

High Availability Stability

DRAFT

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

DRAFT

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2023 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

Contents	3
Split-Brain Detection and Prevention	4
Activate High Availability	6
Advanced Options: HA Failover Detection Time Multiplier	7
Monitoring High Availability Edges	8
Overview	8
High Availability specific information on the Monitor > Edge	8
High Availability Failover Bars on the System tab	9
The Monitor > Edge > HA Standby Tab	11
HA Event Details	14
Enterprise-Level Orchestrator Alerts and Events	15
Request Packet Capture Bundle with New Orchestrator UI	18
Packet Capture for Edges configured for High Availability	19

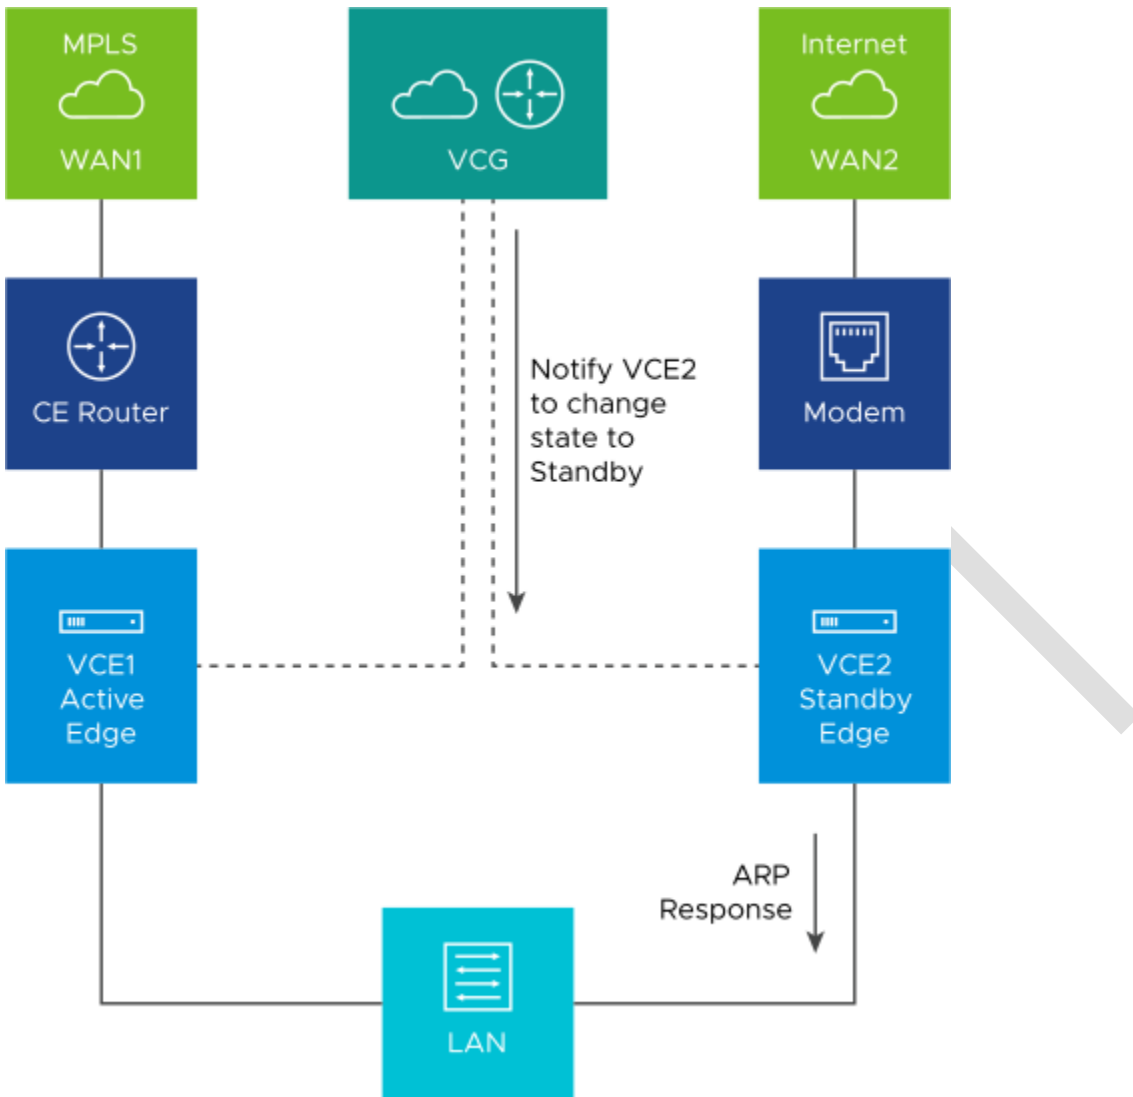
Split-Brain Detection and Prevention

The primary Gateway is used to prevent split-brain conditions.

The Gateway has a pre-existing connection to the Active Edge (VCE1). In a split-brain condition, the Standby Edge (VCE2) changes state to Active and tries to establish a tunnel with the Gateway (VCG). The Gateway will send a response back to the Standby Edge (VCE2) instructing it to move to Standby state, and will not allow the tunnel to be established. Gateway will always have tunnels only from the Active Edge.

As soon as the HA link fails, the VCE2 moves to the Active state and enables the LAN/WAN ports, and tries to establish tunnels with the Primary Gateway. If the VCE1 still has tunnels, the Primary Gateway instructs the VCE2 to revert to the Standby state and thus the VCE2 blocks its LAN ports. Only the LAN interfaces remain blocked (as long as the HA cable is down). As illustrated in the following figure, the Gateway signals VCE2 to go into the Standby state. This will logically prevent the split-brain scenario from occurring.

Note The normal failover from Active to Standby in a split-brain scenario is not the same as the normal failover. It could take a few extra milliseconds/seconds to converge.



Note Beginning in Release 5.2.0, the **HA Failover Detection Time Multiplier** feature can be used to set a longer High Availability failover threshold. The timer represents how long a Standby Edge will wait for a heartbeat packet from the Active Edge before becoming active. In some instances, where a lower model Edge is under high traffic load, the Active Edge's heartbeat packet may take longer than the default threshold time to be delivered to the Standby Edge. As a result the Standby Edge triggers a failover and is promoted to Active, resulting in a Split-Brain state.

Setting the HA Failover Detection Time Multiplier to a value higher than the default can lessen the risk of a Split-Brain state in this scenario. The default value is 700 milliseconds (ms), and this value can be increased up to a value of 7000 ms. For more information, see [Chapter 2 Activate High Availability](#).

Activate High Availability

You can activate High Availability (HA) on a pair of Edges to ensure redundancy.

- 1 In the Enterprise portal, click **Configure > Edges**.
- 2 Select the SD-WAN Edge from the list and click the **Device** tab.
- 3 Scroll down to the **High Availability** section and click **Active Standby Pair**.

High Availability

HA: Active Standby Pair Segment Agnostic

High Availability is enabled at the Edge level. When using Active/Standby Pair HA, enable HA prior to connecting the Standby SD-WAN Edge. To learn more, please consult our HA documentation

Select Type

None

Active Standby Pair

Cluster

VRRP with 3rd party router

HA Interface GE1

⚠ The VLAN value configured for the switched access port is reset to the value derived from the associated profile before moving to an HA Interface

Deploy with Unique LAN MAC Address ⓘ Enable Graceful Switchover (require Graceful Restart in routing protocol)

⌵ Advanced Settings

- 4 Click **Save Changes** at the bottom of the **Device** window.

By default, the HA interface to connect the pair is selected as follows:

- For Edges 520, 520v, and 540: The LAN1 port is used as HA interface and DPDK is not enabled on these platforms.
- For Edges 510, 610, 620, 640, 680, 840, 2000, 3400, and 3800: The GE1 port is used as HA interface and DPDK is enabled on these platforms.

Note The above HA interfaces are selected automatically but beginning with Release 5.2.0 you can configure any LAN interface to be the HA interface with the **HA Interface** option.

By default, High Availability uses a common virtual MAC address to support seamless failover between devices. If you need to use a unique MAC address in certain virtual environments, instead of generating a common or shared virtual MAC address, you can select the **Deploy with Unique LAN MAC** checkbox, which is deactivated by default. This option will use the physical MAC address for hardware Edges and the assigned MAC address for virtual Edges. The LAN and Routed LAN use physical MAC address, while the WAN links would still use virtual MAC address.

You can activate or deactivate the **Deploy with Unique LAN MAC** option only when you enable High Availability by choosing **Active Standby Pair**. Once High Availability is enabled, you cannot activate or

deactivate **Deploy with Unique LAN MAC** at a later point of time.

If you need to activate or deactivate the option, turn off High Availability as follows:

- 1 In the **High Availability** section, click **None**.
- 2 Click **Save Changes** at the top of the **Device** window.

Enable the High Availability again and then click the **Deploy with Unique LAN MAC** checkbox to activate or deactivate the option.

Advanced Options: HA Failover Detection Time Multiplier

Beginning in Release 5.2.0, a user can manually configure the time threshold before the Active Edge is marked as non-responsive which would trigger a failover to the Standby Edge. On some Edge platforms an Edge may experience a high amount of traffic sufficient to delay sending out a heartbeat response to the Standby Edge indicating that it is still functioning. This delay may exceed the default 700 millisecond threshold and trigger the Standby Edge to become active and results in an Active-Active (Split-Brain) state. With this feature, the user can increase the time threshold before the Active Edge is declared down and trigger a failover and prevent a potential split-brain state.

The value is changed under the **Advanced Options** section where a user configures the **HA Failover Detection Time Multiplier**. This multiplier is a number that is multiplied by 100 milliseconds (ms). The default value is 7 (700 ms) and be configured up to 70 (7000 ms).

The screenshot shows the configuration page for High Availability. At the top, there is a dropdown menu for 'High Availability' set to 'None'. Below this, there is a section for 'HA: None' with a 'Segment Agnostic' button. A note states: 'High Availability is enabled at the Edge level. When using Active/Standby Pair HA, enable HA prior to connecting the Standby SD-WAN Edge. To learn more, please consult our HA documentation'. Under 'Select Type', there are four radio button options: 'None', 'Active Standby Pair' (which is selected), 'Cluster', and 'VRRP with 3rd party router'. Below this is an 'HA Interface' dropdown menu. There are two checkboxes: 'Deploy with Unique LAN MAC Address' (unchecked) and 'Enable Graceful Switchover (require Graceful Restart in routing protocol)' (unchecked). A red box highlights the 'Advanced Settings' section, which contains the 'HA Failover Detection Time Multiplier' set to 7. At the bottom right, there are buttons for 'DISCARD CHANGES' and 'SAVE CHANGES'.

Monitoring High Availability Edges

The Orchestrator includes special monitoring for a site deployed with a High Availability topology which are outlined in this section.

Overview

Beginning in Release 5.2.0, the Orchestrator includes improved monitoring for sites deployed in a High Availability topology:

- The **Monitor > Edge > Overview** tab now includes two HA specific improvements:
 - WAN link information includes which HA Edge the link is associated with by serial number, which is especially important with Enhanced HA deployments.
 - HA Interface Status.
- HA specific failover bars on the **Monitor > Edge > System** tab denoting where an HA Edge site failed over.
- A new **Monitor > Edge > HA Standby** tab for Standby Edge monitoring information.

High Availability specific information on the Monitor > Edge > Overview tab

When you navigate to the **Monitor > Edge > Overview** tab, a site deployed in High Availability has added information specific to an HA deployment:

The screenshot shows the VMware Orchestrator interface. The top navigation bar includes 'vmw Orchestrator', 'Customer 5-site', and 'SD-WAN'. The main navigation menu on the left includes 'Monitor', 'Configure', 'Diagnostics', and 'Service Settings'. The current view is 'Edges / B2 - 610', which is 'Offline'. The 'Overview' tab is selected, showing a 'Links Status' table and an 'HA Interface Status' section.

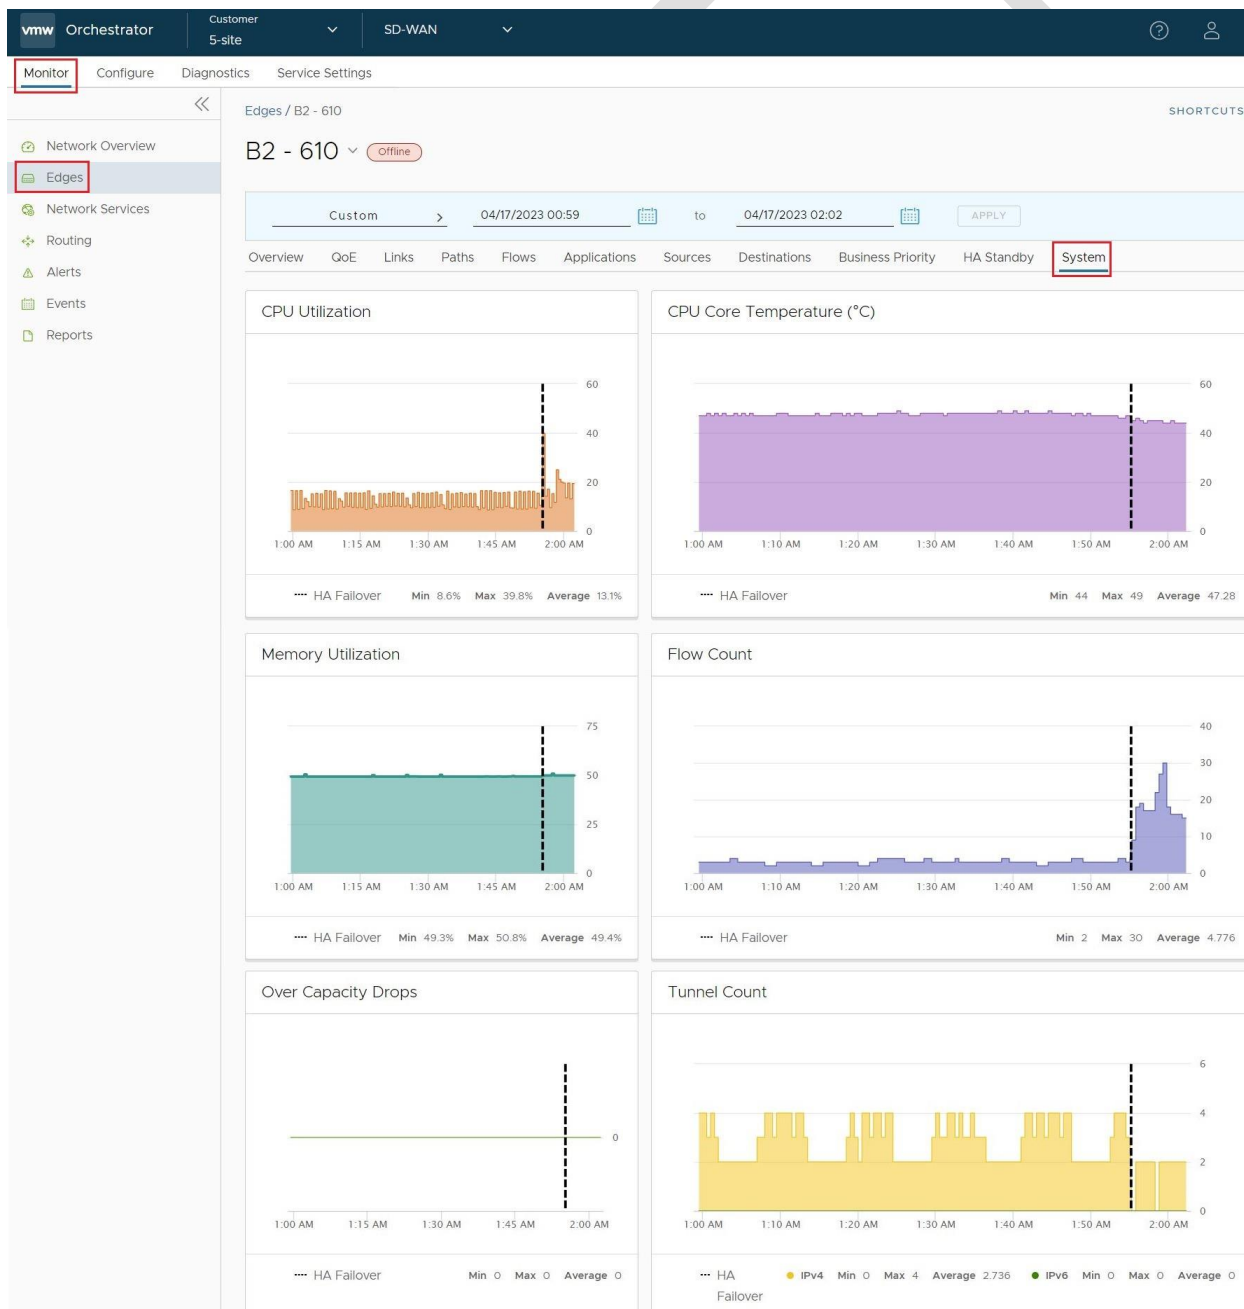
Links	Auto Dual-Mode SIM	Device Serial No (Device State)	Link Status	Interface (WAN Type)	Throughput Bandwidth	Pre-Notifications	Alerts	Sign
169.254.6.42 169.254.6.42	N/A	GYKJV43 (ACTIVE)	● Stable	GE3 (Ethernet)	0 bps ↑ 254.919 Mbps 0 bps ↓ 257.777 Mbps	✔ Edit	✔	N/A

Name	Interface Status	Interface (HA Type)
169.254.2.1 169.254.2.1	● Up	GE1 (standard)

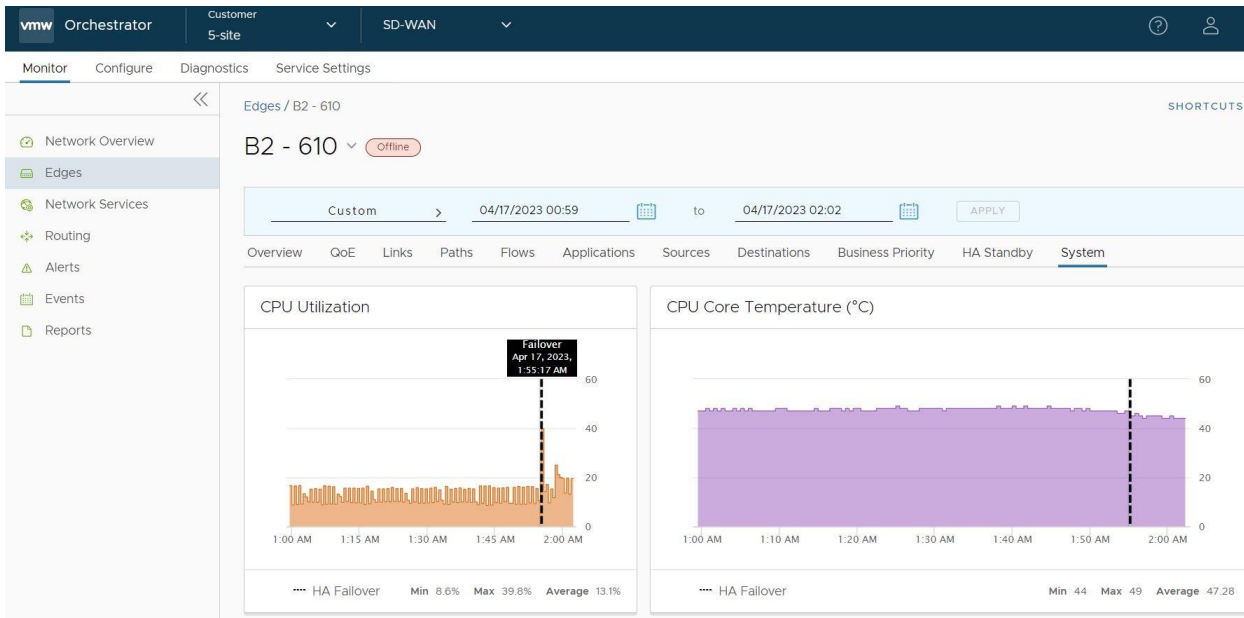
- In the **Link Status** section, each connected WAN link includes a column **Device Serial No (Device State)** that includes the Edge serial number associated with that WAN link and the HA status of that Edge (Active or Standby). This information is valuable in Enhanced HA deployments where WAN links are uniquely associated to different HA Edges and allows you to see the status of WAN links on the Standby Edge just as you see them on the Active Edge.
- The **Overview** tab adds an **HA Interface Status** section which includes the IP Address, Interface Status (Up or Down), and Interface (HA Type).

High Availability Failover Bars on the System tab

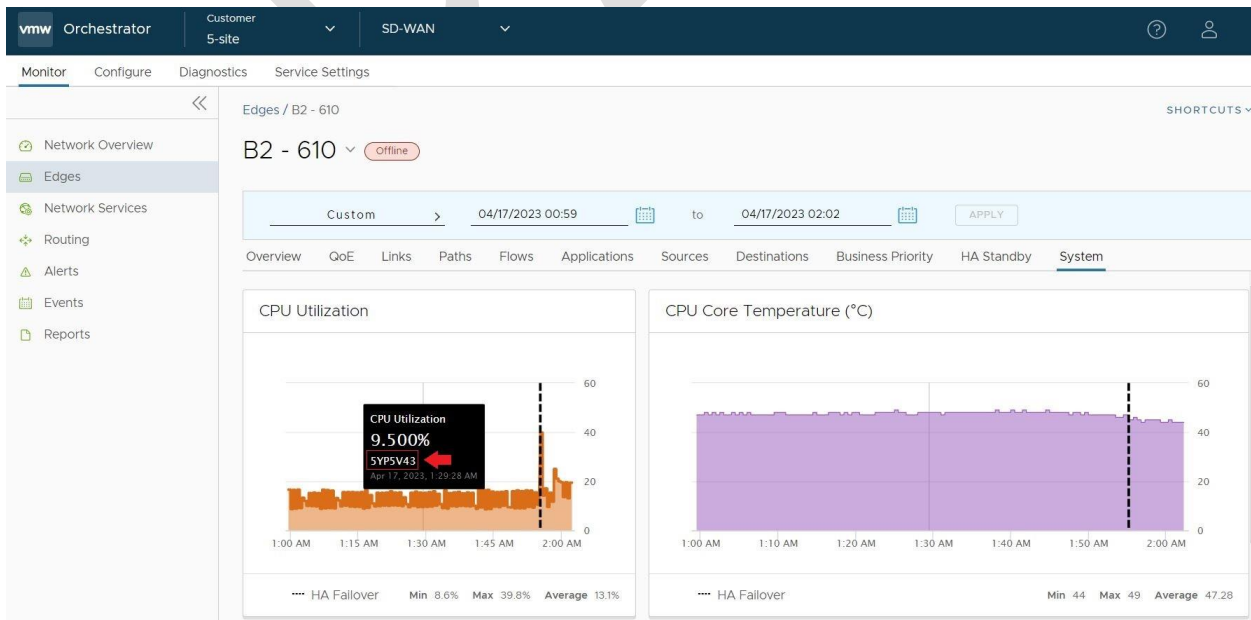
When you navigate to the **Monitor > Edge > System** tab, a site deployed in High Availability has additional functionality which is best seen when a HA failover has occurred. When an HA failover occurs, the Orchestrator renders a vertical bar marking the point of the failover.



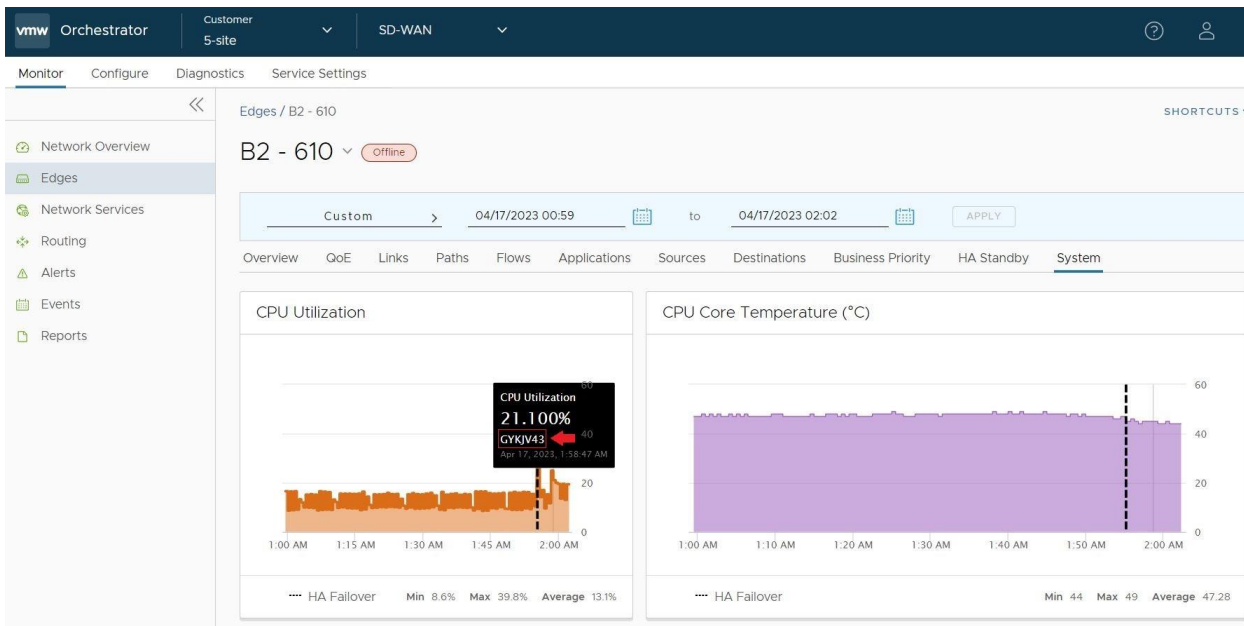
Focusing on the CPU Utilization graph, if a user hovers their mouse on the bar, the bar indicates when the failover occurred. This failover bar represents a boundary marking how the System statistics apply to each respective HA Edge when they serve as the Active Edge.



If you click on the graph to the left (earlier) side of the failover bar, the Orchestrator will indicate which Edge was Active at the time of those statistics. In this example, the Active Edge prior to the failover was Edge serial # SYP5V43.

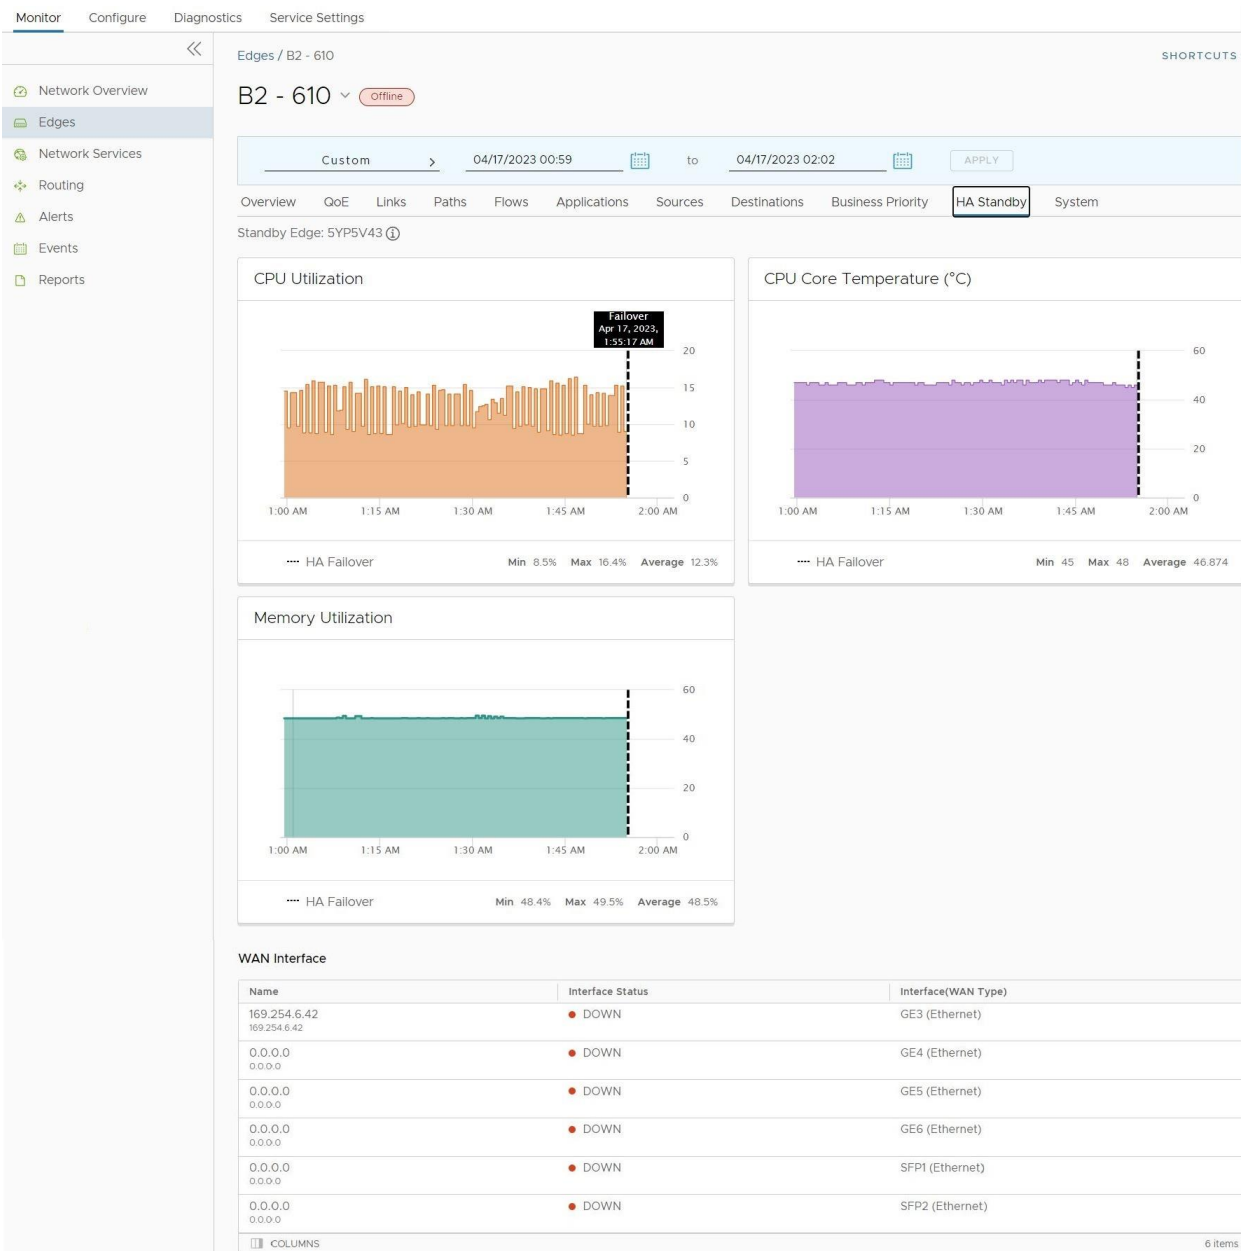


If you click on the graph to the right (later) side of the failover bar, the Orchestrator indicates that the Active Edge for the post-failover statistics is Edge serial # GYKJV43. In this way you can always know which Edge applies to which System values.



The Monitor > Edge > HA Standby Tab

You can now view the System Health statistics for the HA Edge when it is serving a standby role by clicking on the **Monitor > Edge > HA Standby** tab. This page also includes a failover bar indicating when an HA site has triggered a failover and there is now a different Edge in the role of Standby.



The **HA Standby** tab also provides live **WAN Interface** information which is especially helpful in Enhanced HA deployments where the Standby Edge is using unique WAN links.

Click on the information button to get Standby Edge information at the same level as what a user sees on the **Monitor > Edge > Overview** tab for the Active Edge.

- Network Overview
- Edges
- Network Services
- Routing
- Alerts
- Events
- Reports

Edges / B2 - 610

B2 - 610 Offline

Custom > 04/17/2023 00:59 to

Overview QoE Links Paths Flows Applications Sources

Standby Edge: 5YP5V43 i ←

CPU Utilization

1:00 AM 1:15 AM

... HA Failover

Memory Utilization

Description

Location
US

Edge Info
 Activation: Activation pending
 Act. Key: 7XBV-JS8C-Z9D5-Q93J
 Last Contact: Apr 17, 2023, 2:05:30 AM
 System Up Since: Apr 11, 2023, 1:04:09 AM
 Service Up Since: Apr 14, 2023, 8:48:16 AM
 Operator Alerts: Enabled
 Authentication: Certificate Deactivated

High Availability
 HA Status: Standby ready
 Active Device: GYKJV43
 Standby Device: 5YP5V43

Device Hardware
 Model: 5YP5V43
 Serial Number: EDGE610LTE-RW
 LTE Region:

Device Software
 Software Version:
 Factory Version:

Device Firmware
 Platform Version: **NA** [**NA** (BIOS_3.43.0-9-13_CPLD_0x33_PIC_v20N), HASupported Upgradable]
 Modem Version: **NA** [**NA**]

Configuration Profile
 Profile: Quick Start Profile

HA Event Details

This section describes HA events. **Note**, Release 5.2.0 adds the bottom 4 events.

HA Event	Description
HA_GOING_ACTIVE	A standby SD-WAN Edge is taking over as Active because it has not heard a heartbeat from the peer.
HA_STANDBY_ACTIVATED	When a new Standby is detected by the Active, the Active tries to activate the Edge by sending this event to the SD-WAN Orchestrator. On a successful response, the Active will sync the configurations and sync data.
HA_FAILED	Typically happens after the HA pair has formed and the Active SD-WAN Edge no longer hears from the Standby SD-WAN Edge. For example, if the Standby SD-WAN Edge reboots, you will receive this message.
HA_READY	Means the Active SD-WAN Edge now hears from the Standby SD-WAN Edge. Once the Standby SD-WAN Edge comes back up and reestablishes the heartbeat, then you will receive this message.
HA_TERMINATED	When the HA configuration is deactivated, and it is successfully applied on the Edges, this Event is generated.
HA_ACTIVATION_FAILURE	If the SD-WAN Orchestrator is unable to verify the HA activation, it will generate this Event. Examples include: <ul style="list-style-type: none"> ■ the SD-WAN Orchestrator is unable to generate a certificate ■ the HA has been deactivated (rare)
VCO_IDENTIFIED_HA_FAILOVER	Event message reads: Edge HA Failover Detected The SD-WAN Orchestrator has detected that a High Availability failover has occurred on the Edge.
VCO_IDENTIFIED_HA_FAILURE	Event message reads: Edge HA Failure Detected The SD-WAN Orchestrator has detected that the Standby Edge has gone down. This event will include the serial number of the Edge.
HA_UPDATE_FAILOVER_TIME	Event message reads: Updating HA Failover time from ####ms to ####ms A user changed the failover time for when an HA Edge will failover based on how long the Edge will wait to receive a heartbeat from the Active Edge. Increasing this value can prevent an Active-Active "Split Brain" state for HA Edges under high load. This is done through the HA Failover Detection Time Multiplier located at Configure > Edge > Device > High Availability on the Orchestrator.
HA_RESET_FAILOVER_TIME	Event message reads: Updating HA Failover time from ####ms to ####ms When an HA Edge's system has been stable for 60 seconds, the process reduces the failover threshold time by 50%.

Enterprise-Level Orchestrator Alerts and Events

Describes a summary of alerts and events generated within the VMware SD-WAN Orchestrator at the Enterprise level.

The document provides details about all Enterprise-level Orchestrator events. Although these events are stored within the SD-WAN Orchestrator and displayed on the Orchestrator UI, most of them are generated by either an SD-WAN Edge or an SD-WAN Gateway and/or one of its running components (MGD, EDGED, PROCMON, and so on) with the exception of a few which are generated by the Orchestrator itself. You can configure notifications/alerts for events in Orchestrator only.

The following table provides an explanation for each of the columns in the "Enterprise-level Orchestrator Events" table:

Column name	Details
EVENT	Unique name of the event
DISPLAYED ON ORCHESTRATOR UI AS	Specifies how the event is displayed on the Orchestrator.
SEVERITY	The severity with which this event is usually generated.
GENERATED BY	The VMware SD-WAN component generating the notification can be one of the following: <ul style="list-style-type: none">■ SD-WAN Orchestrator■ SD-WAN Edge (MGD)■ SD-WAN Edge (EDGED)■ SD-WAN Edge (PROCMON)
GENERATED WHEN	Technical reason(s) and circumstances under which this event is generated.
RELEASE ADDED IN	The release this event was first added. If not specified, this event existed prior to release 2.5.
DEPRECATED	Specifies if the event is deprecated from a specific release.

Enterprise-level Orchestrator Events (**Note:** The Final Four Events are New for 5.2.0)

EVENT	DISPLAYED ON ORCHESTRATOR UI AS	SEVERITY	GENERATED BY	GENERATED WHEN	RELEASE ADDED IN	DEPRECATED
HA_STANDBY_ACTIVATED	HA Standby Activated	INFO	SD-WAN Edge (EDGED)	When active Edge detects standby peer send this event to SD-WAN Orchestrator to activate standby Edge.		
HA_INTF_STATE_CHANGED	HA Interface State Changed	ALERT	SD-WAN Edge (EDGED)	HA interface went down/up.		
HA_GOING_ACTIVE	High Availability Going Active	INFO	SD-WAN Edge (EDGED)	Standby Edge transition to Active Edge after detecting no heartbeat for more than 700ms.		
HA_FAILED	High Availability Peer State Unknown	INFO	SD-WAN Edge (EDGED)	Active Edge detects no heartbeat or activity from standby Edge for more than 700 milliseconds.		
HA_READY	High Availability Ready	INFO	SD-WAN Edge (EDGED)	Active Edge detects activated standby peer.		
VCO_IDENTIFIED_HA_FAILOVER	Edge HA Failover Detected	ALERT	SD-WAN Edge (EDGED)	Orchestrator has detected that a High Availability failover has occurred on the Edge.		
VCO_IDENTIFIED_HA_FAILURE	Edge HA Failure Detected	ALERT	SD-WAN Edge (EDGED)	Orchestrator has detected that the Standby Edge has gone down.		

HA_UPDAT E_FAILOVE R_TIME	Updating HA Failover time from ####ms to ####ms	INFO	SD-WAN Edge (EDGED)	User changed the failover time for when an HA Edge will failover due to a lack of heartbeat response. This time is measured in milliseconds (ms).		
HA_RESET_ FAILOVER_ TIME	Failover time reset from ####ms to ####ms.	INFO	SD-WAN Edge (EDGED)	When an HA Edge's system has been stable for 60 seconds, the process reduces the failover time by 50%.		

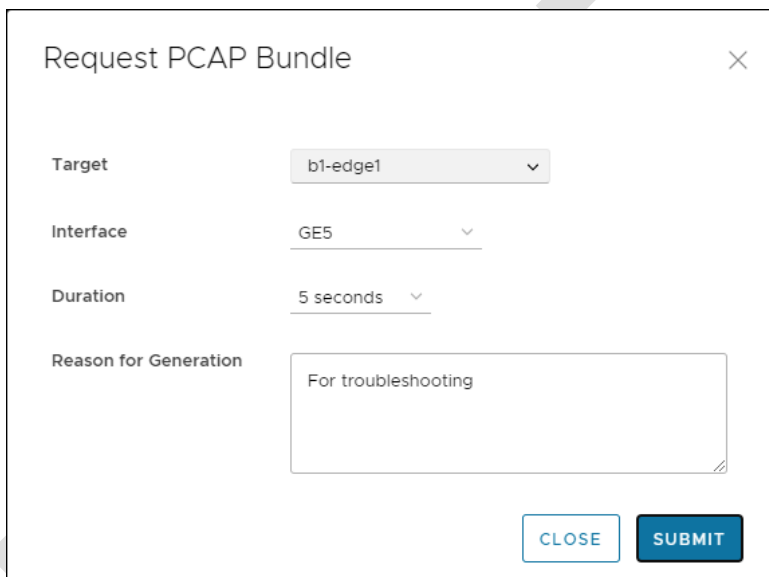
DRAFT

Request Packet Capture Bundle with New Orchestrator UI

The Packet Capture bundle collects the packets data of a network. These files are used in analyzing the network characteristics. You can use the data for debugging an Edge device.

To generate a PCAP bundle using the new UI:

- 1 In the Enterprise portal, click the **Diagnostics** tab.
- 2 Click **Diagnostic Bundles > Request PCAP Bundle**.
- 3 In the **Request PCAP Bundle** window that appears, configure the following:



The screenshot shows a window titled "Request PCAP Bundle" with a close button (X) in the top right corner. The window contains the following fields:

- Target:** A dropdown menu with "b1-edge1" selected.
- Interface:** A dropdown menu with "GE5" selected.
- Duration:** A dropdown menu with "5 seconds" selected.
- Reason for Generation:** A text input field containing "For troubleshooting".

At the bottom right of the window, there are two buttons: "CLOSE" (light blue) and "SUBMIT" (dark blue).

Table 6-1.

Option	Description
Target	Choose the target Edge from the drop-down list. The packets are collected from the selected Edge.
Interface	Choose an Interface or a VLAN from the drop-down list. The packets are collected on the selected Interface.
Duration	Choose the time in seconds. The packets are collected for the selected duration.
Reason for Generation	Optionally, you can enter your reason for generating the bundle.

The window displays the details of the bundle being generated, along with the status.

Packet Capture for Edges configured for High Availability

In Release 5.2.0 and later, a user can request a packet capture for the Standby Edge's HA interface, the interface that connects the Standby Edge to the Active Edge. This option appears at the bottom of the menu and reads: **Standby Edge**, and then lists the HA interface.

