



Monitoring VMware HCX with VMware Aria Operations for Networks

Table of contents

Introduction	3
VMware Aria Operations for Networks support for VMware HCX	4
Built-in dashboards and alerts	4
Useful searches for network extension flows	4
Get flows traversing the network extension tunnels	4
Get flow metrics per network extension	4
Get flow metrics per network extension appliance pair	5
Get top talking VM pairs over network extension tunnels	5
Get top talking L2 networks over network extension tunnels	5
Get tromboning flows over network extension tunnels	6
Get flows to the internet over network extension tunnels	7
VMware HCX appliance and tunnel metrics	7
Configuring threshold alerts for your VMware HCX environment	7
Appliance VM bandwidth usage exceeds threshold	8
Appliance VM CPU usage exceeds threshold	8
Appliance VM memory usage exceeds threshold	8
Average TCP RTT for VMware HCX L2E flows exceeds threshold	8
VMware HCX tunnel packet drops exceed threshold	9
Summary	10
About the authors	10
Acknowledgements	10
Resources	10

Introduction

VMware HCX® is the application mobility platform by VMware. It is heavily used to migrate workloads from legacy VMware vSphere® or VMware NSX® environments to modern VMware NSX-T™ or VMware Cloud™ environments. VMware HCX also supports workload migration from on-premises to hyperscalers, such as Azure VMware Solution, Google Cloud VMware Engine, and Oracle Cloud VMware Solution, and is also the leading choice for cloud migration.

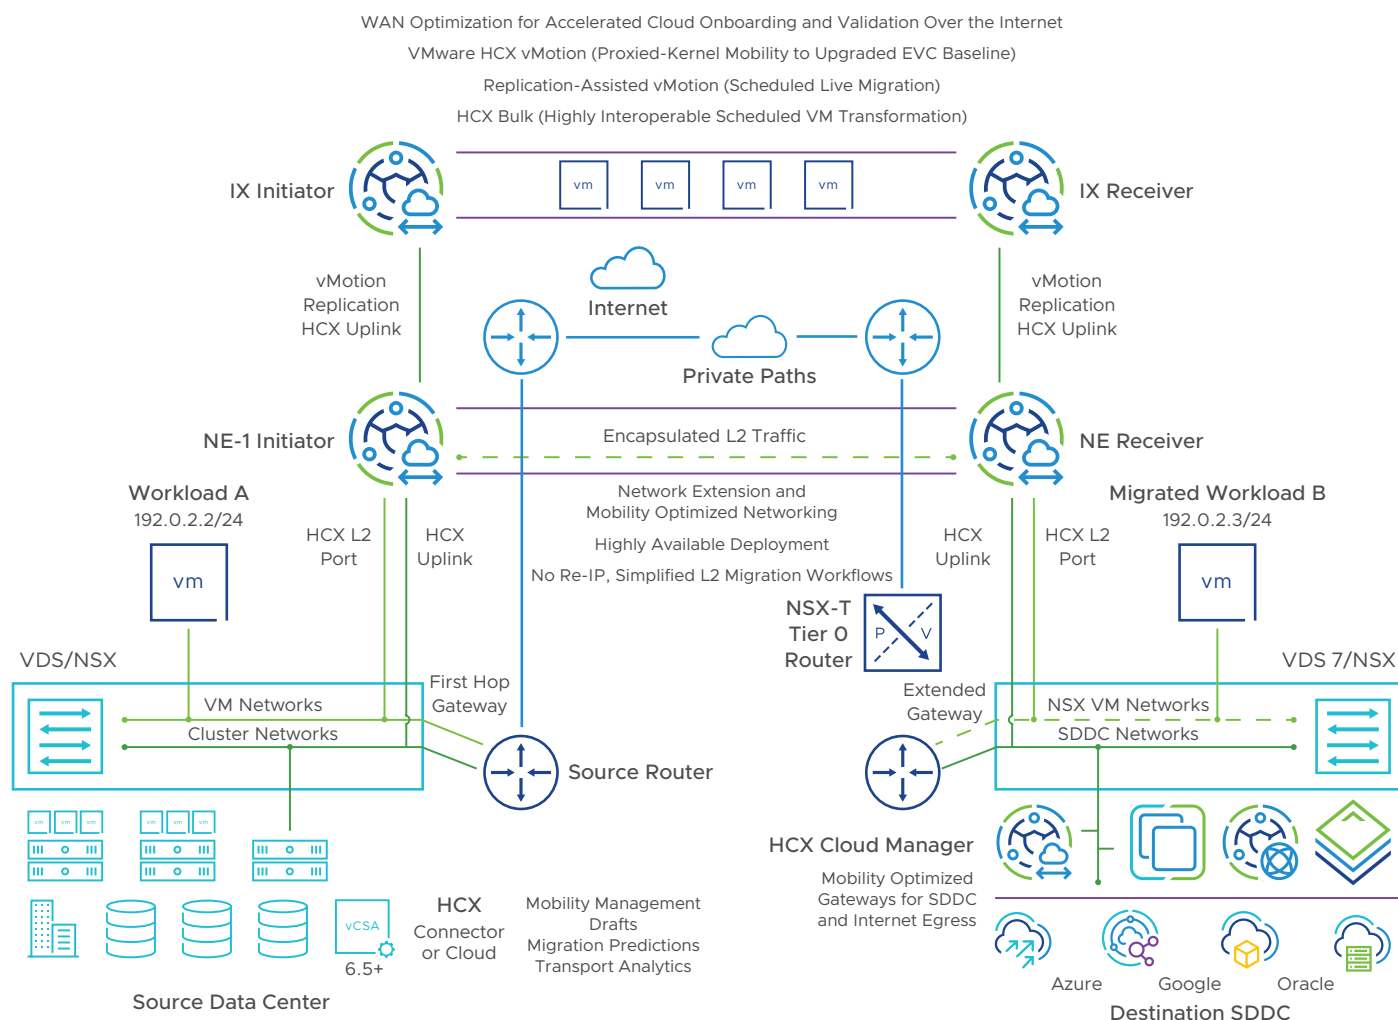


Figure 1: Example of workload migration with VMware HCX.

VMware HCX deploys its interconnect service as part of its service mesh configuration. The interconnect service uses interconnect (IX) appliance peers between local and remote sites. A transport tunnel is established between these appliances. These tunnels are used with multiple migration services to move workloads between paired sites enabled by VMware HCX.

One migration approach is to re-IP the workloads as they reach the remote site. This introduces risk and additional dependencies in an already complex workload migration process.

To simplify this, VMware HCX has another service known as network extension (NE). When the NE service is enabled on a service mesh, VMware HCX will deploy NE peer appliances between the local and remote data centers, and establish a transport tunnel between them.

VMware HCX network extensions allow users to extend their on-premises Layer 2 networks (VLANs/VXLANs/GENEVE) as NSX-T segments in their remote sites. When a network is extended using VMware HCX network extensions, a new network with the same subnet is created on the remote site as a disabled standby gateway address, leveraging the on-premises gateway until the customer is ready to migrate the full subnet, or enables mobility optimized networking.

This allows any virtual machines (VMs) from the extended network, which are migrated to the remote site, to retain their IP addresses and keep communication with other non-migrated workloads on-premises intact.

Workload migration should be done carefully by analyzing the impact of migration on the ingress/egress traffic of the destination site gateway. Metrics such as bytes, packets, bytes per second, and packets per second across different dimensions (such as intra-VLAN, inter-VLAN, VLAN to internet, and internet to VLAN) need to be analyzed before deciding on the candidate VMs for a migration wave. The VMware Fling for migration planning helps in getting all these metrics. See the [VMware Aria Operations™ for Networks Toolkit](#) for more information.

VMware Aria Operations for Networks support for VMware HCX

VMware Aria Operations for Networks (formerly VMware vRealize® Network Insight™) supports the addition of VMware HCX as a data source. It collects configuration information about service meshes, appliances, Layer 2 extensions, and the like. VMware Aria Operations for Networks also collects metrics, such as receive/transmit (Rx/Tx) packets, bytes, and packet drops, for the tunnels between appliances.

VMware Aria Operations for Networks receives IPFIX flow data from the vSphere Distributed Switch™ (VDS) in the case of pure VMware vCenter Server® environments and from the NSX Distributed Firewall™ in the case of modern software-defined data centers (SDDCs) based on NSX Data Center for vSphere or NSX-T. These IPFIX flows are enriched with information about the VMs, the host, the data center, security groups, and the like.

VMware Aria Operations for Networks identifies flows traversing the VMware HCX network extension transport tunnels and enriches them with the information about the VMware HCX L2 extension they are part of. This helps you understand the network traffic in your VMware HCX environment deeply.

You need to add vCenter Server, NSX Manager™, and VMware HCX Manager to VMware Aria Operations for Networks for end-to-end visibility of your VMware HCX multisite topology.

Built-in dashboards and alerts

After adding the VMware HCX Managers of all your sites as data sources to VMware Aria Operations for Networks, the initial dashboards provide an immediate overview of the VMware HCX environments. The starting point to check this can be the VMware HCX Manager dashboard. Read the [Viewing VMware HCX Manager Details](#) documentation to find out more about the VMware HCX dashboards built in to VMware Aria Operations for Networks.

VMware Aria Operations for Networks also raises alerts if the health statuses of the VMware HCX services, appliances or tunnels change. The [Alerts List](#) provides a full list of supported alerts; find the keyword “HCX” to see alerts related to VMware HCX.

Useful searches for network extension flows

VMware Aria Operations for Networks has a powerful search query mechanism to extract information from its vast collection of configuration, inventory, metrics and network flow data. See the [Working with Search Queries](#) documentation to learn more about the search capabilities of VMware Aria Operations for Networks. The right search queries can give you interesting insights into your environment.

The following search queries can be used to get insights into the flow traffic over VMware HCX network extension tunnels.

Get flows traversing the network extension tunnels

flows where HCX L2 Extension is set

This query will give a list of all flows going over the network extension tunnels.

Get flow metrics per network extension

series(sum(bytes)), series (avg (byte rate)), series(sum(packets)) of flows where HCX L2 Extension is set group by HCX L2 Extension

Get flow metrics per network extension appliance pair

series(sum(bytes)), series (avg (byte rate)), series(sum(packets)), series(avg(Average TCP RTT)) of flows where HCX L2 Extension is set group by HCX L2 Extensions.Source Appliance, HCX L2 Extensions.Remote Appliance

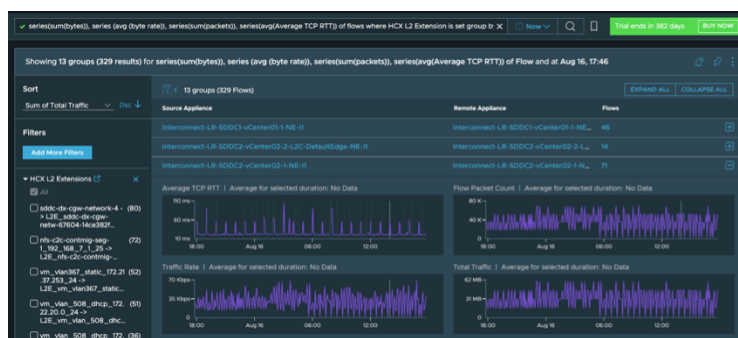


Figure 2: Example of flow metrics per network extension appliance pair.

This shows you the load on each NE appliance pair. It can be used to assess whether you need to begin using a new appliance pair, or plan to move any of the L2 extensions to a dedicated appliance pair to avoid performance bottlenecks.

Identify which network extension is taking the most bandwidth of any appliance pair

series(sum(bytes)), series (avg (byte rate)), series(sum(packets)), series(avg(Average TCP RTT)) of flows where HCX L2 Extension is set and HCX L2 Extensions.Source Appliance = 'Interconnect-LrSDDC1-LrSDDC2-1-NE-I' and HCX L2 Extensions.Remote Appliance = 'Interconnect-LrSDDC1-LrSDDC2-1-NE-R' group by HCX L2 Extensions



Figure 3: Example of network extensions taking the most bandwidth in an appliance pair.

Get top talking VM pairs over network extension tunnels

series(sum(bytes)), series (avg (byte rate)), series(sum(packets)), series(avg(Average TCP RTT)) of flows where HCX L2 Extension is set group by Source VM, Destination VM

This can be filtered on different VMware HCX L2 extensions by adding additional filters, such as *HCX L2 Extension = <>*.

Get top talking L2 networks over network extension tunnels

series(sum(bytes)), series(avg(byte rate)), series(sum(packets)), series(avg(Average TCP RTT)) of flows where HCX L2 Extension is set group by Source L2 Network, Destination L2 Network

Get tromboning flows over network extension tunnels

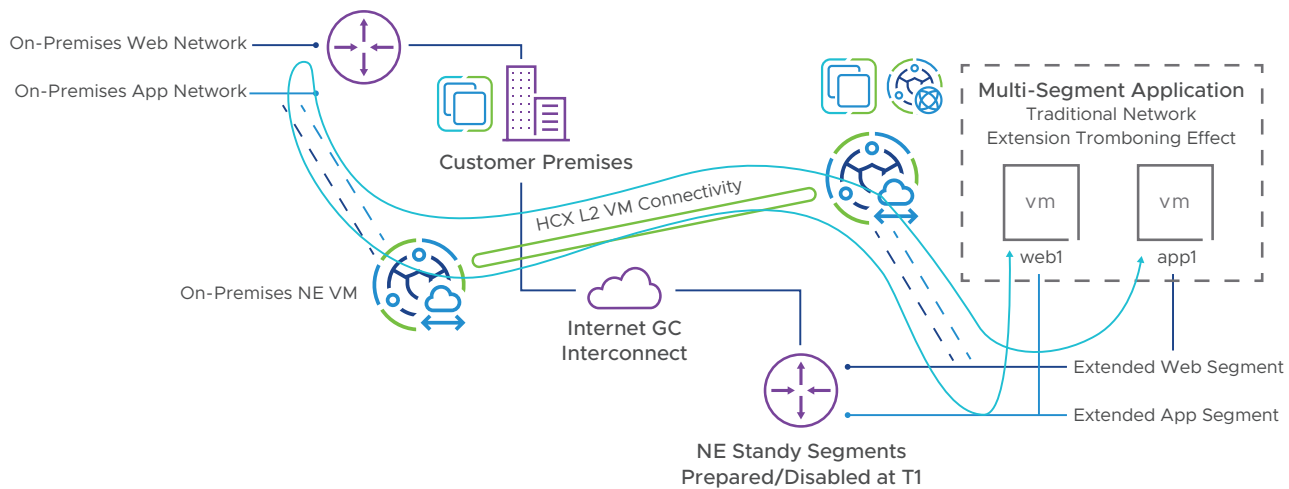


Figure 4: Tromboning flow over NE tunnels.

When VMs on two different extended networks present on the cloud data center want to communicate with each other, they do it via the on-premises data center as the configured gateways for both the extended networks are on the on-premises data center.

This hairpinning introduces additional latency in communication between the VMs even though those VMs are present in the same data center. This can adversely affect the performance of the application.

The following query will give the list of tromboning flows:

flows where HCX L2 Extension is set and flow type = 'Same DC'

This works because the VMware HCX L2 extension is set only if the flow is going over the tunnel. So, if the flow is going over the tunnel and the endpoint VMs are in the same data center, then we can say that it is tromboning.

You will find two VMware HCX L2 extensions enriched in these kinds of flows. Figuring out these flows is important to make decisions about enabling mobility optimized networking (MON) on the participating VMware HCX L2 extensions so that the flow does not hairpin and the VMs on the cloud site can talk with each other via the cloud gateway.

To evaluate if you need to enable MON on the participating VMware HCX L2 extensions, you need to understand how much traffic is tromboning. Use the following query to see how much traffic is going over the tunnel per extension and decide based on the traffic whether MON is required or not.

Tromboning traffic per VMware HCX L2 extension

series(sum(bytes)) of flows where HCX L2 Extension is set and flow type = 'Same DC' group by HCX L2 Extension

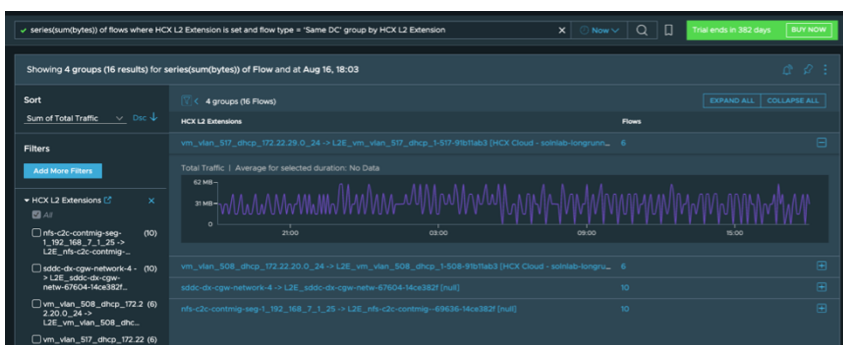


Figure 5: Example of tromboning traffic per L2 extension.

VMware HCX also allows you to configure MON at an individual VM level instead of at the L2 extension level. To find out per VM tromboning traffic, use the following query.

Tromboning traffic per VM pair

series(sum(bytes)) of flows where HCX L2 Extension is set and flow type = 'Same DC' group by Source VM, Destination VM

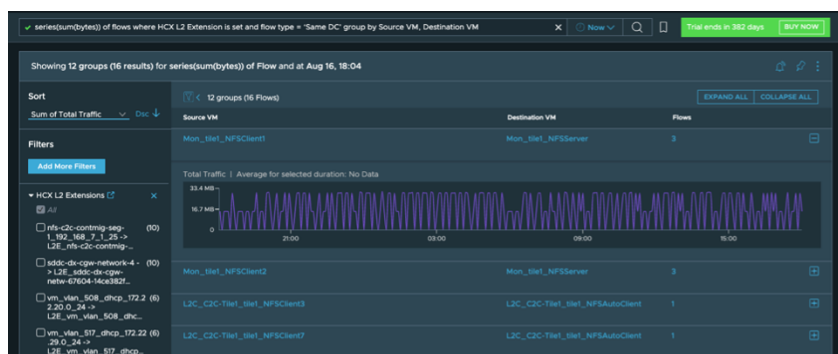


Figure 6: Example of tromboning traffic per VM pair.

Get flows to the internet over network extension tunnels

flows where HCX L2 Extension is set and flow type = 'Destination Is Internet'

If any of the migrated workloads are talking to the internet, then the flow will traverse over the network extension tunnel to the on-premises data center and egress to the internet via the on-premises gateway. This can be avoided by enabling MON on the VMware HCX L2 extension so the VMs can talk to the internet via the cloud gateway.

VMware HCX appliance and tunnel metrics

VMware Aria Operations for Networks collects various metrics related to the VMware HCX appliance VMs and the tunnels that exist between the IX and NE appliance pairs. Monitoring the health of these two entities is important to ensure the smooth functioning of your VMware HCX workload migrations and network extensions.

If the bandwidth, CPU or memory usage of the VMware HCX IX/NE appliances reaches the maximum capacity of the appliance, it can lead to slowness in migration, and severely impact the performance of the workloads that have flows over the NE tunnels, which in turn leads to a bad user experience for the end-user of the application.

To keep a tab on the appliance and tunnel health, configure some thresholds in VMware Aria Operations for Networks that will alert you if these metrics exceed the configured thresholds.

Configuring threshold alerts for your VMware HCX environment

Set up the following alerts under **Analytics > Thresholds** for monitoring different metrics for VMware HCX entities. Read the [Configure Thresholds and Alerts](#) documentation for more details on how threshold alerts can be configured in VMware Aria Operations for Networks

Appliance VM bandwidth usage exceeds threshold

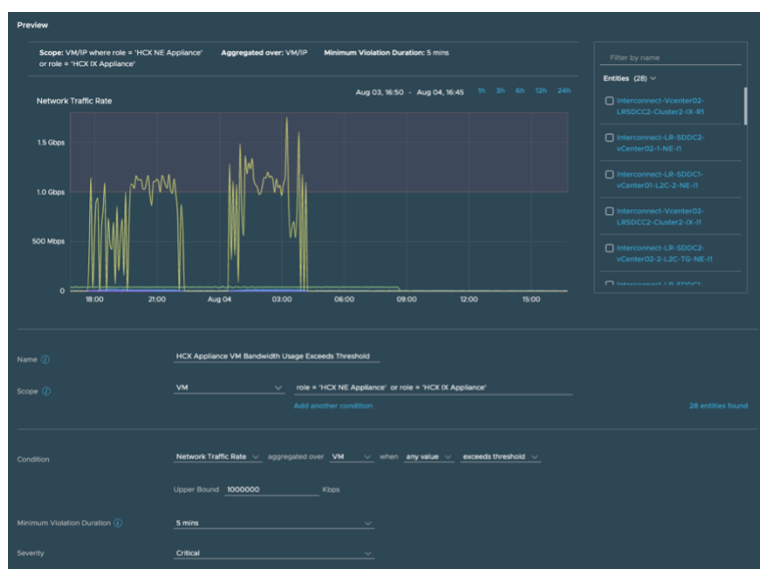


Figure 7: Example of configuring threshold alerts in VMware Aria Operations for Networks.

First, give a unique meaningful name to our threshold configuration. This name is shown in the alert when it is raised.

Select **VM** as the Scope and add **role = 'HCX NE Appliance' or role = 'HCX IX Appliance'** as a condition on the role of the VM.

Now, select **Network Traffic Rate** as the metric to be used in Condition.

After this, you should see a list of VMs in the Entities section on the top right and its Network Traffic Rate time series on the left.

Now, set the Upper Bound for the Network Traffic Rate around 4,000,000 Kbps as VMware HCX appliances support 4–6Gbps. This ensures you will get alerted before the bandwidth usage climbs to its peak capacity.

You can choose to set up other forms of notification, such as email or push notification.

Appliance VM CPU usage exceeds threshold

The scope in this event will be the same as Figure 7, but the metric will change. Select **CPU Usage Rate** as the Condition and set the Upper Bound to around 90 percent.

Appliance VM memory usage exceeds threshold

The scope in this event will be the same as Figure 7, but the metric will change. Select **Memory Usage Rate** as the Condition and set the Upper Bound to around 90 percent.

More such alerts can be set on other metrics of the appliance VMs using the same scope with different metrics selected.

Average TCP RTT for VMware HCX L2E flows exceeds threshold

VMware Aria Operations for Networks collects the IPFIX flows exported by NSX-T. These flows have the **Average TCP RTT** metric, which indicates the average round-trip time (RTT) required for a particular flow. Usually, the RTT values will be higher for cases where packets traverse over the VMware HCX network extension tunnels. To keep track of such flows and their RTT metrics, set a threshold alert on it as follows:

- Select **Flow** as the Scope, and add **HCX L2 Extension is set** as a condition of the scope.
- Add conditions to narrow down the scope to a particular application or set of VMs.
- Select **Average TCP RTT** as the Condition, aggregated over **Entire Scope** when **the average** for any **30 mins** interval **exceeds threshold**.

- Set the Upper Bound as per your TCP RTT requirements. (Note that this value is in microseconds.)

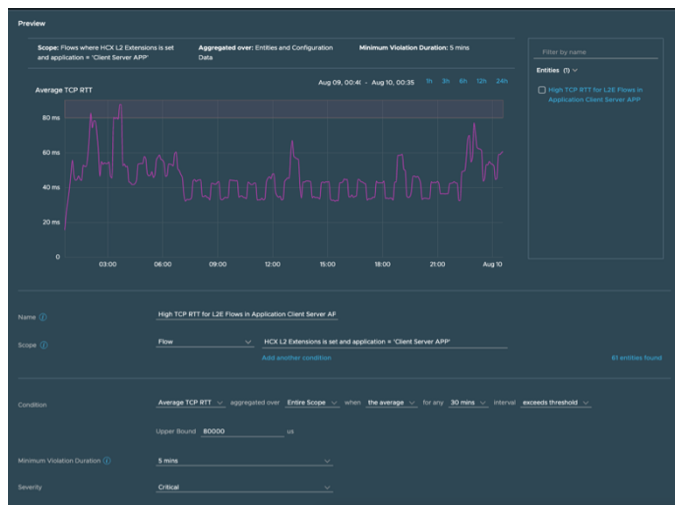


Figure 8: Example of a threshold alert for average TCP RTT.

VMware HCX tunnel packet drops exceed threshold

VMware Aria Operations for Networks collects various metrics for the IPIP_FOU tunnels created between VMware HCX IX/NE appliance pairs. These tunnels are used for migrating workloads or for carrying VM-VM traffic over network extensions. To monitor the health of these tunnels, set a threshold alert on these metrics:

- Select **HCX Tunnel** as the Scope, and add **Source Appliance = <> and Remote Appliance = <>** as a condition of the scope. This will allow you to set the threshold for all tunnels between two appliances.
- Select **Dropped Rx Packets** as the Condition, aggregated over **HCX Tunnel** when **the total** for any **30 mins** interval **exceeds threshold**.
- Set the Upper Bound as per your TCP RTT requirements.



Figure 9: Example of a threshold alert for packet drops.

Summary

Monitoring the hybrid cloud environment is important to make sure the cloud migration process goes smoothly. VMware Aria Operations for Networks gives you the desired visibility and alerting required to successfully navigate your multi-cloud journey.

About the authors

Ambarish Pande ([@ambarishpande](#)) is a member of technical staff for VMware Aria Operations for Networks in the Modern Applications and Management Business Group at VMware. He works on VMware Aria Operations for Networks integrations with VMware HCX, VMware Cloud on AWS, VMware Cloud on Dell, and hyperscalers such as Azure VMware Solution, Google Cloud VMware Engine, and Oracle Cloud VMware Solution.

Gabe Rosas is a staff technical product manager for VMware HCX in the Networking and Advanced Security Business Group at VMware. He is experienced in data center networking and server virtualization, and has written technical papers about data center migration and multi-cloud technologies, such as VMware HCX. He is an active vExpert and VMware Communities (VMTN) Community Manager. You can find him on Twitter [@gabe_rosas](#) or follow his blog at [hcx.design](#).

Acknowledgements

The author would like to thank Martijn Smit ([@smitmartijn](#)) and Matt Just ([@h1_tek_matt](#)) for their guidance and valuable feedback on this paper.

Resources

For more information, visit the [VMware Aria Operations for Networks Search Exchange](#).

Try [VMware Aria Operations for Networks](#) with a free trial.

