platforms, MAC address entries may not be inserted in the CPU's copy of the forwarding database after being learned on the underlying NPU or ASIC. This is most likely to occur under heavy periods of learning or aging. This does not affect the normal operation of the switch.
- EFOS software implements IEEE 802.1s running in IEEE 802.1w-compatibility mode, rather than the IEEE 802.1w protocol itself. As such, ANVL RSTP test suites are not supported. ANVL MSTP Test cases are supported.
- An 802.1X port that is up and authenticated may not return to the authenticated state after a switch reboot when spanning tree is enable. The recommended workaround is to set the RADIUS server timeout to 15 seconds or higher.
- The PVSTP/PVRSTP feature has issues handling a large number of BPDUs by virtue of the number of interfaces participating in VLANs that are PVSTP/PVRSTP enabled. The failure scenarios include, but are not limited to:
  - Unexpected topology changes in *x*STP
  - Very long times for convergence
  - (Very stressful conditions only) System reboots and crashes.

  The present solution scales reasonably well up to a cumulative total of around 200 <i*nterface, VLAN*> tuples participating in PVSTP/PVRSTP, beyond which system behavior is unpredictable.
- FP-41237. When applying a saved configuration with a large VLAN configuration (such as all ports in 4K VLANs), attempting to create new port-channels as soon as the user prompt appears may cause the CLI to be unresponsive. The user must wait until all the VLANs are created before attempting to configure the device.
- FP-61836. A console break-in feature is enabled by default. To disable it by default, set "`static L7_uint32 consoleBreakinFlag = L7_FALSE;`".
- FP-75340. LLDP MED application should not allow configuration of location and inventory transmit TLVs as the underlying application is not present. The switch allows configuring transmission of these location and inventory TLVs even though the underlying application to support location and inventory are not supported. There is no operational impact from this setting, although these TLVs will not be transmitted.
- FP-78109. Maximum MAC entries are not learned as there are few default entries present in the hardware table.
- FP-91249. For Denial of Service ICMP command, EFOS adds 8 bytes to the value configured by the user, to accommodate the ICMP header size.
- FP-98290. The IP address for a configured trap receiver host name gets resolved at configuration time and is then stored as a resolved address. If the DNS resolution fails at this time then a null (e.g., 0.0.0.0) address is stored.
- FP-98517. The DiffServ policy name received from the RADIUS server as a filter-id attribute in the RADIUS accept message fails to apply on the authorized port when the port control mode is Auto (i.e., port-based dot1x mode).

- FP-98879. Automatic configurations (e.g., vlan ingressfilter, vlan acceptframe, vlan participation, etc.) done by dot1AD commands can be overwritten by the administrator, but upon save and reload those overwritten commands might not be retained.
- FP-101015, FP-100897. LAGs take a long time (2 to 5 min.) to come up or down after they are administratively enabled or disabled when the L2 multicast table is full (1K).
- FP-104201. When the primary TPID is changed to something other than 0x8100, the packet tag is not getting stripped at the egress even though VLAN tagging is disabled.
- FP-110640, FP-116224. When there are large numbers of MSTP instances and large numbers of connected ports to another switch and auto edge (auto portfast) is enabled, there is a chance of small loop when one of the switches reboots causing traffic disruption.
- FP-112703. Occasionally, when dot1x is enabled on a port and a VLAN participates in this dot1x-enabled unauthorized port, the port egresses the traffic on that VLAN.
- FP-128631. If the connected device sends both IEEE and CEE DCBX TLV(s) in the same packet, the EFOS behavior is as follows:
  - If the configured mode is **auto**, then the operational version is set to IEEE.
  - If the configured mode is **IEEE**, then the operational version is set to IEEE.
  - If the configured mode is **CEE**, then the operational version is set to CEE.
- FP-149428. When 4094 VLANs are created, the show running-configuration command takes 8 to 10 seconds to respond.
- FP-149479, FP-103674. When RADIUS or Syslog uses a DNS host name for configuring server entries, they are unaware of changes in the DNS cache host mapping. When a RADIUS or Syslog server entry is added, the DNS host address is resolved and is persistent until the server entry times out; therefore, the RADIUS or Syslog application attempts to send the packet to the RADIUS or Syslog server configured on the switch.
- FP-158505. When maximum MSTP instances are configured and a large number of ports are connected to another switch, continuous log messages are observed and the switch becomes nonresponsive.
- FP-158745. The clear config command maintains the previous Ethernet counters after a clear counters all command.
- FP-171812. For port-based MGMD snooping, dropping unregistered multicast data and flooding of unregistered multicast data only to mrouter ports is not supported.
- FP-183177. The administrator cannot remove a physical port from the LAG if the admin key configured on the LAG interface is different from the physical port’s admin key.
- FP-183323. When there are 128K MAC addresses learned on a VPC interface, the FDB entries are not in sync between the primary and secondary VPC devices. For entries that are not configured on one of the devices, traffic will be broadcast on all interfaces in the VLAN.
- FP-205769. The RA packet counters are per rule and are incremented whenever a packet matches the rule (rather than when an action corresponding to a matching rule is taken). Hence, there is a possibility that the counters get incremented even when the action corresponding to the RAGUARD rule was not taken. It is not recommended to attach conflicting ACL rules on the interface on which RAGUARD is configured.
- FP-224481. The MAC entry database on the VPC interface goes out of sync with the peer’s when the partner’s LAG interface is flapped. Subsequent MAC entries that are learned are in sync
- FP-237397. If the MAC-limit is set by the user while continuously sending traffic, MAC addresses learned are more than the configured limit, due to simultaneous Flushing and Learning operations. It is recommended to stop traffic while configuring the MAC-address limit in VLAN MAC Locking.
- FP-242609: A jump in system time is accommodated and handled appropriately for deducing monotonic time for PTP-related operations. A message `Time jumped by ~130` is displayed on the console when the switch is configured to synchronize time with the SNTP; indicating a time jump was detected and accommodated. This log can be ignored.
- FP-245676: When the PTP sync interval on a slave is configured to 1, the servo (which has a granularity of 1s) only receives sync messages/timestamps every 2 seconds causing this issue. This can be avoided by keeping the sync interval to ≥ 1 pkt/sec.

- FP-247476: In BVCI, if an EFOS plugin bundle is installed after the switch gets detected in ODL, then BST does not work. The user needs to make sure that the EFOS plugin bundle is installed in ODL before the switch gets detected in ODL.
- FP-247534. When either an ingress port or egress port is configured as a routing port, and untagged packets are sent on the ingress port, it is observed that packets egressing out of the output port are tagged with VLAN 1/4093.
- FP-248744. BroadView™ Agent: We may observe some statistics slightly more than 100% on some realms. We need to treat it as 100%.
- FP-250057. BroadView Agent: When trigger-rate-limit interval and trigger-rate-limit parameters together are configured to generate more than one trigger per second, then trigger reports may get generated for some duration even after clearing BST thresholds.
- FP-250140. BroadView: BroadView HTTP requests tend to stall when continuous triggers are generated on the system.
- FP-252898. BroadView: Async reports are not generated when the user disables BST and then re-enables BST with the same config parameters.
- FP-253287. Two StrataXGS V RSX devices connected back-to-back act as Roots when the STP mode is PVST.FP-257825. Config migration for "sflow flow-based rules" from previous EFOS releases will not work if the original configuration has more ACL rules that are supported on that silicon.
- FP-263599. MAC movement feature for an 802.1X unaware client is not supported if the new port's PVID happens to be the same as the VLAN on which the client was initially authorized on the first port.
- FP-264342. "soc_sand function returned error" logs on the console while clearing the FDB.
- FP-267963: "ERR Failed to map asic index" traces observed on console are non-critical and does not impact the functionality
- FP-268396: An EMCQ trigger is received despite configuring a ECPUQ threshold.

2.3.3 EFOS Switching Configuration

- Not all commands in the EFOS *CLI Command Reference* are available on all platforms.
- FP-44881. When creating 4K VLANs and including all ports in 500 VLANs before enabling spanning tree and other configurations, a NIM timeout can occur (dot1q may not return a completion of a port not forwarding event in time). This situation can be avoided by using interface-based commands to configure ports in VLANs, rather than global configuration commands.

2.3.4 EFOS Switching SNMP

- MIB II object if Speed does not list LAG ports.
- The following MIB objects are labeled Not Supported:
  - RFC 1493 Bridge MIB: dot1dTpLearnedEntryDiscards, dot1dStaticTable
  - RFC 1643 Ethernet MIB: dot3CollTable
  - RFC 2233 Interfaces MIB: ifCounterDiscontinuityTime, ifStackTable, ifRcvAddressTable, ifMIBObjects Group
  - RFC 2674 VLAN MIB: dot1qFdbTable, dot1qTpFdbTable, dot1qTpGroupTable, dot1qForwardAllTable, dot1qForwardUnregisteredTable, dot1qStaticUnicastTable, dot1qStaticMulticastTable, dot1qConstraintSetDefault, dot1qConstraintTypeDefault, dot1qLearningConstraintsTable.
- In the private MIB, the object agentLagSummaryName cannot be set.
- Set and Get operations on object dot3adAggPortActorAdminState occurs in reverse order.
- The VLAN port option AdmitUntaggedOnly is missing in SNMP.

2.3.5 EFOS QoS

The following features are included in the QoS package:

- Access Lists
- Class of Service (CoS)

- Differentiated Services (DiffServ)

NOTE: Not every feature within the QoS is supported on every platform. In addition, the individual features within the QoS may not always be simultaneously enabled on the same platform.

NOTE: Contact your sales representative at https://support.broadcom.com for information concerning your platform. Also, refer to the release notes for your specific platform for that platform's restrictions or limitations concerning that feature.

2.3.6 EFOS QoS Configuration

The following items are applicable to DiffServ configuration:

- It is only intended for use with IP packets.
- The optional `match not` command in Class-Map Config mode, when specified, is used to negate the class match condition. This parameter is not supported for the class-map (for example, reference class) match condition.
- The IP DSCP, IP Precedence, and IP TOS (with mask) class match conditions are alternative methods to specify classification based on the contents of the IP Service Type (TOS) octet in the IP packet header:
- IP DSCP compares the high-order six bits of the IP Service Type octet and ignores the remaining bits, or
- IP Precedence compares the high-order three bits of the IP Service Type octet and ignores the remaining bits, or
- The IP TOS (with mask) is intended for use as a free-form match specification of the IP Service Type octet.
- A class of type all or any may reference at most one other class, but it must be of the same class type. Class referencing is not supported for class type ACL. Class references may be chained, but the total number of class match conditions in a chain is limited to twice the maximum number of rules normally allowed for a single class (actual value is platform specific).
- The Service Table Operational Status denotes the current up/down state of the DiffServ operation on the directional service interface. For the DiffServ Operational Status to be up, all of the following conditions must be true:
  - A policy is successfully attached to the service interface in the appropriate direction.
  - The policy contains one or more policy-class instances (it need not have policy attributes defined when best-effort service is desired).
  - Each policy-class instance refers to a valid class.
  - A valid class consists of at least one class match condition.
  - Each class match condition is supported by the policy type (direction) based on platform features and limitations.
  - All MIB table rows representing each of the preceding items must have a row status of active (of particular importance when using SNMP).
  - The port must be up, administratively enabled, and generally able to forward traffic.
- FP-79708. The bucket count for EFP meters is only modified when bytes are actually egressed on the port. If the conforming action is drop, the conforming packets are not transmitted, and the bucket counter is not changed. Therefore, all packets end up being treated as conforming packets and are discarded in this configuration.

2.3.7 EFOS QoS SNMP

- The following items pertain to the DiffServ standard MIB support (RFC 3289):
  - This MIB is supported as read-only. All SNMP configurations for DiffServ are handled through a private MIB (FASTPATH-QOS-DIFFSERV-PRIVATE-MIB).
  - Set operation is not supported on diffServAlgDropQThreshold and get would always read 16384(16K).
  - The IP Multifield Classification Table (diffServMultiFieldClfrTable) is not used. Instead, an Auxiliary Multifield Classification Table (agentDiffServAuxMfClfrTable) is defined in a Broadcom extension to the standard MIB (FASTPATH-QOS-DIFFSERV-EXTENSIONS-MIB) and is used for all DiffServ classifier definitions. This extension's MIB is also supported as read-only.

  - By default, the DiffServ Standard MIB is not compiled in standard flex-all builds. To include the DiffServ Standard MIB support in a flex-all build, modify the package.cfg file to include diffserv_std in the L7_FLEX definition.
  - When the DiffServ Standard MIB is compiled in a flex-all build, more than 256 MB of memory may be required.
- Regarding the DiffServ private MIB:
  - If a platform only supports Service Table actions for all interfaces in a particular direction (that is, does not allow individual <slot.port> specification), then any SNMP set operation performed on an individual object instance in the agentDiffServServiceTable is automatically applied to all supported DiffServ service interface instances in that direction.

2.3.8 EFOS Routing

- The EFOS Routing OSPF implementation was tested only with broadcast and point-to-point interfaces. NBMA and point-to-multipoint interfaces are not supported.
- The router may fail to calculate an OSPF route to a /32 network if the IP address in that network is used as the router ID of another router in a common OSPF area. This occurs under the following circumstances. Say an area border router is configured with an IP address with a 32-bit netmask and uses that IP address as its router ID. Further, say that this address is part of area 0. The router will then issue a type 3 summary LSA for this network for each of its non-backbone areas. A EFOS router that receives the summary LSA will fail to compute a route for this network. Routing to other destinations is unaffected. The problem can be avoided by configuring the area border router with a router ID that is not an interface address from a 32-bit subnet.
- Source IP address checking is not performed for packets routed in hardware. For IP packets sent to the CPU for forwarding, the source IP address is checked and those packets containing a net-directed broadcast or any value 224.0.0.0 or higher (including the limited broadcast) address are silently discarded.
- OSPF can only store a limited number of Link State Advertisements (LSAs). This number (OSPF_MAX_LSAS) is a function of the forwarding table size. When OSPF receives more LSAs than it can store, it enters the stub router state. OSPF may also enter stub router state when the routing table is full or when memory allocation fails. In the stub router state, the router sets the link cost of non-stub links to LSInfinity so that other routers prefer alternate routes. See RFC 3137. OSPF logs messages when the Link State Database (LSDB) becomes 90% or 95% full and when the router enters the stub router state. The router will not automatically exit the stub router state when the resource limitation is resolved. The user must manually disable and reenable OSPF to exit the stub router state.
- FP-108571. The DHCPv6 client does not send a RELEASE message when a routing interface is disabled. This behavior does not impact the user, but the DHCPv6 server reserves the address for this client until the binding entry times out.
- FP-154991. While the system boots up, a couple of link flaps are observed on fiber ports. Because of this, the routing interface is suppressed after reload, even though the restart-penalty is less than suppress threshold.
- FP-172759. Any control traffic with a match on *match term* in a route-map that is applied on an interface is affected if the action specified through *set terms* routes that traffic differently. This behavior occurs because if the route-map specifies the *set term* to policy-route the packet to a different next-hop, errors such as *neighbor-ship not getting formed* happen.
- FP-172866. When an ACL rule contains a rate-limit action or any action that requires metering in hardware, that particular ACL is not a candidate to be included in a route-map statement as a match condition. This is because for each route-map statement, a counter is added. Meters and counters are mutually exclusive in hardware, so applying a route-map with such an ACL in the match condition results in a failure.
- FP-173871. The IP address assigned from the DHCP server is not completely removed after disabling and enabling DHCP on the routing management port.
- FP-185220. When using the CLI, the configured client identifier cannot be seen in the `show ip dhcp pool configuration all` command output. The `show running-config` command shows the configured client-identifier value.
- FP-190596. When OSPF is enabled on unnumbered interface, the MIB object ospfAddressLessIf points to the ifIndex that is linked to unnumbered interface.

- FP-206129. The command **show ip ospf neighbor** takes 10-15 seconds to display all the entries when DUT has maximum routes installed.
- FP-212474. VLAN routing interfaces can be configured only from the EFOS CLI and not from the Linux shell interface.
- FP-222483. When traceroute is initiated with large packet sizes, some of the fragments may get dropped due to ICMP rate limiting.
- FP-250220. The number of VRF instances that can be created depends on the available free memory in the switch, which in turn depends on the package/feature combination in the switch software.
- FP-257331. In certain complex BGP configurations, the running configuration is not completely cleared after a reload.
- FP-257483: Directly connected routes of a configured routing interface are not removed if the routing interface is disabled. When the routing interface is enabled again, an attempt is made to add the already-present connected route. This results in a log message with text `Error message of type 20 received on netlink-request, sockid - 9, due to error File exists`. This error message can be ignored as it causes no functional impact.
- FP-259545. When the route learning for large number of routes in the order of 128K is in progress, the `show ip bgp summary` command responds with a delay of a couple of minutes.

NOTE: Once the routes are learned, there is no delay observed while executing the `show ip bgp summary` command.

- FP-267886: Sometimes, the first DHCP Release packet sent on a Network port is not sent out of the hardware. Need to re-disable DHCP to send out the Release packet.
- FP-268357: Platforms using ALPM template scaled routes and with 64 VRFs need at least 16 GB RAM.
- FP-268696: Only one Static-Route BFD session [one per family IPv4/IPv6] is supported on an interface.

2.3.9 EFOS IPv6 Routing

- FP-24150. IPv6 site-local multicast addresses are not supported.
- FP-24936. EFOS will fail ANVL tunneling ipv6ov4 test 8.2, 83, and 8.6 due to the system not sending a Neighbor Solicitation packet in response to an Echo Request. RFC 4213 states in section 3.8 that a device is not required to send a Network Solicitation in this situation since a tunnel is not considered to have a link address that must be discovered.
- FP-53694. When an IPv6 packet (with a valid link local SA and global DA) is sent to the router and when the next hop neighbor is resolved, the packet is dropped but no ICMPv6 Destination Unreachable message is returned to the sender.
- FP-170990. An IPv6 prefix-list cannot be deleted if a description is configured on that prefix-list.
- FP-214632. Console debug tracing for IPv6 ping packets is not supported.
- FP-215176. There is no CLI command to clear the IPv6 DHCP client statistics on an interface operating in client mode.
- FP-246222. A 6to4 tunnel has to be unconfigured and configured again, if the reachability to the tunnel endpoint is disrupted in the presence of an active 6to4 traffic flow(s).

2.3.10 EFOS Security

- When enabling SSH on a system that does not have the private keys, the CLI will be unresponsive while creating the RSA and DSA private keys. There is no indication on the CLI that the keys are being generated. The amount of time that the system remains unresponsive is system dependent.
- For SSL/TLS, the device will not generate its own certificates, and no default certificates are present. In order to use SSL/TLS, certificates must be obtained (in PEM format) from a certificate granting authority (commercial or otherwise) and downloaded to the device. Tools for establishing a certificate granting authority are not provided with EFOS software.
- FP-151600. When the login authentication method is configured as enable or line to authenticate users using the enable or line password, then the user name is not required. The “User:” prompt is not displayed for console and Telnet users; however, SSH users are still prompted to provide user details.

2.3.11 EFOS Multicast

- The IP multicast feature is not available on some platforms.

NOTE: Contact your sales representative for information concerning your platform.

- IP multicast statistics are not supported.
- The following PIMDMv4 ANVL test suites are labeled Not Supported:
  - EFOS does not support a single Hello timer for all the interfaces.
  - IP Address Change is not easy to simulate with current EFOS configuration and ANVL scripts.
  - A Prune Echo message will be sent only if an interface has more than one neighbor.
  - DUT supports only the largest active holdtime and not the most recently used. There is a bug in ANVL for this. Step 15 should be DUT and should reset the PT to prune-holdtime.
  - DUT does not record the TTL value for each data packet. DUT records the TTL value only when first data packet arrives and maintains that throughout.
- The following PIMSMv4 ANVL test suites are labeled Not Supported:
  - Joins with unknown RP are discarded
  - Change the IP address on DUT interface requires to delete the old `ip-addr` and reconfigure the new `ip-addr`. Scripts as not easy to write to simulate this test.
- The following PIMSMv6 ANVL test suites are labeled Not Supported:
  - Joins with unknown RP are discarded
  - The user is not allowed to change link local address of an interface
  - C-RP can only be configured for global scope groupaddress (ff[0-f]e::)
- Tests in ANVL suites for IGMPv2, IGMPv3, MLDv1 and MLDv2 related to Host-behavior are labeled Not Supported.
- MLDv1 ANVL test suites 7.20 and 7.21 are labeled Not Supported because of echo forwarding not supported.
- IGMPv3 ANVL test suites 10.20 and 10.21 are reported as inconclusive and 10.23 as failed, but these pass if verified manually. The rationale is ANVL sends v1 report and expects the DUT to send v1 query packets when the DUTs interface is configured in v3.RFC does not clearly say that when lower version reports are received, query packets of that version only should be sent on that interface of the router. CISCO® also does not send lower version query packets when lower version reports are received. All that matters is what version is configured on that interface. So, to be compatible with CISCO, it has been decided that the FASTAPTH router will also behave the same way as CISCO and not send lower version query packets when lower version reports are received.
- DVMRPv4 ANVL test suite 5.8 fails because of timing issue. The test basically passes if verified manually, With a tolerance time of 10 (setting in.prm file), the ANVL receives all the expected prunes, but with slight variation in timing for in between prunes. The prunes are expected to receive at 0, 3, 6, 12, 24 seconds, but they are received at 0, 3.5, 7.9, 14, 24 seconds, respectively. So, the test case fails ANVL-wise, but manual observation is required to see that protocol passes in sending all the 5 expected prunes. The dumps of received prunes and the timing can be clearly monitored from ANVL log to judge it as passed.
- MLDv2 ANVL test suite 3.4 and 7.21 fails because of a timing issue. The test passes if verified manually. Even though the ANVL receives the expected two queries within the listen time (2*query-interval + tolerance time), for example, 257 seconds, the interval between the two queries is more than 125 seconds, which the ANVL is complaining about.
- FP-18687, FP-18690. Multiple DVMRP ANVL tests fail due to forwarding entries not being updated quickly enough after ANVL sends a control frame. Data frames are sent immediately following the control frame, before the entry has been updated. When viewing the forwarding table after running these tests, it can be seen that the forwarding entry has been properly formed.
- FP-54028. Multicast traffic is not forwarded while performing negative testing with PIMSM as the MRP.
- FP-76345, FP-180295. When PIMSM is used as the Multicast Routing Protocol with IGMPv2/MLDv1 for 2048 Groups, it is observed that the steady state CPU utilization is around 45% – 50% and is peaked at 80%. This CPU overload exists only in the combination of PIM-SM and IGMPv2/MLDv1: extra processing associated with the lack of source-specific information in IGMPv2/MLDv1.
- FP-173779. CPU utilization on a PIM non-DR router is high when the non-DR is in the shortest path to the host that is receiving data.
- FP-251125. Beyond 1200 multicast route entries, the PIMDM protocol is not able to prune few entries due to the flood of PIM State Refresh messages.

2.3.12 EFOS Data Center

- FP-210877. Tunnels of different tenants with matching {source TEP, remote TEP} will share the same tunnel in the hardware. Each tunnel in the hardware is created with individual unicast packet and byte counters. Even though the application tunnel handles are different, tenant tunnels sharing the same hardware tunnel will always have the same tunnel packet/byte counter values since they fetch the counters of the shared hardware tunnel. Similarly, when a clear operation is performed on a tenant tunnel, it clears the counters on the shared hardware tunnel.

2.3.13 EFOS Open API

The Quagga Routing Information Base (RIB) and Quagga routing protocols execute in separate processes. If the routing protocol process is killed or otherwise exits, the learned and configured routes will remain in the RIB. The EFOS RIB also exhibits this behavior. Restarting the routing protocol will clean up the RIB.

2.3.14 NETCONF

- FP-240165 NETCONF: config tags on request XML supports only the format “<config xmlns=“urn:ietf:params:xml:ns:netconf:base:1.0">”.
- FP-245214 NETCONF: “XPath error: Undefined namespace prefix” prints are observed on console. These prints are harmless and can be ignored.
- FP-245628. A message appears in the log file (Missing YIN module “broadcom-common-types”) after clearing the configuration. These messages can be ignored.
- FP-254897 NETCONF: PTP section on NETCONF startup config fails to apply during system startup. The following logs are observed “Failed to perform initial copy of startup to running. Proceeding after failing to apply startup.”
- FP-262150: User cannot fetch PTP state data under "ptp-state" subtree. If such requested is attempted, it would return empty data. However, one can fetch ptp-state subtree data by rather fetching the entire PTP tree.

2.3.15 EFOS Management

- FP-218328. While entering certain modes via the CLI, the Help prompt inadvertently provides information on commands that may not be available. For example, while configuring a physical interface, the autostate command is listed, but is actually applicable for VLAN routing interfaces only.
- FP-218527. The CLI error message is not optimal when an invalid range is entered during configuration of standard or enhanced ACL numbers, will be repaired in later release.
- FP-225389. The user is not being prompted to save the configuration after an 'application install' config operation by the switch automatically. This issue is only seen when users try to install their own applications on the switch. The workaround for this issue is to manually save the configuration after the 'application install' config operation.
- FP-257331. In certain complex BGP configurations, the running configuration is not completely cleared after a reload.

Appendix A: Related Documents

The references in this section may be used in conjunction with this document.

NOTE: Broadcom provides customer access to technical documentation and software through its Customer Support Portal (CSP) and Downloads & Support site. For a CSP account, contact your Sales or Engineering support representative.

For Broadcom documents, replace the “x” in the document number with the largest number available in the repository to ensure that you have the most current version of the document.

Document Name	Number	Source
EFOS Administrator’s Guide	EFOS3.X-SWUM1xx	Broadcom CSP
EFOS CLI Command Reference	EFOS3.X-SWUM2xx	Broadcom CSP
EFOS Functional Specification	EFOS3.X-PG1xx	Broadcom CSP
EFOS Getting Started Guide	EFOS3.X-PG2xx	Broadcom CSP
EFOS Scaling Parameters and Values	EFOS3.X-RM1xx	Broadcom CSP

Revision History

EFOS3.X-RN300; November 26, 2018

Initial release

BROADCOM®