ArubaOS-CX 10.02.0012 Release Notes for the Aruba 8325 Switch Series

a Hewlett Packard
Enterprise company

Part Number: 5200-5880a
Published: March 2019
Edition: 2

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Description

This release note covers software versions for the ArubaOS-CX 10.02 branch of the software.

NOTE: If you run the `show version` command on the 8325, the version number will display GL.10.02.*xxxx*, where *xxxx* is the minor version number.

ArubaOS-CX is a new, modern, fully programmable operating system built using a database-centric design that ensures higher availability and dynamic software process changes for reduced downtime. In addition to robust hardware reliability, the ArubaOS-CX operating system includes additional software elements not available with traditional systems, including the features included in the Features section of this release note.

Version 10.02.0001 is the initial build of major version 10.02 software.

Product series supported by this software:

Aruba 8325 Switch Series

Important information

To avoid damage to your equipment, do not interrupt power to the switch during a software update.

Industry and government certifications

Refer to the Approved Product Lists sites for the Common Criteria, FIPS 140-2 and DoDIN APL to obtain the product certification details. Products should be used as evaluated and defined in the respective configuration guides.

- Common Criteria: **https://www.niap-ccevs.org/Product/**
- FIPS 140-2: **https://csrc.nist.gov/Projects/Cryptographic-Module-Validation-Program/Validated-Modules/Search**
- DoDIN APL: **https://aplits.disa.mil/processAPList.action%3bjsessionid=f2l2uEoOL6g4YYsEyVrFgx4W4f8J-Fgu4DLFZZmPXCvl-7Ft9SGf%21809605859**

Version history

All released versions are fully supported by Hewlett Packard Enterprise, unless noted in the table.

Version number	Release date	Remarks
10.02.0012	2019-03-08	Released, fully supported, and posted on the web.
10.02.0010	2019-01-29	Released, fully supported, and posted on the web.
10.02.0001	2018-12-14	Initial release of ArubaOS-CX 10.02 for the Aruba 8325 Switch Series. Released, fully supported, and posted on the web.

Products supported

This release applies to the following product models:

Product number	Description
JL624A	Aruba 8325-48Y8C 48p 25G SFP/+/28 8p 100G QSFP+/28 Front-to-Back 6 Fans and 2 PSU Bundle
JL625A	Aruba 8325-48Y8C 48p 25G SFP/+/28 8p 100G QSFP+/28 Back-to-Front 6 Fans and 2 PSU Bundle
JL626A	Aruba 8325-32C 32-port 100G QSFP+/QSFP28 Front-to-Back 6 Fans and 2 Power Supply Bundle
JL627A	Aruba 8325-32C 32-port 100G QSFP+/QSFP28 Back-to-Front 6 Fans and 2 Power Supply Bundle

Compatibility/interoperability

The switch web agent supports the following web browsers:

Browser	Minimum supported versions
Edge (Windows)	38
Chrome (Ubuntu)	54 (desktop) 56 (mobile)
Firefox (Ubuntu)	52
Safari (MacOS, IOS Only)	10

NOTE: Internet Explorer is not supported.

The following table provides information on compatibility of the switches found in this release note with network management software:

Management software	Supported version(s)
Airwave	8.2.6
Network Automation	10.10, 10.11, 10.20, 10.21, 10.30, 10.40
Network Node Manager	10.10, 10.20, 10.21, 10.30, 10.40
IMC	7.3 (E0506P05)

NOTE: For more information, see the respective software manuals.

Enhancements

This section lists enhancements added to this branch of the software.

Software enhancements are listed in reverse-chronological order, with the newest on the top of the list. Unless otherwise noted, each software version listed includes all enhancements added in earlier versions.

Version 10.02.0012

No enhancements were included in version 10.02.0012.

Version 10.02.0010

No enhancements were included in version 10.02.0010.

Version 10.02.0001

6in4 Tunnels

Support for tunneling IPv6 traffic in an IPv4 network has been added.

BGP connections over GRE tunnels

This feature enables BGP peering and routing through a GRE tunnel.

Control plane ACLs

Control plane ACLs control access to the control plane. This is primarily used to control access to services like SSH, SNMP, NTP, or the web server. However, this can also be extended to control plane network services like BGP. Control plane ACLs support IPv4 and IPv6 type access lists and can be applied per VRF.

Egress queue shaping

Egress queue shaping limits the amount of traffic transmitted per output queue. The buffer associated with each egress queue stores the excess traffic in order to absorb bursts and smooth the output rate according to the configuration. For example, an administrator might limit the strict-priority queue traffic to prevent low-priority queue starvation in the event of a device inappropriately sending too many higher-priority packets. Egress queue shaping can be configured on an ethernet port or on a link aggregation group (LAG).

IPSLA

IPSLA enables network performance monitoring. IPSLA is supported for ICMP Echo, TCP Connect, UDP Echo, HTTP, and VOIP jitter. Monitoring is enabled using the Network Analytics Engine to enable robust history and ability to capture additional information when anomalies are detected.

IPv6 multicast routing

IPv6 multicast routing supports MLD, MLD snooping, and PIM-SM v6 routing, providing the capability to enable routing of IPv6 multicast traffic.

Mirror to CPU

This feature adds the capability to mirror dataplane packets to the CPU for monitoring directly on the switch using Tshark.

Multi protocol BGP

Multi protocol BGP with IPv6 address family, also known as BGP4+, enables sharing of IPv6 routes using BGP.

Multicast routing

Loopback for RP and BSR is now supported for both IPv4 and IPv6.

NAE encrypted credentials

The Network Analytics Engine (NAE) now supports encrypted credentials for connecting to external services. These credentials are encrypted securely using the TPM contained on the switch.

NAE periodic callback actions

This feature introduces a new condition syntax to periodically execute a callback function for a given period of time. Using the Network Analytics Engine (NAE) python API, users can set callbacks to be called in regular intervals. This allows a script writer to create conditions based on a fixed time interval.

NAE time series for external APIs

Using Network Analytics Engine (NAE) period callback actions, an NAE agent can be created using an external API from another device or services. Monitoring an API in a connected access switch, or even the Mobility Master, is now possible. When an anomalous event is discovered (like excessive errors on a trunk port from access to distribution), NAE can collect additional information from the distribution switch or access switch automatically.

NTP master

This feature allows the switch to act as the NTP master in the network.

Object groups for ACLs

This feature enables the creation of named groups representing sets of IPv4 or IPv6 addresses and L4 port ranges. Object groups allow administrators to simplify their configurations of ACLs. By defining a few rules with address or port groups, users can potentially effect hundreds of hosts and services in a clear and simple manner.

Policy Based Routing

Policy Based Routing (PBR) is a flexible feature for creating various routing decisions based on additional information in the packets. One common use of PBR is to implement source based routing, for example, routing all guest VLAN traffic out a separate WAN link.

Remote mirroring

The remote mirroring capability uses GRE encapsulated mirrored frames to a destination network device.

Rx flow control

Frames received on a port will pause sending egress packets. When the pause timer expires, the transmission of packets will proceed. This is commonly used in legacy scenarios or with services like iSCSI that cannot handle any packet loss very well.

Security

RADIUS accounting, PKI for syslog, and ServiceOS console password have been added to enhance security on the switch.

Syslog over TLS

This feature enables secure configuring of a syslog server with TLS security.

VLAN ACLs/Policies/Classifiers

ACLs, policies, and classifiers can now be applied to a VLAN interface, simplifying the application of these elements on VLANs facing clients or other networks.

VSX

- VSX and Spanning Tree interoperability
  - VSX interoperability with MSTP.
  - Each VSX member shares a virtual bridge ID, ensuring switches connected to the VSX pair see the pair as a single Spanning Tree entity.
- VSX active/active multicast routing enables active/active control plane with PIM dual-DR for multicast routing, further limiting traffic that must traverse the ISL and ensuring optimal performance.
- New VSX sync features, including CoPP, PBR, QoS, VLAN ACL, VLAN classifier/policy, AAA/users, DNS, NTP, sFlow, SNMP, SSH, and static route.
- VSX static LAG enables VSX lags without LACP.

VXLAN

Support for static L2 VXLANs.

Fixes

This section lists released builds that include fixes found in this branch of the software. Software fixes are listed in reverse-chronological order, with the newest on the top of the list. Unless otherwise noted, each software version listed includes all fixes added in earlier versions.

The Symptom statement describes what a user might experience if this is seen on the network. The Scenario statement provides additional environment details and trigger summaries. When available, the Workaround statement provides a workaround to the issue for customers who decide not to update to this version of software.

NOTE: The number that precedes the fix description is used for tracking purposes.

Version 10.02.0012

LACP
CR_45192

Symptom: LACP failover takes more than five seconds to recover/redirect traffic through the backup link.

Scenario: When doing a manual failover for LACP by manually removing the line card and redirecting traffic to the backup link can take longer than five seconds.

OSPF
CR_45094

Symptom: OSPF link failover takes more than five seconds to recover traffic with 5k or more routes.

Scenario: In an OSPF failover scenario, you may see failover taking more than five seconds to recover traffic with 5k or more routes.

CR_45207

Symptom: OSPF adjacency is down.

Scenario: OSPF adjacency goes down after a burst of traffic causes the best effort queue on the port to be oversubscribed.

Workaround: Remove or fix the cause of the traffic burst.

Version 10.02.0010

Classifier
CR_41870

Symptom: In certain conditions, the switch fails to honor the drop policy action for iIPv4 and IPv6 classes.

Scenario: While an IPv6.v6 class is updated while being applied to exiting traffic, the switch may fail to honored the drop policy action.

Workaround: Disable the interface where the policy is applied to allow for all class entries to be reprogrammed.

IP SLA
CR_44958

Symptom: The switch incorrectly floods the log messages with IPSLA related messages.

Scenario: When NTP is configured in the device, the switch may incorrectly flood the log messages with IPSLA related messages.

Multicast
CR_44064

Symptom: In certain conditions, the switch experiences loss of multicast traffic.

Scenario: In a VSX configuration, after restoring a configuration from a checkpoint that includes VRF and IGMP configurations, the multicast traffic may fail for some VLANs.

Workaround: Resend IGMP joins in all VLANs where the multicast traffic is failing.

Issues and workarounds

The following are known open issues with this branch of the software.

The Symptom statement describes what a user might experience if this is seen on the network. The Scenario statement provides additional environment details and trigger summaries. When available, the Workaround statement provides a workaround to the issue.

BGP

CR_37739

Symptom: When the switch uses route leaking and a BGP peer to learn the same route, the switch may incorrectly install the two routes as ECMP routes.

Scenario: In a multi VRF environment, while performing mutual route leaking on the VRRP peers with BGP neighborship established in between and towards the upstream network, the switch installs both routes as ECMP instead of preferring the leaked route.

Workaround: Use OSPF routing as the interconnect between VRRP peers instead of BGP.

CR_43864

Symptom: BGP Sessions configured under non-default VRF fails to come up.

Scenario: When restoring a configuration from a checkpoint, BGP sessions under non-default VRF fail to come up.

Workaround: Use SFTP to restore the configuration or remove and re-configure BGP sessions under non-default VRF after a configuration restore from a checkpoint.

Fans

CR_42515

Symptom: Various LEDs alternate colors every few seconds and the fan status in the `show environment fan` command alternates between `ok` and `fault` every few seconds.

Scenario: When a fan assembly that does not match the current system airflow direction is inserted into the switch, the fan assembly LED on the rear of the switch alternates between green and red, the Fan LED on the front of the switch alternates between green and amber, and the Global LED on the front of the switch alternates between green and amber every few seconds. The status of the fan in the `show environment fan` command output will also alternate between `ok` and `fault`. Additionally, the event log will display corresponding messages regarding fan faults and recoveries.

Workaround: Replace the fan assembly with one that matches the system airflow direction.

OSPF

CR_08491

Symptom/Scenario: OSPFv2 and OSPFv3 do not support detailed LSA `show` commands.

Workaround: Use the `diag` command, instead.

CR_35544

Symptom: OSPFv3 neighor is not formed when the area type is changed on the fly.

Scenario: In a scaled setup with a large number of interfaces, when the area type is changed from Normal to NSSA, the OSPFv3 Neighborship may get stalled in Exchange state.

Workaround: Shut down the OSPF peering interface or stop the current traffic on this interface and then make the OSPF area changes.

VRRP

CR_24910

Symptom: Unable to configure same IPv6 link local address as primary virtual IP address under different VRFs.

Scenario: Unique virtual link local addresses have to be configured for all VRRP IPv6 instances irrespective of VRF.

Workaround: Do not use the same virtual link local address across different VRFs.

VLAN

CR_42560

Symptom: In certain conditions, protocol VLAN mapping does not work as expected.

Scenario: When protocol VLAN mappings are applied on a port and then if a native VLAN configuration is applied or changed on some other port, the protocol VLAN mappings do not work as expected. For example, the packets are either dropped or accepted on a wrong VLAN.

Workaround: Re-configure the protocol-based VLAN on the ingress port, after completing all VLAN configurations.

VSX

CR_37202

Symptom: The switch boots up with an incorrect startup configuration file.

Scenario: When executing `write memory` or `copy running-config startup-config` simultaneously or in very close succession (few seconds) on both the primary and secondary VSX switches, the VSX switch may end up with an incorrect startup configuration after system reboot.

Workaround: Erase current startup configuration then save to memory using `write memory`.

CR_43419

Symptom: A VSX switch reports a missing-reference error and stops syncing further configurations through VSX-Sync.

Scenario: When VSX-Sync is enabled for MCLAG interfaces at the global level, if an MCLAG interface is created on the secondary VSX switch without creating the same on the primary switch first, the VSX-Sync will report a mismatch or missing-reference between VSX pairs.

Workaround: When VSX-Sync is enabled for MCLAG interfaces at the global level, always create MCLAG interfaces on the primary VSX switch before configuring the same on the secondary VSX switch.

CR_43542

Symptom: Traffic passing through the secondary VSX switch is lost.

Scenario: After deleting a VLAN (with vsx-sync enabled) and VLAN interface in the primary VSX switch, the VLAN is also deleted on the secondary VSX peer, but the VLAN interface remains on the peer. When the same VLAN (with vsx-sync enabled) and VLAN interface is added back to the primary switch, there is still a traffic loss on the secondary switch.

Workaround:

1. Enable the VLAN with vsx-sync enabled on the primary switch and verify the VLAN is synced to the secondary.

NOTE: Removing the VLAN from the primary initially removed he VLAN membership on both peers. Make sure to re-configure the VLAN back to the client interfaces of both peers.

2. Delete the VLAN interface on the secondary switch.
3. Add the deleted VLAN interface back to the secondary with (with the proper IP address).

NOTE: If `vsx-sync active-gateway` is enabled on the primary, wait for it to get synced to the secondary.

Traffic will resume right after re-configuring the initial configuration of the VLAN interface to the secondary switch.

Feature caveats

Feature	Description
MVRP and VSX	MVRP is mutually exclusive with VSX.
VSX and RPVST+	RPVST+ is mutually exclusive with VSX.
RPVST+ and MSTP	Spanning Tree can only run in MSTP or RPVST+ mode.
RPVST+ and MVRP	RPVST+ is mutually exclusive with MVRP.
VRRP and Proxy ARP	VRRP is mutually exclusive with Proxy ARP on the same interface.
IGMP/PIM on Loopback and GRE interfaces	PIM and IGMP cannot be enabled on Loopback and GRE interfaces.
Counters	Layer 3 Route-only port counters are not enabled by default. Enabling them will reduce ipv4 route scale to 80K.
Network Analytics Engine (NAE)	Agents monitoring a resource that has column type enum with a list of strings (as opposed to a single string enum) is not supported.
Network Analytics Engine (NAE)	The following tables are not supported for NAE scripts: OSPF_Route, OSPF_LSA, OSPF_Neighbor, BGP_Route.
Network Analytics Engine (NAE)	Network Analytics Engine (NAE) agents execute Command Line Interface (CLI) actions as 'admin' user, so they have permission to run any command by default. However, when the authentication, authorization and accounting (AAA) feature is enabled, the same restrictions applied to 'admin' will also apply to NAE agents. Keep that in mind when configuring the AAA service, e.g. TACACS+, and make sure to give admin user permission to run all commands needed by enabled agents. Otherwise, some CLI commands may be denied and their outputs won't be available. Actions other than CLI won't be affected and will execute normally. Also, NAE agents won't authenticate, thus the AAA service configuration must not block authorization for unauthenticated 'admin' user. ClearPass doesn't support such configuration, so it cannot be used as a TACACS+ server.
Classifiers	IPv4 egress ACLs can be applied only to route-only ports.
Classifiers	Classifier policies, IPv6 and MAC ACLs are not supported on egress.
Classifiers	DSCP remarking is performed only on routed packets.

Table Continued

Feature	Description
Classifiers	For security ACLs, HPE strongly encourages modifications be done as a two step process: Bring down the port and then modify.
Classifiers	Policies containing both MAC and IPv6 classes are not allowed.
Classifiers	Egress ACL logging is not supported.
REST	REST supports the 'admin' and 'operator' roles but does not work with TACACS+ command authorization.
REST	With the exception of ACLs and VLANs, REST APIs using POST/PUT/DELETE are not validated before performing the function. Therefore, to avoid unintended results or side effects, HPE recommends testing the API write action first.
VSX	VSX active-forwarding only works when the L3 interface is IPv4. Enabling it on an interface that has dual-stack or IPv6 may result in traffic losses.
VSX and MSTP	MSTP Inter Region is not supported.
VSX and MSTP	L2 link connected parallel to ISL link will be blocked by MSTP.
VSX and MSTP	MSTP config sync between VSX peer switches is not supported.
VSX and MSTP	All port-specific spanning tree configurations are not recommended to configure on ISL.
VSX and MSTP	Topology change for VSX lags are accounted on active multi-chassis lag role only.
VSX and MSTP	It is recommended to use MSTP only in a VSX environment replacing Loop protect.
VSX and MSTP	It is recommended to synch time using NTP for `show spanning-tree vsx-peer` commands for the last topology change field.
VSX and MSTP	Common Bridge ID will continue to be used even after VSX split brain scenario is identified.
VSX and MSTP	Interop with other STP flavors are not supported. VSX pair configured as non-root switch.
VSX and MSTP	L2 links parallel to ISL link in VSX will be blocked by MSTP; Peer-Keep-Alive as SVI over L2 is not supported.
VSX and MSTP	Configured/default system-mac+1 and system-mac-1 is used in VSX-pair and should not be used by any STP bridge.
VXLAN and tunneling	VXLAN is mutually exclusive with any other tunneling (GRE and 6in4).
VXLAN and VSX	VXLAN is mutually exclusive with VSX.

Upgrade information

Version 10.02.0012 uses ServiceOS GL.01.03.0002.

IMPORTANT: Do not interrupt power to the switch during this important update.

IMPORTANT: If you are upgrading from any version of 10.00 or any version of 10.01 earlier than 10.01.0020 and the switch is configured with MCLAG, VSX reconfiguration is required after the upgrade. Network downtime is required for this upgrade. For more information on upgrading from MCLAG to VSX, see the *ArubaOS-CX Virtual Switching Extension (VSX) Guide*.

IMPORTANT: The device running configuration may be lost if you downgrade from 10.02.0012 back to a version of 10.01 earlier than 10.01.0040 or 10.00 (any version). Prior to upgrading to 10.02.0012, save the current working configuration backup using a checkpoint or external servers (TFTP/SFTP, refer to **Upgrade information** on page 12). After the upgrade, if you must downgrade from 10.02.0012, erase startup-config. Reboot into the previous version you are downgrading to then restore the working configuration from the backup you created before the upgrade.

NOTE: Some Network Analytics Engine (NAE) scripts may not function properly after an upgrade. HPE recommends deleting existing NAE scripts before an upgrade and then reinstalling the scripts after the upgrade. For more information, see the *Network Analytics Engine Guide*.

File transfer methods

The switches support several methods for transferring files to and from a physically connected device or via the network, including TFTP, SFTP, and USB. This section explains how to download and run new switch software.

File transfer setup

TFTP

Before using TFTP to transfer the software to the switch, make sure:

- A software version for the switch has been stored on a TFTP server accessible to the switch via management port. (The software file is typically available from the Switch Networking website at **http://www.hpe.com/networking/support**.)
- The switch is properly connected to your network via the management port and has already been configured with a compatible IP address and subnet mask.
- The TFTP server is accessible to the switch via IP. Before you proceed, complete the following:
  - Obtain the IP address of the TFTP server in which the software file has been stored.
  -

NOTE: If your TFTP server is a UNIX workstation, ensure that the case (upper or lower) that you specify for the filename is the same case as the characters in the software filenames on the server.

SFTP

For some situations you may want to use a secure method to issue commands or copy files to the switch. By opening a secure, encrypted SSH session and enabling IP SSH file transfer, you can then use a third-party software application to take advantage of SFTP. SFTP provide a secure alternative to TFTP for transferring information that may be sensitive (like switch configuration files) to and from the switch. Essentially, you are creating a secure SSH tunnel as a way to transfer files with SFTP channels.

Before using SFTP to transfer the software to the switch, make sure:

- A software version for the switch has been stored on a computer accessible to the switch via management port. (The software file is typically available from the Switch Networking website at **http://www.hpe.com/networking/support**.)
- The switch is properly connected to your network via the management port and has already been configured with a compatible IP address and subnet mask.
- The computer containing the software image is accessible to the switch via IP. Before you proceed, complete the following:
  - Obtain the IP address of the computer on which the software file has been stored.
  -
- Establish a secure encrypted tunnel between the switch and the computer containing the software update file (for more information, see the *Fundamentals Guide* for your switch).

  NOTE: This is a one-time procedure. If you have already setup a secure tunnel, you can skip this step.
- Enable secure file transfer using the `ssh server vrf <VRF-name>` command (for more information, see the *Command-Line Interface Guide* for your switch).

  ```
  switch(config)# ssh server vrf mgmt
  ```

USB

Before using USB to transfer the software to the switch, make sure to:

- Store a software version on a USB flash drive.
-
- Determine the name of the software file stored on the USB flash drive.
- Enable USB on the switch:

```
switch(config)# usb
switch(config)# do usb mount
switch(config)# do show usb
Enabled: Yes
Mounted: Yes
```

Copying the software and rebooting the switch

Procedure

1. Copy the software to the secondary flash on the switch using the `copy <remote-URL> {primary | secondary} [vrf <VRF-name>]` command (for more information, see the *Command-Line Interface Guide* for your switch).
   - For TFTP:
   - For SFTP:
   - For USB:

   When the switch finishes downloading the software file, it displays this progress message:

```
Verifying and writing system firmware…
```

2. When the installation finishes, confirm the version and the file saved to disk are what was transferred. Do this using the `show images` command (for more information, see the *Command-Line Interface Guide* for your switch).
3. You must reboot the switch to implement the newly downloaded software image using the `boot system [primary | secondary | serviceos]` command (for more information, see the *Command-Line Interface Guide* for your switch).
4. Upon successful reboot, execute the `show system` command and verify the correct firmware revision.

A Security Bulletin is the first published notification of security vulnerabilities and is the only communication vehicle for security vulnerabilities.

- Fixes for security vulnerabilities are not documented in manuals, release notes, or other forms of product documentation.
- A Security Bulletin is released when all vulnerable products still in support life have publicly available images that contain the fix for the security vulnerability.

Finding Security Bulletins

Procedure

1. Go to the HPE Support Center - Hewlett Packard Enterprise at **www.hpe.com/support/hpesc**.
2. Enter your product name or number and click **Go**.
3. Select your product from the list of results.
4. Click the **Top issues & solutions** tab.
5. Click the **Advisories, bulletins & notices** link.

Security Bulletin subscription service

You can sign up at **http://www.hpe.com/support/Subscriber_Choice** to initiate a subscription to receive future Hewlett Packard Enterprise Security Bulletin alerts via email.