

VirtualWisdom User Guide 6.7

[Download PDF version](#)

Table of Contents

- VirtualWisdom Platform Architecture 6
- VirtualWisdom User Interface and Navigation 10
 - User Settings 12
- Notifications 14
- Data and Entities 19
 - Entity Overview 20
 - Entity Types by Category 22
 - ITL Examples 47
 - Entity Creation 51
 - Entity Management 53
 - Auto-Defined Entity Management 53
 - User-Defined Entity Management 56
 - Entity Management Best Practices 75
 - Applications and Tiering 76
 - Application Entities 76
 - Application Tiering 90
 - Create Tiers 91
- VirtualWisdom Metrics 94
 - Metric Data Aggregation 95
 - Data Persistence 97
- Infrastructure Monitoring Concepts 100
 - Health, Utilization, Capacity, and Performance 100
 - Using a Dashboard to Identify Application Infrastructure Issues 102
- Inventory 122
 - Inventory List View 123
 - Application Inventory Page 125
 - Tier Inventory View 132
 - Other Entity Types Inventory View 134
- Topology 138
 - Selecting a Topology View and Entity 143
 - Saving Your Topology View 150
 - Excluding (Filtering) Nodes 150
 - Entity Grouping 153
 - Using the Topology Map Controls 156
 - Understanding Entity Status 157
 - Viewing Entity Data 158
 - Topology Use Case 161
- Reports 163
 - Report Page 165
 - Report Templates 165
 - Chart Types 167
 - Basic Chart Types 169
 - Top N Charts 175
 - SLA Status Charts 180
 - Special Charts 183

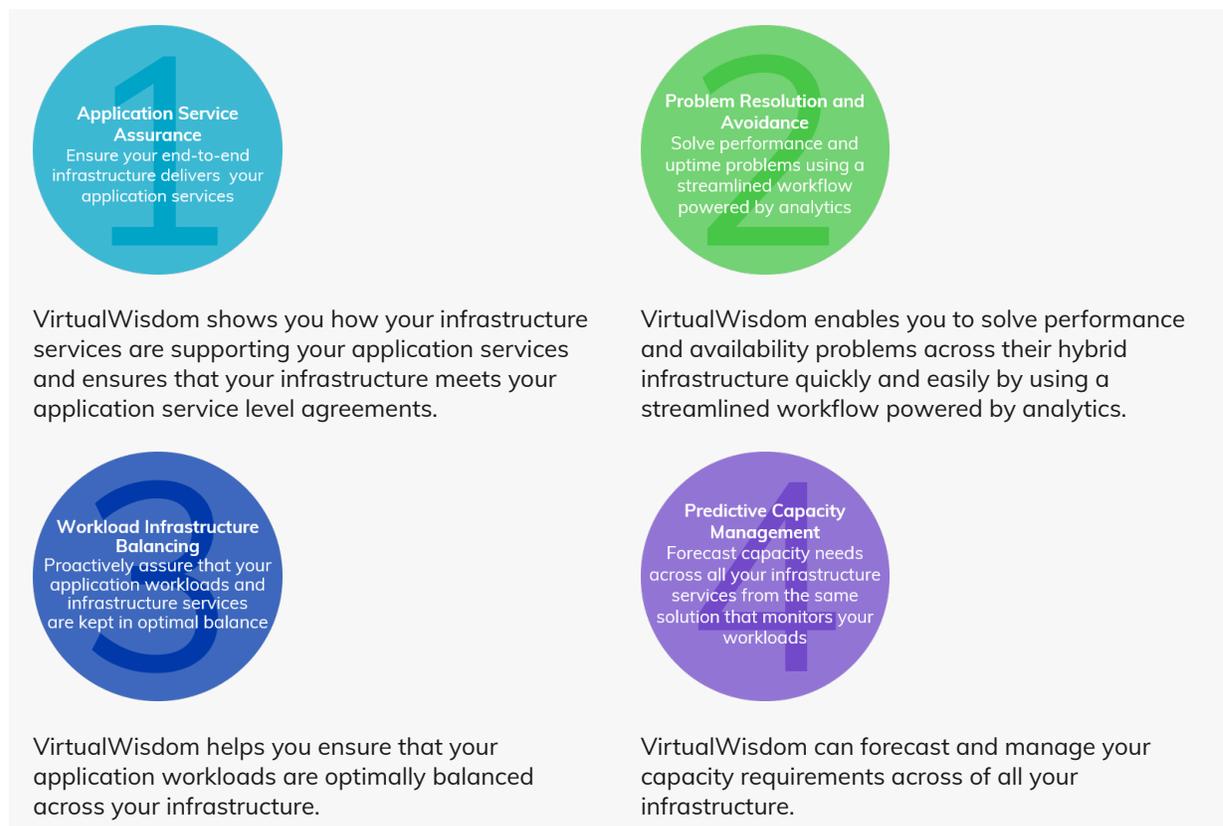
Using Reports	187
Using Report Variables	187
Zoom and Fetch and Set Report Time	189
Report Callouts	192
Exporting a Chart or a Report	193
Sharing a Report Template	194
NPIV Indicator Message	197
Errors in Reports	198
Creating Reports	199
Report Filtering	208
Creating Report Variables	216
Report Snapshots	221
Scheduling Snapshot Cleanup	222
Event Integration	223
Enabling Event Polling	224
Viewing Events	225
Disabling (Stopping) Event Polling	227
Standard Services Reports	227
Dashboards	229
Visibility Dashboards	232
Executive Dashboard	232
Application Dashboard	237
Admin Dashboards	240
Storage Admin Dashboards	241
Compute Admin Dashboard	246
Operations (NOC) Dashboard	250
Analytics	256
Analytics Home Page	257
Problem Resolution and Avoidance Analytics	258
Event Advisor	259
Trend Matcher	261
Workload Infrastructure Balancing Analytics	266
VM Coordinator	268
Storage Port Balancer	274
Workload Right Sizer	278
Queue Solver	283
Migration Analysis	289
Workload Analysis	292
Predictive Capacity Management Analytics	295
Capacity Forecast	297
VM Deployment Advisor	301
Capacity Auditor	305
Application Service Assurance Analytics	309
Seasonal Trend	310
Balance Finder	313
Saved Analytics	319
Analytics Output	319
Deleting Analytics Output	320

Alarms and Cases	322
Alarm Rule Templates	324
Standard Alarm Rule Templates	328
Single Metric Alarms	330
Viewing Your Configured Alarms	335
Configuring an Alarm Rule Template	337
Cases	340
Open Case Page	344
Archived Cases	346
Investigations	350
Infrastructure Assessment Workflow	360
Starting from Topology	361
Starting From Alarms	364
Starting From Analytics	365
Starting From Dashboards	369
Contact Information	372
Legal	374

VirtualWisdom Overview

Our customers leverage Virtana's expertise and experience to bring greater visibility and faster resolution of data center and network-related issues. To enable this success, we offer the VirtualWisdom infrastructure monitoring and AI-powered analytics platform to solve for some of the IT industry's most challenging problems.

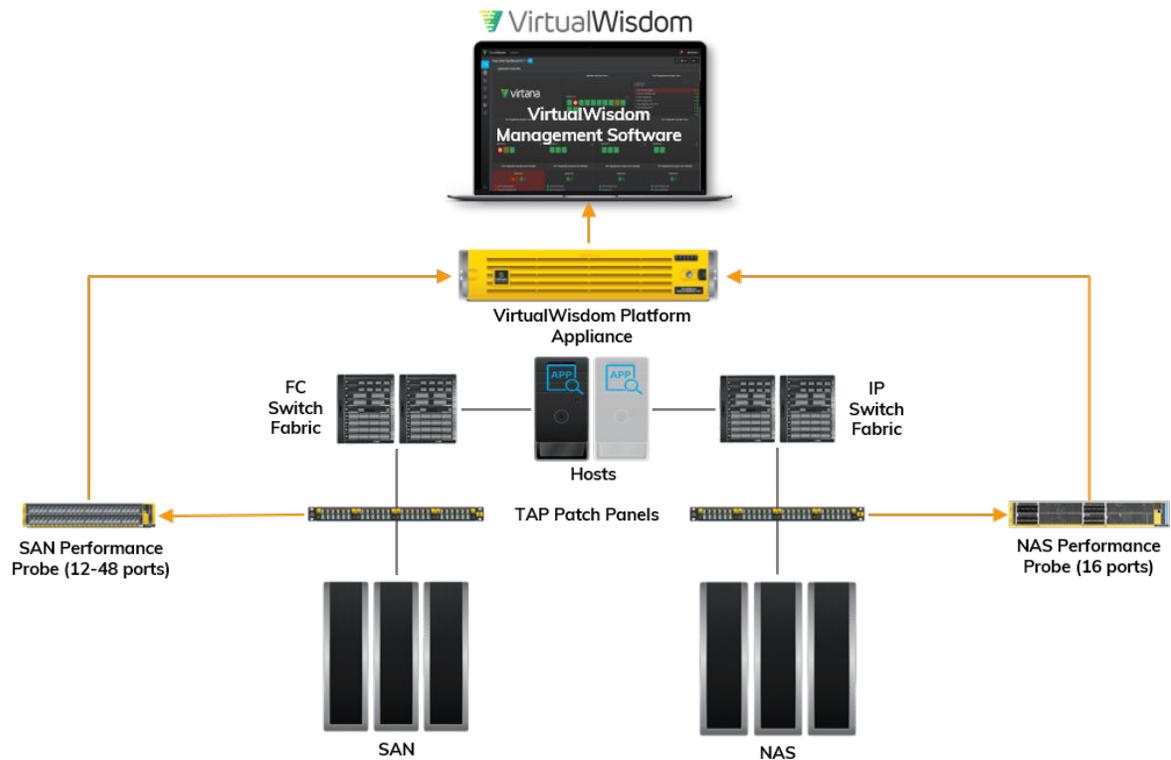
Here are four important value areas where VirtualWisdom provides significant value. We call these our “value pillars” and our training and operationalization activities are designed to help you use VirtualWisdom to address each of these areas.



VirtualWisdom Platform Architecture

VirtualWisdom offers both a hardware-based and a software-based architecture. Our hardware-based solution is comprised of a platform appliance that can be integrated with your infrastructure using hardware probes and software-based integrations. Alternately, you can install our software-based platform that uses a software-based appliance and software integrations only.

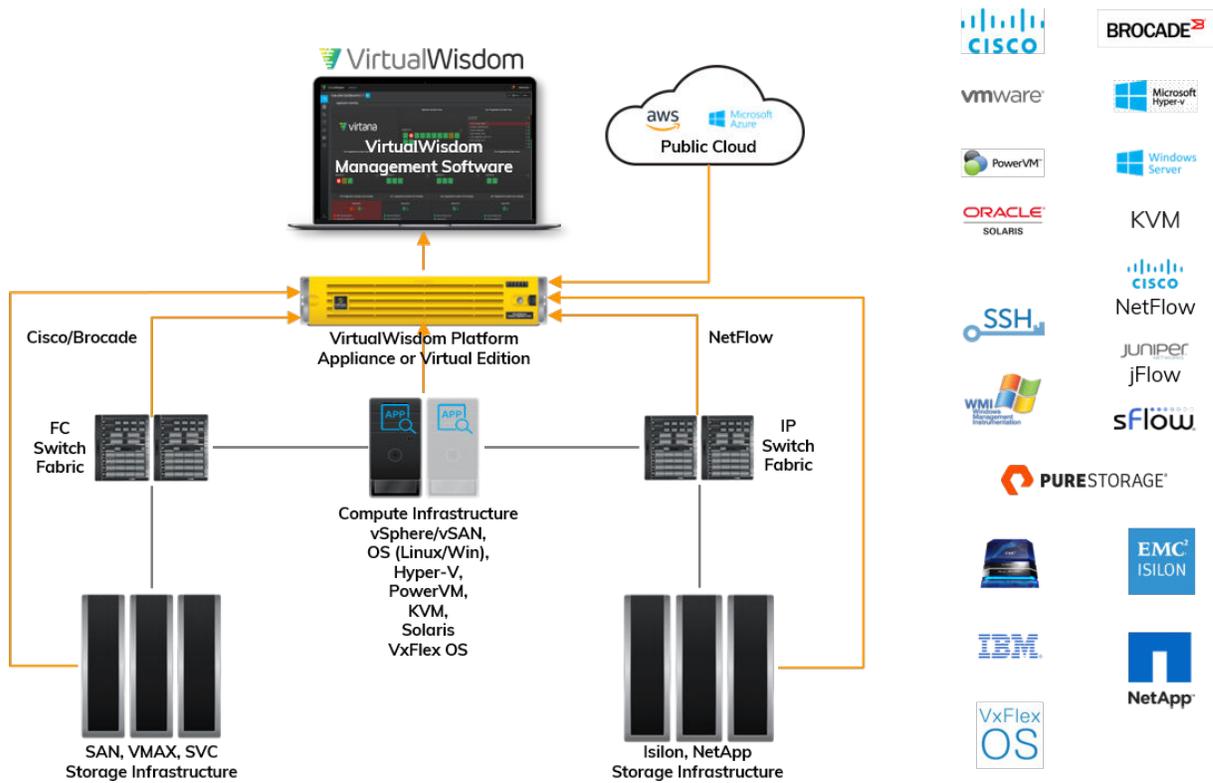
Hardware Probes for Fibre Channel SAN and IP NAS Infrastructure Monitoring



The VirtualWisdom hardware Performance Probes connect to your SAN or NAS infrastructure using Traffic Access Points (TAPs). The probes collect live data from the SAN or NAS infrastructure and send it to the VirtualWisdom Platform Appliance.

The VirtualWisdom appliance persists and correlates the data collected by the probes and presents it in the browser-based VirtualWisdom user interface using the VirtualWisdom Management Software. The hardware appliance must be used with the hardware probes.

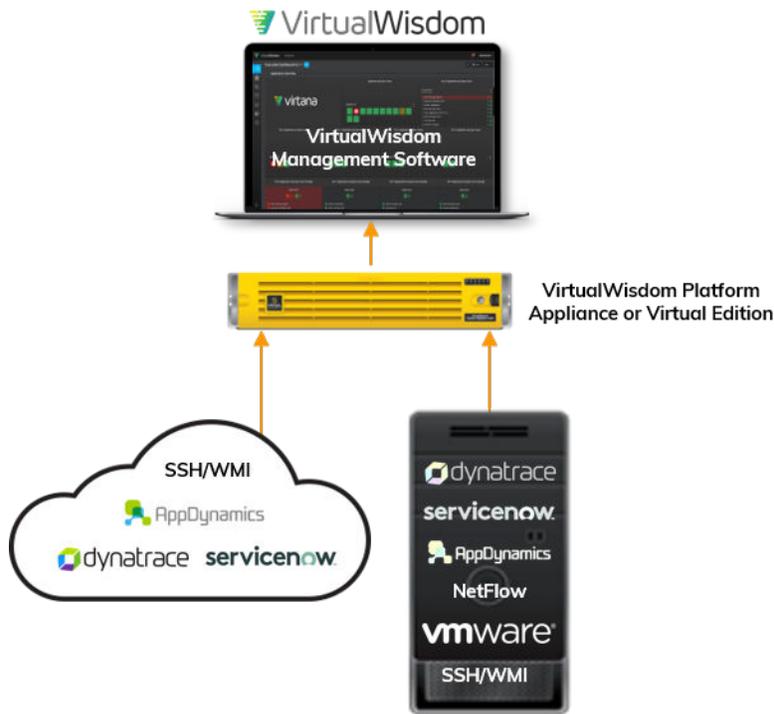
Software Integrations for Network, Compute, and Storage Infrastructure Monitoring



The software integrations collect data from your network, compute, and storage infrastructure and send it to the VirtualWisdom appliance.

The VirtualWisdom appliance persists and correlates the data collected by the integrations and presents it in the browser-based VirtualWisdom user interface using the VirtualWisdom Management Software. You can use the hardware or virtual edition of the appliance with the software integrations.

Software Integrations for Application Discovery and Event Monitoring



The software integrations collect data from your cloud or local APM or CMDB instances and hosts.

The VirtualWisdom appliance persists and correlates the data collected by the application integrations and presents it in the browser-based VirtualWisdom user interface using the VirtualWisdom Management Software. You can use the hardware or virtual edition of the appliance with the software integrations.

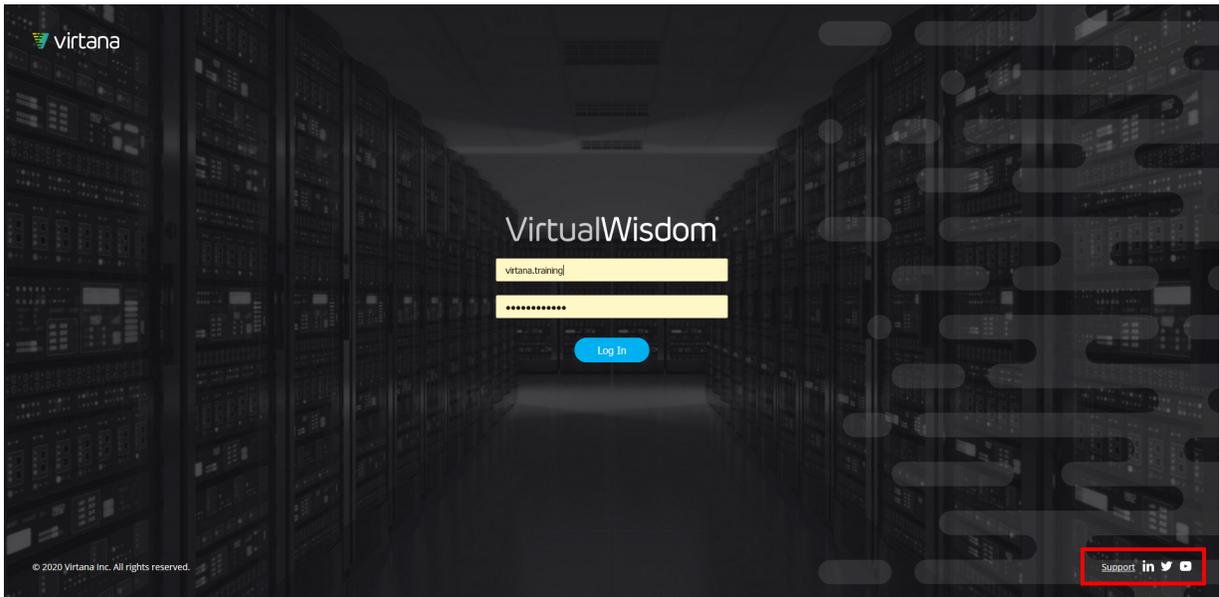
VirtualWisdom User Interface and Navigation

The VirtualWisdom Management Software's user interface is a browser-based graphical user interface that is compatible with the most recent version of any of these supported browsers: Chrome, Firefox, Edge, and Safari.

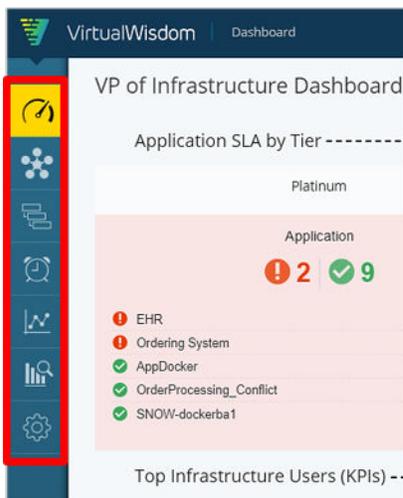
Native browser capability is supported for navigation:

- Forward/Back buttons
- Saved links
- Browser history

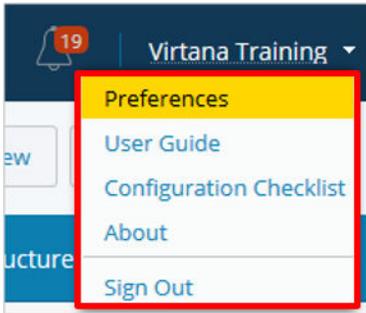
The login screen presents username and password fields and links to the Customer Support Portal and Virtana social media pages.



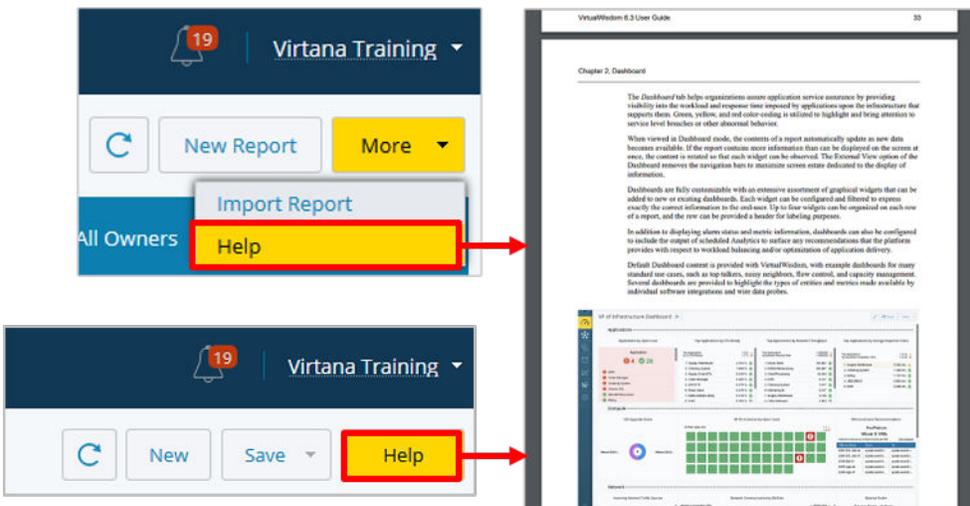
After logging in, links to the different modules are displayed in tabs down the left side of the page. You can use these buttons to return to a previous page after navigation to another module.



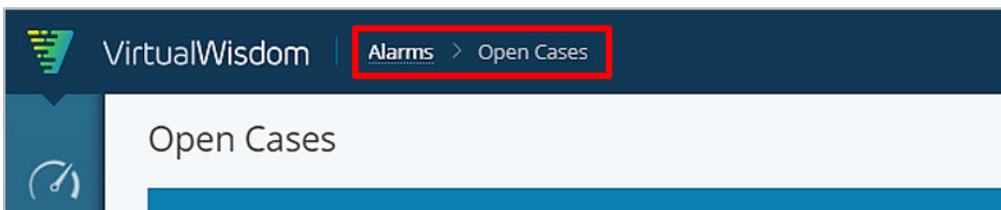
Clicking on the down arrow next to the username, or on the username itself, provides access to common functions across the UI, for example, Preferences, User Guide, Configuration Guide, and About.



To access the VirtualWisdom User Guide, click More, then the Help button, or the Help button, to open a new browser window to display the User Guide in PDF format.

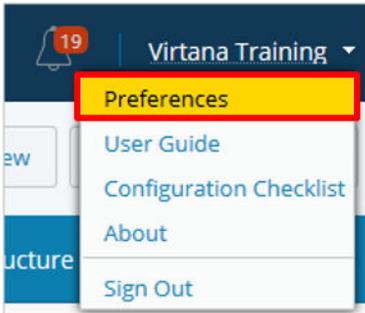


You can use the breadcrumb navigation to easily return to a previous screen.

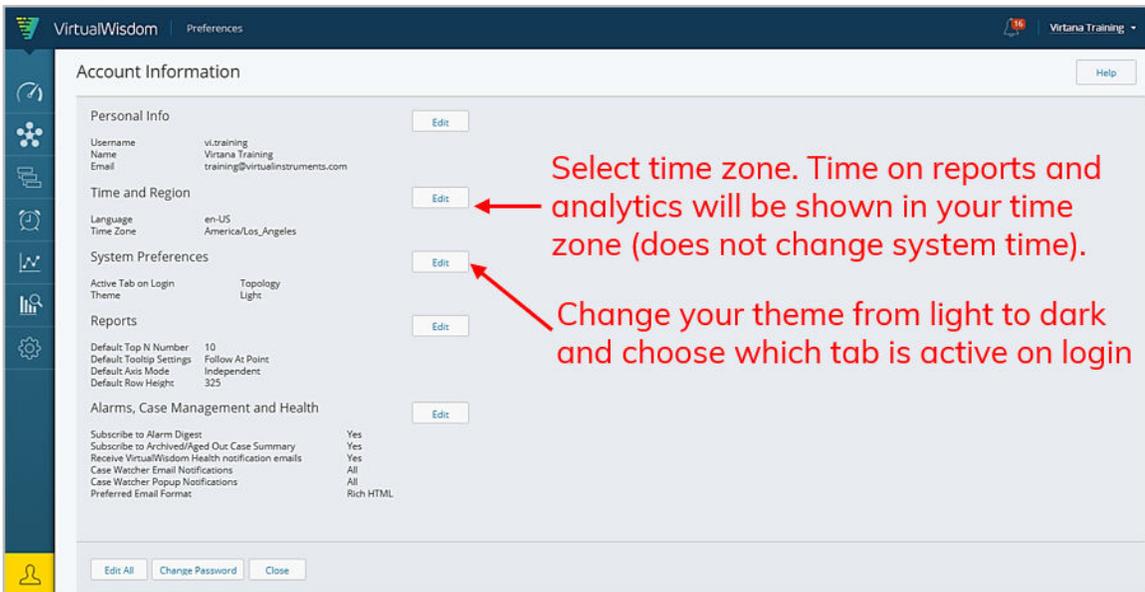


User Settings

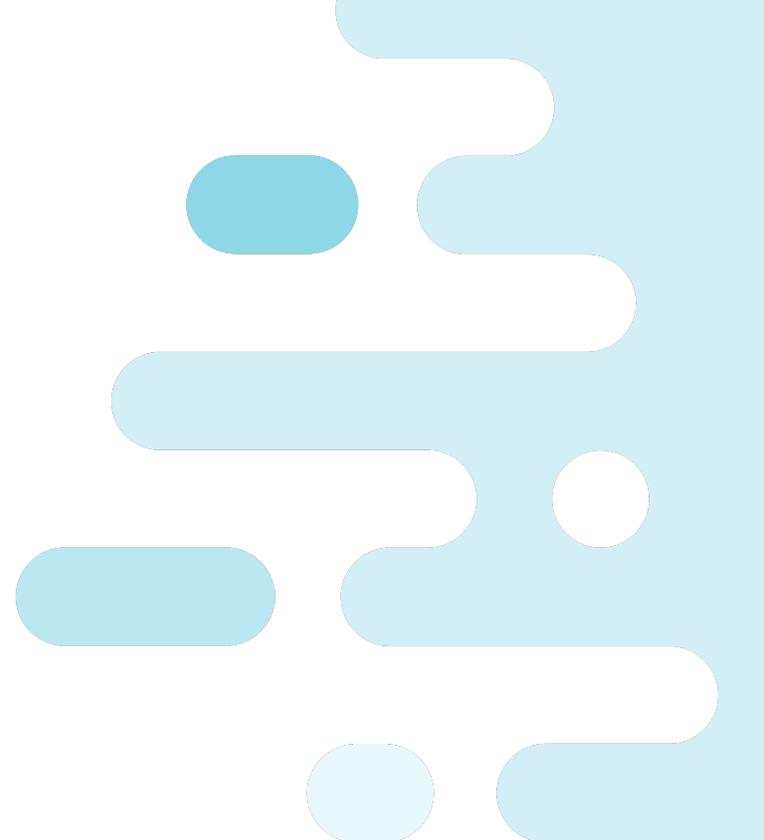
1. To set user preferences, select Preferences from the drop down menu at the top right of the page or click the person icon at the bottom left of the page.



- From the Account Information page, you can edit your account preferences such as local time and region, inactivity timeout, reporting defaults and alarm digest subscription.



Notifications



Banner Notifications

Ongoing activities, such as log collection and backup in progress, display a pop up banner notification. Banner notifications display until you click Close, or until the activity associated with the banner completes.

Banner notifications are visible to all users.

Other Notifications

An alarm bell with a red oval is displayed if there are user notifications available. These notifications include direct messages, cases/alarms/investigations, health, and system notifications. Click on the alarm bell to display a list of notifications.



There are four types of notifications available:

Notification Type	Description
Direct Messages	Messages sent directly to you regarding and associated with an open case.

Notification Type	Description
Cases & Alarms, Investigations	Indicate open cases in Cases & Alarms, Investigations.
VirtualWisdom Health	Visible only to Administrators, indicate that there are open cases in Settings & VirtualWisdom Health Notifications, which have not been closed. These notifications are often related to integrations and data collection. If there are no VirtualWisdom Health open cases, or if all the open cases have been closed, there are no notifications.
System	<p>Notifications related to your VirtualWisdom system. An email server configuration issue could be an example of a system notification.</p> <p>Notifications in the Notification Drawer are marked with a severity icon of Info, Warning, or Critical. The severity icon on the notification corresponds to the case with which they are associated. Notifications in the System section might be marked with the blue circular icon, indicating that the associated notification is unread.</p>

Note that the type of notifications displayed depends on the user's role. Administrators only will see VirtualWisdom Health and System notifications.

Viewing Notifications

1. Click the bell icon.



2. The Notifications Drawer opens. The drawer is divided into four sections: Direct Messages, Cases & Alarms, Investigations, VirtualWisdom Health, and System. The number associated with each section corresponds to the number of notifications in that section. When you clear a notification the number on the bell goes down.

Notifications ✕

- > **Direct Messages** ⋮
- > **Inventory Updates** ⋮
- > **Cases & Alarms, Investigations** ⋮
- ∨ **VirtualWisdom Health**
! 8 | ! 4 | ⋮

! **DNS Reverse** ⋮

04/29/2020 02:15:37 PM PDT

Reverse-lookups of DNS servers are failing on VirtualWisdom Server (VW-64-Demo-28)

! **NTP Issues** ⋮

04/29/2020 02:15:08 PM PDT

VirtualWisdom Server (VI-Appliance.vi.local) is not in sync with the configured NTP server

! **NTP Issues** ⋮

04/29/2020 02:15:07 PM PDT

The NTP server on VirtualWisdom Server (VI-Appliance.vi.local) is not configured/reachable

! **VM Host Issue** ⋮

03/29/2020 03:57:01 PM PDT

Invalid VM properties from host : one or more domain names are not valid:

! **VM Host Issue** ⋮

03/29/2020 03:56:55 PM PDT

Invalid VM properties from host : one or more IP addresses are not valid:

! **DNS Reverse** ⋮

12/29/2019 05:20:32 AM PST

Reverse-lookups of DNS servers are failing on VirtualWisdom Server (VI-Appliance)

! **Host Metrics Connectivity Issue** ⋮

12/20/2019 09:22:08 AM PST

Failed to connect to host 10.20.10.41

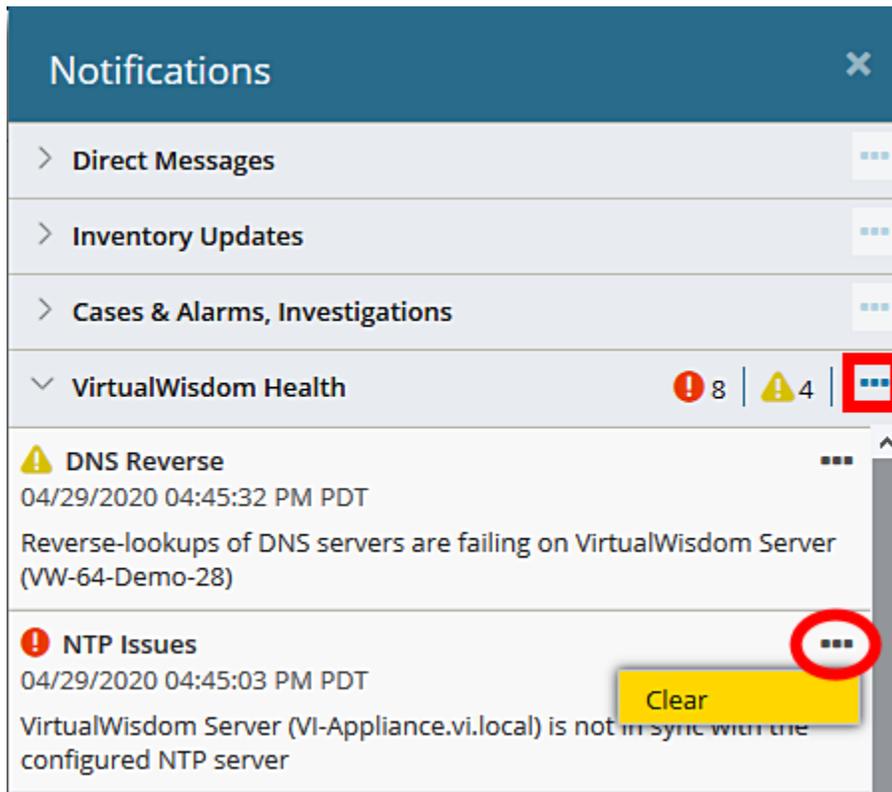
! **Host Metrics Connectivity Issue** ⋮

> **System**
! 2 | ! 5 | ⋮

- Click a section of the Notification Drawer to expand the section and view its notifications. Click the section again to close the section.
- Drill down on a notification to view more data. The Notifications pane will remain open after navigation. You can move this pane by dragging and dropping it to another location on the screen or close it by clicking on the x in the top right corner of the pane.

The screenshot shows the VirtualWisdom interface. The main content area is titled 'Network Setup' and is divided into three sections: 'Virtual Edition', 'Network Ports', and 'DNS and Host Name'. The 'Virtual Edition' section shows 'Appliance Type: vm', 'Version: 6.4.0', and 'Build: 16'. The 'Network Ports' section shows 'NIC0' as the primary interface, with 'Enabled: true', 'DHCP: false', 'IPv4 Address: 10.1.0.28', 'IPv4 Mask: 255.255.255.0', 'IPv4 Gateway: 10.1.0.1', 'RemoteWisdom?: Yes', and 'RemoteWisdom gw: 10.1.0.1'. The 'DNS and Host Name' section shows 'DNS Servers: 10.1.0.1', 'Domains: Hostname', and 'Hostname: VW-64-Demo-28'. On the right side, the 'Notifications' drawer is open, displaying a list of notifications under the 'VirtualWisdom Health' category. The notifications include 'DNS Reverse' (04/29/2020 04:45:32 PM PDT), 'NTP Issues' (04/29/2020 04:45:03 PM PDT), 'NTP Issues' (04/29/2020 04:45:02 PM PDT), 'VM Host Issue' (03/29/2020 03:57:01 PM PDT), 'VM Host Issue' (03/29/2020 03:56:55 PM PDT), 'DNS Reverse' (12/29/2019 09:20:32 AM PST), and 'Host Metrics Connectivity Issue' (12/20/2019 09:22:08 AM PST). Each notification has a status icon (warning or error) and a 'Close' button.

- Click the three dots in the notification to clear it from the list. You can clear all notifications from a category by clicking the three dots in the notification category header.



Data and Entities

What kind of data does VirtualWisdom collect?

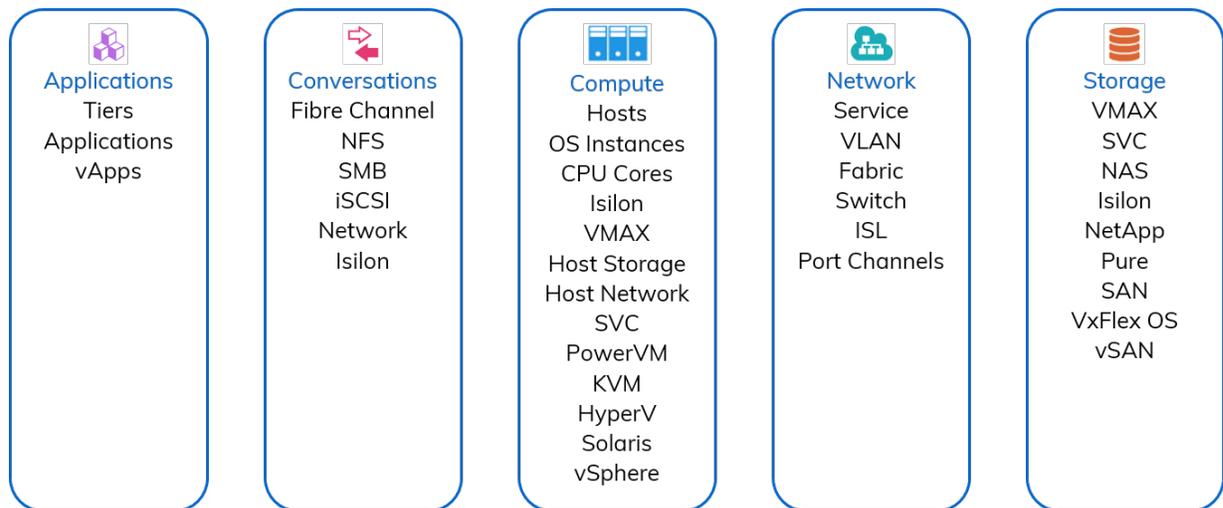
<i>Infrastructure Properties</i>	<i>Events and Metrics</i>
 NETWORK	
 COMPUTE	
 STORAGE	
 APPLICATIONS	
 CLOUD	
 HYPER-CONVERGED	
<ul style="list-style-type: none">• Names• System Properties• Relationships	<ul style="list-style-type: none">• Health Events and Metrics (including AppDynamics events)• Utilization and Capacity Metrics• Performance Metrics• ServiceNow Tickets

VirtualWisdom collects data about the components that make up your hybrid infrastructure and their relationships. VirtualWisdom discovers and stores information about these components including their name, their system properties, and their relationships.

VirtualWisdom also monitors and collects data on the events occurring in your infrastructure, for example, data on infrastructure, health, utilization and capacity, and performance.

Entity Overview

What is an entity?



The entity is the fundamental and most atomic element in VirtualWisdom. Entities allow VirtualWisdom to group resources based on their function, correlation, and inter-dependencies. Entities are logical groupings of the physical and virtual components of your infrastructure and include all the infrastructure components monitored by VirtualWisdom.

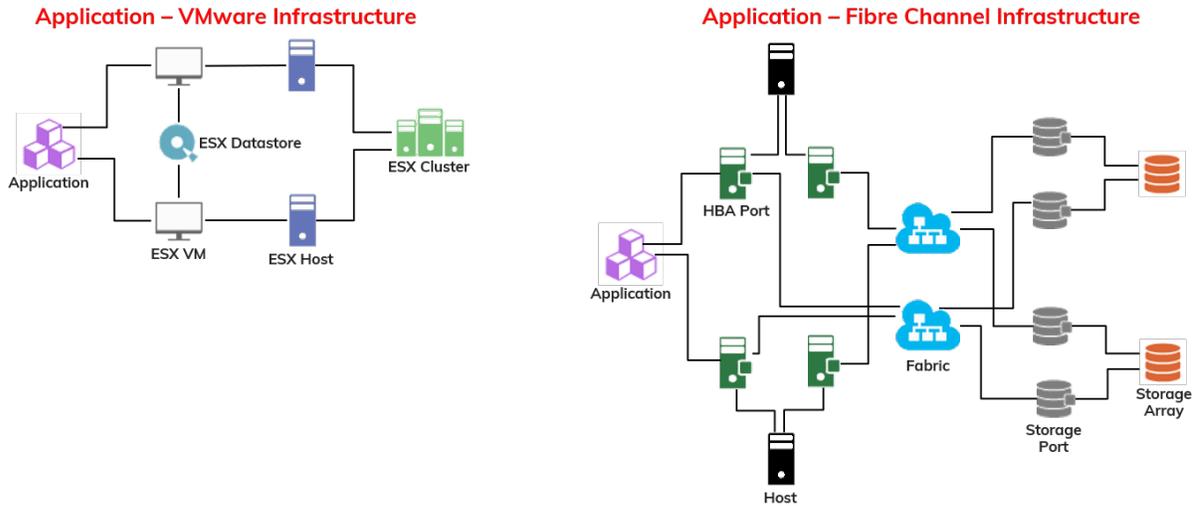
Entities can be linked to other entities. The VirtualWisdom user can build groups of entities to display the end-to-end infrastructure of an application in a meaningful fashion in the VirtualWisdom software.

All entities have associated metrics with built-in aggregation rules for each metric type. The metrics, which measure the data flow in the environment, are collected, accessed and analyzed through Entities. Data can be viewed in the context of hosts, arrays and applications, for example, the top 10 hosts in an application.

The VirtualWisdom software is entity-centric. Entities are used to view topology, set alarm rules, create reports, and run analytics.

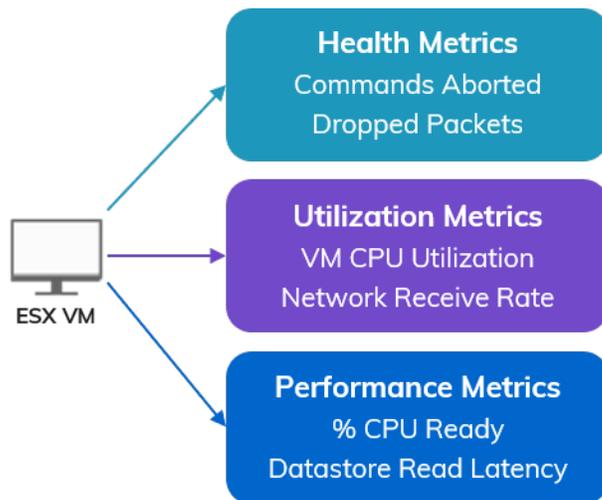
The Value of Entities

Entities provide visibility into the end-to-end infrastructure supporting your application. In the image below, we see two application-centric views of the infrastructure supporting an application: the VMware infrastructure and the fibre channel infrastructure.



Entities are also associated with collected metrics and use built-in aggregation rules to enable:

- Reporting - "What are the top 10 ESX VMs by VM CPU Utilization for an ESX cluster?"
- Alerting - "Alert me when VM CPU Utilization exceeds 95% for the VMs on an ESX cluster?"
- Troubleshooting - "Display events from the last 24 hours where the VM CPU Utilization was high on an ESX cluster."



Entity Types by Category

The Entity Types page displays tables of entity types and their properties. All entity types include the name, tags, and created on properties, plus various additional entity-specific properties.

The following tables list entity types that are included with a standard VirtualWisdom Basic License. Entity types for optional (non-core) integrations are detailed in the relevant Integration User Guide.

Application



Table 1. Entity Types - Applications

Entity Type	Icon	Properties	Created By
Tier		<ul style="list-style-type: none"> • Application • Application Count • Created On • Discovered Name • Entity Type • External ID • Name • Rank • Tags • Tier Key • VW UID 	User or Discovery
Application		<ul style="list-style-type: none"> • Conflict Key • Created On • Entity Type • Name • Number of Hosts not Imported • Tags • Tier • Tier Id • Unread • VW UID 	User or Discovery

Entity Type	Icon	Properties	Created By
vApp		<ul style="list-style-type: none"> • Annotation • Child vApps • Created On • Discovered Name • Entity Type • Inventory Path • Name • Overall Status • Owner • Parent vApp • Tags • VW UID • Virtual Machines 	Discovery

Compute



Table 2. Entity Types - Compute

Sub-Category	Entity Type	Icon	Properties	Created By
Hosts	All Host Types		<ul style="list-style-type: none"> • Created On • Entity Type • Name • Tags • VW UID 	
Hosts	Host		<ul style="list-style-type: none"> • Components • Created On • Domain Name • Entity Type • Last Discovered Role • Name • OS Version • Role • Role Updated By • Tags • VW UID 	User or Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Hosts	OS Instance		<ul style="list-style-type: none"> • Created On • Domain Name • Entity Type • Hypervisor Type • Name • OS Release • OS Type • OS Version • State • Tags • UUID • VW UID 	Discovery
Hosts	CPU Core		<ul style="list-style-type: none"> • Core ID • Created On • Entity Type • Model • Name • OS Instance • Speed (MHz) • Tags • VW UID 	Discovery
Host Storage	Volume Group		<ul style="list-style-type: none"> • Capacity (KB) • Created On • Device Name • Entity Type • Name • OS Instance • Tags • VW UID 	Discovery
Host Storage	Logical Volume		<ul style="list-style-type: none"> • Capacity (KB) • Created On • Device Name • Entity Type • Logical Device Name • Name • OS Instance • Tags • VW UID • Volume Group 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Host Storage	Physical Volume		<ul style="list-style-type: none"> • Capacity (KB) • Created On • Device Name • Entity Type • Created On • Device Name • Entity Type • Logical Device Name • Name • OS Instance • Storage Device • Tags • VW UID 	Discovery
Host Storage	Storage Device		<ul style="list-style-type: none"> • Capacity(KB) • Created On • Device Name • Entity Type • Logical Device Name • Name • OS Instance • Tags • VW UID 	Discovery
Network	HBA Card		<ul style="list-style-type: none"> • Created On • Driver • ESX Host • Entity Type • Host • Model • Name • Node WWN • Tags • VW UID 	User or Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Network	HBA Port		<ul style="list-style-type: none"> • Attached Ports • Created On • Device Type • Discovered Name • Entity Type • FCID • HBA Card • Host • Is Virtual • Logical Fabric • Name • Nickname • Port Speed • Proxy FC Port • Proxy FCID • Proxy Fabric Name • Tags • VW UID • WWN 	Discovery
Network	Virtual Ethernet Port		<ul style="list-style-type: none"> • Created On • Discovered Name • ESX Host • ESX VM • Entity Type • IP Address • Name • Role • Tags • VW UID 	Discovery
Network	Source Ethernet Port		<ul style="list-style-type: none"> • Created On • DHCPv4 Enabled • DCPv6 Enabled • Device Type • Entity Type • MAC Address • Name • Storage Array • Storage Controller • Storage I/O Module • Tags • VW UID 	

Sub-Category	Entity Type	Icon	Properties	Created By
Network	IP Address		<ul style="list-style-type: none"> • Bonded Network Interface • Created On • Device Type • Domain Name • Entity Type • Ethernet Port • Host • IPv4 Long Value • Name • Network Interface • Prefix Length • Tags • VLAN • VW UID • Value • Version 	Discovery
Network	Source IP Address		<ul style="list-style-type: none"> • Bonded Network Interface • Created On • Device Type • Domain Name • Entity Type • Ethernet Port • Host • IPv4 Long Value • Name • Network Interface • Prefix Length • Tags • VLAN • VW UID • Value • Version 	
Network	Network Interface		<ul style="list-style-type: none"> • Bonded Network Interface • Created On • Entity Type • IP Address • Interface Name • MAC Address • Name • OS Instance • Speed (Mbps) • Status • Tags • VW UID 	<ul style="list-style-type: none"> • Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Network	Bonded Network Interface		<ul style="list-style-type: none"> • Bonding Mode • Created On • Entity Type • IP Address • Interface Name • MAC Address • Name • OS Instance • Speed (Mbps) • Status • Tags • VW UID 	Discovery
IBM PowerVM	PowerVM Host		<ul style="list-style-type: none"> • Actual Cores • Actual Memory GB • Created On • Current Available Cores • Current Available Memory GB • Deconfigured Cores • Deconfigured Memory GB • Dedicated Cores • Entity Type • Firmware Memory GB • Model • Name • Pending Available Cores • Pending Available Memory GB • Pool Size • Sample Rate • Serial • Status • Tags • Total Cores • Total Memory GB • VW UID • Virtual Processors 	Discover

Sub-Category	Entity Type	Icon	Properties	Created By
IBM PowerVM	PowerVM Partition		<ul style="list-style-type: none"> • Active Memory Expansion Factor • CPU Mode • CPU Pool Maximum • CPU Pool Name • CPU Pool Reserved • CPU Sharing Mode • CPU Uncapped Weight • Components • Created On • Current CPU • Current Memory GB • Current Paging VIOS • Domain Name • Entity Type • IP Address • LPAR Env • Last Discovered Role • Mac address • Maximum CPU • Maximum CPU Entitled Capacity • Maximum Memory GB • Memory Mode • Memory Weight • Minimum CPU • Minimum CPU Entitled Capacity • Minimum Memory GB • Name • OS version • Power VM Host Name • Primary Paging VIOS • Processor Compatibility Mode • RMC state • Role • Role Updated By • Secondary Paging VIOS • Status • Tags • Using NPIV • VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
IBM PowerVM	PowerVM VIOS Partition		<ul style="list-style-type: none"> • Active Memory Expansion Factor • CPU Mode • CPU Pool Maximum • CPU Pool Name • CPU Pool Reserved • CPU Sharing Mode • CPU Uncapped Weight • Created On • Current CPU • Current Memory GB • Current Paging VIOS • Entity Type • IP Address • LPAR Env • Mac address • Maximum CPU • Maximum CPU Entitled Capacity • Maximum Memory GB • Memory Mode • Memory Weight • Minimum CPU • Minimum CPU Entitled Capacity • Minimum Memory GB • Name • OS version • Power VM Host Name • Primary Paging VIOS • Processor Compatibility Mode • RMC state • Secondary Paging VIOS • Status • Tags • VW UID 	Discovery
Microsoft Hyper-V	Hyper-V Cluster		<ul style="list-style-type: none"> • Created On • Discovered Name • Domain Name • Entity Type • Hyper-V Hosts • Name • Tags • VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Microsoft Hyper-V	Hyper-V Host		<ul style="list-style-type: none"> • Created On • Discovered Name • Domain Name • Entity Type • HBA Cards • Hyper Visor Present • Hyper-V Cluster • Hyper-V VMs • Inventory Path • Logical Processors • Name • Power State • Public IP Address • Tags • Total Physical Memory (GB) • VW UID • Version • Windows GUID 	Discovery
Microsoft Hyper-V	Hyper-V VM		<ul style="list-style-type: none"> • Components • Created On • Discovered Name • Domain Name • Entity Type • FC Ports • Hyper-V Host • Inventory Path • Last Discovered Role • Name • OS Version • Power State • Role • Role Updated By • Tags • Total Memory (GB) • VW UID • Virtual CPUs • Windows GUID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Microsoft Hyper-V	Hyper-V VHD		<ul style="list-style-type: none"> • Created On • Discovered Name • Entity Type • Filename • Hyper-V VM • Name • Tags • VW UID 	Discovery
VMware vCenter	ESX Cluster		<ul style="list-style-type: none"> • Created On • Discovered Name • ESX Hosts • Entity Type • Inventory Path • Name • Tags • VW UID 	Discovery
VMware vCenter	ESX Datastore		<ul style="list-style-type: none"> • CIFS User • Created On • Data Store Type • Discovered Name • Disk Groups • Entity Type • Inventory Path • Is Accessible • NAS Host • NAS Host IP • NAS Mount Path • Name • Overall Status • Tags • VW UID • Virtual Machines 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
VMware vCenter	ESX Host		<ul style="list-style-type: none"> • CPUMhz • Cache Disks • Capacity Disks • Connection State • Created On • Discovered Name • Disk Groups • ESX Cluster • ESX Datastores • Entity Type • Ethernet Ports • HBA Cards • Hyper-Threading Enabled • Inventory Path • Is Supported Version • Logical Processors • MemorySize • Mounted File Systems • Name • Number of CPU Packages • Overall Status • Power State • Tags • VW UID • Version • Virtual Ethernet Ports • Virtual Machines 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
VMware vCenter	ESX VM		<ul style="list-style-type: none"> • Capacity • Components • Connection State • Created On • Datastores • Discovered Name • Domain Name • ESX Host • Entity Type • Ethernet Ports • FC Ports • Free Space • Inventory Path • Last Discovered Role • MemorySizeMB • Name • OS Version • Overall Status • Power State • Role • Role Updated By • Tags • VW UID • Virtual CPUs 	Discovery

Conversations



NOTE

If you are working on a software VirtualWisdom Edition then only Network Conversations and Isilon Conversations* can be viewed in Inventory. To view other conversation types (FC, NFS, SMB, iSCSI) the VirtualWisdom hardware probes must be installed.

*Requires the Isilon integration to be installed and configured. See the Isilon Integration User Guide for a list of Isilon entities.

Contact [Virtana Sales](#) for more information.

Table 3. Entity Types - Conversations

Entity Type	Icon	Properties	Created By
FC Conversation		<ul style="list-style-type: none"> • Created On • Entity Type • Initiator FCID • Initiator Name • Initiator WWN • LUN • Name • Tags • Target FCID • Target Name • Target WWN • VW UID • initiatorId • targetId 	Discovery
NFS Conversation		<ul style="list-style-type: none"> • Created On • Destination • Entity Type • FSID • Name • Source • Tags • VLANID • VW UID 	Discovery
SMB Conversation		<ul style="list-style-type: none"> • Created On • Destination • Entity Type • Name • Share Name • Source • Tags • VLANID • VW UID 	Discovery
ISCSI Conversation		<ul style="list-style-type: none"> • Created On • Destination • Entity Type • LUN • Name • Source • Tags • VLANID • VW UID 	Discovery

Entity Type	Icon	Properties	Created By
Network Conversation		<ul style="list-style-type: none"> Created On Discovered Name Entity Type Name Tags VW UID 	Discovery

Network



Table 4. Entity Types - Network

Sub-Category	Entity Type	Icon	Properties	Created By
IP Network	Ethernet Port		<ul style="list-style-type: none"> Created On DHCPv4 Enabled DHCPv6 Enabled Device Type Entity Type MAC Address Name Storage Array Storage Controller Storage I/O Module Tags VW UID 	
IP Network	Network Service		<ul style="list-style-type: none"> Created On Discovered Name Entity Type Name Tags VW UID Version 	Discovery
IP Network	VLAN		<ul style="list-style-type: none"> Created On Entity Type ID Name Tags VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Storage Network	Physical Fabric		<ul style="list-style-type: none"> • Created On • Discovered Name • Entity Type • Name • Tags • VW UID 	Discovery
Storage Network	SAN Switch		<ul style="list-style-type: none"> • Created On • Discovered Name • Entity Type • IP Address • Manufacturer • Model • Name • Physical Fabrics • Serial Number • Tags • VW UID • Version • WWN 	Discovery
Storage Network	Switch Blade		<ul style="list-style-type: none"> • Created On • Discovered Name • Entity Type • Module Number • Name • SAN Switch • Tags • VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Storage Network	Switch Port		<ul style="list-style-type: none"> Attached ISL Port Attached Ports Created On Device Type Discovered Name Entity Type FCID Is Virtual Logical Fabric Logical Switch Name Nickname Port Speed Port Type SAN Switch Switch Blade Tags VW UID WWN 	Discovery
Storage Network	Inter-Switch Link		<ul style="list-style-type: none"> Attached Ports Created On Discovered Name Entity Type Name Tags VW UID 	Discovery
Storage Network	LAN		<ul style="list-style-type: none"> Created On Entity Type Name Tags VLAN VW UID 	
Logical Network	Logical Fabric		<ul style="list-style-type: none"> Created On Discovered Name Entity Type Fabric ID Name Physical Fabrics Tags VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
Logical Network	Logical Switch		<ul style="list-style-type: none"> Created On Discovered Name Entity Type Fabric ID Logical Fabric Name SAN Switch Tags VW UID WWN 	Discovery
Logical Network	Port Channel		<ul style="list-style-type: none"> Attached ISLs Created On Entity Type Is Virtual Name Port Speed Tags VW UID WWNs 	Discovery

Storage



Table 5. Entity Types - Storage

Sub-Category	Entity Type	Icon	Properties	Created By
NAS>File System	NFS File System		<ul style="list-style-type: none"> Created On Entity Type Ethernet Port FSID NAS File System Key NFS Conversation Name Storage Array Tags VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
NAS>File System	SMB File System		<ul style="list-style-type: none"> Created On Entity Type Ethernet Port Name SMB Conversation SMB File System Key Share Name Storage Array Tags VW UID 	Discovery
NAS>File System	Link Aggregation Group		<ul style="list-style-type: none"> Created On Entity Type LAG Key LAG Number Name Tags VW UID 	Discovery
NAS>File System	Monitored Link		<ul style="list-style-type: none"> Created On Entity Type Link Aggregation Group NAS Probe Port Key Name Port Number Tags VW UID 	Discovery
NAS>File System	Destination Ethernet Port		<ul style="list-style-type: none"> Created On DHCPv4 Enabled DHCPv6 Enabled Device Type Entity Type IP Addresses MAC Address Name Storage Array Storage Controller Storage I/O Module Tags VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
NAS>File System	Destination IP Address		<ul style="list-style-type: none"> Bonded Network Interface Created On Device Type Domain Name Entity Type Ethernet Port Host IPv4 Long Value Name Network Interface Prefix Length Tags VLAN VW UID Value Version 	Discovery
NAS>NetApp	NetApp Cluster		<ul style="list-style-type: none"> Cluster Location Created On Discovered Name Entity Type Manufacturer Model Name Serial Number Tags UUID VW UID 	Discovery
NAS>NetApp	NetApp Storage Node		<ul style="list-style-type: none"> Asset Tag Created On Discovered Name Entity Type Manufacturer Model Name NetAppCluster DisplayLabel NetAppCluster UUID Node Location Serial Number Tags UUID VW UID Version 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
NAS>NetApp	NetApp SVM		<ul style="list-style-type: none"> Created On Discovered Name Entity Type Name NetAppCluster DisplayLabel NetAppCluster UUID Tags UUID VServer Type VW UID 	Discovery
NAS>NetApp	NetApp LIF		<ul style="list-style-type: none"> Created On Discovered Name Entity Type IP Address Name NetApp SVM Role Tags VW UID 	Discovery
SAN	Storage Array		<ul style="list-style-type: none"> Created On Entity Type Name Tags VW UID 	User or Discovery
SAN	Storage Controller		<ul style="list-style-type: none"> Created On Entity Type Name Storage Array Tags VW UID 	User or Discovery
SAN	Storage I/O Module		<ul style="list-style-type: none"> Created On Entity Type Name Storage Controller Tags VW UID 	User or Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
SAN	Storage Port		<ul style="list-style-type: none"> Attached Ports Created On Device Type Discovered Name Entity Type FCID Is Virtual Logical Fabric Name Nickname Port Speed Proxy FC Port Proxy FCID Proxy Fabric Name Storage Array Storage Controller Storage I/O Module Tags VW UID WWN 	Discovery
SDS>VxFlex OS	VxFlex OS System		<ul style="list-style-type: none"> Cluster Mode Cluster State Created On Entity Type Good Nodes Num Good Replicas Num Name Perf Profile VxFlex OS System Key System Version Name Tags VW UID 	Discovery
SDS>VxFlex OS	VxFlex OS Metadata Manager		<ul style="list-style-type: none"> Created On Entity Type Name Role VxFlex OS Mdm Node Key VxFlex OS Mdm Node Name Status Tags VW UID Version 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
SDS>VxFlex OS	VxFlex OS Protection Domain		<ul style="list-style-type: none"> Created On Entity Type Name VxFlex OS Protection Domain Key VxFlex OS Protection Domain Name State Tags VW UID 	Discovery
SDS>VxFlex OS	VxFlex OS Storage Pool		<ul style="list-style-type: none"> Created On Entity Type Name VxFlex OS Storage Pool Key VxFlex OS Storage Pool Name Tags VW UID 	Discovery
SDS>VxFlex OS	VxFlex OS Data Server		<ul style="list-style-type: none"> Created On Entity Type Maintenance State Mdm Connection State Membership State Name Perf Profile Rmcache Size In Kb VxFlex OS Data Server Key VxFlex OS Data Server Name State Tags Use Rmcache VW UID Version 	Discovery
SDS>VxFlex OS	VxFlex OS Network Interface		<ul style="list-style-type: none"> Created On Entity Type IP Address Name Role Tags VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
SDS>VxFlex OS	VxFlex OS Device		<ul style="list-style-type: none"> • Capacity Limit In Kb • Created On • Device State • Entity Type • Error State • Max Capacity In Kb • Name • Path Name • VxFlex OS Device Key • VxFlex OS Device Name • Tags • VW UID 	Discovery
SDS>VxFlex OS	VxFlex OS Volume Tree		<ul style="list-style-type: none"> • Created On • Entity Type • Name • VxFlex OS Volume Tree Key • VxFlex OS Volume Tree Name • Tags • VW UID 	Discovery
SDS>VxFlex OS	VxFlex OS Volume		<ul style="list-style-type: none"> • Capacity In KB • Created On • Entity Type • Name • VxFlex OS Volume Key • VxFlex OS Volume Name • Tags • Use Rmcache • VW UID • Volume Type 	Discovery
SDS>VxFlex OS	VxFlex OS Fault Set		<ul style="list-style-type: none"> • Created On • Entity Type • Name • Role • VxFlex OS Fault Set Key • VxFlex OS Fault Set Name • Status • Tags • VW UID • Version 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
SDS>VxFlex OS	VxFlex OS Data Client		<ul style="list-style-type: none"> • Created On • Entity Type • Mdm Connection State • Name • Perf Profile • VxFlex OS Data Client IP • VxFlex OS Data Client Key • VxFlex OS Data Client Name • Tags • VW UID • Version 	Discovery
vSAN	Disk Group		<ul style="list-style-type: none"> • Cache Disks • Capacity Disks • Created On • Disk Group Key • ESX Datastore • ESX Host • Entity Type • Name • Node UUID • Tags • UUID • VW UID 	Discovery
vSAN	Cache vSAN Disk		<ul style="list-style-type: none"> • Created On • Device Type • Discovered Name • Disk Group • Entity Type • LUN • Name • PowerVM Partition • Tags • VW UID 	Discovery
vSAN	Capacity vSAN Disk		<ul style="list-style-type: none"> • Created On • Device Type • Discovered Name • Disk Group • Entity Type • LUN • Name • PowerVM Partition • Tags • VW UID 	Discovery

Sub-Category	Entity Type	Icon	Properties	Created By
vSAN	SCSI Disk		<ul style="list-style-type: none"> • Created On • Device Type • Discovered Name • Disk Group • Entity Type • LUN • Name • PowerVM Partition • Tags • VW UID 	Discovery

ITL Examples

- Divide a host to specify one ITL (*initiator:target:lun*) combination responsible for the application and another ITL combination responsible for the backup operations.
- For a backup operation, select only the target port and LUNs to which all back-ups are sent. VW discovers which ports and LUNs are communicating with each other and build the entity for you.
- You want a quick way to configure an application entity but do not know its downstream resources. You can create an application entity by specifying only the initiator port WWN.
- If you create a combination with a * notation but do not have a performance probe monitoring the traffic, no metric information is be stored. This example could occur when an administrator wants to provision application entities before the performance probe is deployed.

FC Conversations



NOTE

The hardware SAN Performance Probe (ProbeFC) must be installed in your environment to discover and collect data for FC Conversation entities.

Contact [Virtana Sales](#) for more information.

Fibre Channel conversations are communications between Initiators, Targets, and LUNs (ITLs):

- Initiator (HBA Port)
- Target (Storage Port)
- LUN (Logical Unit Number)

As conversations occur and are discovered by ProbeFC, the metrics are captured and stored, and the *FC Conversation* entities are named with an *initiator:target:lun* convention. These entities can be modified like any other discovered entity, with a Name and Description. The FC Conversation entity type is available in the following areas of VirtualWisdom:

- Entity list
- Entity editing (name, description)
- Alarms (Exchange Performance only)
- Reports (all types)
- Analytics (Trend Matcher, Event Advisor)

When you discover FC Conversations, you must apply filter criteria. Otherwise, depending on the size of your network, an unmanageably large number might be displayed. Narrow the choices by filtering:

- Application
- Initiator - Host, HBA Card, HBA Port
- Target - Storage Array, Storage Controller, Storage I/O Module, Storage Port

An initiator (typically a host) negotiates with a target (typically a storage device) to connect to a LUN (a disk partition, or one or more disk drives).

FC Conversations cannot be searched by name from the FC Conversation page. After filtering, you can sort by the Name column and manually scan the results.

Table 6. Application Entities

Representation	Description
ITL	Initiator-side port, target storage port, and LUN
IT*	Initiator-side port, target storage port, and any LUN
I**	Initiator-side entity, wildcard for target, LUN(s) (replaced by Devices)
*TL	Target storage port (WWN) and LUN

- If you specify a wildcard (*), the performance probe fills in the entity information based on the conversations it observes. Applications configured this way must have a performance probe monitoring the traffic.

- You can configure an application entity by choosing a port entity, such as a target port or host port. When a port type is the starting point in the entity creation process, you have the option to select the exact target port and LUN. You could also allow the system to discover the conversations by selecting I**.
- When you configure an application entity by choosing a configured entity type such as a host or storage processor, VirtualWisdom uses only discovered conversations. For example, if you start configuring an application entity with a host entity (consisting of ports), VW doesn't allow you to then specify the target entity. The configuration wizard uses only the * notation.
- A target port can be discovered if a ProbeFC monitors the traffic or zoning is configured.

Discovering FC Conversation Entities

As conversations occur and are discovered by ProbeFCs, the metrics are captured and stored, and the FC Conversation entities are named with an *initiator:target:lun* convention. As with any other discovered entity, these entities can be customized with a Name and Description.

1. Click **Inventory** and select FC Conversation from the Conversations group. The FC Conversation Filter displays.

The screenshot shows a dialog box titled "FC Conversation Filter" with a close button (X) in the top right corner. The dialog contains three rows of dropdown menus:

- Application:** A single dropdown menu that is currently empty.
- Initiator:** A dropdown menu with "Host" selected.
- Target:** A dropdown menu with "Storage Array" selected.

At the bottom of the dialog, there are two buttons: a yellow "Apply" button and a white "Cancel" button with a blue border.

2. Specify one or more filter criteria, and click the **Apply** button. At least one entity type must be specified to filter FC Conversations; if you specify multiple entity types, an AND filter is created. A search box is available for each type. The filtered FC Conversation list is displayed, in which the grayed-out discovered Name column is the combination of the *Initiator*, *Target*, and *LUN* columns. These columns can be sorted in ascending (default) or descending order.
3. Select an FC Conversation by clicking the row.

The row is highlighted. You can display the properties of a conversation with the *Show Properties* menu item. Properties of the *Initiator*, *Target*, and *LUN* are displayed.

4. Click the *Edit* menu item.

The *FC Conversation* page is displayed, showing the original discovered name of the conversation. Properties are also displayed, in a read-only area.

5. Specify the information you choose in the *Name* and/or *Description*, fields, and click the **Save** button.

The revised information in the newly-created conversation entity is now shown in the *FC Conversation* list, and the *Name* field is no longer grayed out.

The revised detail information of the newly-created conversation entity is now shown on the edit page.

NFS/SMB Conversation

NFS/SMB conversations are communications between:

- Source
- Destination
- Filesystem

As conversations occur and are discovered by ProbeNAS, the metrics are captured and stored, and the *NFS Conversation* and *SMB Conversation* entities are named with an *source:destination.filesystem* convention. These entities can be modified like any other discovered entity, with a *Name*, and/or *Description*.

When you discover NFS or SMB conversations, all of them are displayed, and no filter criteria are applied.

The following child entities are supported for the NASProbe, by manual creation, entity management, or entity import:

- Applications - NFS Conversation, SMB Conversation
- Hosts - Source IP Address
- Storage Arrays - Destination Ethernet Ports
- Storage Controllers - Destination Ethernet Ports
- I/O Modules - Destination Ethernet Ports

Alarms based on NAS metrics:

- Performance
 - Histogram Performance
 - Average Performance
 - Procedure Rate

- Procedure Limit
- Errors
 - Link Errors
 - Packet Errors
 - Flow Control

Discovering NFS or SMB Conversation Entities

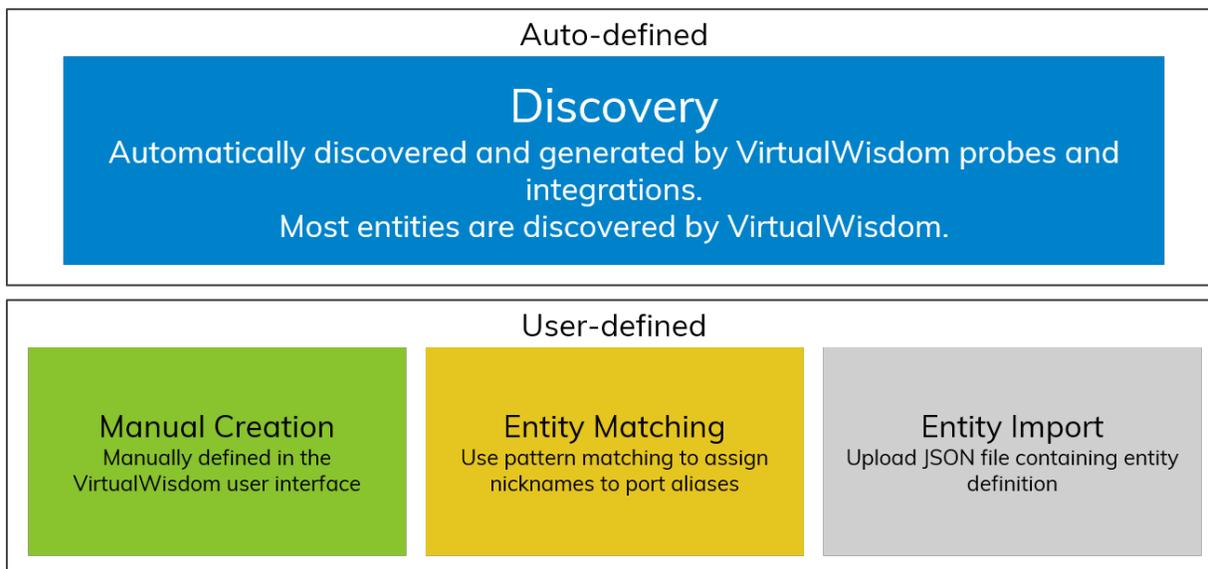
As conversations occur and are discovered by ProbeNAS, the metrics are captured and stored, and the *NFS Conversation* entities are named with an *source:destination:filesystem* convention. As with any other discovered entity, these entities can be customized with a *Name* and/or *Description*.

1. Click **Inventory** and select *NFS or SMB Conversation* from the *Conversations* group.
2. Select the *NFS/SMB Conversation* entity type.
The list of conversations is displayed.
3. Select an NFS or SMB conversation by clicking the row.
The row is highlighted.
You can display the properties of a conversation with the *Show Properties* menu item from the dropdown.
4. Specify the information you choose in the *Name* and/or *Description* fields, and click **Save**.
The revised information is now shown in the *NFS Conversation* or *SMB Conversation* list.

Entity Creation

How are entities created?

Most entities are automatically created by VirtualWisdom when new infrastructure is discovered by VirtualWisdom's probes and integrations. This is called auto-defined or auto-discovery entity creation.



Entities can also be defined manually through the VirtualWisdom user interface using one of three methods:

1. **Manual Creation:** The user creates each entity with the entity management feature in the VirtualWisdom user interface. User Defined Entities allow the VirtualWisdom user to organize their environment and the collected metrics in a fashion that is familiar to them, for example, by Host, Application, or Storage Array.
2. **Entity Matching:** Entity Matching is a feature that allows the user to assign a meaningful nickname to discovered port Entities. Pattern matching is applied against port alias values based on a nickname scheme to streamline what would otherwise be a tedious process.
3. **Entity Import:** Entities can also be imported using a JavaScript Object Notation (JSON) file. JSON is an open standard format that uses human-readable text to transmit data between a server and a web application.



IMPORTANT

Beginning in VirtualWisdom 6.7, a limit has been placed on the number of conversation entities that VirtualWisdom stores for ProbeFC, ProbeNAS, and NetFlow.

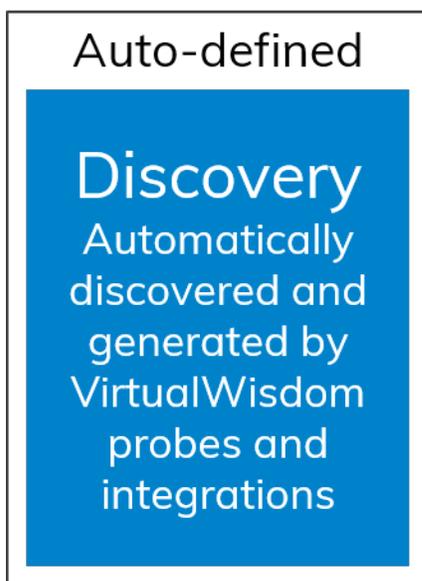
If the system limit of the number of conversations is reached, the least-recently-seen conversations are automatically deleted.

Deletion of these entities is intended to increase performance and reliability for long-running deployments. If you wish to modify or disable this feature, contact [VirtualWisdom Support](#).

Entity Management

When new integrations or probes are configured with VirtualWisdom, most entities associated with the integrations or probes are automatically discovered. However, some entities might not be identified. In that case, the VirtualWisdom administrator (**vw-admin** role) must manually create the missing entities.

Auto-Defined Entity Management



Most entities in VirtualWisdom are discovered automatically. Auto-discovery is performed when new integrations and probes are configured, on a defined schedule, e.g., every 24 hours, or when new infrastructure components are installed and discovery is started manually after installation.

Discovery identifies topology, zoning, and various system properties that depend on the integration and entity type. These properties are called system properties and can be viewed on the entity's Inventory page. VirtualWisdom also tracks meta data on entities, e.g., type, created by/on, tags, etc.

Configuring Auto-Discovery Schedules for Software Integrations

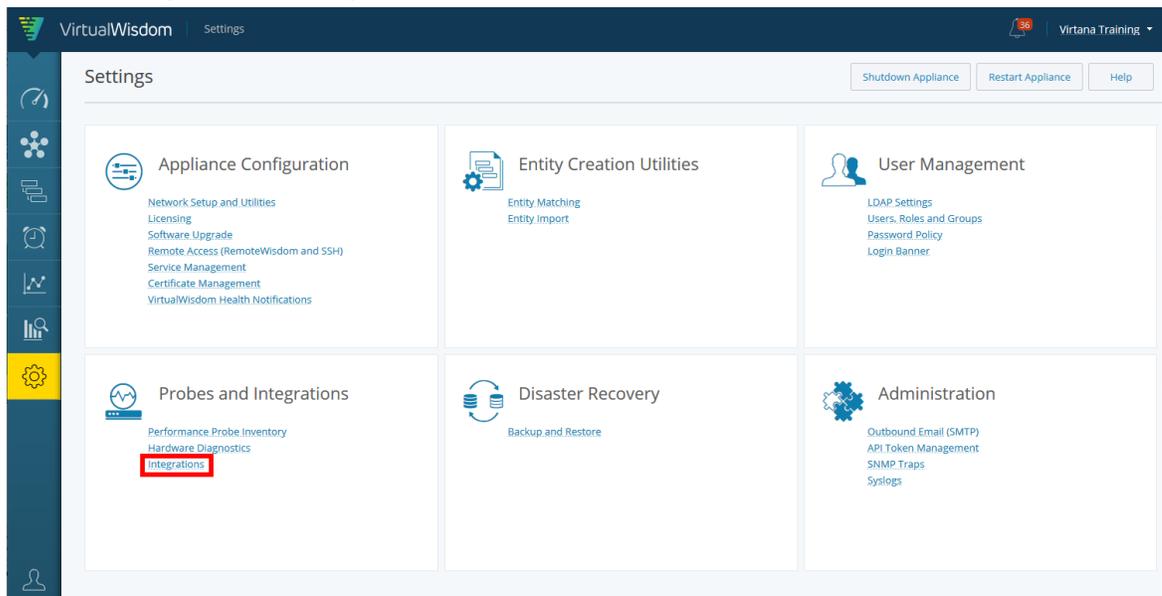
Discovery can be configured to occur automatically on a recurring schedule.



NOTE

This task is available only to users with the VirtualWisdom Administrator (**vw-admin**) role.

1. From Settings, select Integrations.



2. Locate the desired integration and select View.

The screenshot shows the 'Integrations' page with the following cards:

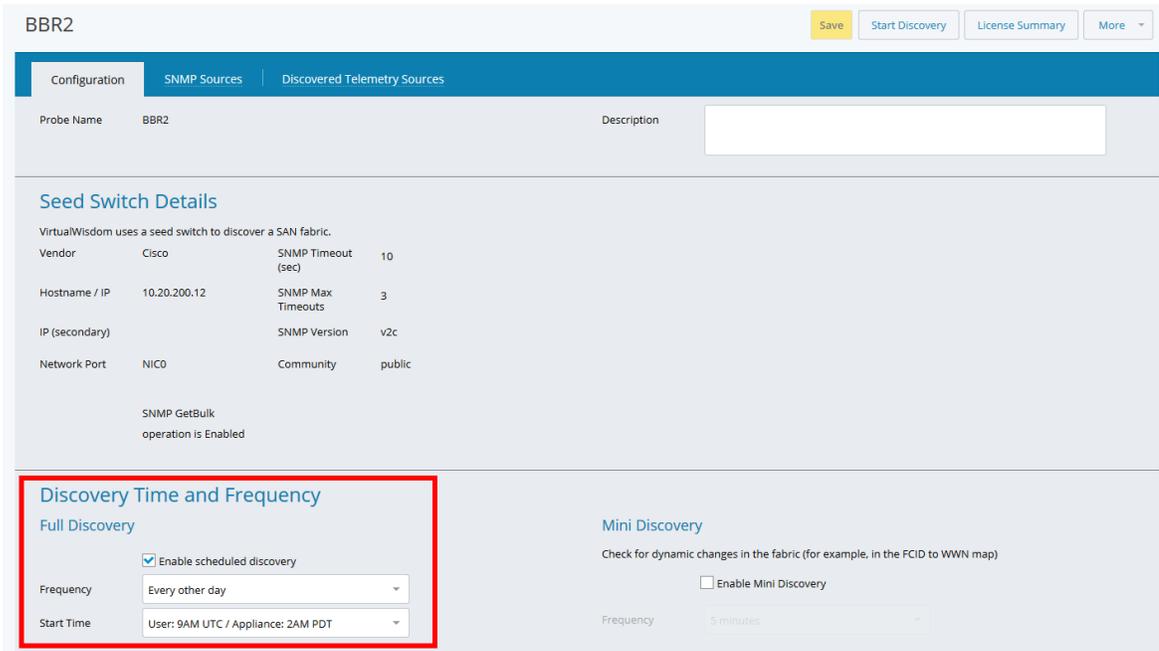
- U.12.1.16 (4): Discover and monitor KVM environments
- Discover and Monitor Microsoft's Hyper-V environment
- U.14.0.1 / 5 (2): Subscribe to Solaris OS instances to enable metric collection using SSH
- Virtana VirtualWisdom - 1.0.2 (1): Monitors the health of the VirtualWisdom Appliance
- VMware vSphere (2): Discover and Monitor VMware's vSphere environment
- Brocade SAN (2): Discover and Monitor Brocade's Fibre Channel switch environment
- Cisco SAN (4): Discover and Monitor Cisco's Fibre Channel switch environment (View button highlighted)
- Dell EMC Isilon Integration - 2.1.4.8 (1): Discover and Monitor DELL EMC Isilon
- Dell EMC VxFlex OS (1): Discover and Monitor Dell's VxFlex OS (ScaleIO) environment
- EMC VMAX Integration - 2.4.0.2 (1): Discover and Monitor DELL EMC's VMAX Storage Arrays through Unisphere storage manager
- IBM SVC - 2.2.2.1 (0): Discover and Monitor IBM's SVC storage virtualizer
- NetApp FAS (1): Discover and Monitor NetApp C-mode filers (NFSv3)
- NetApp Storage - 0.12.0-beta.16 (2): Discover and Monitor NetApp ONTAP environment
- Pure FlashArray - 1.5.1.11 (2): Discover and Monitor Pure FlashArray storage arrays

3. A list of configured subscriptions is displayed, with data on when the last discovery occurred, and any integration or metrics collection warnings or errors. Select a row to review the configuration.

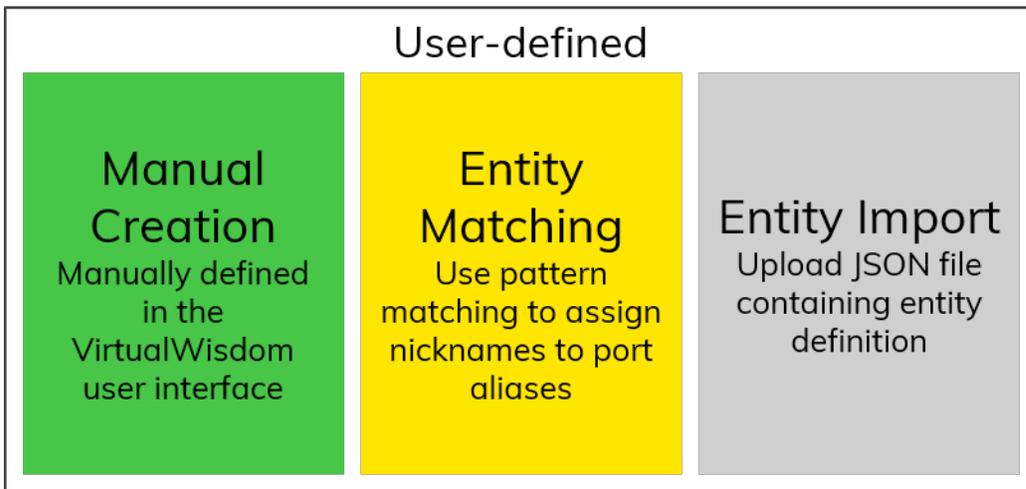
The screenshot shows the 'Cisco SAN' page with a table of subscriptions:

Name ↑	Subscription	Last Discovery	Last Metrics Collection
BBR2	Subscribed (1/1)	Warning - 05/29/2020 09:00:05 AM UTC	05/29/2020 10:45:01 PM UTC
FCOE-C5548	Subscribed (8/9)	Warning - 05/29/2020 09:03:56 AM UTC	Collecting Metrics...
mds01-bbr1.lab.vi.local	Subscribed (1/1)	05/29/2020 09:00:15 AM UTC	05/29/2020 10:45:01 PM UTC
mds9148-qe-1	Subscribed (1/1)	Warning - 05/29/2020 09:00:00 AM UTC	Collection Failed (1/1)
mds9148-qe-2	Subscribed (1/1)	Warning - 05/29/2020 09:00:00 AM UTC	Collection Failed (1/1)
MDS9706-QE	Subscribed (3/3)	Warning - 05/29/2020 09:00:51 AM UTC	05/29/2020 10:45:39 PM UTC

4. Select the desired frequency and start time. Each integration has an option to configure the frequency and start time for automatic discovery. Note that the option may vary slightly based on the integration.

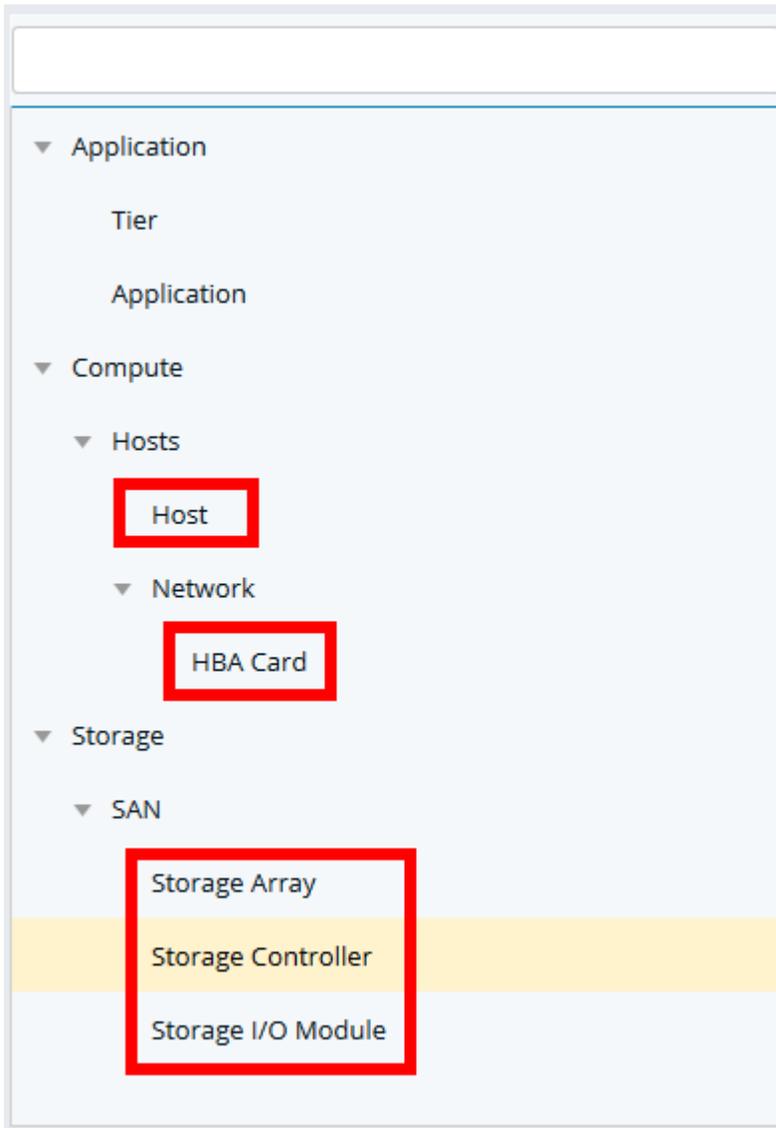


User-Defined Entity Management



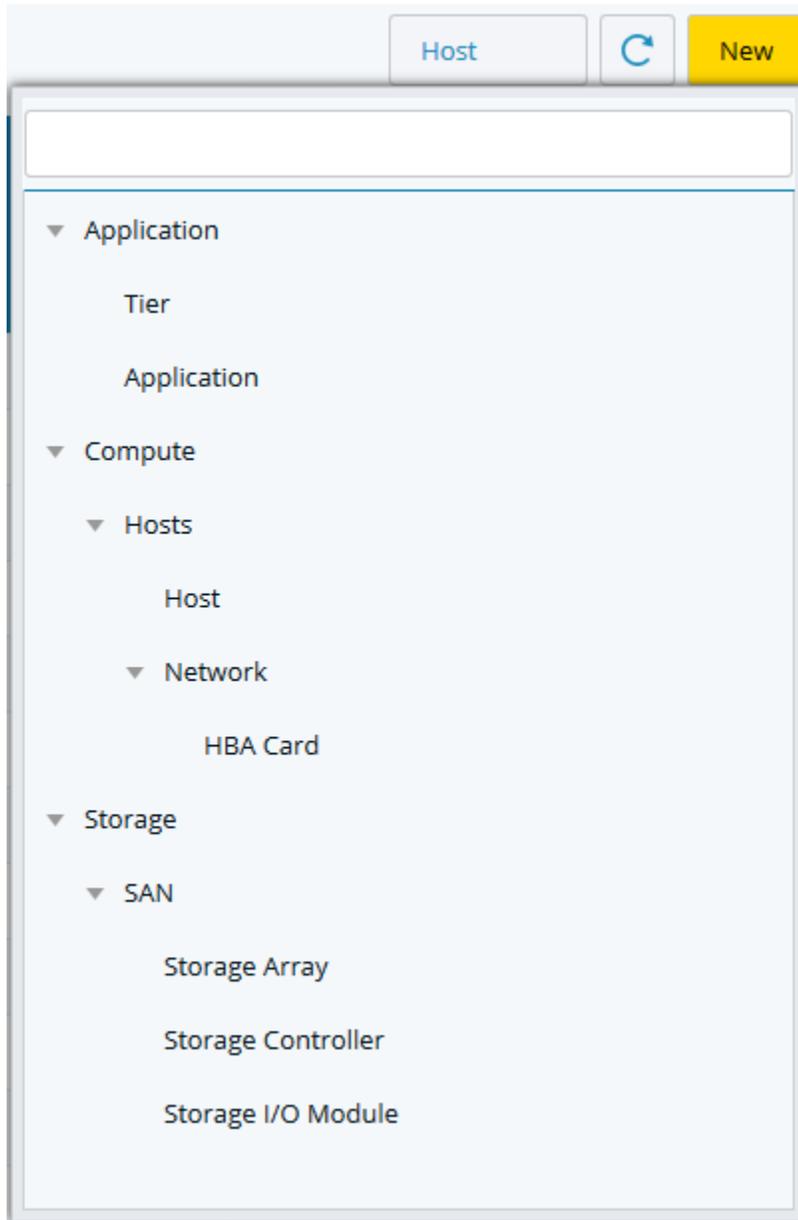
A handful of entity types can (or must) be defined using the VirtualWisdom administrator (**vw-admin** role). Manual entity creation can occur after initial auto-discovery is completed, or after new components have been added.

The highlighted entity types can only be created manually. Tiers and Applications can be created automatically through VirtualWisdom's APM integrations.



Manual Entity Creation

1. Navigate to the Inventory module and select New, then select the entity type (Host, HBA Card, Storage Array, Storage Controller, Storage I/O Module).



2. Name is the only required field but you can define custom properties such as description and tags. Creator and creation date are also tracked.
 - a. Enter the new entity's name and any other desired information.

Host

Name * SVCS_Test_1 Description Services test host

Tags Test

Role Test

System Properties

Custom Properties Add Property

Data Center Services

Entities Add Bulk Delete

Name Type

Click Add to associate entities



IMPORTANT

The following characters cannot be entered in entity names: ~ % & *
= [] \ ; , | < > ?

- b. Click Add to associate the entity with sub-entities, e.g., HBA ports in the case of a host or HBA card.

Add Entity

HBA Card (Items: 61)

Compute

Hosts

Network

HBA Card

HBA Port

IP Address

Source IP Address

p770-9117-MMD-SN849C567:1E450041...

p770-9117-MMD-SN849C567:1E450040...

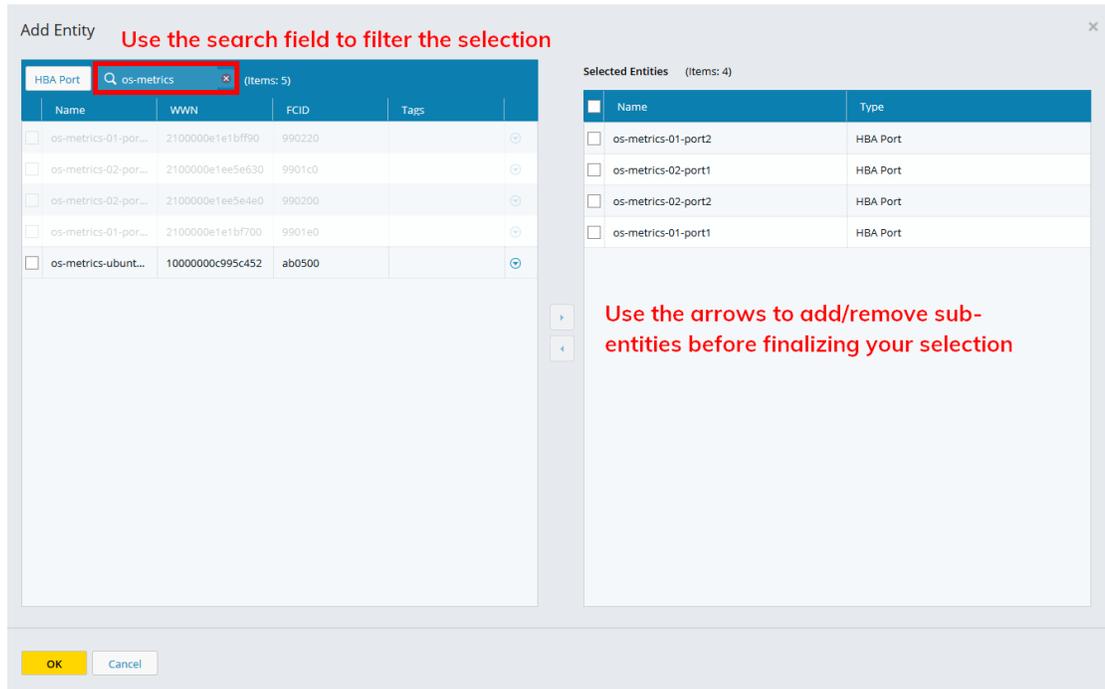
HYPERV-01:hba0

Selected Entities

Name Type

OK Cancel

You will see a different menu depending on the entity type you are creating.



For example, entities that can be comprised of more than one sub-entity types will require you to select the sub-entity type first. Click the entity button at the top to select an entity type. You can use the search field to filter the entities by name. For entity types for which there is only one possible sub-entity type, you will see a list of sub-entities that you can select from.

Creating SAN Fabric Entities

Creating host and storage array entities from discovered HBA and storage ports is recommended to support advanced troubleshooting capabilities.



Entities should be created after vSphere discovery is completed (ESX Hosts are automatically discovered and do not need to be created).

The [Entity Matching \[61\]](#) or [Entity Import \[73\]](#) utilities are used to create these entities.



Entity Matching

The Entity Matching utility creates entities based on pattern matches using discovered port information, e.g., WWN and nickname (alias).



NOTE

Entity Matching works really well when your organization uses a highly regimented approach to providing aliases (human-readable nicknames for HBA and storage port WWNs) for their devices. Here are some examples:

Host	HBA Ports
SJPEXWIN23	SJPEXWIN23_HBA0 SJPEXWIN23_HBA1

This name breaks down as follows: SJ = San Jose, P = Production, EX = an abbreviation of the primary application (Microsoft Exchange), WIN = Windows, 23 = the 23rd of its kind.

This host has one or more HBA ports. A common naming convention of these ports is shown above.

With this kind of convention in place, using Entity Matching to create the Host entity (SJPEXWIN23) is a trivial exercise.

Storage Array	Storage Ports
VMAX0589	VMAX0589_10E0 VMAX0589_10F0 VMAX0589_10G0 VMAX0589_10H0 VMAX0589_9E0 VMAX0589_9F0 VMAX0589_9G0 VMAX0589_9H0

This name breaks down as follows: VMAX = storage array model, 0589 = last four digits of the array's serial number.

This storage array has a number of storage ports associated with it, with names as shown above.

Using Entity Matching to create the storage array entity is very simply done.

The Entity Matcher uses parse rules to group discovered port-level entities into higher level entities like hosts and storage arrays. Parse rules are regular expressions: sequences of characters that form search patterns.

You can use the REGEX tester at this link to test your pattern matches before using the Entity Matching Utility: <https://regex101.com/#pcre>

If your organization uses a consistent naming strategy for hosts, storage arrays, and ports, using the Entity Matching Utility is straightforward.

If your organization does not use a consistent naming strategy (this is common in companies that have undergone mergers), you can request assistance from Virtana Services to design parse rules and assist you with using the Entity Matching Utility.

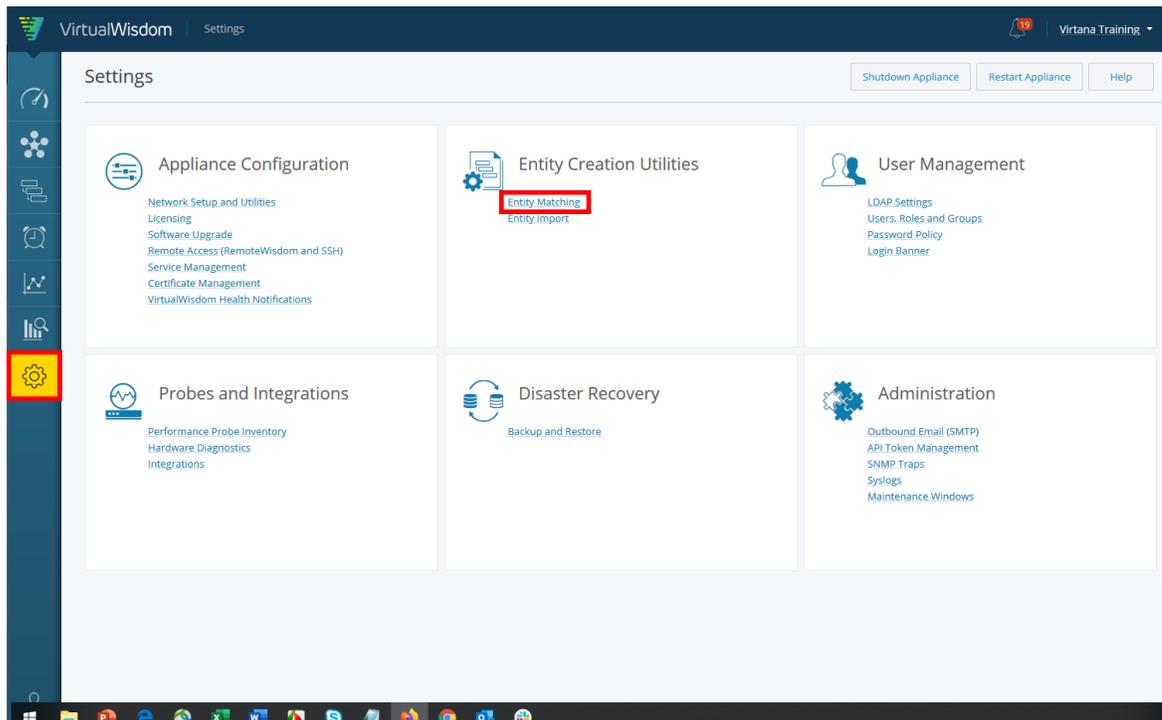
Using the Entity Matching Utility



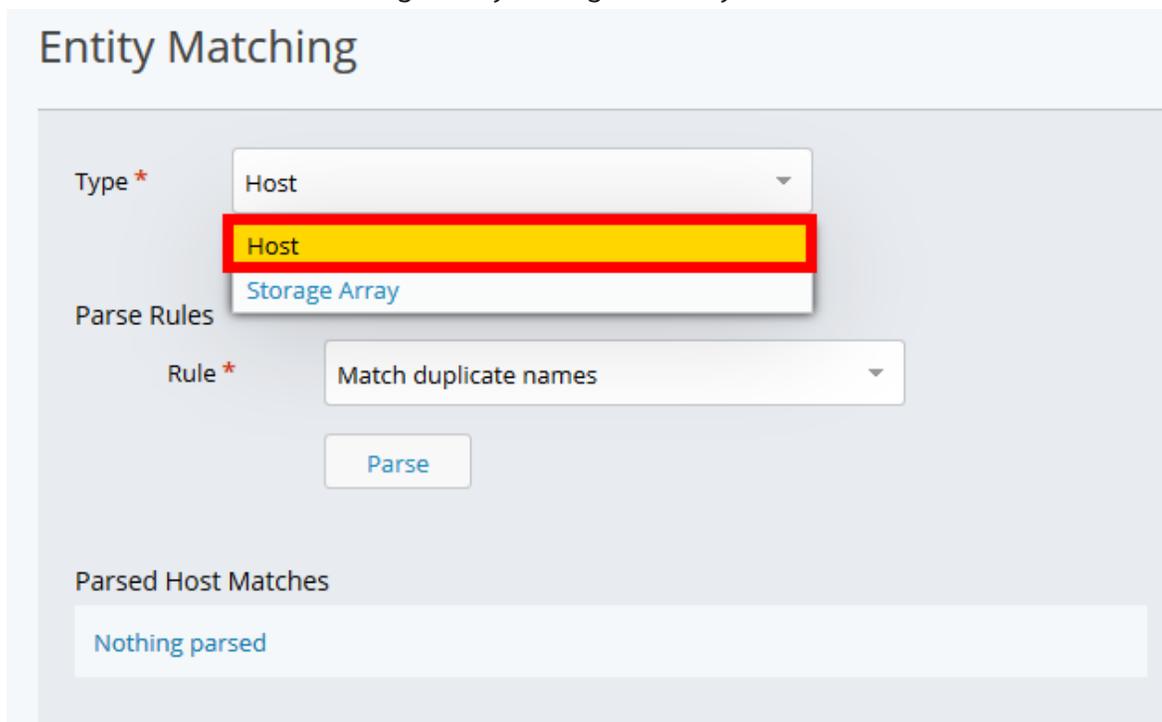
NOTE

This task is available only to users with the VirtualWisdom Administrator (**vw-admin**) role.

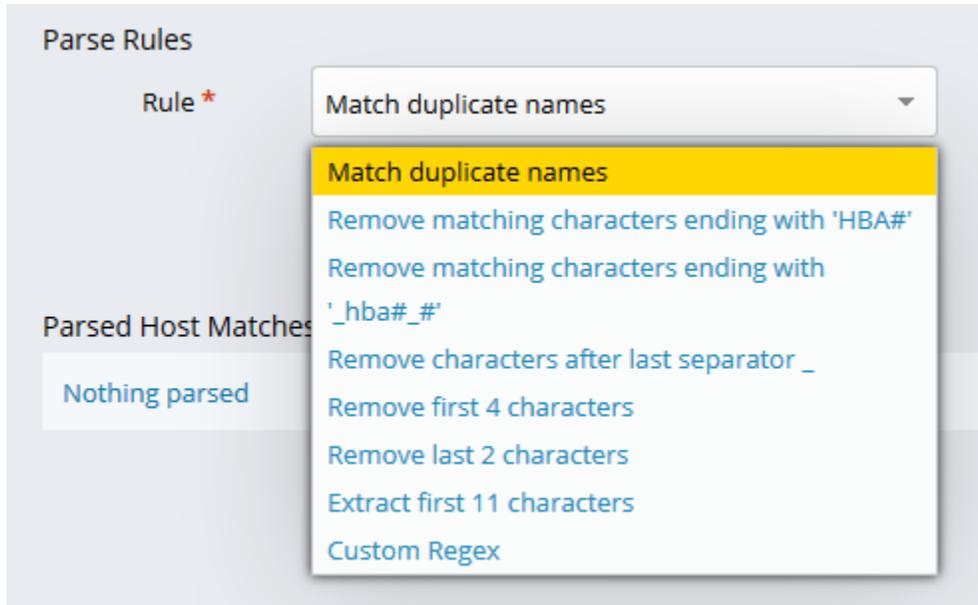
1. To use the utility, navigate to Settings > Entity Matching.



2. You can create hosts or storage arrays using the utility.



3. Select a parse rule to match the entities' aliases. You can use any of the existing parse rules or create your own custom regex.



4. Use these recommended parse rules to perform entity matching.

Remove matching characters ending with 'HBA#'
 Remove matching characters ending with
 '_hba#_#'
 Remove characters after last separator _
 Remove first 4 characters
 Remove last 2 characters
 Extract first 11 characters
 Custom Regex

- a. **Remove characters after last separator (_)**

This parse rule will remove all characters from the alias after the last separator. The base rule uses an underscore, `_`, as the separator but you can create a custom rule to change the separator.

After selecting the standard "last separator" rule, select the Custom Regex rule.

Remove first 4 characters
 Remove last 2 characters
 Extract first 11 characters
 Custom Regex

Change the underscore to any other symbol you want to use for matching. Click Parse to view the matches.

Parse Rules

Rule *

Regex *

The utility returns the number of matches found. Click the down arrow next to the rule to view all matches in the target group.

Parsed Host Matches

Custom Regex (^(.*)_.*\$) (2)

All entities matching the rule are displayed. Hover over the matched ports to see the complete list of target hosts. You can remove any entities you do not wish to include in the creation process.

Target Group : Custom Regex (^(.*)_.*\$)

<input type="checkbox"/>	Host	Ports
<input type="checkbox"/>	sb10-1	sb10-1-INIT
<input type="checkbox"/>	sblaze2-4-init0	sblaze2-4-init0-virtual

Once you are satisfied with the target group, select Create (Host) Entities to create the entities.

b. **Extract first 11 characters**

Another useful parse rule is the "Extract first 11 characters" rule. Using the same process that was outlined above, you can create a custom regex to change the number of characters to extract.

We recommend that you start with the longest port names first.

- Match duplicate names
- Remove matching characters ending with 'HBA#'
- Remove matching characters ending with '_hba#_#'
- Remove characters after last separator _
- Remove first 4 characters
- Remove last 2 characters
- Extract first 11 characters**
- Custom Regex

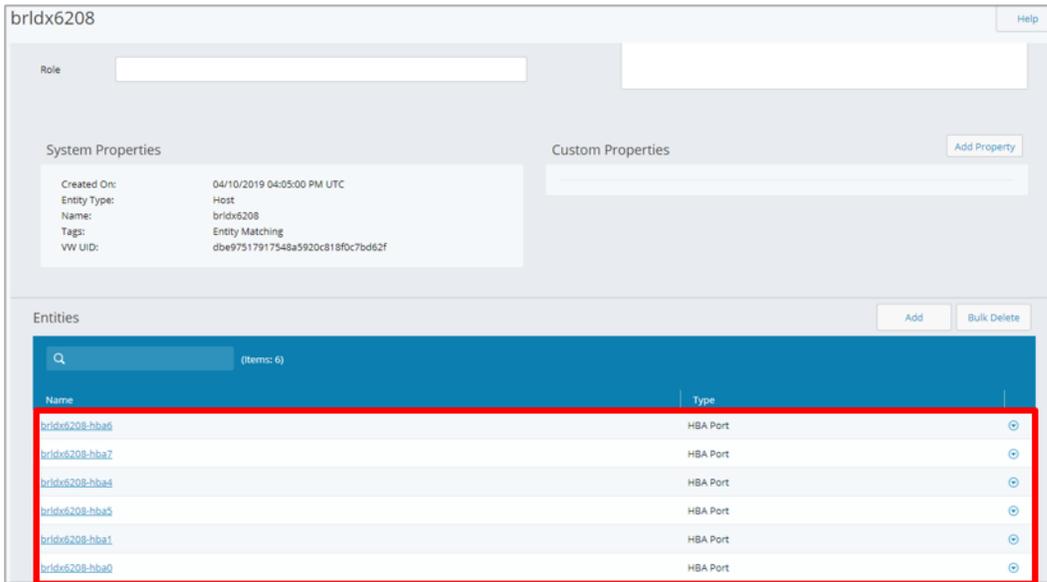
Confirming Entity Creation Using Inventory

Use the Inventory module to confirm entity creation.

The entity is tagged with "Entity Matching" as part of its system properties.



The sub-entities used to create the entity are displayed on the entity's inventory page.

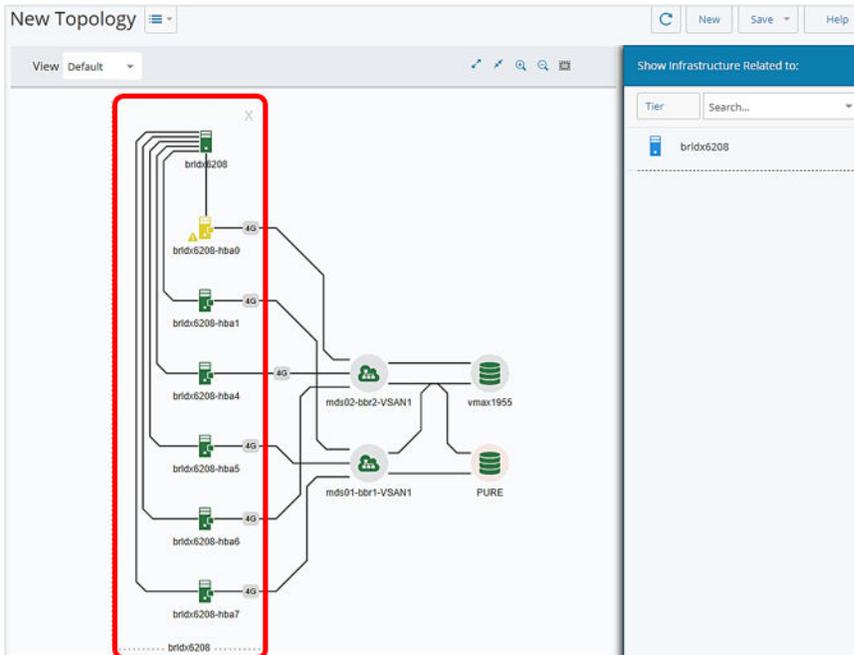


Confirming Entity Creation Using Topology

You can also use the Topology module to confirm entity creation, and to view the entity's relationships within the infrastructure.



Expand the host's topology to view the HBA ports and the relationships within the infrastructure.



Entity Matching Example

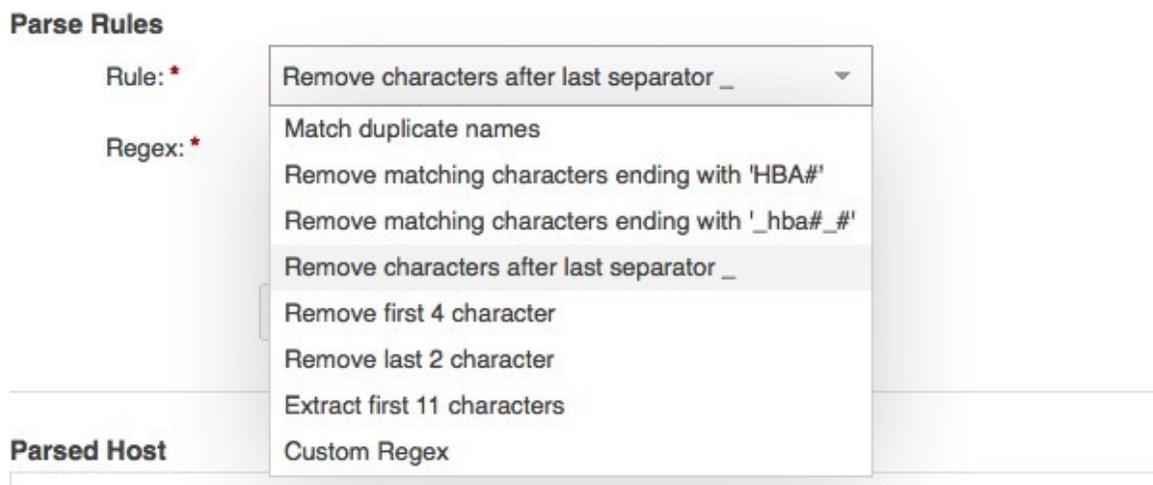
Here's an example of how the Entity Matching utility was used to create a host from HBA port aliases.

1. The Entity Matcher displays a list of unassigned HBA ports. Let's focus on the two HBA ports highlighted in the image below.

Unassigned HBA Ports (1328)

Name	Nickname	WWN
e50580		2022002a6a0cce80
e508a0		24fa002a6aaa8a00
esxAAA01094p01_hba0	esxAAA01094p01_hba0	20000025b511a1ba
esxAAA01094p01_hba1	esxAAA01094p01_hba1	20000025b511b2ba
esxAAA01094p02_hba0	esxAAA01094p02_hba0	20000025b511a1ca
esxAAA01094p02_hba1	esxAAA01094p02_hba1	20000025b511b2ca

2. It's clear from the naming conventions used that these HBA ports belong to a service named **esxAAA0109p01**. The objective is to use the provided parse rules to select these two HBAs to create the parsed Host name. Luckily, there's a rule that does just that. Select the rule called **"Remove characters after last separator _"** from the Rule pull-down as shown below.



Choosing this rule results in the regular expression (abbreviated Regex in VW):

`^(.*)_.*$`

Let's translate this rule:

^ Starting from the beginning of the line

(.*) Match any arbitrary sequence of characters: . matches any single character, * matches zero or more of what precedes

Until an underscore character is found

.* Followed by any arbitrary sequence of characters

\$ Until end of line is encountered

3. After clicking on the **Parse** button, the rule is added to the **Parsed Host** list under the **Parse Rule** selector area. Click on the **View Target Group** to expand the group.

Type

Parse Rules

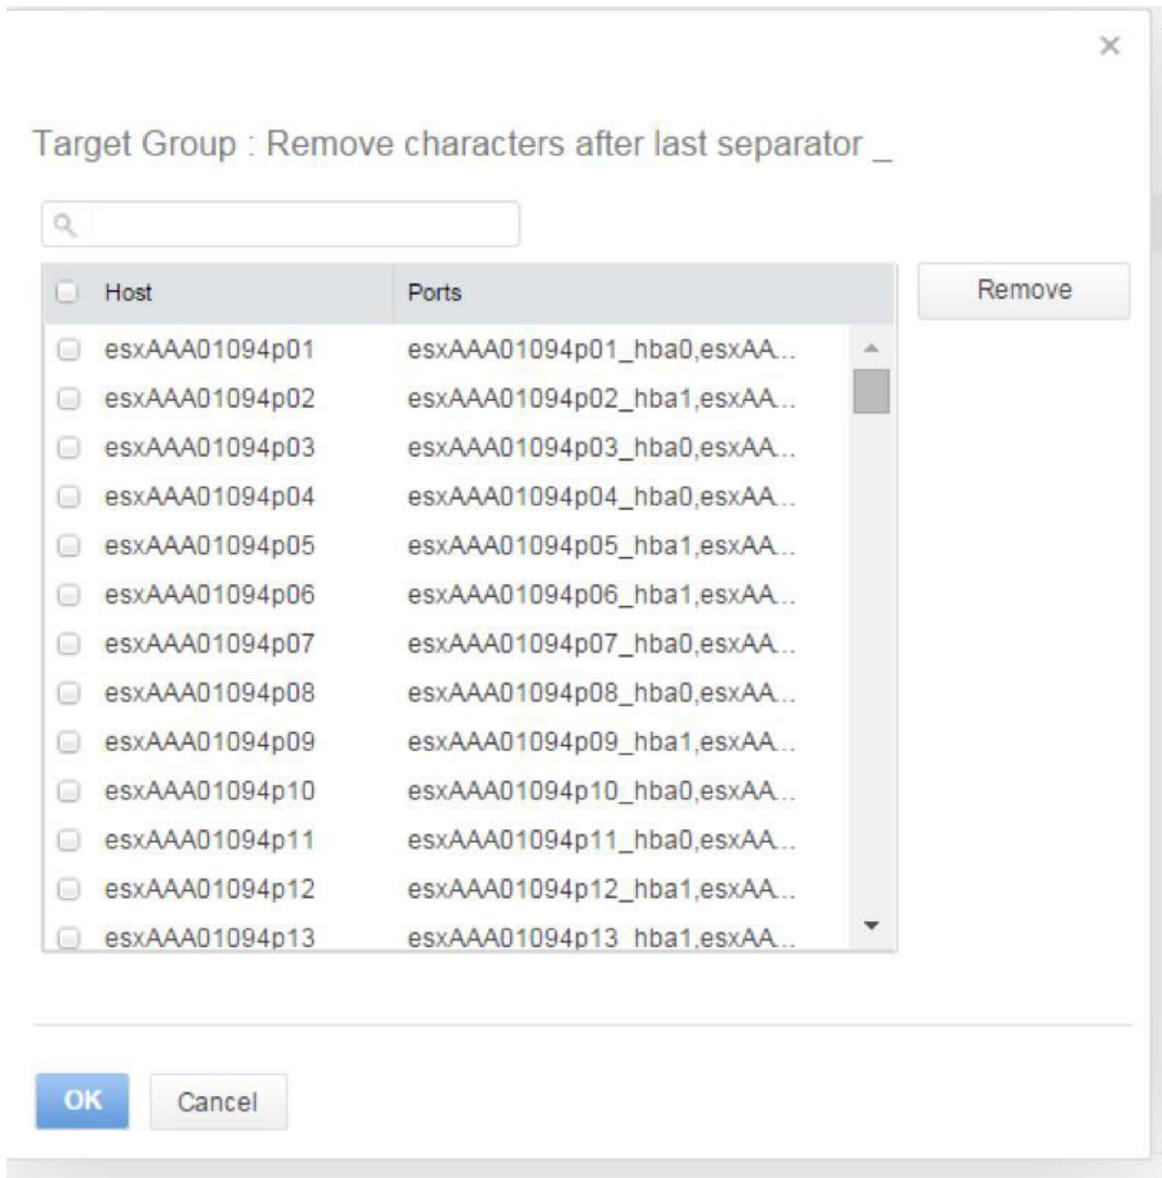
Rule

Regex

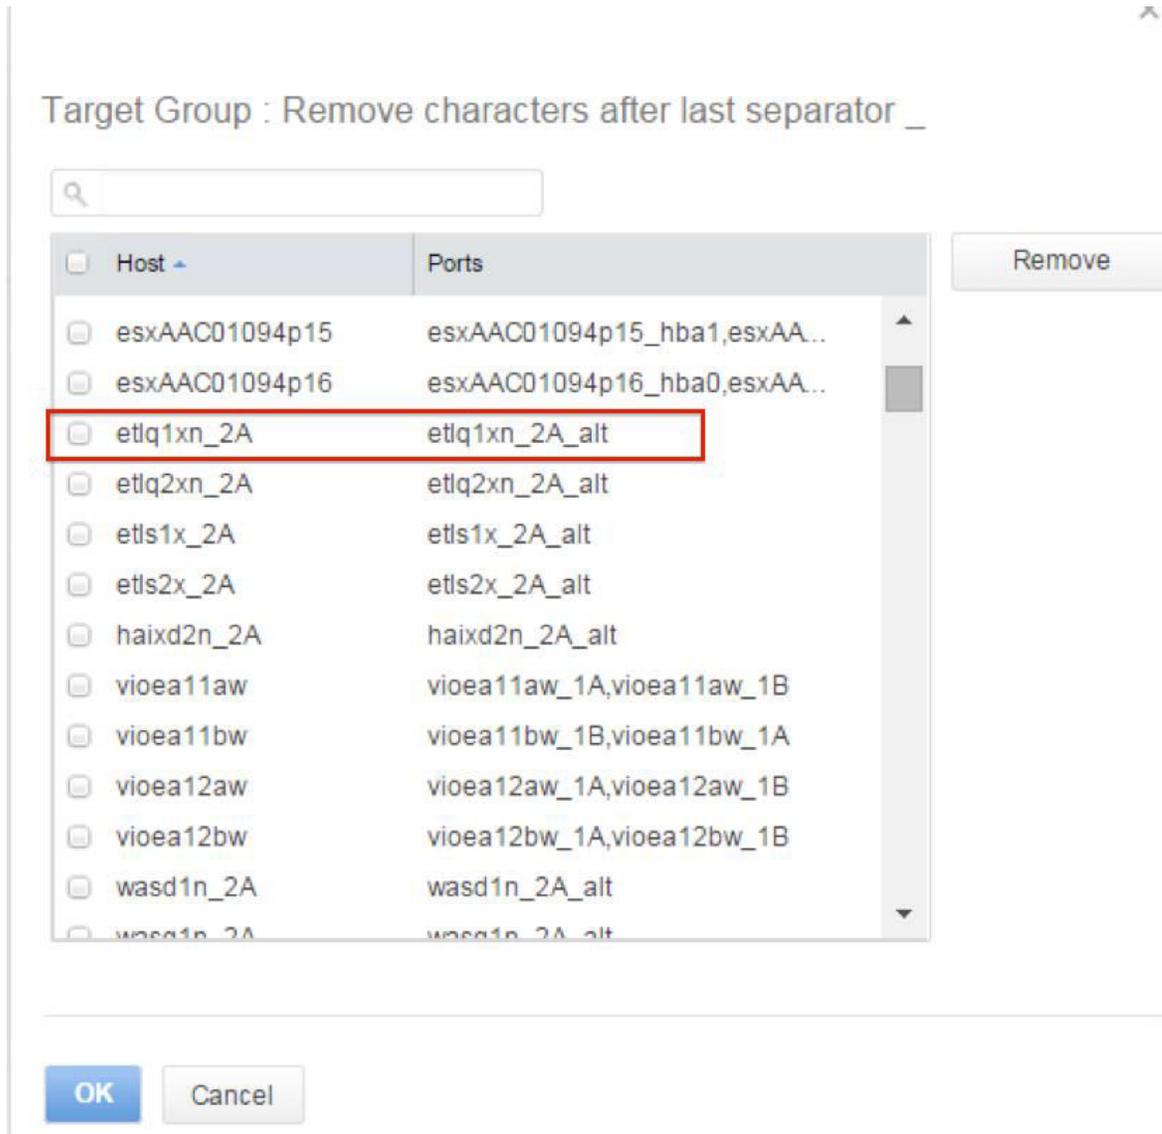
Parsed Host

Remove characters after last separator _ (603) - **Errors Found:** [425](#)

4. A list consisting of Host entity names, each with a list of HBA ports that the entity will contain, is displayed.



- Review the generated list and confirm that each entry is correct. It's likely that there may be a few host entities that were constructed incorrectly, usually resulting from inconsistent nicknaming. If you go ahead and click OK without reviewing, there may be a set of host entities created that are not correct and that will have to be removed. Here's an example: look at the host entity named **etlq1xn_2A** highlighted in the image below.



It contains a single HBA nicknamed **etlq1xn_2A_alt** which is odd. It's unlikely that any host in this day and age has only a single HBA port in it. There are numerous reasons why VirtualWisdom may not have discovered a second port:

- It didn't have a nickname so it appears in the list as an FCID
- It isn't connected to anything; hence VirtualWisdom couldn't discover it

This line should be excluded until you can resolve the anomaly. Remove the line by clicking on the checkbox next to the name.

6. Once you've finished reviewing the list, you click the **Remove** button to remove suspicious entries. Now you can click the **OK** button and then the **Create Host Entities** button and all of the host entities in the target group(s) will be created.

Custom Regular Expressions

If your organization uses a less rigorous approach to naming your ports, you will have to work a bit harder to use Entity Matching to create Host or Storage entities. Here's an example of two HBA ports belonging to the same virtual server:

```
Sto_sjctxesx1_PROD-fc-HBA-1-lab_HBA0_New Sto_sjctxesx1_PROD-fc-HBA-1-lab_HBA1_New
```

The rule we used above won't work here. It would create two hosts named **Sto_sjctxesx1_PROD-fc-HBA-1-lab_HBA0** and **Sto_sjctxesx1_PROD-fc-HBA-1-lab_HBA1**, each with a single port, which is not correct. We want a single host called **Sto_sjctxesx1_PROD-fc-HBA-1-lab** with two HBA ports (...HBA0_New and ...HBA1_New) in it.

What we need to do is modify the Regex for that rule.

1. Start by selected the rule as before, but instead of applying it to the list of unassigned HBA ports, simply copy the Regex field into your edit buffer (ctrl-c or cmd-c).
2. Select Custom Regex and paste in the copied Regex.
3. Edit as shown below and the rule will now match the HBA ports.

Parse Rules

Rule: *

Custom Regex

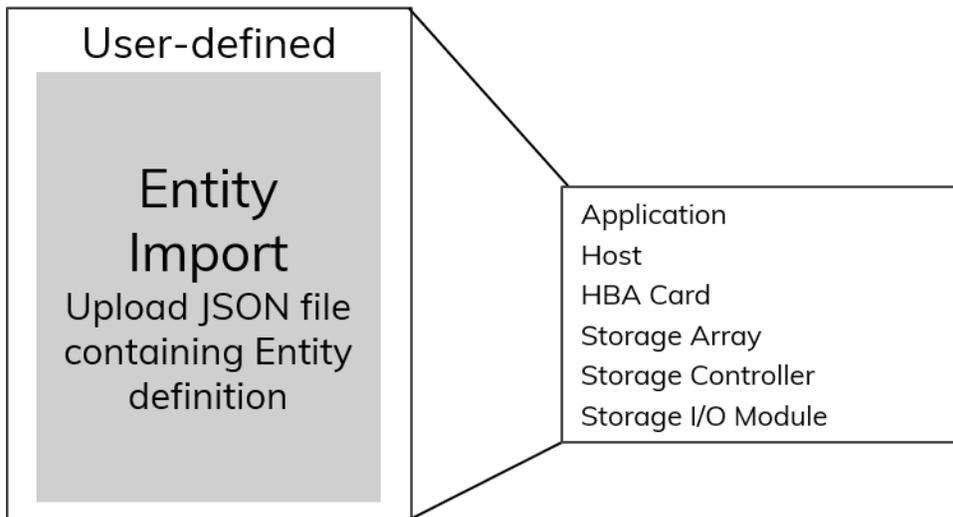
Regex: *

$^(.*)_.*\boxed{.}.*\$$

Parse

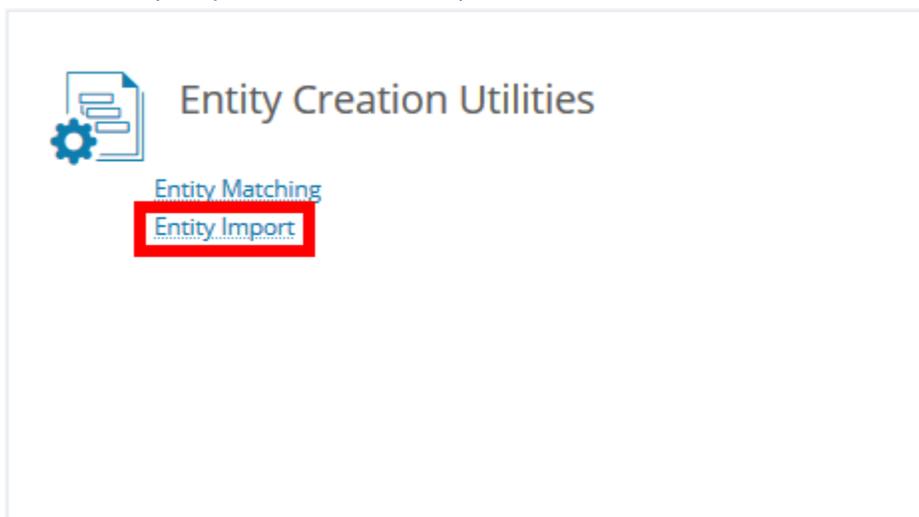
Entity Import

You can use an imported JSON file to create entities in bulk. You can also use this method to create Application entities and other entity hierarchies:

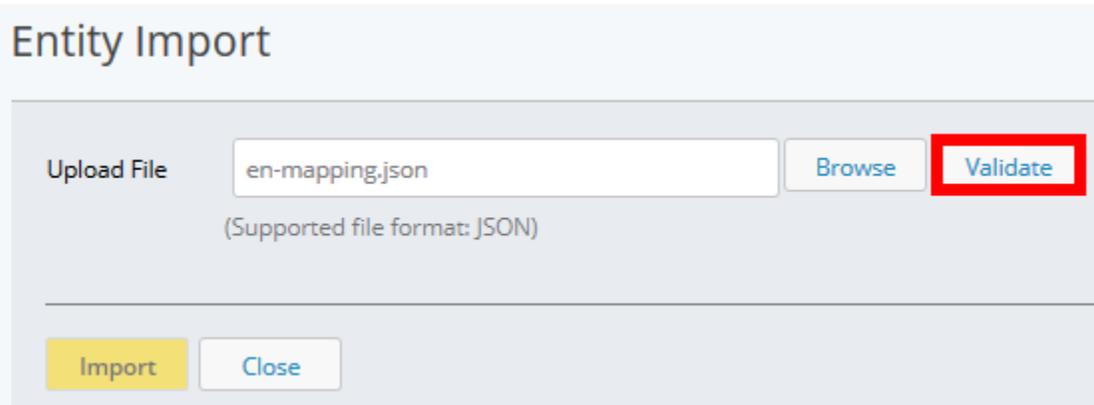
**NOTE**

This task is available only to users with the VirtualWisdom Administrator (**vw-admin**) role.

1. Select Entity Import from the Entity Creation Utilities section on the Settings page.



2. Upload your JSON file and click Validate to check it for errors. Select Import to create the entities.



Entity Import

Upload File

(Supported file format: JSON)

See the VirtualWisdom Administrator Guide, [Entity Import](#) section, for more information on this feature.

Entity Management Best Practices

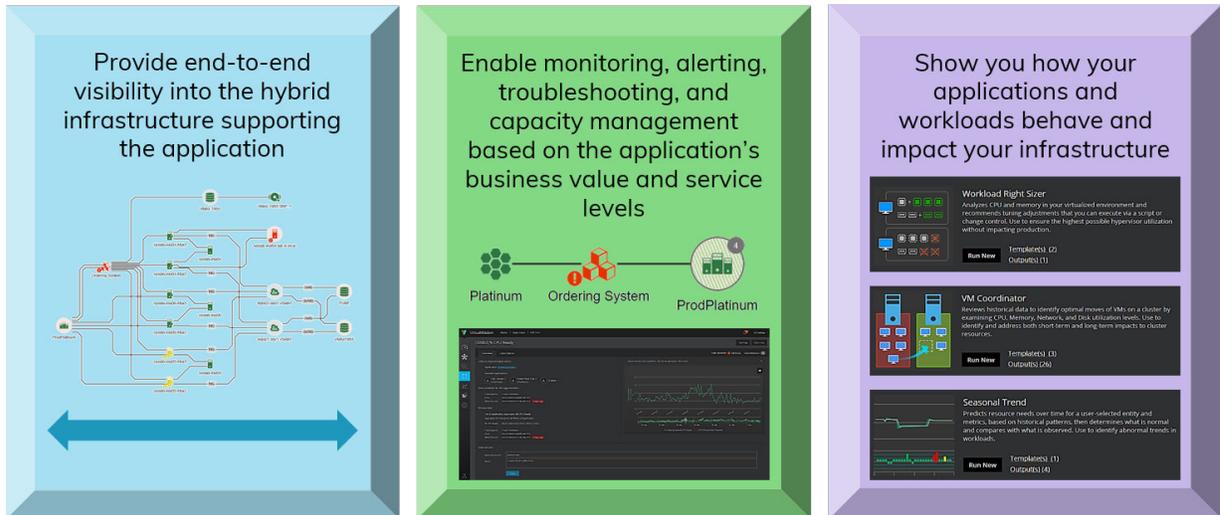
1. Use the entity matching utility to capture newly discovered devices.
2. Perform entity matching at least weekly.
3. Schedule a 10-minute task on Mondays to capture any weekend changes.
4. Perform post-switch discovery if discovery updates take multiple days to complete.
5. Validate entities using reports:
 - a. All hosts are presenting data as expected.
 - b. All conversations are captured.
 - c. Metrics are displayed (use report metrics appropriate for the integration, e.g., Consumed Bandwidth for switch integration).

Which Method Should I Use to Create Entities?

When to Use Entity Matching	When to Use Entity Import
The most commonly used expressions and examples apply to your environment.	Your organization uses a well-defined CMDB or other source that has export to .csv file capabilities.
Your organization used well-defined naming conventions and zoning practices.	When aliases are not automatically discovered (Cisco example FC Alias, zoned by interface).
Your storage ports are clearly defined.	Storage Virtualizer is used for the initiator defined ports (depending on architecture).
	Your organization uses complex naming conventions or aliasing on switches that Regex wouldn't be efficient to use.

Applications and Tiering

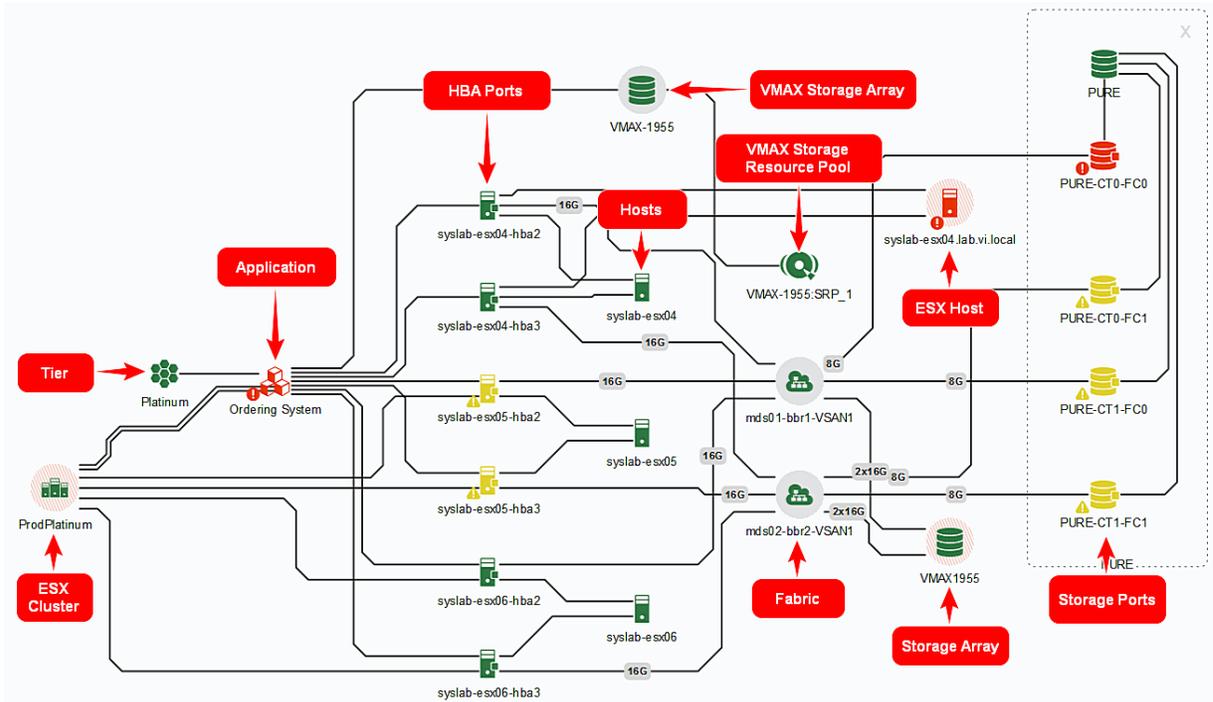
Application Entities



Application entities are groupings of the infrastructure components that comprise and support the application.

Discovering or creating application entities is of critical importance in achieving the most value from the VirtualWisdom platform.

Application entities provide end-to-end visibility into the infrastructure where the application lives: the compute, network, and storage infrastructure.



Application entities also enable VirtualWisdom to provide monitoring, alerting, case management, and troubleshooting based on the application's business value and SLA tier.

Finally, application entities enable you to understand how your applications behave and impact the infrastructure.

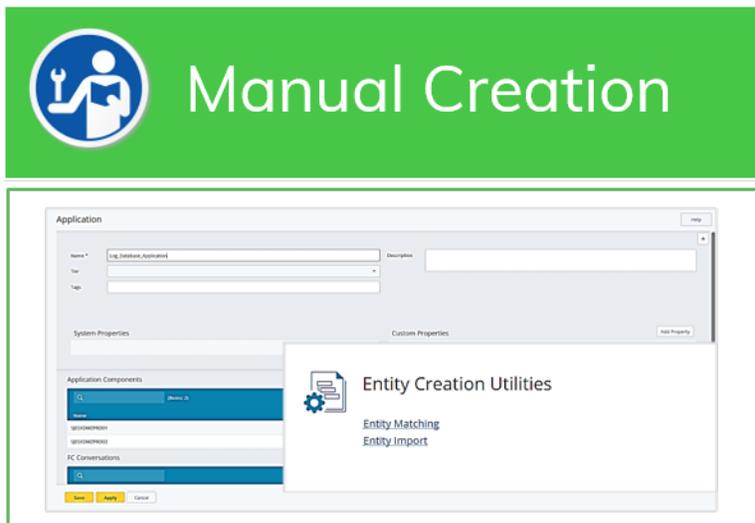
Application Entity Creation

Applications can be discovered and created automatically from an **AppDynamics**, **Dynatrace**, or **ServiceNow** instance, or through suggestions populated by VirtualWisdom through the **VMware** integration, the **Host OS** integration, and through **NetFlow**.

The entities that make up the application must already be part of the infrastructure discovered and monitored by the VirtualWisdom probes and integrations.



Alternately, application entities can be created manually using the VirtualWisdom interface or through import using a JSON file.



Manual Application Entity Creation



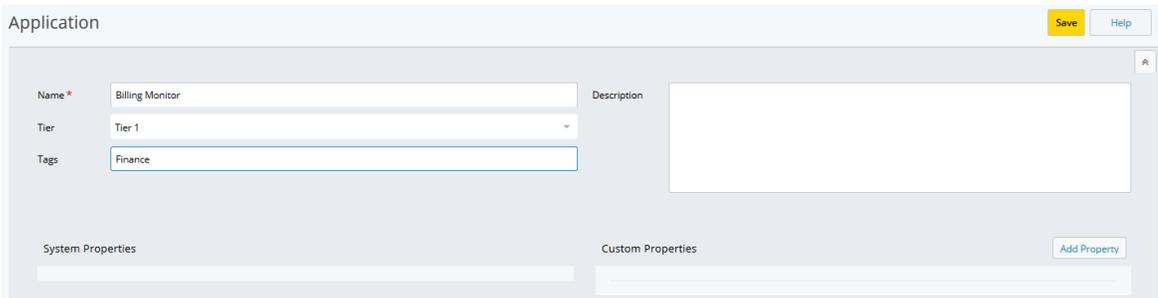
NOTE

Application entities are created top-down. Other entity types are created bottom-up, and child entities are created before their parents. If an entity is created before it is populated with children, it shows an empty topology. When you add children to an entity, the metric associated with those children are applied retroactively. This capability is useful for reports and charts because you can see the historical metrics of the child entities after the creation of the parent.

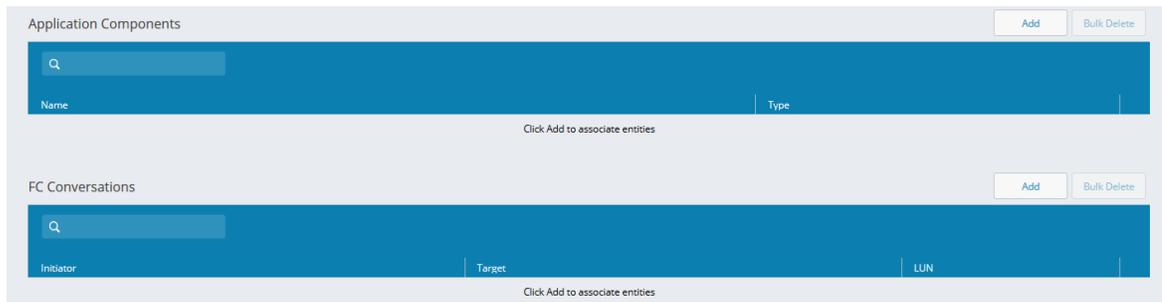
1. From the **Application** list in the **Inventory** module, select **New** then **Application**.



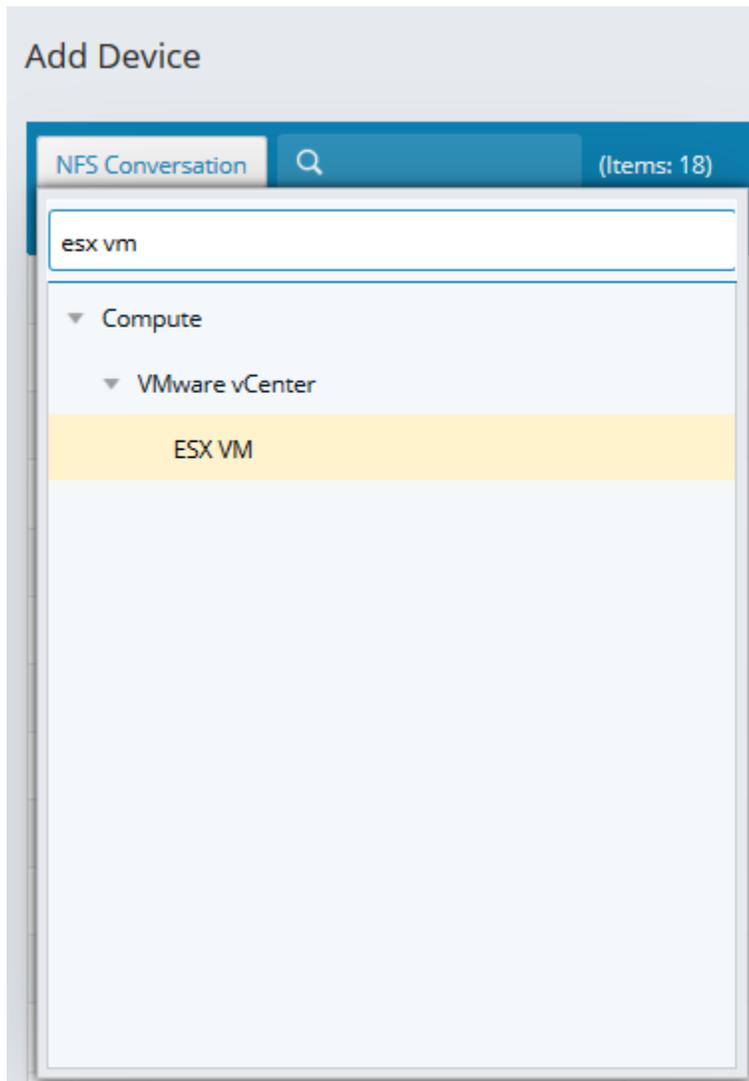
2. Complete the application properties fields.



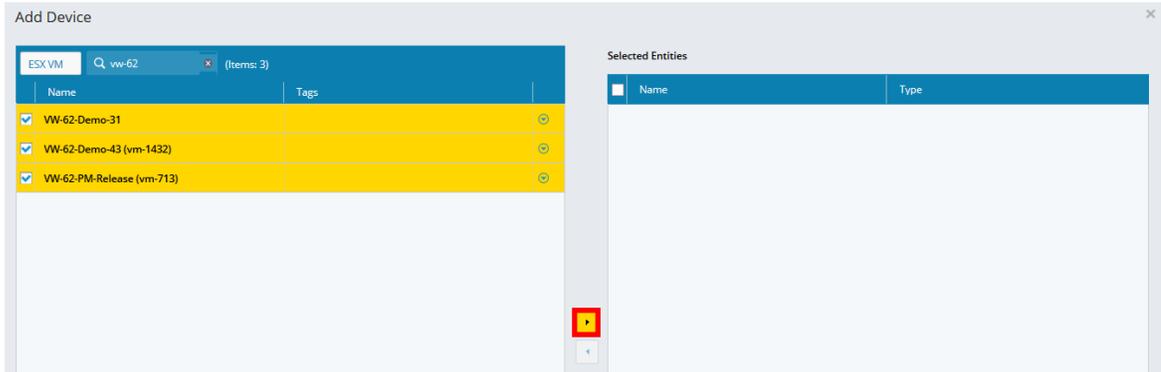
3. Select the application components or conversations that comprise the application.



4. Choose the entity type.

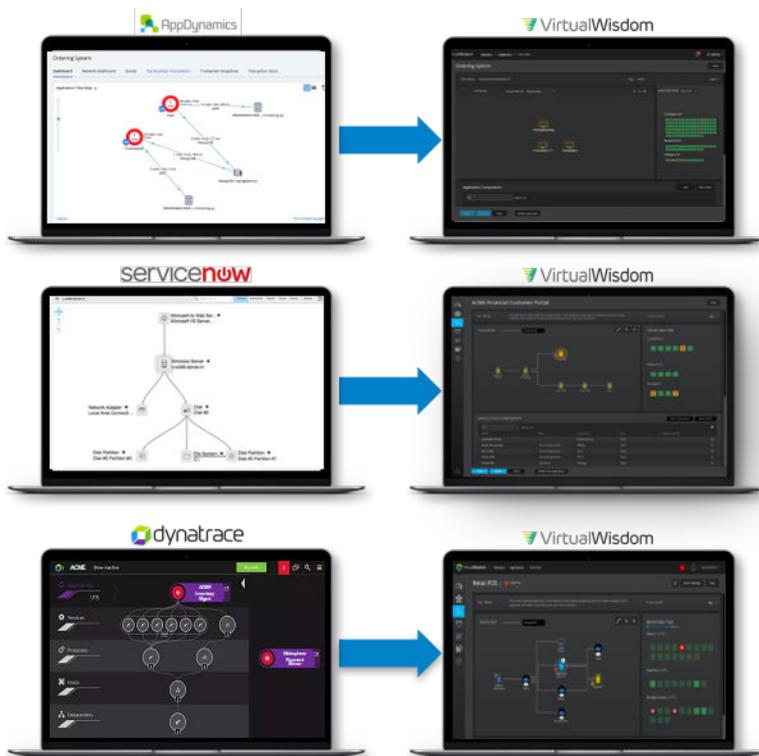


5. Select named entities by checking the boxes next to their names. Use the drop down menu to view their properties or topology. Move the selected entities to the right panel using the right arrow button. Click OK when you are satisfied with your selection.



Application Entity Discovery

Auto-Discovered Application Entities



The AppDynamics, ServiceNow, and Dynatrace integrations automatically discover applications. VirtualWisdom application entities are created automatically for these integrations. VirtualWisdom application entities are created automatically for these integrations.

VirtualWisdom also interprets AppDynamics tiering strategies, allowing you to map the AppDynamics tiers to VirtualWisdom tiers.



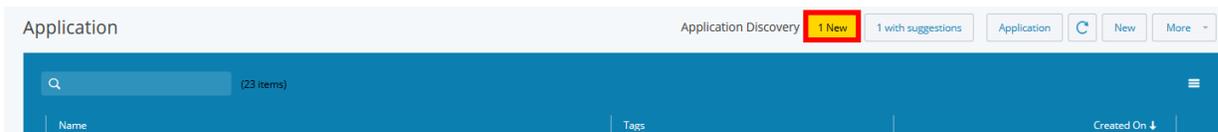
NOTE

Application entities (IP address, host, ESX VM, Hyper-V VM, PowerVM partition) must already exist in VirtualWisdom prior to application discovery through the application, OS, or NetFlow integrations.

Application Entity Suggestions

The Host OS and NetFlow integrations suggest new applications.

A button on the Application inventory page tells you when new suggestions are available.



Drill down on the suggested application to view its properties.

1 New Application Discovered

Entity Name	Created On	Components
Application--(569659952)	06/26/2020 09:00:00 AM UTC	3

[Done](#)

Suggested Application Page

Properties discovered by the integration are displayed under System Properties.

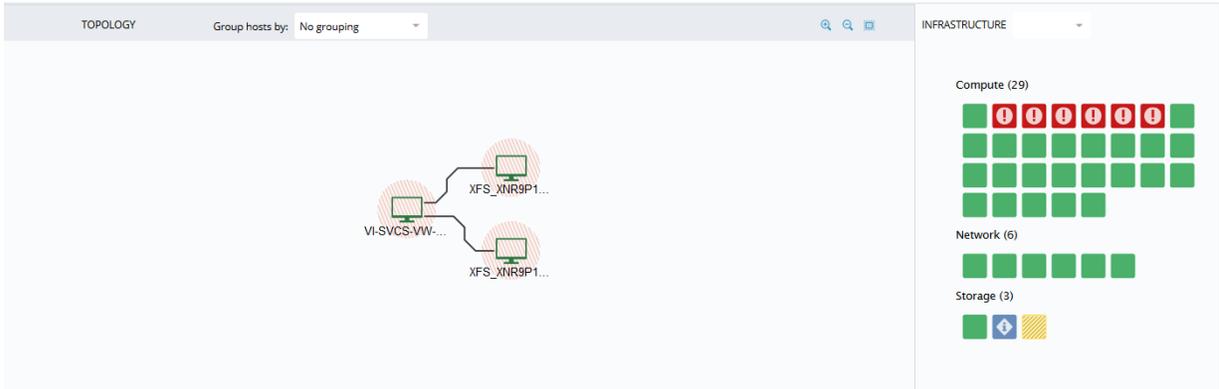
Suggested Application Create Application Ignore Suggestion Help

Name * Description

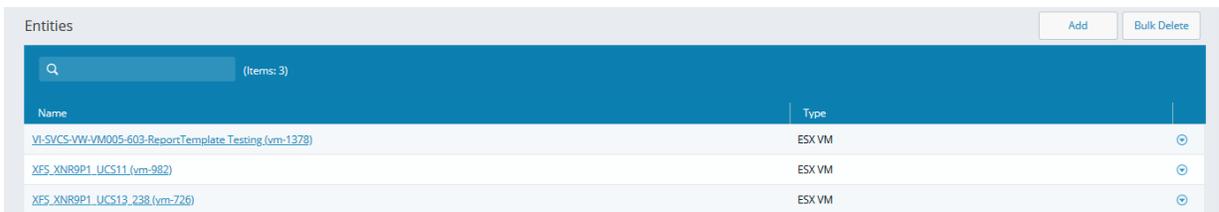
Tier

Tags

The topology of the suggested application is shown with the ability to zoom in or out. An Infrastructure map is also included to the right. Highlighting an entity in the topology map filters the infrastructure map for that entity.



The sub-entities that comprise the application are shown along with their type and a drop down menu to view more info. You can add or delete entities using the buttons at the top right of the list.



Use the Create Application or Ignore Suggestion buttons at the top right of the page to create the app or ignore this suggestion. The suggestion is no longer shown.



Suggested Application Changes

A button on the Application inventory page tells you when changes have been made to existing applications.



Drill down on the application to view its suggested changes.

Found Suggestions for 1 Application ✕

Application	Source	Discovered On
Service Mgmt	VirtualWisdom	07/08/2020 09:01:00 AM UTC

Service Mgmt

Done

Review the suggested changes, then select which changes to apply or ignore all changes.

Service Mgmt - Suggested Changes ✕

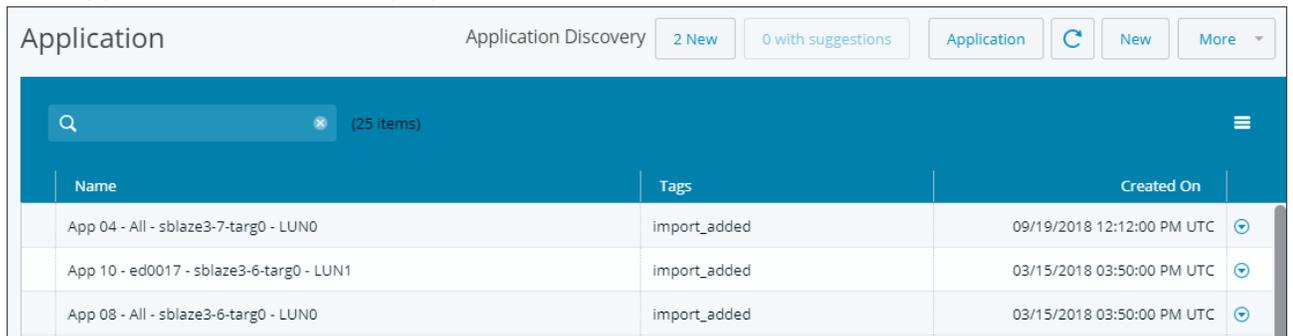
<input type="checkbox"/>	Entity Name	Type	Suggested Change
<input checked="" type="checkbox"/>	TEST-PA-mfg-ova-vw.6.6.0-13.x86_64 (vm-1717)	VirtualMachine	Add to Application
<input checked="" type="checkbox"/>	VI-SVCS-Prod-VW-221d (vm-1493)	VirtualMachine	Add to Application
<input type="checkbox"/>	UCS-11_IP_Only	Host	Add to Application
<input type="checkbox"/>	VI-SVCS-SCRIPTING-VW640 (vm-1641)	VirtualMachine	Add to Application
<input type="checkbox"/>	SVCS-W10-VM05-VitooS (vm-1179)	VirtualMachine	Add to Application

Apply Selected Changes
Cancel
Ignore All Changes

Resolving Application Overlap with Conflict Management

Conflicts can occur only between a ServiceNow-discovered app and an application discovered by either the Operating System Integration or NetFlow integration. If a conflict occurs between Operating System integration and NetFlow Integration applications, the conflicts are merged into a single suggested application. If a suggested application is ignored, it is ignored for 30 days; however, that time period is configurable.

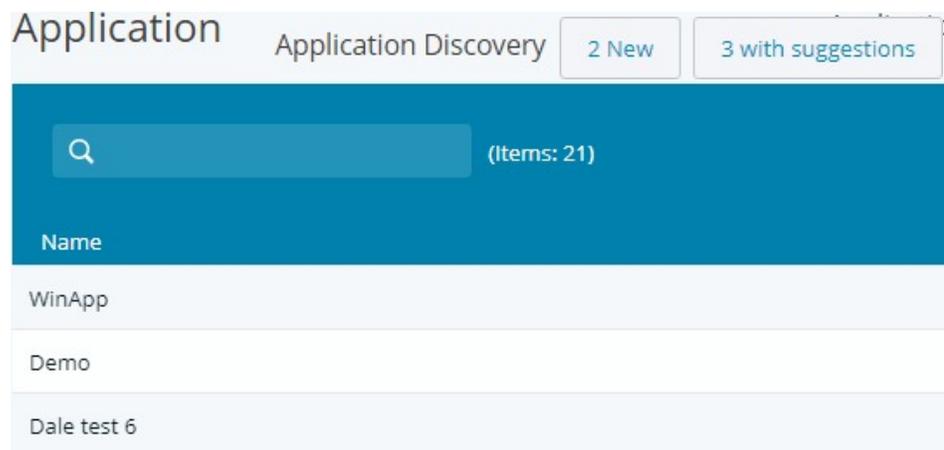
1. Click **Inventory** and then *Application*.
The *Application* screen is displayed.



The screenshot shows the 'Application' screen with 'Application Discovery' selected. It features a search bar with '(25 items)' and a table with the following data:

Name	Tags	Created On
App 04 - All - sblaze3-7-targ0 - LUN0	import_added	09/19/2018 12:12:00 PM UTC
App 10 - ed0017 - sblaze3-6-targ0 - LUN1	import_added	03/15/2018 03:50:00 PM UTC
App 08 - All - sblaze3-6-targ0 - LUN0	import_added	03/15/2018 03:50:00 PM UTC

Discovered applications are also displayed by clicking either the **<#> New** or **<#> with suggestions** button.



The screenshot shows the 'Application' screen with 'Application Discovery' selected. It features a search bar with '(Items: 21)' and a list of suggested applications:

Name
WinApp
Demo
Dale test 6

2. Click one of the *Application Discovery* buttons.
The **<#> New Applications Discovered** or **<#> Updated Applications Discovered** dialog box is displayed. Suggested applications might have missing or extraneous components.

2 New Applications Discovered

Entity Name	Created On	Components
Application--(MYSQL-APACHE_HTTP)	09/14/2017 10:42:00 AM EDT	4
Application--[MYSQL, GENERIC_WEB_SERVER]	09/14/2017 10:42:00 AM EDT	2

3. Click the suggested application to be examined for conflict resolution.

SuggestedApplication

Search (Items: 5)

- Application--(MYSQL-APACHE_HTTP)
- Application--(TOMCAT-POSTGRES-GENERIC_WEB_SERVER-REDIS)
- Application--(TOMCAT-POSTGRES-REDIS)
- Application--[MYSQL, GENERIC_WEB_SERVER]
- Application--(POSTGRES-GENERIC_WEB_SERVER-REDIS)

Detailed information about the application is displayed.

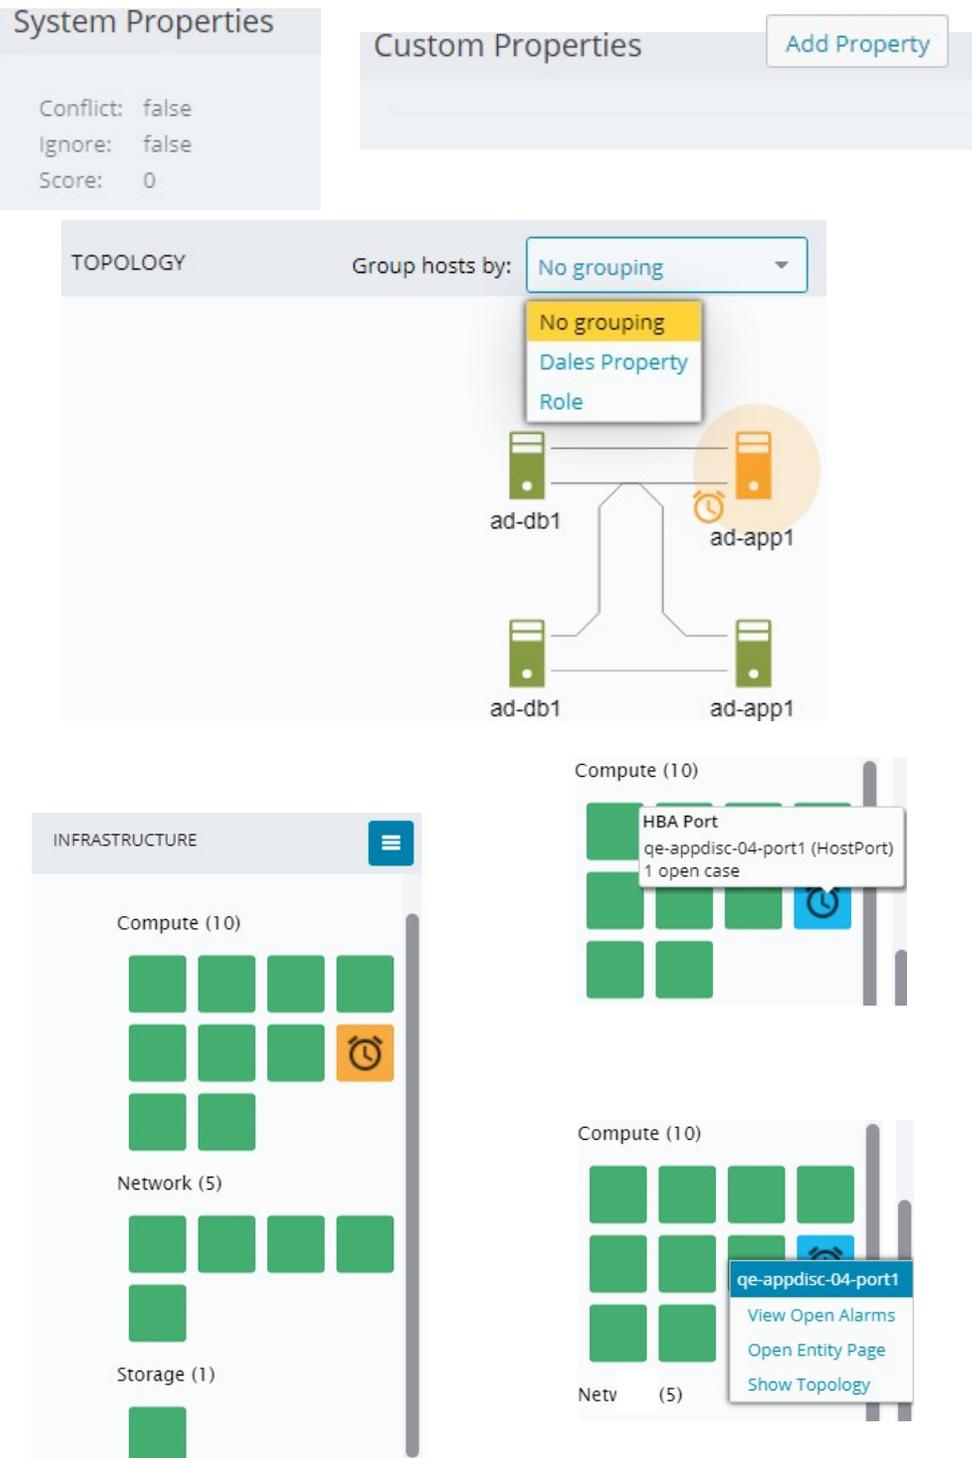
Suggested Application

Name * Application--(MYSQL-APACHE_HTTP)

Tier use tier from ServiceNow

Tags

This information includes system properties, custom properties, topology and infrastructure.



Entity information is also displayed, and entities can be added or deleted.

Entities Add Bulk Delete

(Items: 4)

Name	Type
ad-app1	ESX VM
ad-app1	Host
ad-db1	Host
ad-db1	ESX VM

- You can create or ignore a suggested application, apply all or selected changes, or ignore all changes.

Create Application
Cancel
Ignore

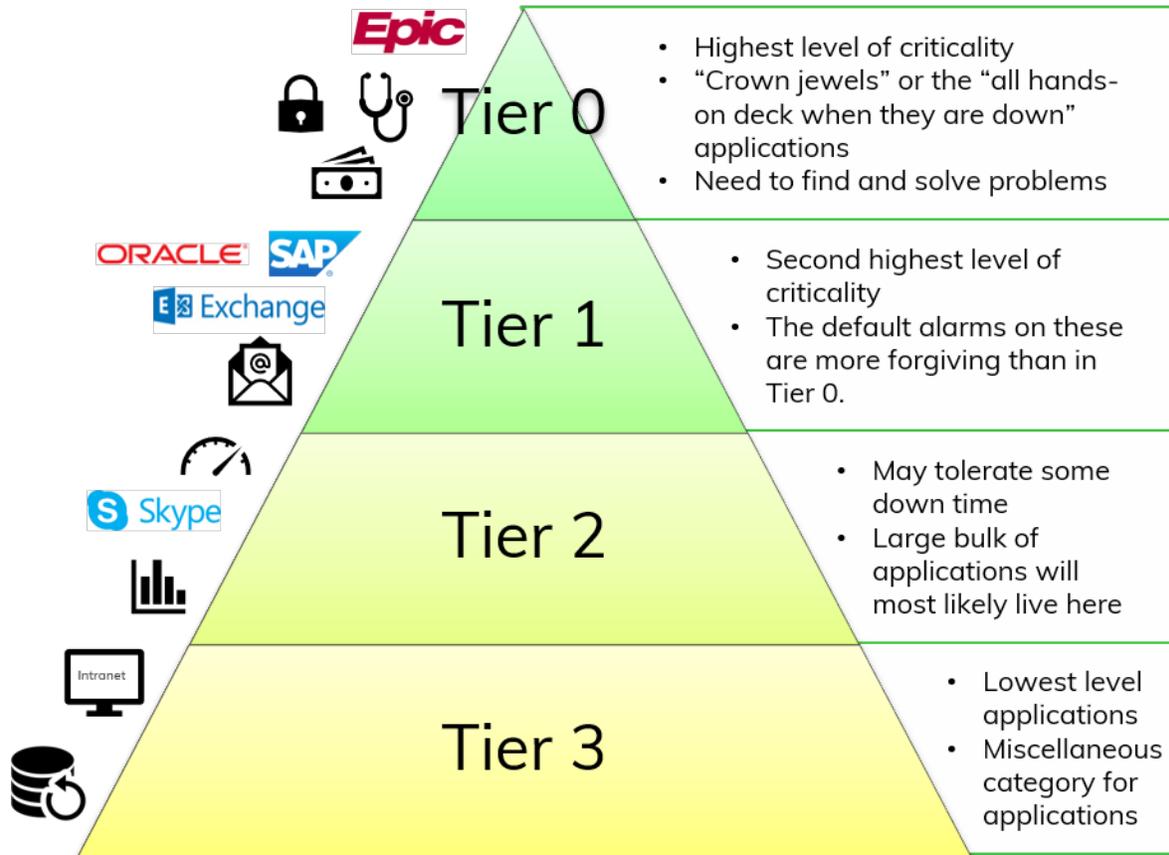
WithITLPattern - Suggested Changes ×

<input type="checkbox"/>	Entity Name	Type	Suggested Change
<input type="checkbox"/>	ubuntu-16-worker3	Host	Add to Application
<input type="checkbox"/>	QE-Swarm-03	VirtualMachine	Add to Application
<input type="checkbox"/>	ubuntu-16-worker4	Host	Add to Application

Apply Selected Changes
Cancel
Ignore All Changes

Apply Selected Changes
Cancel
Ignore All Changes

Application Tiering



Application Tiering is a Business Continuity / Disaster Recovery (BCDR) concept: In the event of a failure, how long can the application be down (RTO), and how much data can be lost (RPO)?

For example, Tier-0 apps cannot be down and cannot lose data so Active/Active failover with sync replication is required; whereas, a Tier-2 app may be able to sustain 8 hours down and 1 hr of data loss.

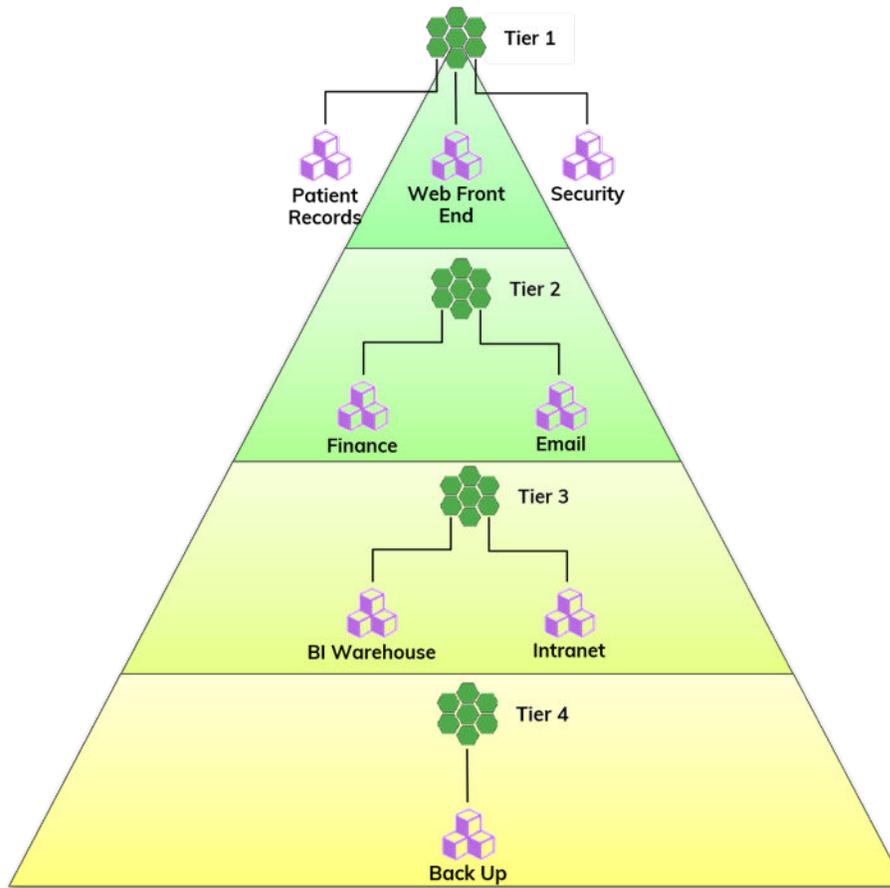
Applications can be added to tiers in VirtualWisdom. When a new application is created in the user interface, the user can select a tier right after choosing a name.

Applications discovered from ServiceNow will have a tier assigned to them based on the Business Criticality value associated with the Business Service. This mapping can be modified in the ServiceNow integration configuration.

Alarms are associated with tiers, allowing you to quickly pinpoint issues affecting your most critical applications first.

Application tiers in VirtualWisdom

Tiers are entities in VirtualWisdom.



Applications are assigned to VirtualWisdom tiers automatically and manually.

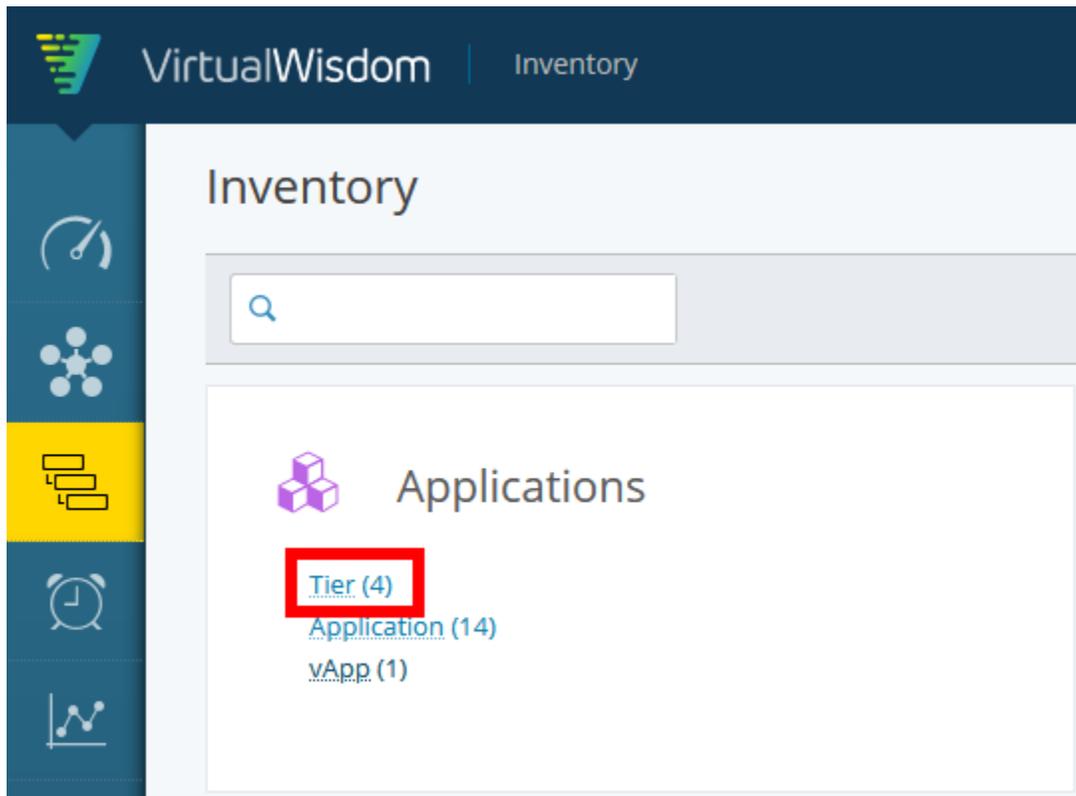
There are four default tiers in VirtualWisdom, with assigned rankings from 1 to 4, where 1 is the highest business priority and 4 is the lowest.

VirtualWisdom also includes pre-configured alarms and dashboards that are based on the application tiers.

Create Tiers

You can create hierarchical tiers that specify the importance of your applications.

1. Click **Inventory** and select *Tier* from the *Applications* group.



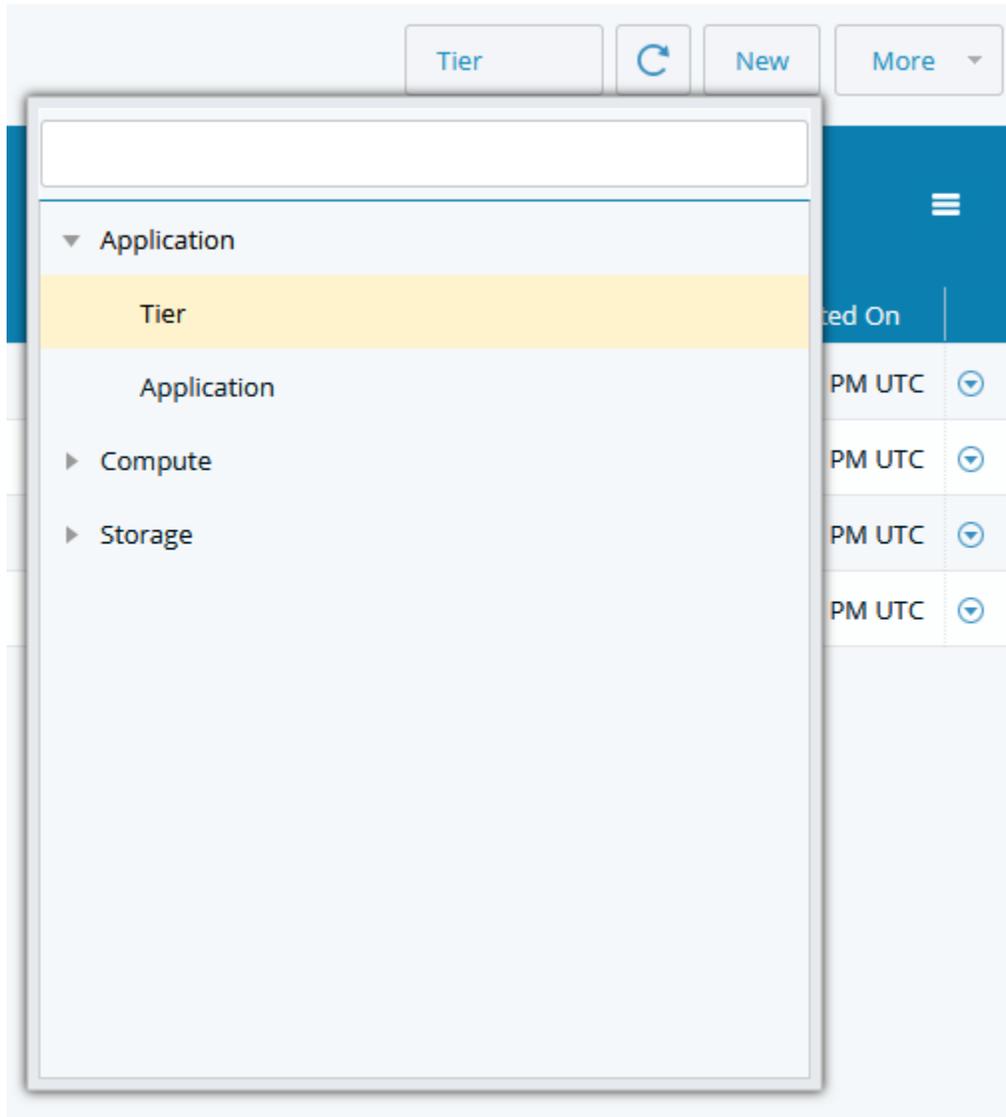
The list of existing tiers is displayed: the four VW-defined tiers are Tier 0, 1, 2, and 3.

Tier Tier

(4 items) ☰

Rank ↑	Name	Tags	Application Count	Created On
☰ 1	Tier 0		4	03/25/2020 11:46:00 PM UTC <input type="button" value="⌵"/>
☰ 2	Tier 1		5	03/25/2020 11:46:00 PM UTC <input type="button" value="⌵"/>
☰ 3	Tier 2		3	03/25/2020 11:46:00 PM UTC <input type="button" value="⌵"/>
☰ 4	Tier 3		2	03/25/2020 11:46:00 PM UTC <input type="button" value="⌵"/>

2. Click **New** and select *Tier* from the Applications menu.



The create entity screen is displayed.

Tier

Save Help

Name* Platinum Description

Tags

Tier Ranking Ahead of all other tiers

Create new tier from scratch
 Clone this existing tier:

System Properties

Custom Properties Add Property

Applications Add Bulk Delete

Name	Type
------	------

- a. You can create a brand-new tier or clone an existing tier.
 - b. You can also specify the *Tier Ranking*.
 - c. You can add *Custom Property / Value* pairs to the tier, which can then be used as identifying factors: for example *Server / Chicago*.
 - d. You can *Add Applications* to the tier.
3. Click **Save**.

VirtualWisdom Metrics

VirtualWisdom collects metrics related to health, performance, and capacity/utilization. Every entity type has a pre-defined set of metrics that are collected from your infrastructure.

These metrics are used throughout VirtualWisdom to provide insight into infrastructure health, utilization and performance and to alarm when problems are detected or are imminent.

Chapter 2, Integrations Brocade and Cisco SAN Integrations (ProbeSW)

Brocade and Cisco SAN Integrations (ProbeSW)

Path	Metric	Description
Storage/SAN/FC Switch Integration	Receive Utilization	Link utilization expressed as a percentage of total link capacity.
	Transmit Utilization	Data transmitted by the switch port expressed as a percentage of available link capacity.
	ISL Channel-A Utilization	Link utilization for an ISL expressed as a percentage of total link capacity.
	ISL Channel-B Utilization	Link utilization for an ISL expressed as a percentage of total link capacity.
	PC Channel-A Utilization	Port Channel utilization expressed as a percentage of total channel capacity.
	PC Channel-B Utilization	Port Channel utilization expressed as a percentage of total channel capacity.
	Receive Consumed Bandwidth	Data transfer rate in bytes/second.
	Transmit Consumed Bandwidth	Transmit throughput rate in bytes/second observed on switch port.
	ISL Channel-A Bandwidth	Throughput rate in bytes/second observed on an inter-switch link.
	ISL Channel-B Bandwidth	Throughput rate in bytes/second observed on an inter-switch link.
	PC Channel-A Bandwidth	Throughput rate in bytes/second observed on a port channel.
	PC Channel-B Bandwidth	Throughput rate in bytes/second observed on a port channel.
	Transmitted CRC Errors	Count of frames detected by the switch port with Cyclic Redundancy Check (CRC) errors. Detected inbound to the switch, CRC errors generally indicate a physical layer issue. In the case of Brocade (post FOS 6.3) these errors represent the crc_err_counter.
	Class 3 Discards	Count of Class 3 Discards observed. Fibre Channel Class 3 transport does not guarantee delivery of frames and if the switch cannot deliver a frame due to timer expiration or destination becoming unavailable, the frame is discarded by the switch without notifying the sender or receiver. The count is a sum of frames discarded at ingress and egress from the port, so no inference can be made as to the direction the discarded frame was travelling.
	% B2B Credits	Normalized buffer credit starvation metric for both Cisco and Brocade switch ports.

How are metrics collected?

Metric collection depends on the integrations and probes configured in your VirtualWisdom portal. Metric collection intervals may be different for each integration.

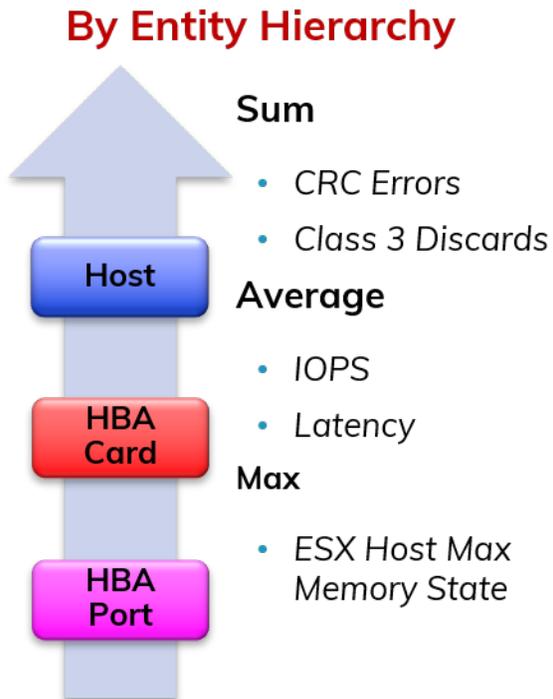
The following resources are available for reviewing and understanding metrics:

- Integration Guides
- Metrics Guide

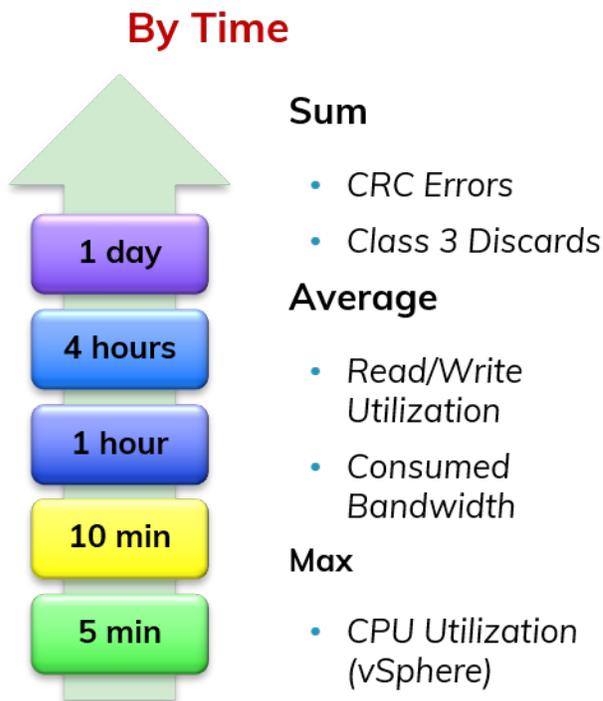
Metric Data Aggregation

VirtualWisdom metric data is aggregated using entity hierarchies and by time.

We show you data rolled up from a group of HBA ports for a single host.



We also show you data rolled up using time intervals, e.g., 5 min, 1 hour, 1 day.



Data Persistence

VirtualWisdom metric data is collected at defined intervals, based on the probe or integration and the configured interval. The length of time metric data is kept in the VirtualWisdom portal depends on its granularity. For example, 1 to 10 second data is kept for 8 days, after which it is aggregated. This is important to understand in reporting. Report ranges that include a long period will show the most granular data (1 sec, 10 sec) for only the period of time shown here.

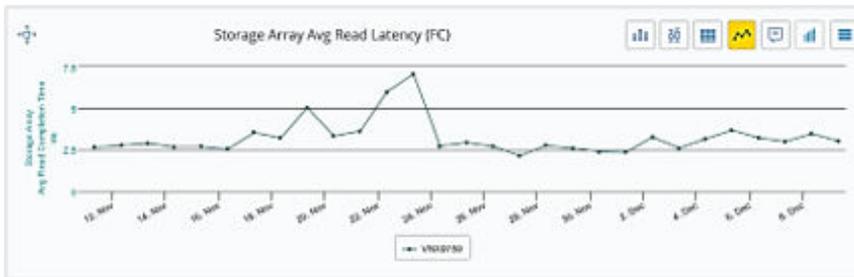
The table below shows the retention policy for the different summary levels.

Summary Level	Days of Retention
1 second	8 days
10 seconds	8 days
1 minute	16 days

Summary Level	Days of Retention
5 minutes	35 days
10 minutes	65 days
1 hour	184 days
4 hours	400 days
1 day	3660 days

Data Persistence Example

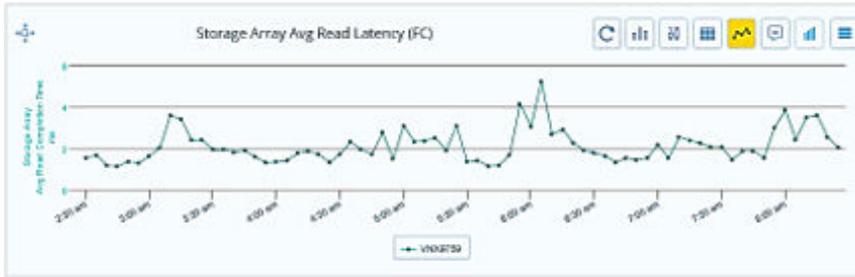
This chart shows Avg Read Latency metric data for the last 30 days (daily averaged value). We can drill down to view more granular data, but only in the context of granularity and the defined data persistence.



We can view 1 minute data for the last 16 days.



But only 5 minute data for the 19 days before that.



Infrastructure Monitoring Concepts

Health, Utilization, Capacity, and Performance

VirtualWisdom's value lies in its ability to proactively detect deviations from your normal levels of Health, Utilization, Capacity, and Performance and provide tools for investigation and remediation. VirtualWisdom provides visibility into metrics pertaining to infrastructure health, utilization, and performance



Health

The health of your application's infrastructure depends on its configuration, how it communicates with other devices, and the cleanliness of its physical layer. Issues and anomalies that impact any of these items can lead to availability and performance issues.

VirtualWisdom observes and reports on these common health issues:

- **Physical layer issues:** CRC errors, code violation errors, frame errors, loss of sync, loss of signal
- **Communication issues:** Class 3 discards, link resets, link failures, aborts, buffer-to-buffer credit exhaustion
- **Configuration issues:** Zoning, incorrect HBA queue depth settings, multi-pathing failure, cluster balance, bully VMs, zombie VMs

Utilization

Utilization is measured as a rate, e.g., MB/s, or a percentage of the available resources. High and low utilization can be a problem for performance, but it depends on multiple aspects. Measurement of one portion of the environment doesn't take into consideration of other bottlenecks that can be in the environment as well (ISL bandwidth, storage array limitations, bandwidth within the server, available CPU resources, etc.) which all can have an impact on overall performance.

Performance

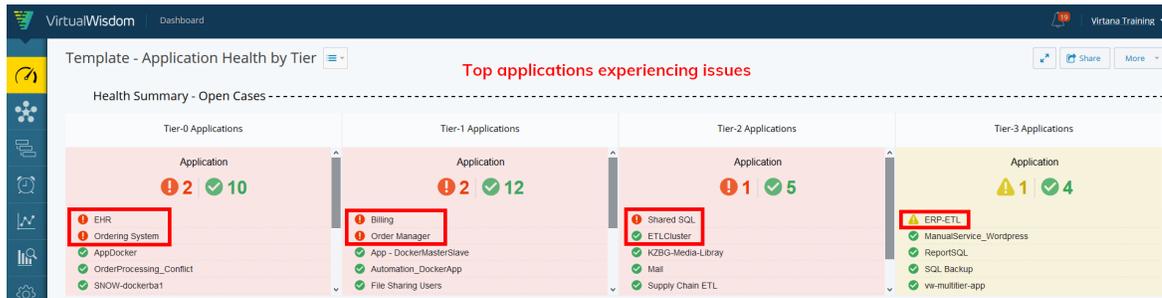
The performance of your application is dependent on the performance of your underlying infrastructure. CPU contention and memory pressure are two factors that can critically impact your application performance.

Using a Dashboard to Identify Application Infrastructure Issues

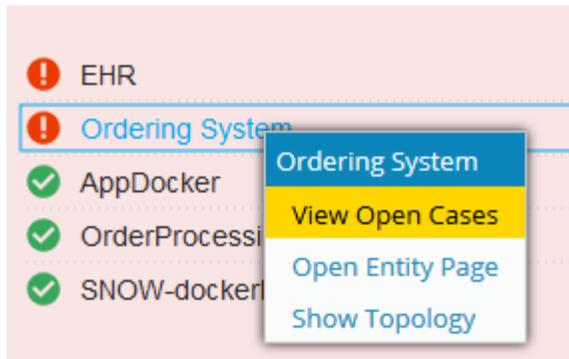
You can use VirtualWisdom to quickly identify issues in your infrastructure. To get you started, here is an example of how you can use a standard dashboard, the **Application Health by Tier** dashboard, to review the health of your application infrastructure, identify issues, and drill down and view related open cases and investigations.

1. Navigate to the Application Health by Tier Dashboard

The Application Healthy by Tier dashboards shows you all the applications and tiers with issues. The Platinum tier has two applications with critical issues, EHR and Ordering System. Let's look at the open cases on the Ordering System application.



2. Drill down to view the Ordering System application's open cases.



3. There are three open cases. Two are based on single metric alarm rules while one is based on the Exchange Performance rule template. Let's look at that open case.

Open Cases - Ordering System ✕

Rule Name	Entity Na...	Type	Occurren...	Last Occurrence
% CPU Ready	Ordering S...	singleMetricAlarm	3	05/05/2020 10:11:00 AM UTC
Exchange Performance	Ordering S...	Exchange Perfor...	13	04/30/2020 02:47:00 AM UTC
STS Avg Read Completion Ti...	Ordering S...	singleMetricAlarm	18	04/30/2020 02:34:00 AM UTC

Close

4. The application experienced some read latency during a one hour period. The Primary Rule shows us the conditions required for the alarm to trigger. Thirteen events were recorded.

[29604] Exchange Performance Topology Close Case

Overview Latest Alarms Case Severity: **CRITICAL** Case Watchers 0

Entity & Impacted Applications:

Application [Ordering System](#)

Impacted Applications:

[SQL Cluster 2 \(Platinum\)](#) | [Snow Proxy Test 1 \(Platinum\)](#) | [10 More](#)

Alarm Statistics for All Triggered Rules:

Total Alarms: 13 over an hour
 First: 04/30/2020 01:22:00 AM UTC
 Most Recent: 04/30/2020 02:47:00 AM UTC 15 days ago

Primary Rule

[Tier 0 Application Read Response Times](#)

Monitors an application's read response time (the aggregated measure of how long each read takes). Triggers when that aggregation exceeds 30ms for more than 6 one-second intervals in any given 1 minute period.

Read Completion Time Alarm when 10% of IOPS have exceeded 30.051 ms within a 1 min interval

Total Alarms: 13 over an hour
 First: 04/30/2020 01:22:00 AM UTC
 Most Recent: 04/30/2020 02:47:00 AM UTC 15 days ago

5. Select the Latest Alarms tab to view the event details.

[29604] Exchange Performance Topology Close Case

Overview Latest Alarms Case Severity: **CRITICAL** Case Watchers 0

Timestamp	Rule	Metric	Threshold	Value ↓
04/30/2020 01:44:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	47.5008361514735
04/30/2020 02:01:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	47.3389355742297
04/30/2020 02:14:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	47.3318629226059
04/30/2020 01:22:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	46.3950900384434
04/30/2020 01:59:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	45.5473880821835
04/30/2020 01:57:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	43.0614827677547
04/30/2020 01:27:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	24.5867970660147
04/30/2020 01:55:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	20.9478041860869
04/30/2020 02:32:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	17.5895946065821
04/30/2020 02:34:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	17.2728875956176
04/30/2020 02:47:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	16.556784461829
04/30/2020 02:30:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	13.3196382406448
04/30/2020 01:40:00 AM UTC	Tier 0 Application Read Response Ti...	NULL	10	10.701390142625

This type of alarm rule does not include an investigation or any trend charts so we need to use other features in the platform to investigate the alarm.

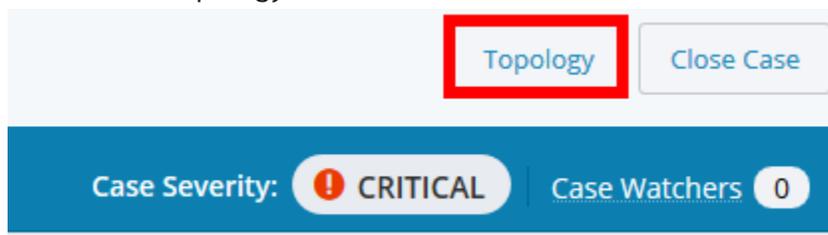
Possible Causes of Application Performance Issues

- Flow control on storage ports
- High CPU utilization on non-ESX hosts
- Incorrectly set HBA queue depth settings
- High utilization on HBA ports
- Speed mismatches between HBA ports and storage ports (FC SAN)

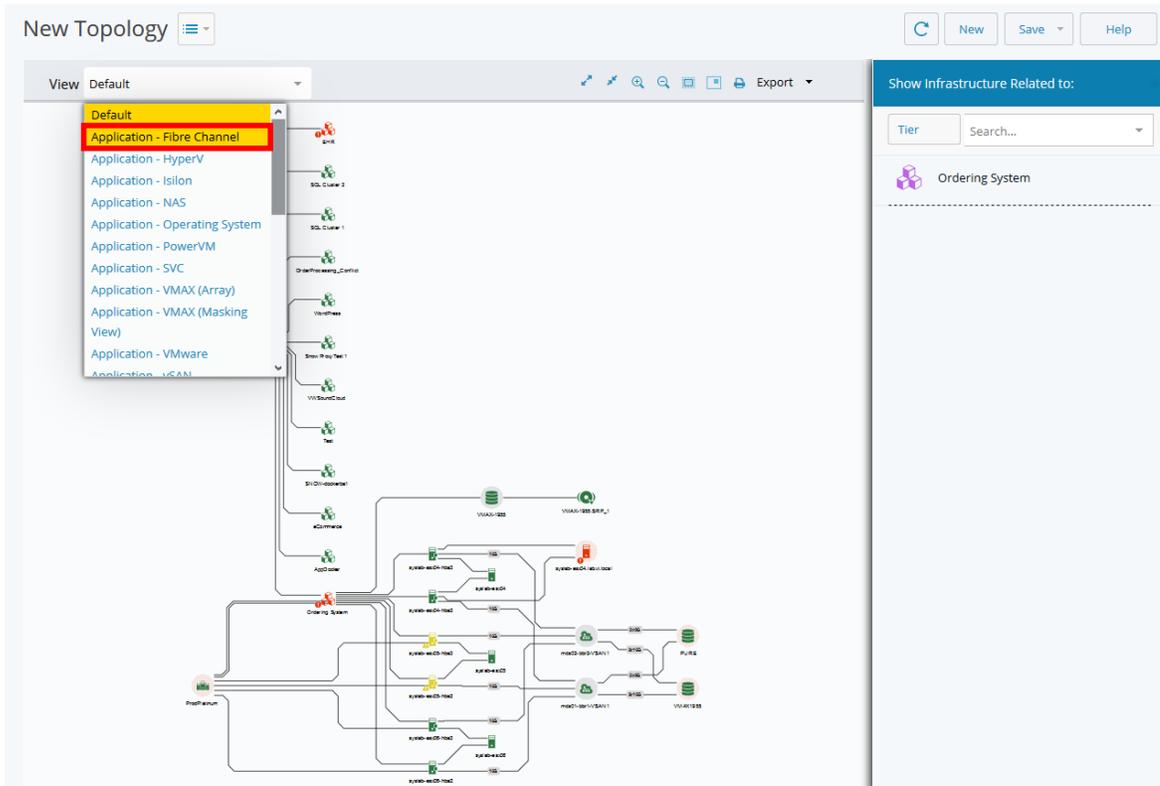
View Issues on Related Application Infrastructure

We can use the Topology feature to determine if there are issues on the related application infrastructure.

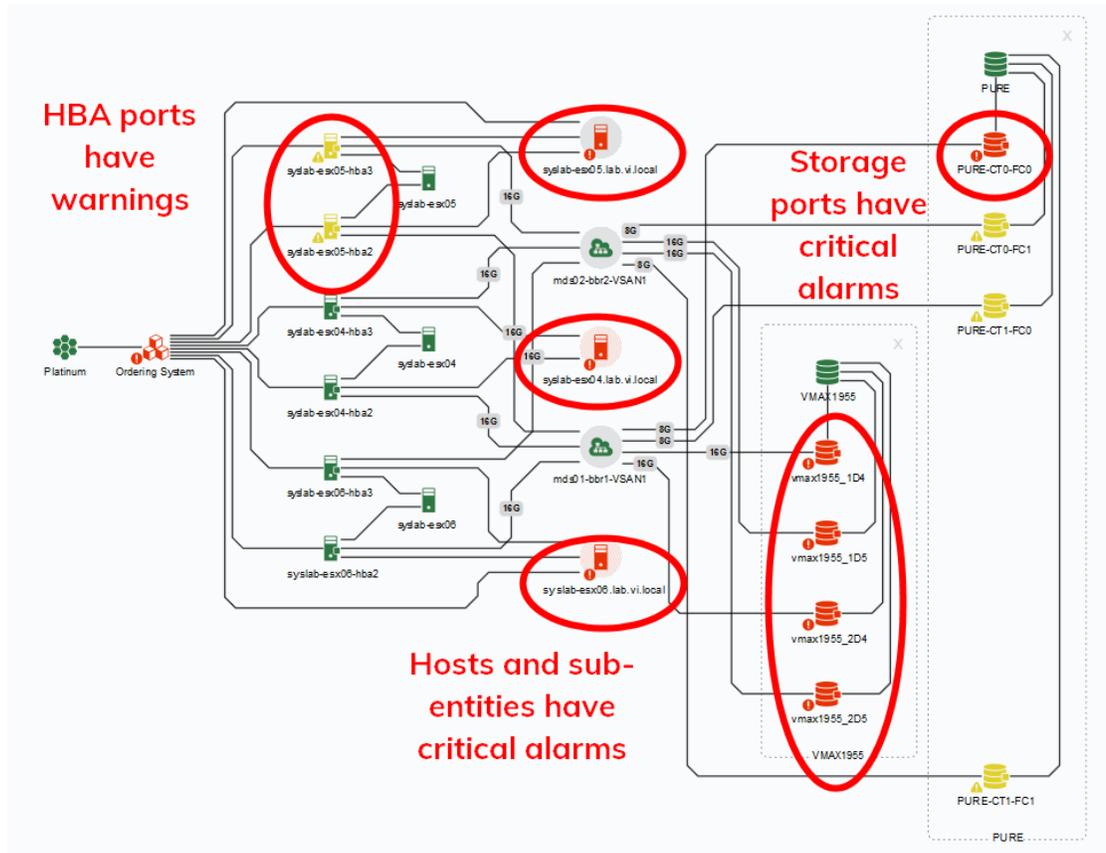
1. Use the Topology button on the open case page to view the impacted application's end-to-end topology.



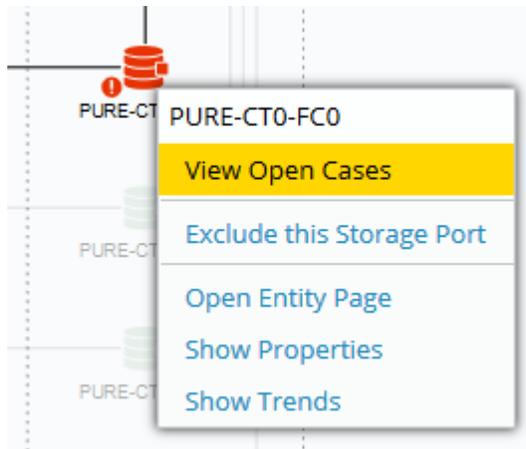
2. The Default view is loaded. All related infrastructure, including other applications and tiers are shown. Switch to the Application – Fibre Channel view to focus on the application.



3. Review the alarms on the Fibre Channel infrastructure
 - a. There are warnings on two of the HBA ports, critical alarms on hosts and sub-entities, and critical alarms on 5 of the 7 storage ports.



- b. Since we are interested in flow control events on the storage ports, let's review the open cases on the impacted storage ports.



- c. The VMAX storage array's ports all show a **Link Buffer-to-buffer Credits** alarm that matches the date and time of the application's performance alarm. We'll investigate these alarms first. Drill down on one of the open cases.

Open Cases - PURE-CT0-FC0

Rule Name	Entity Na...	Type	Occurren...	Last Occurrence
Link Buffer-to-buffer Credits	PURE-CT0-F...	Link Buffer-to-buf...	30,438	05/14/2020 07:00:00 PM UTC
Port Utilization	PURE-CT0-F...	Port Utilization	12,630	05/14/2020 06:55:00 PM UTC

Close

- d. The time frame of the storage port's alarm matches that of the application's alarm. We can use the slider on the Master/Detail trend chart to view the event.

[29595] Link Buffer-to-buffer Credits

Case Severity: **CRITICAL** Case Watchers: 0 Investigations (2)

Overview Latest Alarms

Entity & Impacted Applications:
Storage Port vmx1955_1D4 in VMAX1955

Impacted Applications:
eCommerce (Platinum) | OrderProcessing_Conflict (Platinum) | 22 More

Alarm Statistics for All Triggered Rules:

Total Alarms:	194 over 2 hours
First:	04/30/2020 01:21:00 AM UTC
Most Recent:	04/30/2020 02:59:00 AM UTC 15 days ago

Primary Rule: Tier 1 Perf Probe Storage Port Flow Control
Monitors storage ports (with the VirtualWisdom Performance Probe) for buffer credit starvation. Triggers if the link is in a credit-starved state for more than 30% of the time (18 seconds in 1 minute).

Time with Zero B2B Credits (Switch)
Alarm when more than 30%

Duration
Alarm when threshold is met for 1 min

Time with Zero B2B Credits (Device)
Alarm when more than 30%

Total Alarms: 96 over 2 hours
First: 04/30/2020 01:21:00 AM UTC
Most Recent: 04/30/2020 02:58:00 AM UTC **15 days ago**

External Case

Show Primary Threshold for: Time with Zero B2B Credits (Device) Alarm Threshold

Use the slider to view the event

Use the investigation to troubleshoot

Is there a mismatch in sender and receiver speed?
Closed - Pertinent
Closed 12/06/2019 03:18:26 AM UTC
5 months ago

Is there a mismatch in sender and receiver speed?
Closed - Pertinent
Closed 12/06/2019 03:18:26 AM UTC
5 months ago

- e. We can use the investigation to troubleshoot this issue.

+ Closed - Pertinent Closed 12/06/2019 03:18:26 AM UTC; 5 months ago **X**

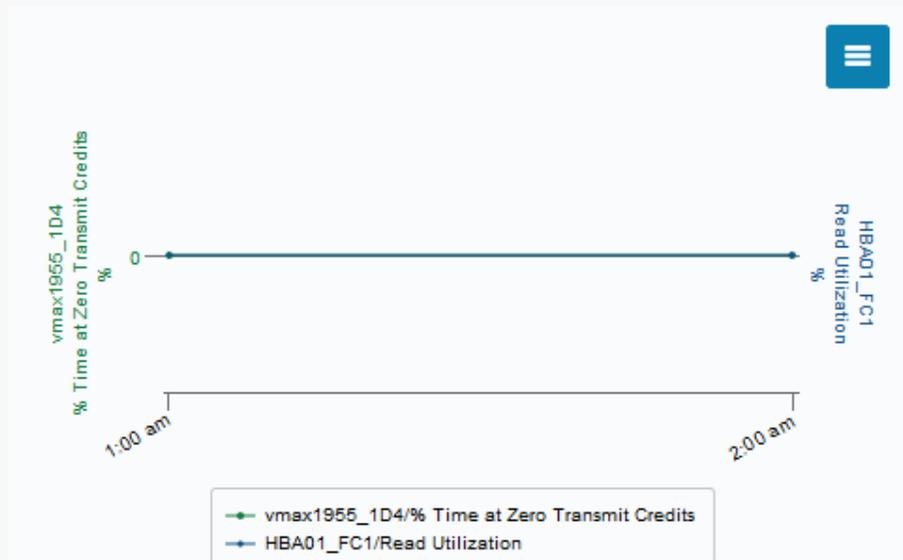
Is there a mismatch in sender and receiver speed?

Case: [29595, Link Buffer-to-buffer Credits](#)
 Storage Port: [vmax1955_1D4](#)

 **VirtualWisdom** 12/06/2019 03:18:23 AM UTC
 Looking for speed mismatches on related HBA ports...

 **VirtualWisdom** 12/06/2019 03:18:26 AM UTC
 Updated status: Closed - Pertinent

Ran [Trend Matcher](#) to find the HBA ports connected to [vmax1955_1D4](#) that match the utilization observed. There were 1 matches. Here is highest match.



[HBA01_FC1](#) port speed, 4G, is lower than [vmax1955_1D4](#), 16G, indicating a Flow Control problem.

Change Recommendation

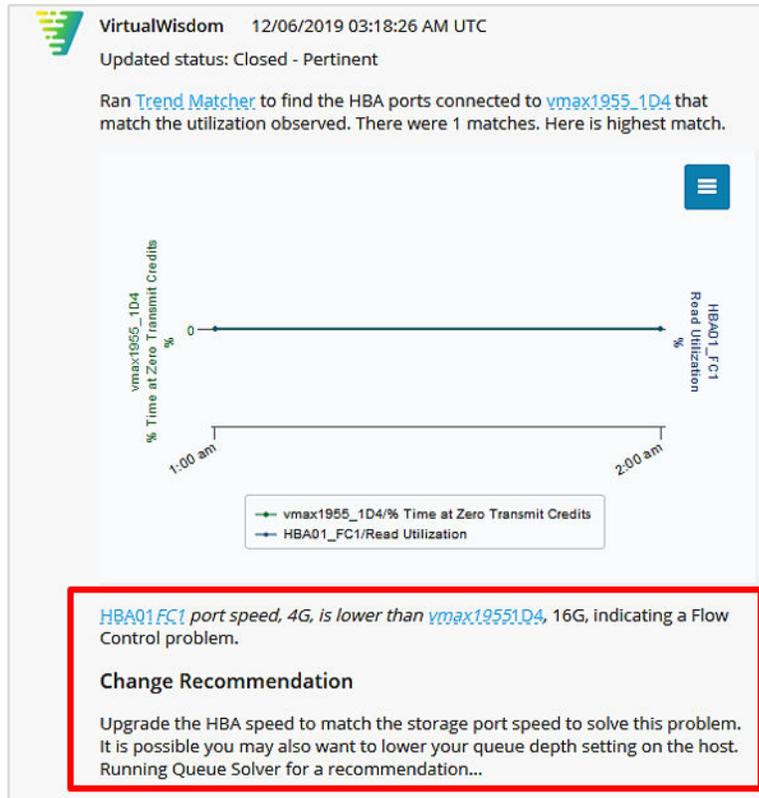
Upgrade the HBA speed to match the storage port speed to solve this problem. It is possible you may also want to lower your queue depth setting on the host. Running Queue Solver for a recommendation...

 **VirtualWisdom** 12/06/2019 03:18:59 AM UTC

[Queue Solver](#) indicates that setting the Queue Depth to 237 on [HBA01_FC1](#) may solve this problem.

 **VirtualWisdom** 12/06/2019 03:19:00 AM UTC

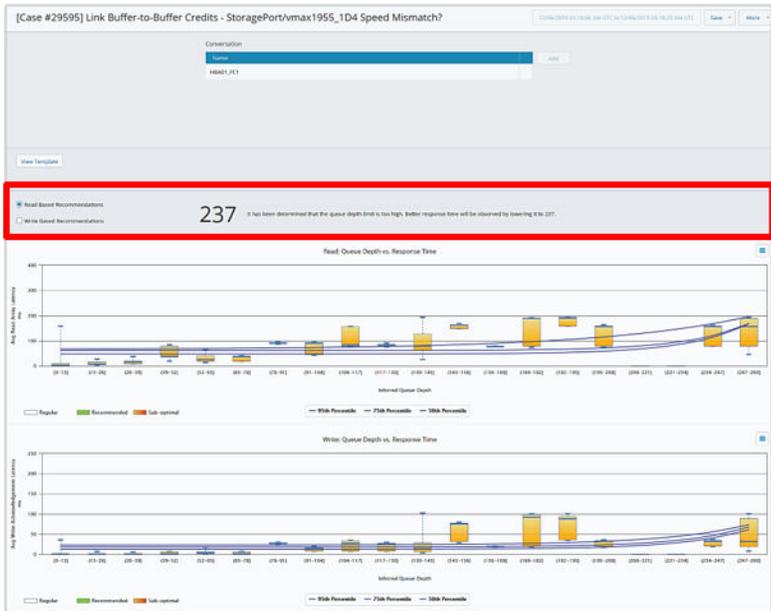
- f. The investigation detected a speed mismatch between the HBA ports connected to the storage port.



HBA and storage port speed mismatch

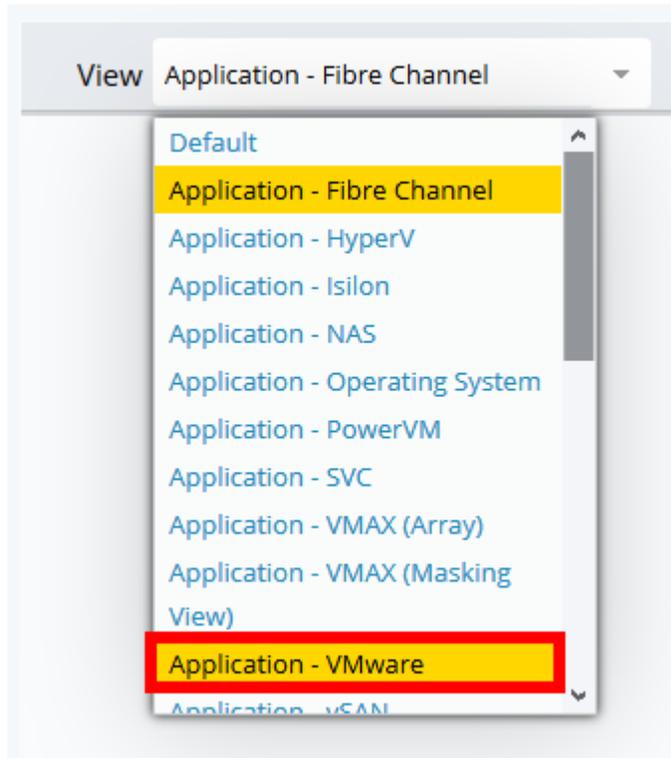
- g. The investigation also ran Queue Solver and determined that changing the queue depth settings on the HBA port may also improve performance.

VirtualWisdom 12/06/2019 03:18:59 AM UTC
Queue Solver indicates that setting the Queue Depth to 237 on **HBA01_FC1** may solve this problem.

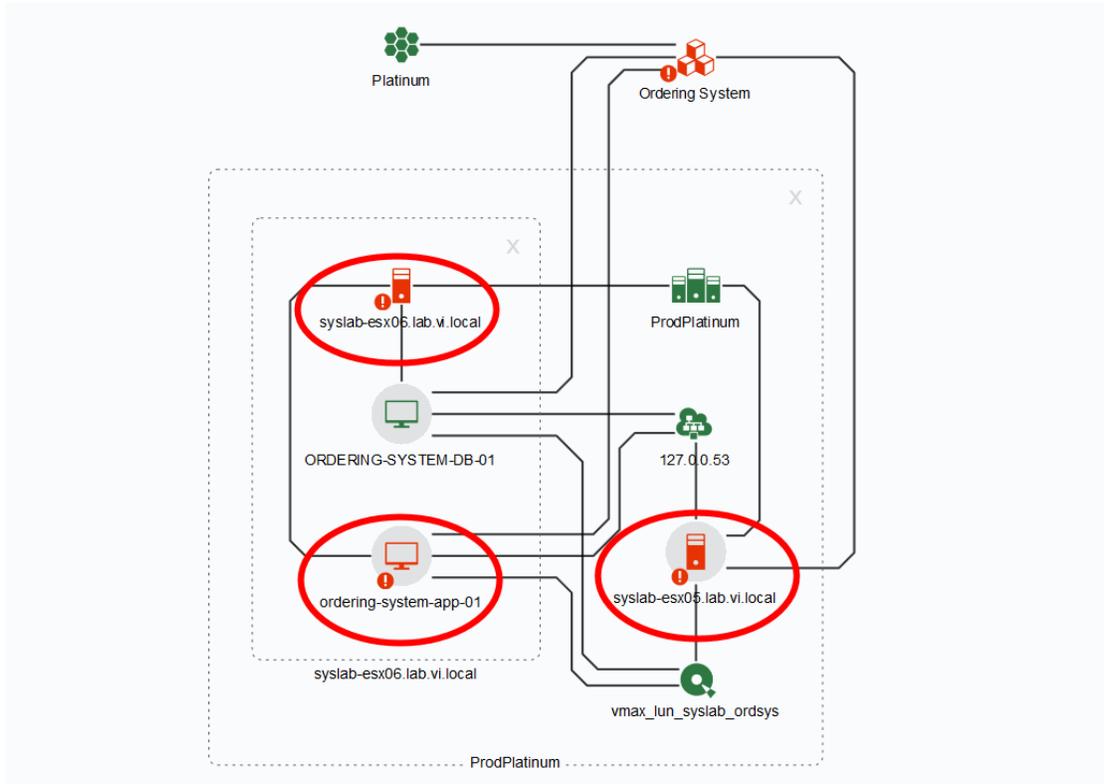


Adjust queue depth settings

4. Review the alarms on the VMware infrastructure
 - a. Switch to the Application - VMware topology view and expand the view.



- b. There are alarms on ESX hosts and VMs.



c. Review the alarms on the ESX hosts

- i. Both ESX hosts have Exchange Performance alarms that occur at the same time as the application's Exchange Performance alarm. These are likely all the same performance event but are shown in VirtualWisdom as separate cases because they are triggered by different alarm rules that were set on the different entity types.

Open Cases - syslab-esx05.lab.vi.local

Rule Name	Entity Na...	Type	Occurren...	Last Occurrence
Exchange Performance	syslab-esx0...	Exchange Perfor...	6	04/30/2020 06:59:00 PM PDT

Close

Open Cases - syslab-esx06.lab.vi.local

Rule Name	Entity Na...	Type	Occurren...	Last Occurrence
CPU Utilization	syslab-esx0...	CPU Utilization	3	05/13/2020 02:16:00 AM PDT
Exchange Performance	syslab-esx0...	Exchange Perfor...	9	04/30/2020 07:32:00 PM PDT

Close

- ii. Review the alarm statistics and investigations.



- iii. **Is the vSphere cluster imbalanced in CPU utilization?**
The automated investigation found that rebalancing the VMs on the cluster would not improve the CPU utilization on the host.

✕ Closed - Unrelated
Closed 12/18/2019 01:23:10 AM PST; 5 months ago ✕

vSphere cluster imbalanced in CPU utilization?

Case: [33339_CPU Utilization](#)
 ESX Host: [syslab-esx06.lab.vi.local](#)

VirtualWisdom 12/18/2019 01:21:09 AM PST

Common causes of High ESX Host CPU Utilization

- An imbalanced vSphere Cluster, where multiple CPU-intensive VMs are hosted on a single ESX Host
- One or more VMs have runaway processes consuming CPU resources

VirtualWisdom 12/18/2019 01:21:09 AM PST

Updated status: Active Investigation

Analyzing allocation of VMs across the cluster ...

VirtualWisdom 12/18/2019 01:23:10 AM PST

Updated status: Closed - Unrelated

Ran [VM.Coordinator](#) to see if better balancing this cluster would alleviate pressure from this host, but the recommended change is not likely to resolve this issue.

Recommendation
4
Based on 7819 iterations (100% completion)

for analyzed cluster ProdPlatinum

VMs to be Moved	From ↓	To
FinanceWeb	syslab-esx06.lab.vi.local	syslab-esx04.lab.vi.local
ERP-ETL-App-01	syslab-esx05.lab.vi.local	syslab-esx06.lab.vi.local
IT-W2K16	syslab-esx04.lab.vi.local	syslab-esx06.lab.vi.local
supply-warehouse-app-01	syslab-esx04.lab.vi.local	syslab-esx05.lab.vi.local

The VM Coordinator recommendation is not likely to improve CPU Utilization on this host. Closing as not related.

VirtualWisdom 12/18/2019 01:23:10 AM PST

Updated external case: INC0586965 with investigation results.

Set Status: [Closed - Unrelated](#) ▼
Save

- iv. **Are there VMs on this ESX host that have a runaway process?**
The investigation does not reveal any runaway processes.

○ Not Started
Created 12/18/2019 01:21:09 AM PST ✕

Are there VMs on this ESX host that have runaway processes?

Case: [33339, CPU Utilization](#)
 ESX Host: [syslab-esx06.lab.vi.local](#)

VirtualWisdom 12/18/2019 01:21:09 AM PST

Common causes of High ESX Host CPU Utilization

- An imbalanced vSphere Cluster, where multiple CPU-intensive VMs are hosted on a single ESX Host
- One or more VMs have runaway processes consuming CPU resources

How to determine if there are there VMs on this ESX host that have runaway processes

Observe the following chart which plots all the virtual machine's CPU utilization on this ESX Host over the last 24 hours

If the CPU utilization trend for any single VM appears to stay fixed at one CPU level over a long period of time (a day or more), this is usually an indication of a runaway process present on this VM.

How to resolve VMs with runaway processes

- Contact the server administrator to ensure the process is truly runaway.
- Terminate the runaway process.

Set Status: [Not Started](#) ▼
Save

d. Review the alarms on the ESX VM

- i. The ESX VM has a number of CPU utilization open cases. Let's review the most recent one.

Open Cases - ordering-system-app-01

Rule Name	Entity Na...	Type	Occurren...	Last Occurrence
CPU Utilization	ordering-s...	CPU Utilization	20	05/14/2020 01:13:00 PM UTC
CPU Utilization	ordering-s...	CPU Utilization	7	04/30/2020 08:19:00 AM UTC
CPU Utilization	ordering-s...	CPU Utilization	38	04/30/2020 06:55:00 AM UTC
CPU Utilization	ordering-s...	CPU Utilization	13	04/29/2020 07:46:00 AM UTC
CPU Utilization	ordering-s...	CPU Utilization	748	04/28/2020 10:55:00 PM UTC

Close

- ii. Review the alarm statistics and investigations.

[33336] CPU Utilization

Case Severity: **CRITICAL** Case Watchers: 0

Investigations (2)

- Runaway process?
 - Not Started
 - Created 12/18/2019 08:56:12 AM UTC
- Insufficient vCPU for the workload on this VM?
 - Not Started
 - Created 12/18/2019 08:56:12 AM UTC

Entity & Impacted Applications: ESX VM ordering-system-app-01 in Prod/Platinum

Impacted Applications: SQL Cluster 2 (Platinum) Show Proxy Test 1 (Platinum) 10 More

Alarm Statistics for All Triggered Rules:

Total Alarms: 20 over 2 days
 First: 05/12/2020 08:49:00 AM UTC
 Most Recent: 05/14/2020 01:13:00 PM UTC 8 hours ago

Primary Rule: Tier 0 Compute ESX VM CPU Utilization
 Monitors ESX Virtual Machine (VM) CPU Utilization. Triggers if the CPU exceeds 98% utilization for a 15-minute interval.
 CPU Utilization Alarm when more than 98% for 15 mins

Total Alarms: 20 over 2 days
 First: 05/12/2020 08:49:00 AM UTC
 Most Recent: 05/14/2020 01:13:00 PM UTC 8 hours ago

Show Primary Threshold for: CPU Utilization Alarm Threshold

- iii. Is there a runaway process?

There is no runaway process on the VM.

○ Not Started Created 12/06/2019 12:17:15 AM PST ✕

Runaway process?

Case: [29746, CPU Utilization](#)
ESX VM: [ordering-system-app-01](#)

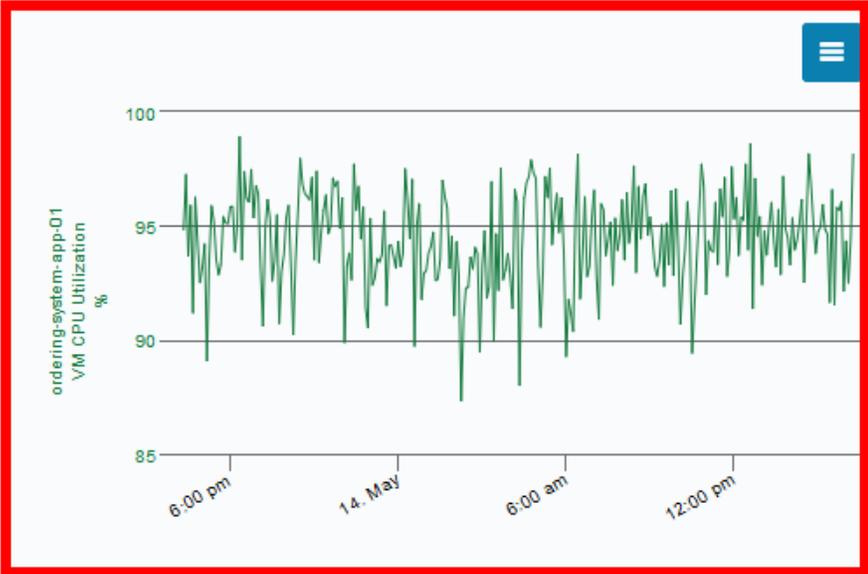
 VirtualWisdom 12/06/2019 12:17:15 AM PST

Common causes of High Virtual Machine CPU Utilization

- A runaway process is consuming an excessive amount of CPU cycles
- Insufficient CPU available for the running processes

How to determine if you have a runaway process

Examine the following chart which plots the virtual machine's CPU utilization over the last 24 hours:



If the utilization trend appears to stay fixed at one CPU level (flat line) over a long period of time (a day or more), this usually indicates a runaway process present on this VM.

How to resolve runaway processes

Contact the server administrator to verify the process is truly a runaway, and then terminate the process.

Set Status: [Not Started](#) ▾

Save

- iv. **Is there insufficient vCPU for the workload on this VM?**
vCPU appears to be sufficient.

○ Not Started
Created 12/06/2019 12:17:15 AM PST ✕

Insufficient vCPU for the workload on this VM?

Case: [29746 CPU Utilization](#)
 ESX VM: [ordering-system-app-01](#)

VirtualWisdom
12/06/2019 12:17:15 AM PST

Common causes of High Virtual Machine CPU Utilization

- A runaway process is consuming an excessive amount of CPU cycles
- Insufficient CPU available for the running processes

How to determine if the VM has insufficient vCPU?

Examine the following chart which plots the VM's CPU utilization over the last 24 hours:

If the utilization appears to spike sporadically, most likely insufficient CPU is allocated to the running processes. More vCPU is needed to support the workload on this host.

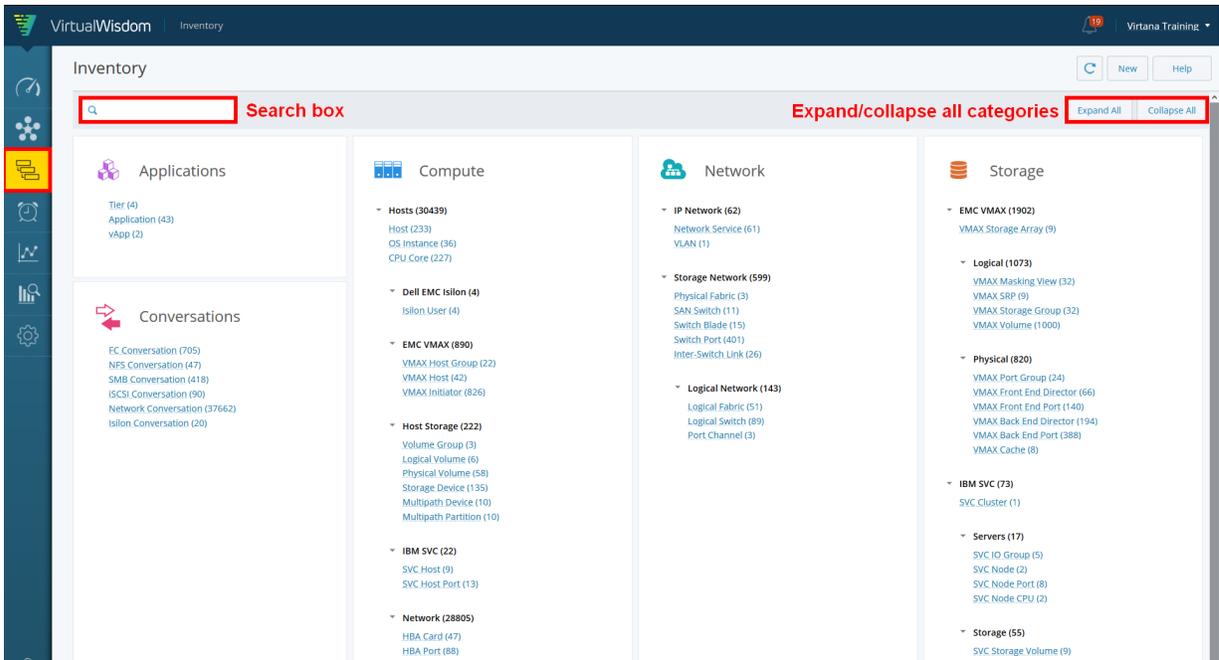
How to resolve insufficient vCPU

Allocate additional vCPU resources to this VM.

Set Status: [Not Started](#) ▼Save

Inventory

Inventory is your "one-stop shop" for viewing and managing entities



After VirtualWisdom has discovered your entities, they are visible throughout the VirtualWisdom user interface. The Inventory module provides a single location for viewing all of your VirtualWisdom entities.

Entities are organized into five categories: Applications, Conversations, Compute, Network, and Storage.

Use the search box to search for an entity using its name.

You can use the expand/collapse all buttons to expand or collapse the categories for easier view. Click on the arrow next to a category to expand it.

Counts on the Inventory landing page are cached and updated every five minutes. You can click the refresh icon on the Inventory landing page to update the counts.

Inventory List View

When you click on any entity type on the Inventory landing page, the Inventory List View page displays with a table of information about the entity. The default table columns vary by entity type, but include Name, Tags, and Created On. You can add or remove columns from the  (Table Options) menu, resize the columns, and move the columns using drag and drop.



NOTE

Not all elements are present for every entity.

The screenshot shows the 'Application Discovery' interface. At the top, there are buttons for '0 New', '1 with suggestions', 'Application', 'New', and 'More'. Below this is a search bar with a magnifying glass icon and the text '(22 items)'. To the right of the search bar are three buttons: 'Application', 'New', and 'More'. Below the search bar is a table with columns: Tier, Name, Tags, and Created On. The table contains several rows of application data. Callouts 1 through 8 point to various UI elements: 1 points to the search bar, 2 points to the '(22 items)' text, 3 points to the 'Application' button, 4 points to the 'New' button, 5 points to the 'More' button, 6 points to the table header, 7 points to a dropdown arrow in the 'Created On' column, and 8 points to a blue dot in the 'Tier' column.

Tier	Name	Tags	Created On
	hg-test-app-no-vm	hg de	04/01/2019 06:32:00 AM EDT
Tier 3	ManualService_Wordpress		02/16/2019 06:00:00 AM EST
Tier 1	App - DockerMasterSlave		03/14/2019 11:00:00 PM EDT
Tier 0	App 05 - All - SB3_0_targ0-targ3 - ...	import_added	02/11/2019 05:35:00 PM EST
Tier 1	App 11 - sblaze2-16-virtual - sblaz...		02/16/2019 06:00:00 AM EST
Tier 3	DockerMasterSlave(Manual Servi...		02/16/2019 06:00:00 AM EST
	App 06 - All - SB3_1_targ1_targ3 - ...	import_added	02/11/2019 05:35:00 PM EST

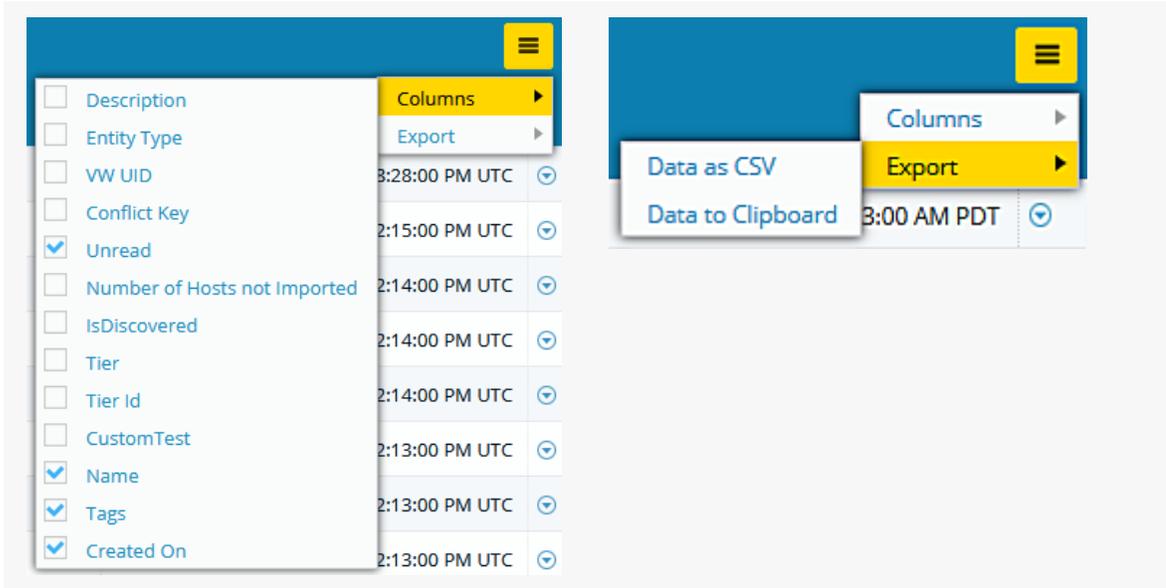
1. Search partial terms or numbers. Not all column content can be searched. To reset the grid to an unfiltered state, click the "X" icon in the search field.



NOTE

If an entity has multiple tags, and you search for an individual entity, the search returns all the tags.

2. Displays the total number of items in the entity list. After searching, this number updates to reflect how many filtered items are shown.
3. Application Discovery actions. Applies only to Application entity types.
 - **New** opens a dialog to create applications based on suggestions from discovery.
 - **with Suggestions** opens a dialog to update applications based on suggestions from discovery.
4. Displays the current type of entities you are viewing in the grid. Clicking the button opens a list of all the entity types, when selecting one, the grid reloads to show you entities of that type.
5. **More** lets you toggle on bulk edit mode or access the VirtualWisdom User Guide.
6. Provides options for showing/hiding list columns and exporting the contents of the list as a CSV file, or copy them to the clipboard.



7. Edit, View Open Cases, Show Properties, Show Topology, and Delete.

If a blue dot is displayed, it means the entity has not yet been inspected.

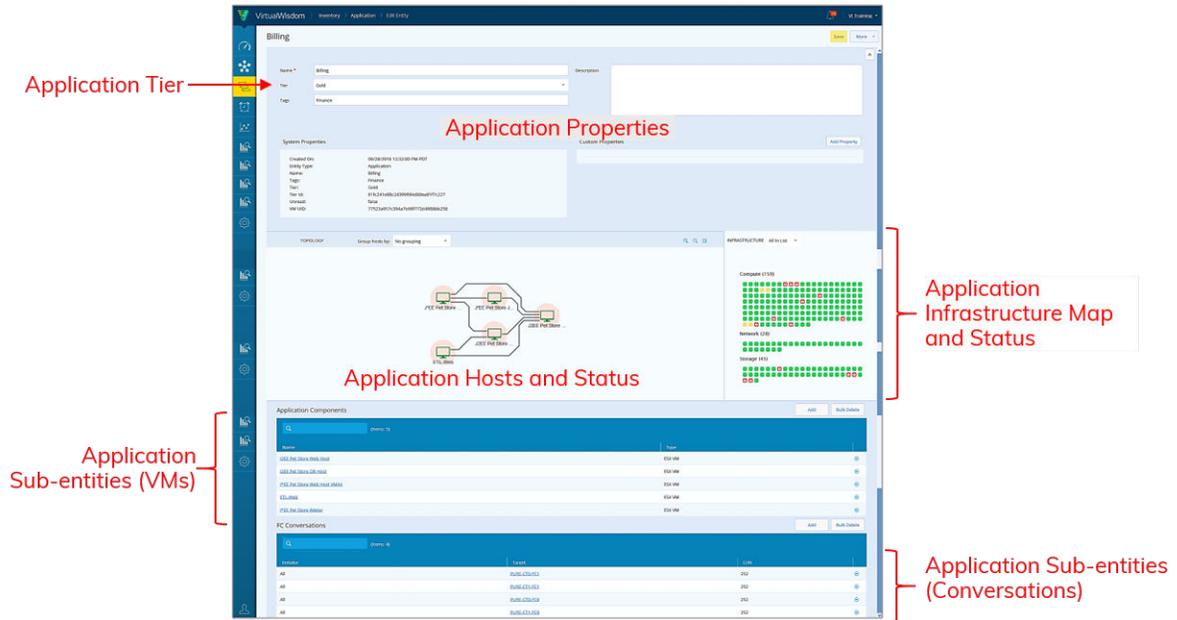


NOTE

Windows Management Instrumentation (WMI) only identifies physical Network Interfaces for which the MAC address of the Bonded NIC is the same as the MAC address of the physical NIC. This results in physical NICs not being reported as children of a NIC bonded to multiple physical NICs. Therefore, the relationship between a Windows Bonded NIC and its children cannot be reported in VirtualWisdom topology views or entity inventory pages.

Application Inventory Page

Application entities have a dedicated page view that displays specialized information about them. At the top of the screen is a field showing you which tier the application belongs to. Description and tags are also displayed. To edit these fields, select the Edit link.



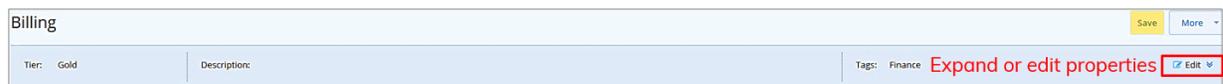
The application’s hosts are displayed below in the Topology section. You can choose whether to group the hosts by their role.

The Infrastructure Map section shows a heat map of the infrastructure supporting the application. These two sections work together to provide insight into applications. Selecting a host from the Topology section highlights the related infrastructure in the Infrastructure Map.

The Application Components section shows a list view of the hosts or containers that comprise the application. Select the down arrow on the right to open the component’s entity page, view its open cases, show its properties or topology.

The Conversations section displays a list of conversations that comprise the application. Click the down arrow to view the menu for the conversation.

Application Properties



The Properties section is hidden by default. Click the Edit button to expand or edit properties.

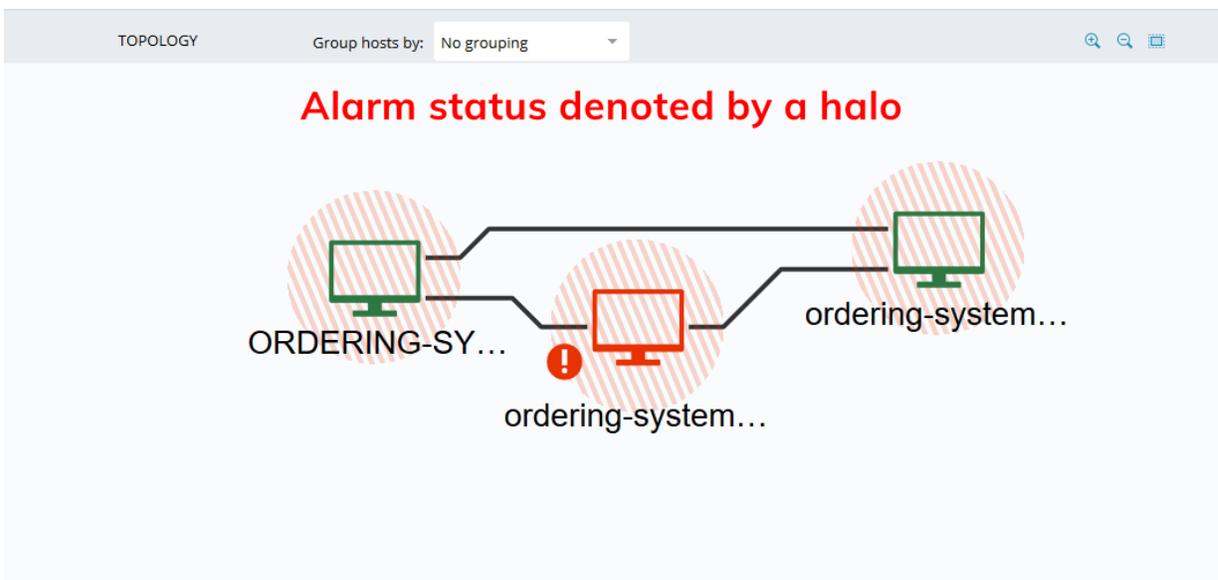
System Properties	
Created On:	09/26/2018 12:32:00 PM PDT
Entity Type:	Application
Name:	Billing
Tags:	Finance
Tier:	Gold
Tier Id:	81f6241e88c24399994d80ea81f7c227
Unread:	false
VW UID:	77523a957c394a7b90f77204988bb258

System properties are properties that are set when the application entity is created. These properties are managed by the VirtualWisdom platform and they cannot be changed.

Custom properties are properties that you can define and manage.

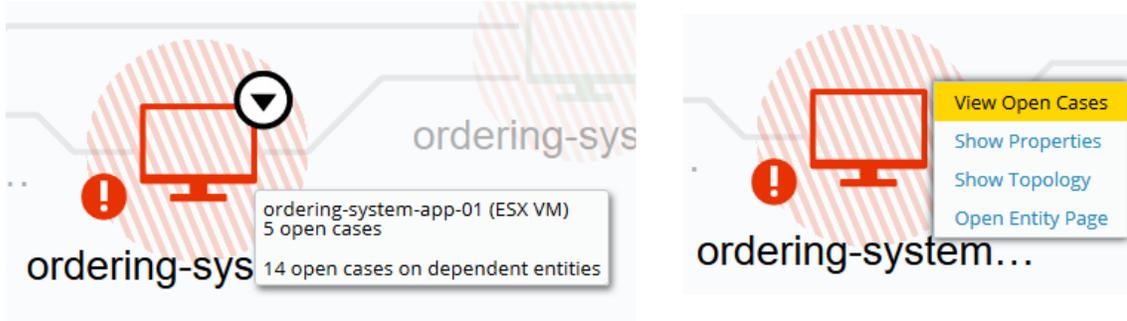
Application Host Alarm Status

The Application Topology view shows the hosts that comprise the application, along with their current alarm status. The host’s status is denoted by the color of its representative icon. A halo surrounding the icon denotes an alarm at a sub-entity level.



Hover over the entity to view more information on the alarms on the entity and its sub-entities. Click on the down arrow to open a menu to view open cases for the host entity.

Hover over or click to view more info



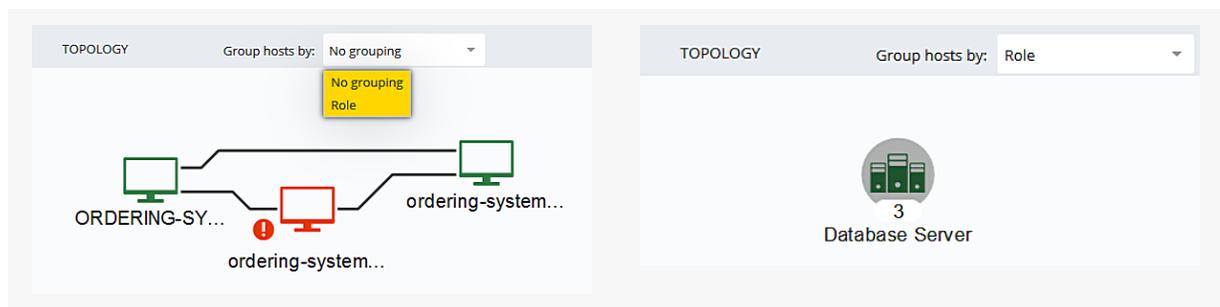
Host entities with red icons and an exclamation point on a red circle have active alarms. Hosts with green icons do not have active alarms.

A halo surrounding a host icon indicates that it has sub-entities. The halo's color indicates the severity of the alarms on the sub-entities:

- Red hatching indicates there is an alarm with critical severity on one or more sub-entities
- Yellow hatching indicates there is an alarm with warning severity on one or more sub-entities
- Green hatching indicates that there are no active alarms on the host's sub-entities.

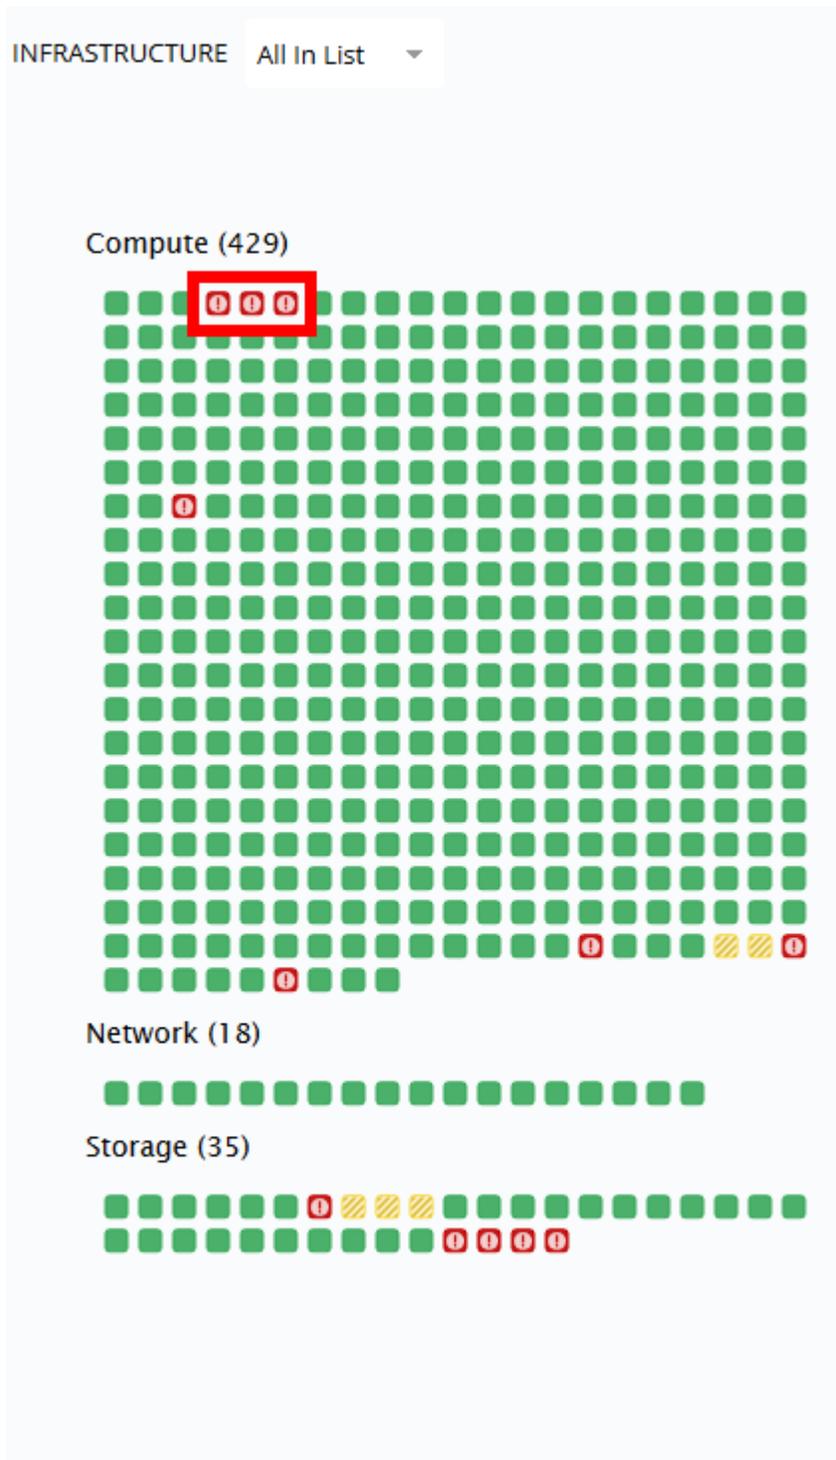
In the application topology shown above, one host has a critical alarm and all three hosts have sub-entities with critical alarms.

You can view the application's hosts grouped by role by selecting Role from the drop down list.

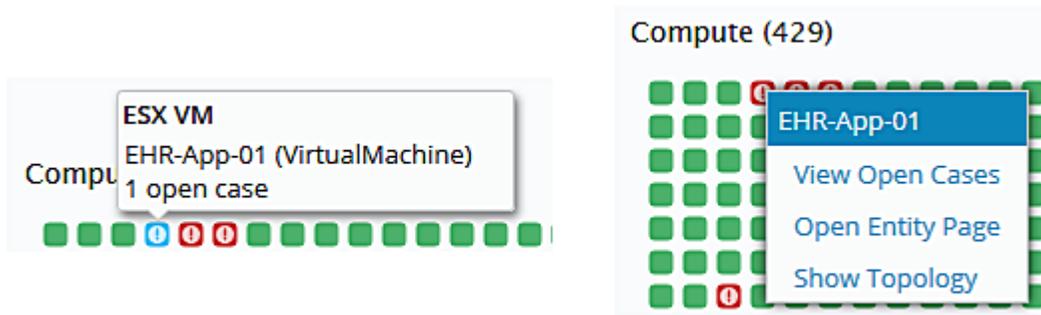


Application Dependency Map

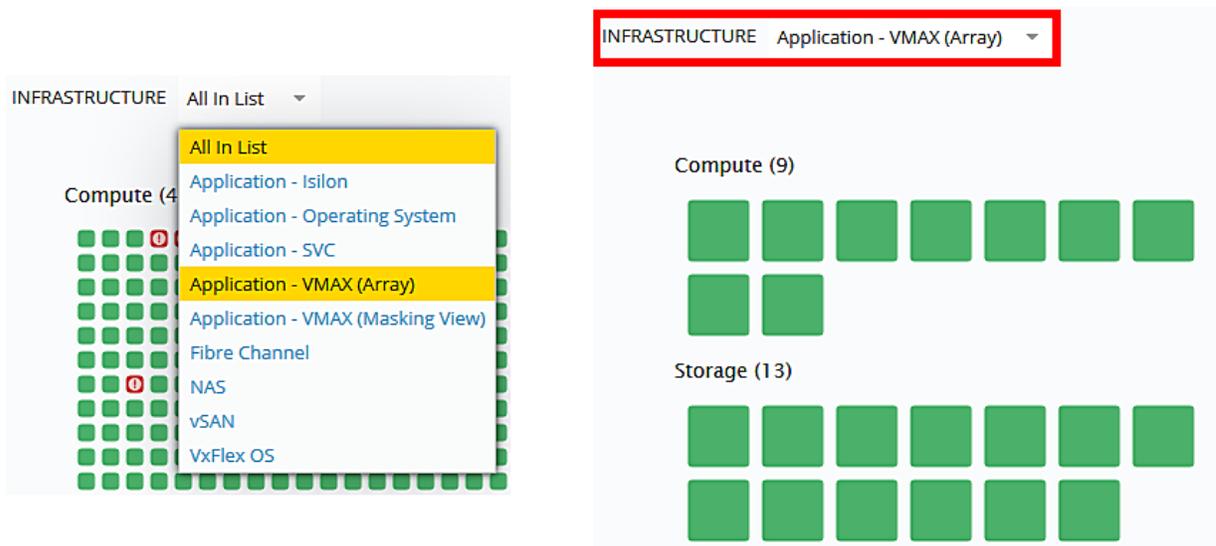
The Application Infrastructure Dependency map is a representation of the entities supporting the application. One box is shown for each entity supporting the application. The boxes show you where there are issues in the infrastructure. They are color coded just like the alarm status shown on the hosts.



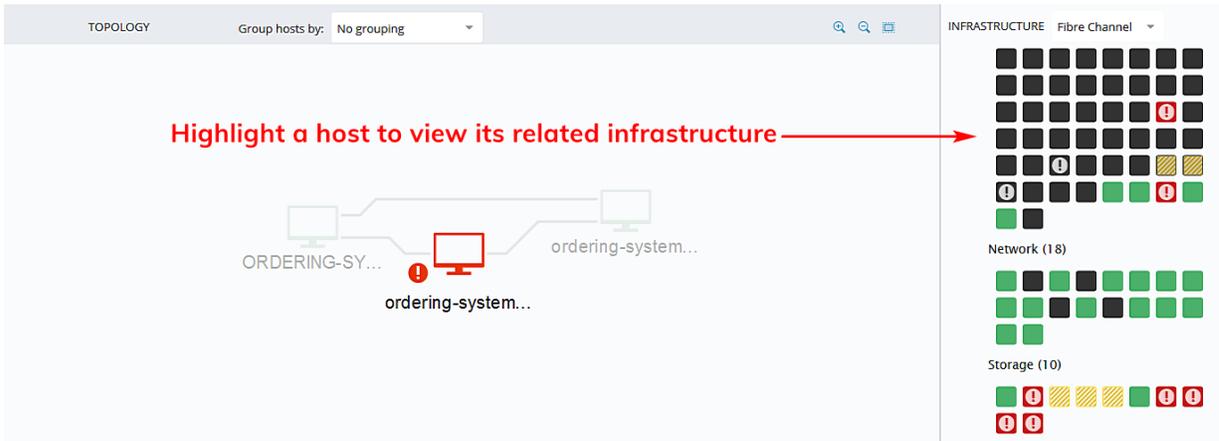
Hover over or click on an entity to view more information about its status.



You can filter the entities by entity type, e.g., Fibre Channel, VMAX, by selecting a view from the dropdown entity type.



You can use the Application Topology view and the Infrastructure map together. By clicking on a host in the Topology view, you can see which of the entities displayed in the Infrastructure map are related to the application host.



Application Sub-Entities

The list views below the application topology map show you the sub-entities that are related to the application. These include entities that comprise the application service and conversations between the application and its data.

Entities that comprise the application service

Conversations between the application and its data

Application Components Add Bulk Delete

Name	Type
J2EE Pet Store Web Host	ESX VM
J2EE Pet Store DB Host	ESX VM
J2EE Pet Store Web Host VMX	ESX VM
ETL Web	ESX VM
J2EE Pet Store JMeter	ESX VM

FC Conversations Add Bulk Delete

Initiator	Target	LUN
All	PURE-CTD-FC1	252
All	PURE-CTL-FC1	252
All	PURE-CTD-FC0	252
All	PURE-CTL-FC0	252

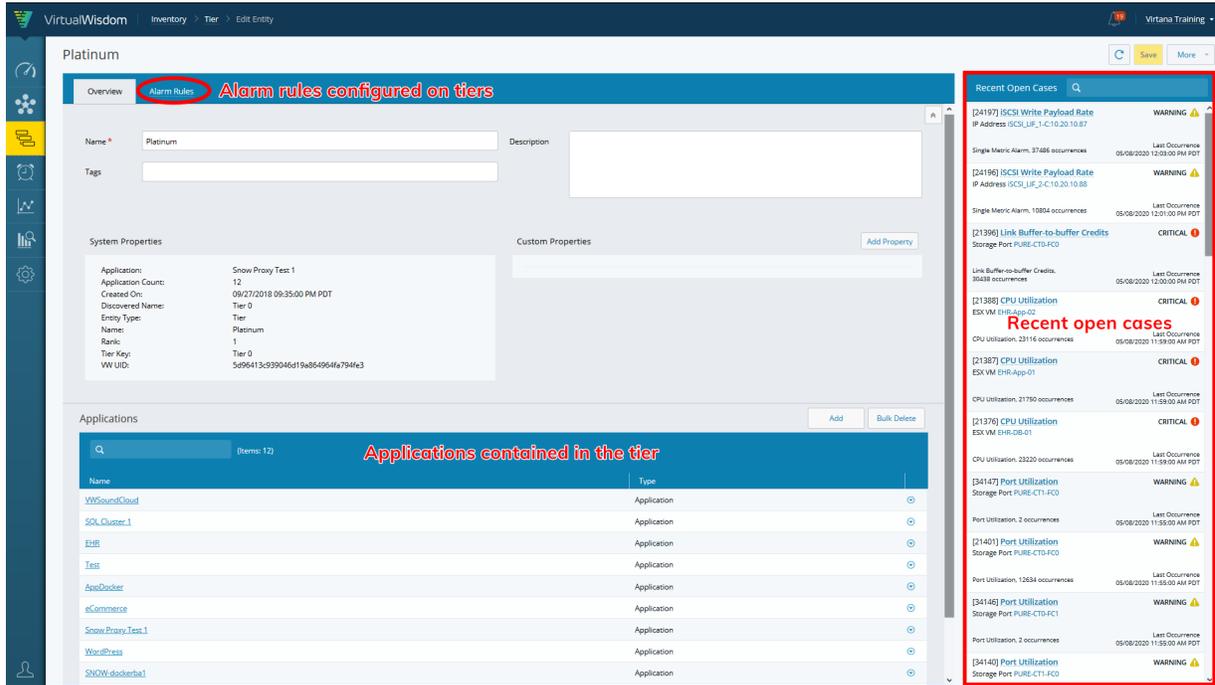
Use the down arrow to view menu options

Tier Inventory View

The Tier Inventory View shows the data and status for an application tier. To view the page, drill down on the name of a tier from the Tier inventory list.

The view shows you the tier's properties and lists the applications that are contained in the tier.

To the right is a panel that lists the most recent open cases on the tier.

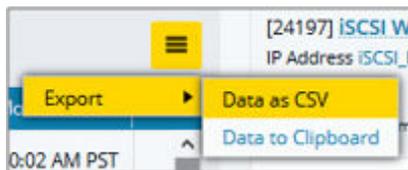


Select the Alarm Rules tab to view the alarms that have been configured for the tier.

Platinum

Rule Name	Description	Enabled	Severity	From Template	Last Modified ↓
Tier 0 Application Workload Drift	Monitors critical applications for changes (drifts) in read...	Yes	Critical	Single Metric	04/15/2020 06:30:42 PM PDT
Tier 0 NAS Controller Packet Errors	Monitors NAS Storage Controller links for packet errors. ...	No	Critical	NAS Packet Errors	04/15/2020 06:30:42 PM PDT
Tier 0 Bad SCSI Status - Queue Full	Monitors non-zero SCSI status messages with a sense co...	Yes	Critical	Bad SCSI Status	04/15/2020 06:30:42 PM PDT
Tier 0 Bad SCSI Status - Busy	Monitors non-zero SCSI status messages with a sense co...	Yes	Critical	Bad SCSI Status	04/15/2020 06:30:42 PM PDT
Tier 0 Bad SCSI Status - Internal Tar...	Monitors non-zero SCSI status messages with a sense co...	Yes	Critical	Bad SCSI Status	04/15/2020 06:30:42 PM PDT
Tier 0 OS Instance Excessive Disk W...	Detects OS Instance excessive disk write latency.	No	Critical	Single Metric	04/15/2020 06:30:42 PM PDT
Tier 0 OS Instance Excessive Disk Re...	Detects OS Instance excessive disk read latency.	No	Critical	Single Metric	04/15/2020 06:30:42 PM PDT
Tier 0 Physical Volume Critical Disk ...	Detects Physical Volume critical disk usage.	No	Critical	Single Metric	04/15/2020 06:30:42 PM PDT
Tier 0 Physical Volume Excessive Dis...	Detects Physical Volume excessive disk usage.	No	Critical	Single Metric	04/15/2020 06:30:42 PM PDT
Tier 0 NAS Controller Link Errors	Monitors NAS Storage Controller links for CRC errors. Tri...	No	Critical	NAS Link Errors	04/15/2020 06:30:42 PM PDT
Tier 0 Perf Probe Storage Port Flow ...	Monitors storage ports (with the VirtualWisdom Perform...	Yes	Critical	Link Buffer-to-Buffer Credits	04/15/2020 06:30:42 PM PDT
Tier 0 NFS Write Performance	Monitors the NFS write performance of Tier 0 applications.	Yes	Critical	Single Metric	04/15/2020 06:30:42 PM PDT

You can export the rule data by selecting the hamburger icon, then selecting Export.



Other Entity Types Inventory View

All entities have inventory pages, which are similar across entity types other than Application and Tier entities.

The Properties section displays the entity's system and custom properties. Below the Properties section is a list of all sub-entities related to this entity. Some entities also have a Conversations section below this.



IMPORTANT

Beginning in VirtualWisdom 6.7, a limit has been placed on the number of conversation entities that VirtualWisdom stores for ProbeFC, ProbeNAS, and NetFlow.

If the system limit of the number of conversations is reached, the least-recently-seen conversations are automatically deleted.

Deletion of these entities is intended to increase performance and reliability for long-running deployments. If you wish to modify or disable this feature, contact [VirtualWisdom Support](#).

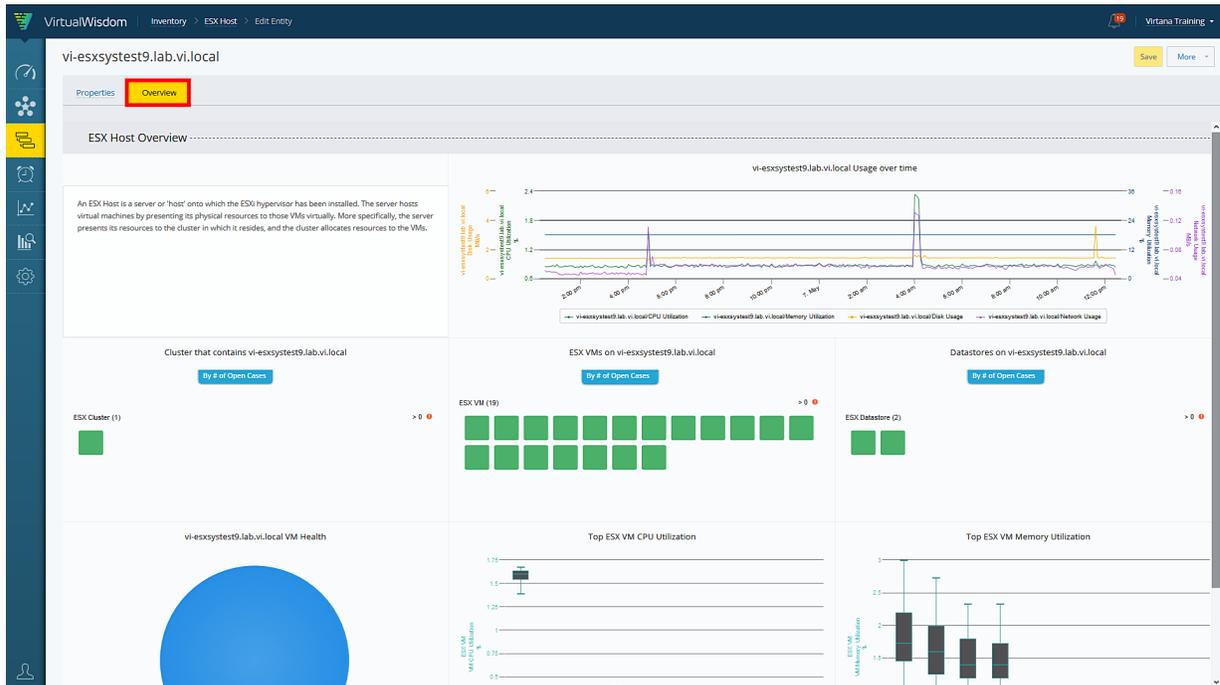
System properties are populated automatically with entity discovery and cannot be changed. Custom properties can be added or changed by selecting the Add Property button and entering your custom property.

The screenshot displays the VirtualWisdom interface for editing an entity named 'vi-esxsystem9.lab.vi.local'. The interface is divided into several sections:

- Entity Properties:** This section, highlighted with a red border, contains fields for Name, Description, and Tags. Below these are System Properties (e.g., Connection State, CPU Info, Created On) and Custom Properties. A red arrow points to an 'Add custom property' button.
- Sub-entities:** This section, also highlighted with a red border, is a table listing related entities. The table has columns for Name and Type.

Name	Type
wmsim-2mit-3	ESX VM
sbsim-paru-4	ESX VM
Dyantrace_Manu-1	ESX VM
vi-esxsystem9.lab.vi.local/vmk0	Virtual Ethernet Port
sbsim-paru-10	ESX VM
sbsim-paru-7	ESX VM
wmsim-paru-5	ESX VM
wmsim-2mit-4	ESX VM
sbsim-paru-8	ESX VM
AppDynamic-2mit	ESX VM
vcsim-paru-2	ESX VM
sbsim-2mit	ESX VM
- FC Conversations:** This section at the bottom shows a table for Fibre Channel conversations, currently empty.

Entity Overview

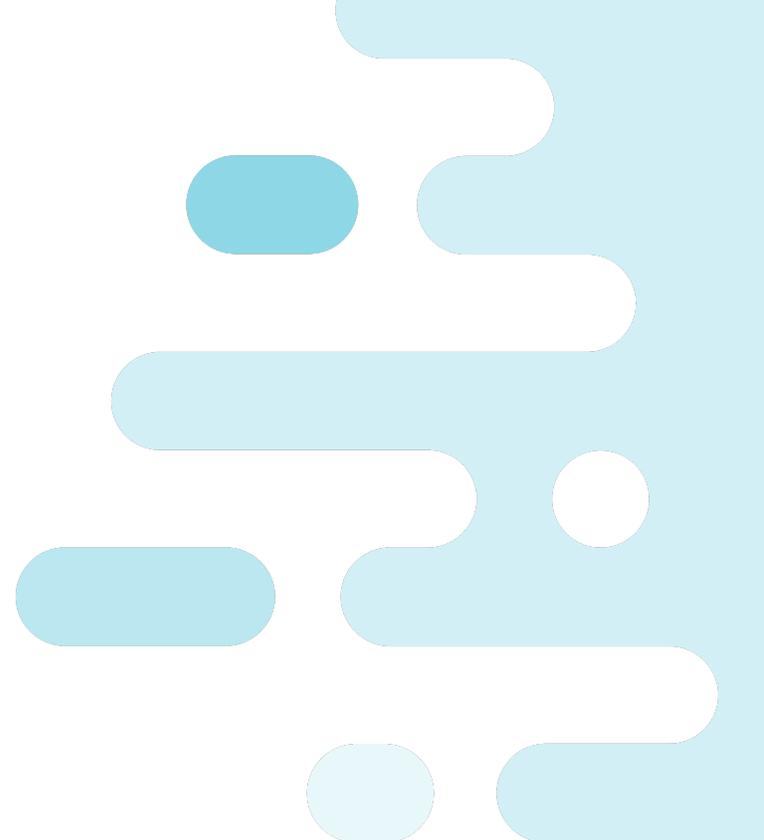


The Overview tab displays a report that provides basic information and metric data for the entity. The data displayed is for the last 24 hours of monitoring and focuses on health, utilization, capacity and performance data. You can use the page to view the current state of the entity.

The data shown differs based on the entity type. Entity overviews are available on the following entity types:



Topology



At the core of the Wisdom AI platform is our approach to app-centric topology that enables you to discover, visualize, diagnose and manage your application's datapaths. You can view the topology from the servers that host the application, to the shared networks that connect it, to the storage that holds its data. This ability to put all the infrastructure in the context of the applications it serves, while understanding the application's business value and the workload it generates against the infrastructure, is what makes VirtualWisdom app-centric in its approach.

VirtualWisdom uses the app-centric topology extensively across the platform, for defining dashboard contents, application or infrastructure topology views, and alarm policies. Topology is also used with the analytic infrastructure advisors for diagnosis and optimization.

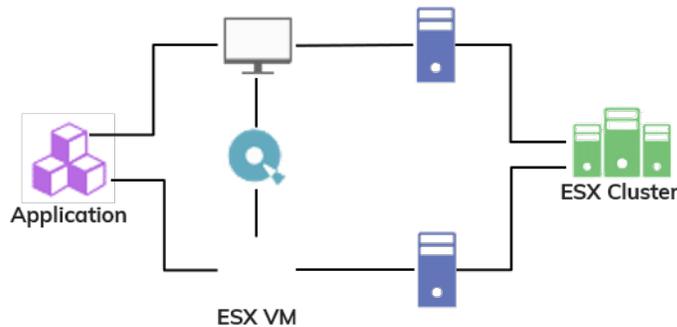


NOTE

Topologies saved in a VirtualWisdom version prior to 6.0 are incompatible with this release. During installation of the current version VirtualWisdom, any existing saved topologies are archived and are no longer accessible from the Select a Saved Topology menu.

What Is Topology?

Topology is a graphical view of the interconnections and activity in your infrastructure.



An entity relationship

Topology shows you the entity relationships in the context of tiers and applications. The end-to-end infrastructure supporting the application is visible in a single view. Topology also shows you the status of each entity.

Topology answers these questions				
What is the end-to-end scope and scale of an application?	Do entities interact with each other and how?	Where do entity dependencies overlap?	Where do problems exist in the infrastructure?	What is the impact to the application?

VirtualWisdom topology enables you to:

- Identify if and how two or more entities interact with each other
- Identify overlap for dependencies of two or more entities
- Understand the end-to-end scope and scale of an application
- Identify and understand potential problems in an entity's data path
- Understand where various entities live in the data center
- Understand the impact of and impact to applications in your environment

Entity Representations in Topology

Topology supports multiple infrastructure views to cover major technology areas (compute, network, and storage). Multiple views provide different perspectives into the end-to-end infrastructure supporting your applications.

More than two dozen icons are used to represent different entity types. Colors vary based upon alarm conditions and node selections.

See the individual user guide for each integration that you install for an equivalent list of entity icons for the integration.

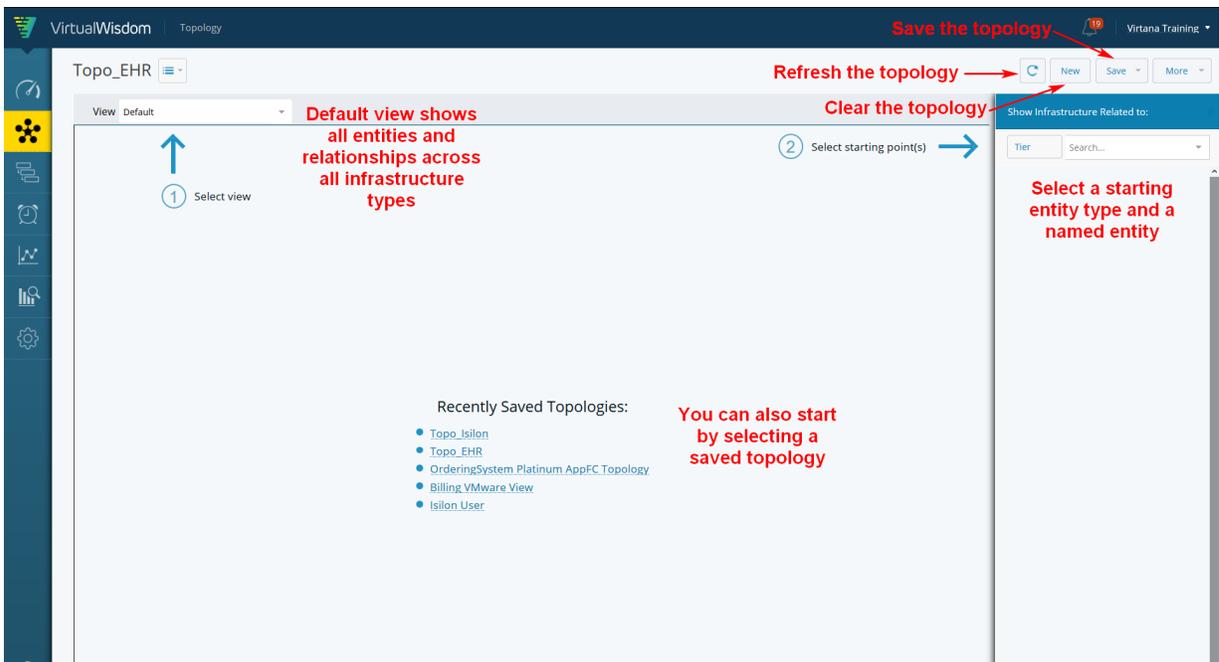
Entity Type	Icon
app	
cache	
conversation	
cpu	
disk	
fabric	
fabric card	
fabric port	
grouped cache	
grouped cpu	
grouped disk	
grouped fabric-card	
grouped fabric-port	
grouped fabric	

Entity Type	Icon
grouped host	
grouped host card	
grouped host port	
grouped storage	
grouped storage card	
grouped storage port	
grouped VM	
host	
host card	
host port	
Netapp	
storage	
storage card	
storage port	
tier	

Entity Type	Icon
VM	

Topology Landing Page

Navigate to Topology by selecting the topology icon in the left-hand navigation panel. You can also navigate to Topology directly from other VirtualWisdom modules, such as Dashboards, Inventory, and Alarms and Cases.



You must select a starting point for the topology. This is done by selecting an entity type and a named entity. You can also start by selecting a saved topology, if one is available.

The default view shows you all of the related entities across all integrations.

The “Show Infrastructure Related to” pane focuses the topology on a specified entity type, showing only entities related to that entity type.

There are two buttons at the top to clear the topology view and start over and to save the topology view so it can be used in the future.

You can refresh the data displayed in the topology by clicking the half circle icon.

Selecting a Topology View and Entity

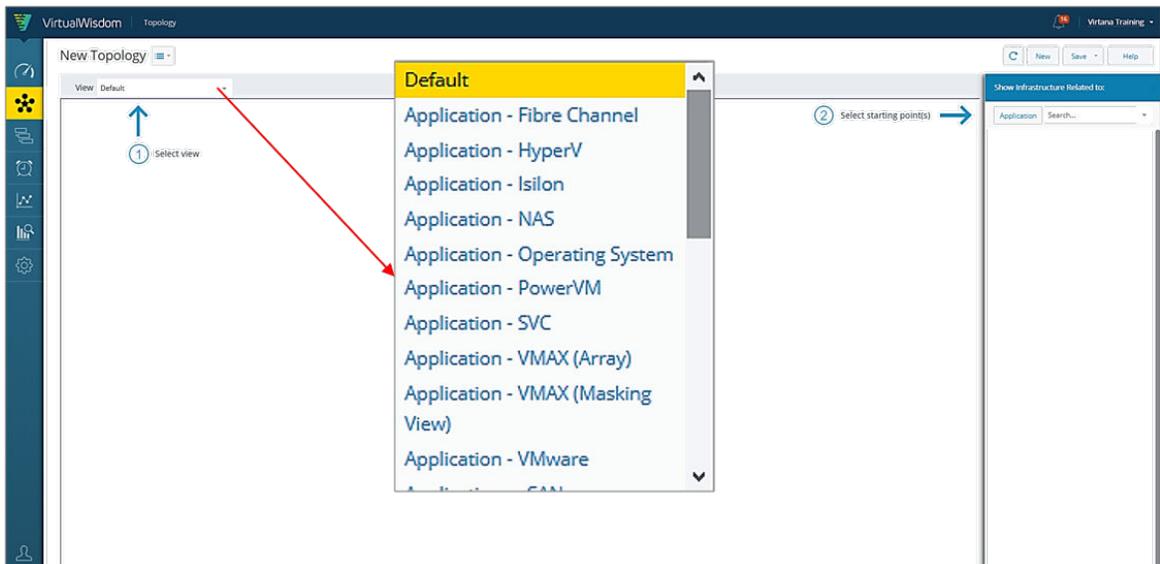
Each topology map is tailored to a given technology, or view, that you select. Views allow you to visualize the datapaths and how various entities interact within that technology, all the way to the storage level. Topologies display hierarchically, from left to right.

As an example, you might view the datapaths of all of your Tier 0 applications from the perspective of your Fibre Channel network, then change the selections to view Tier 0 application datapaths on NAS technology.

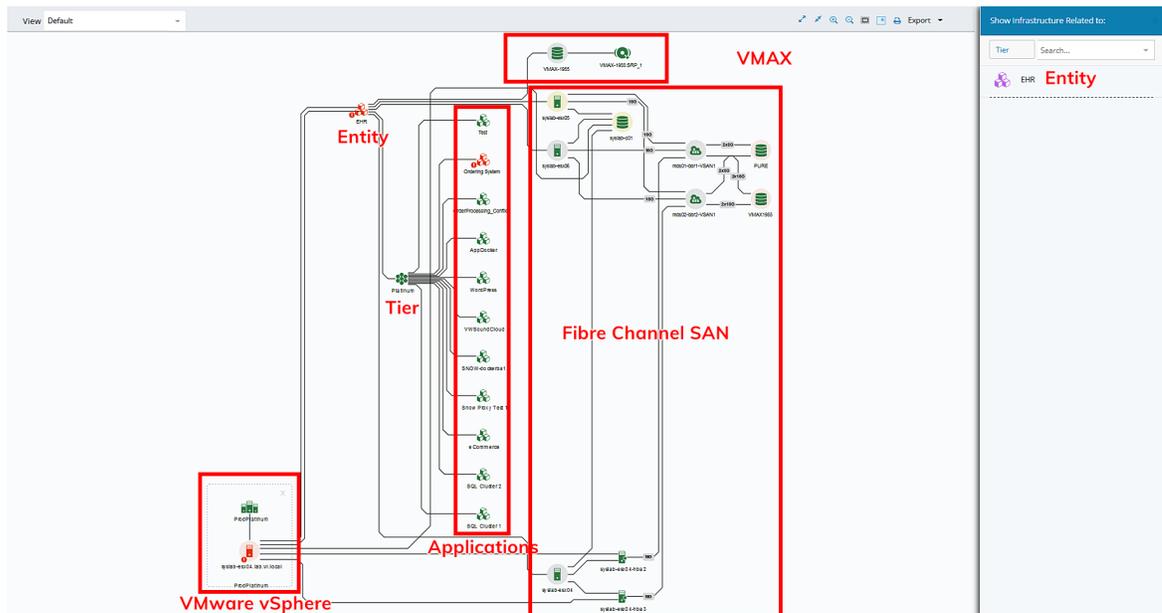
You can also temporarily filter out various entities that display in the topology map. This can provide a simplified view of segments of a complex topology.

You can either create a new topology map to view, or you can choose a topology map that was previously saved.

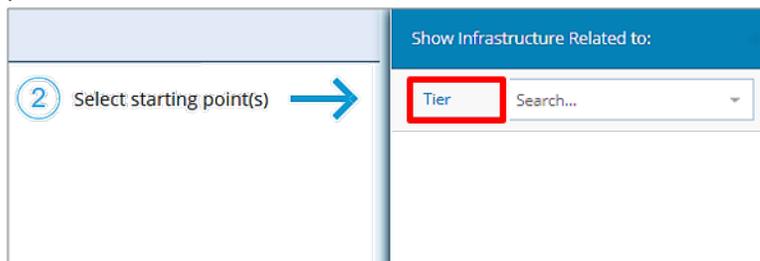
1. Select the arrow on the View field to select a view. The Default view shows you all of the entities and their relationships across all integrations.



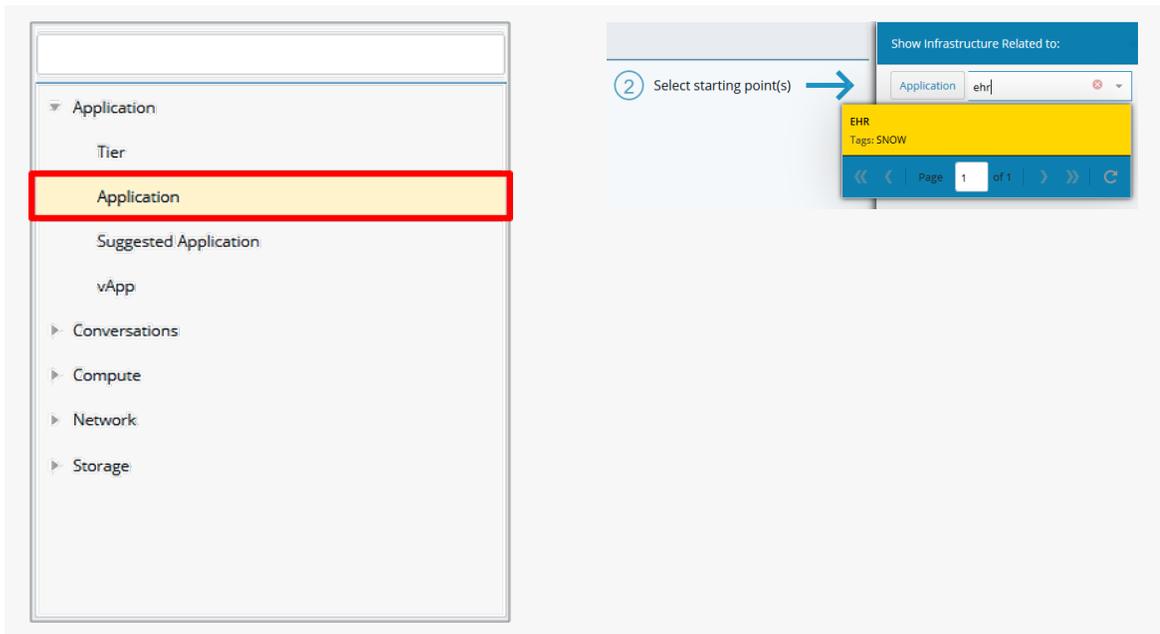
2. All infrastructure types are displayed, e.g., VMware vSphere, FC SAN, and VMAX, as shown in the image above.



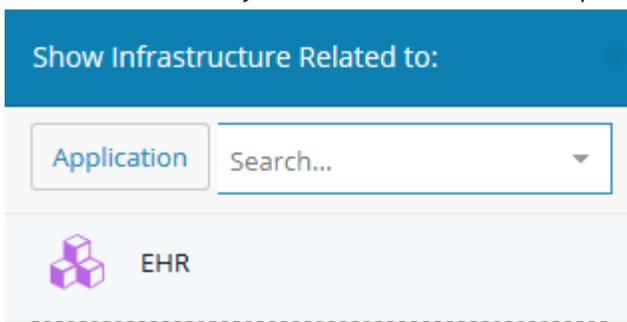
3. Select a starting point for your topology using the **Show Infrastructure Related to** pane.



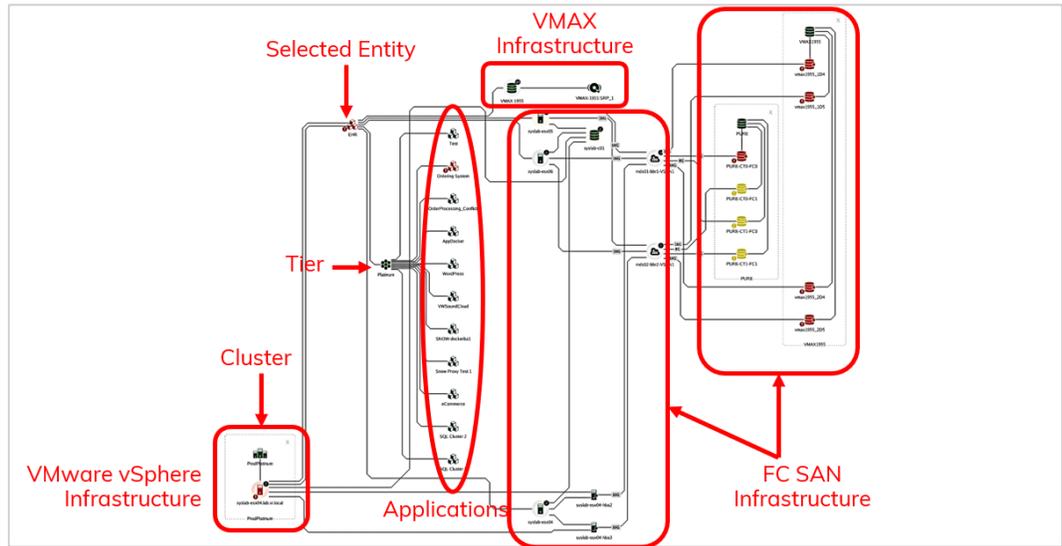
4. Choose an entity type then use the drop down arrow to select an entity from the list of existing entities.



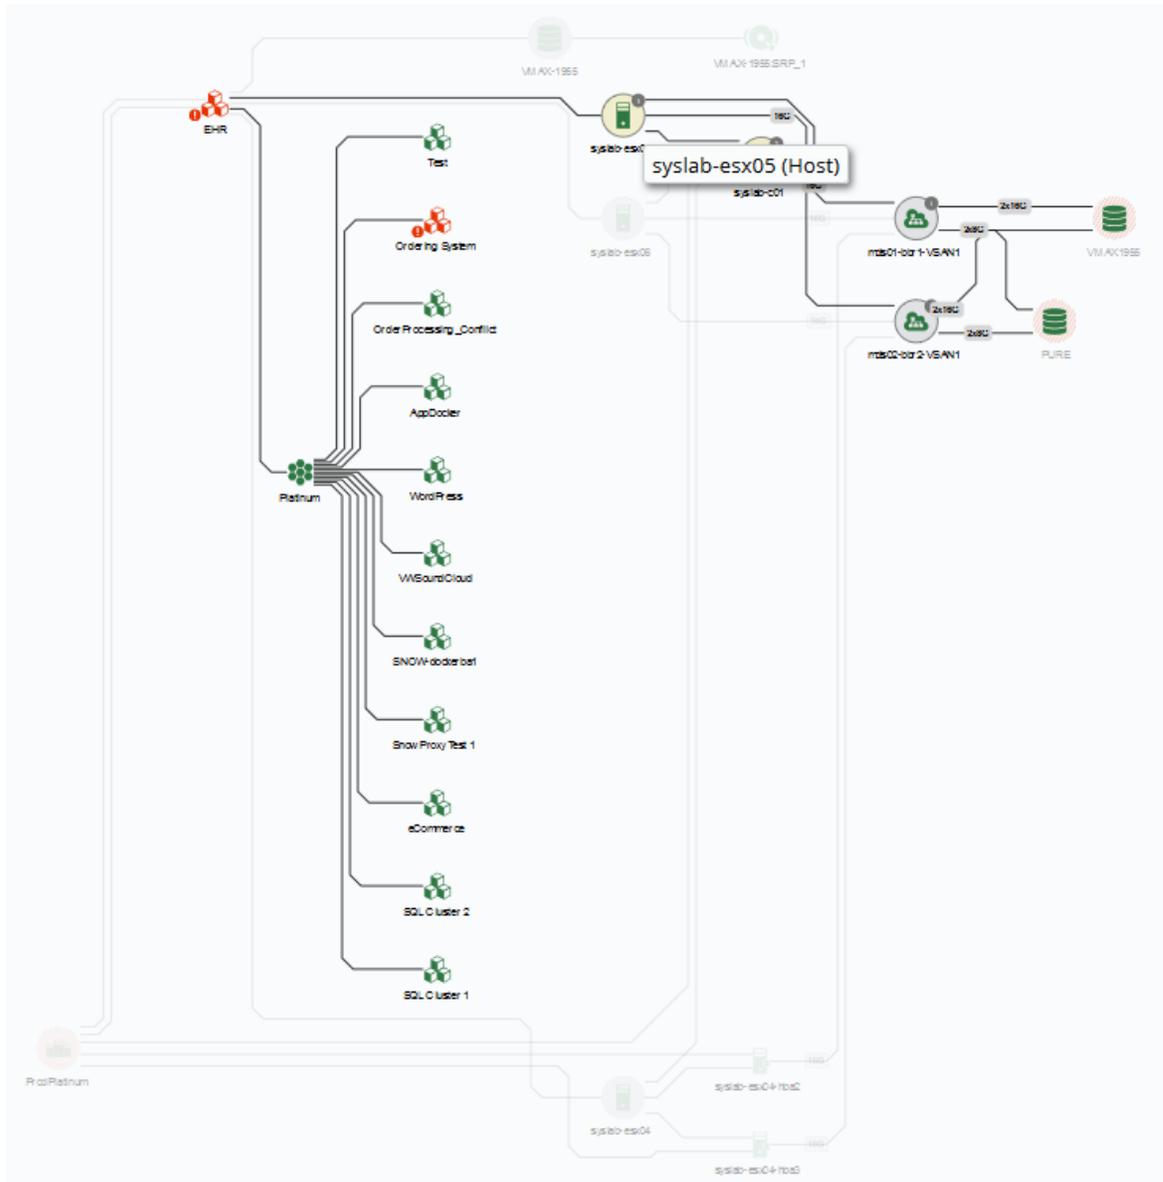
- The selected entity's name is shown in the panel on the right side of the page.



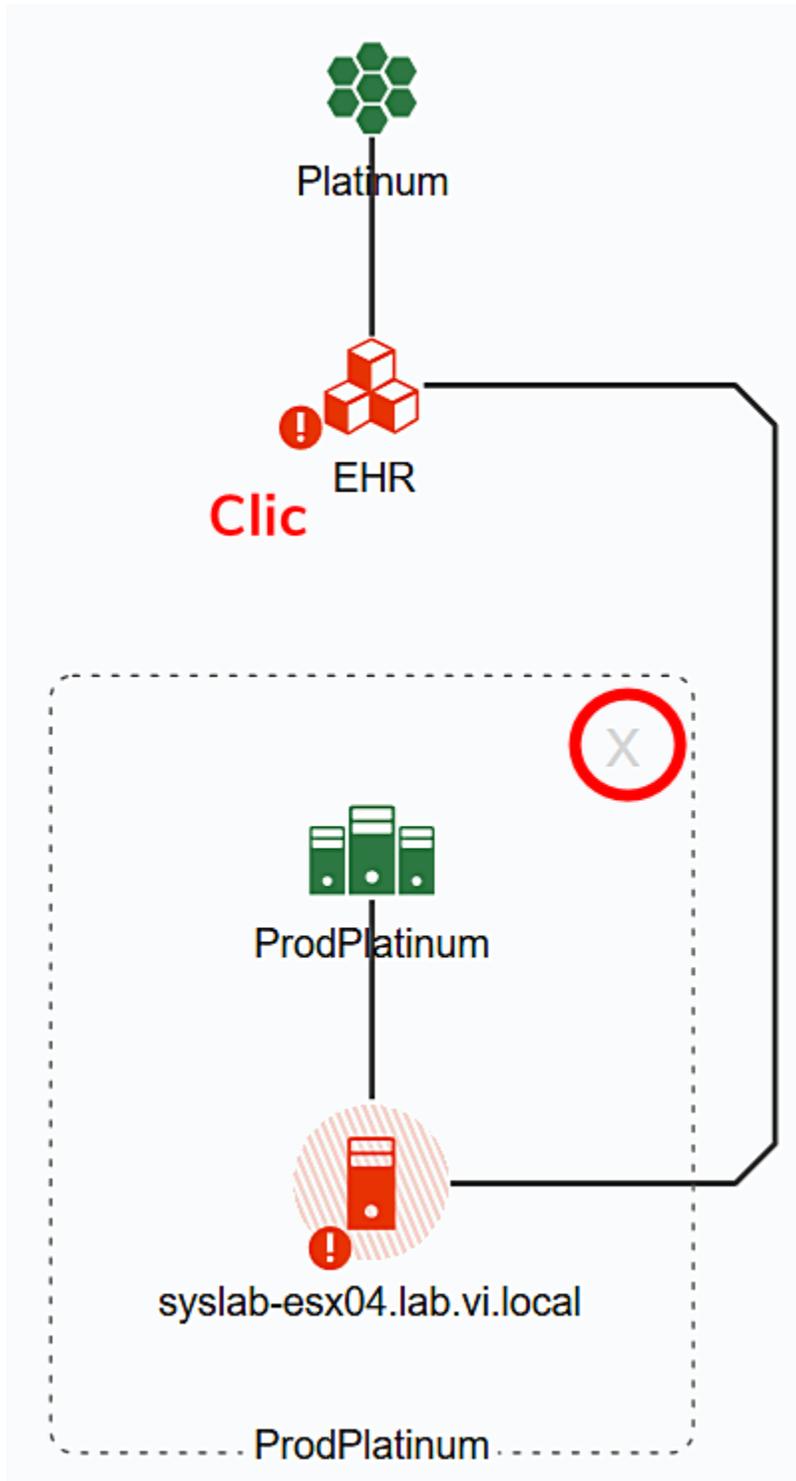
- The window pane on the left shows you the end-to-end infrastructure supporting the application across all infrastructure types. It also shows you the other applications that belong to the tier.



7. Hovering over an entity highlights the entity and its relationships across the end-to-end infrastructure while the unrelated entities fade into the background.

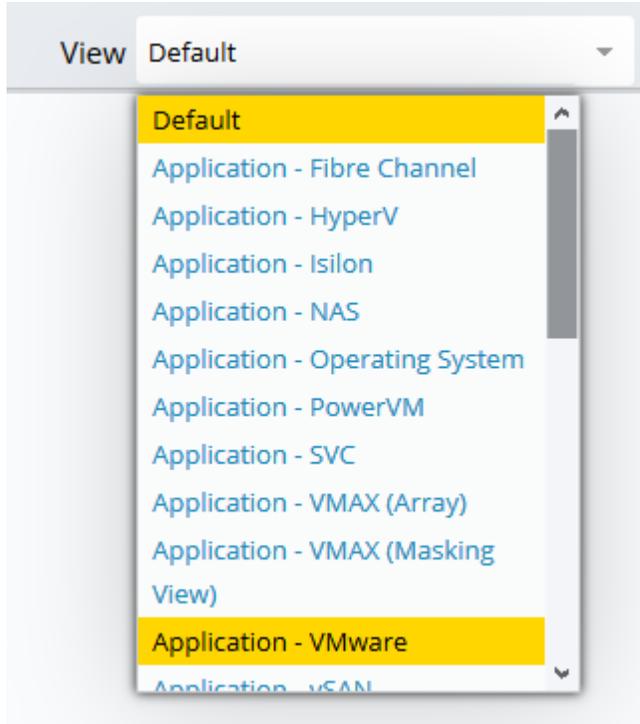


8. Drill down on an entity to expand the view to include its sub-entities. A dotted line denotes the topology for the entity. You may be able to drill down multiple times. Click the x to return to the higher-level view.

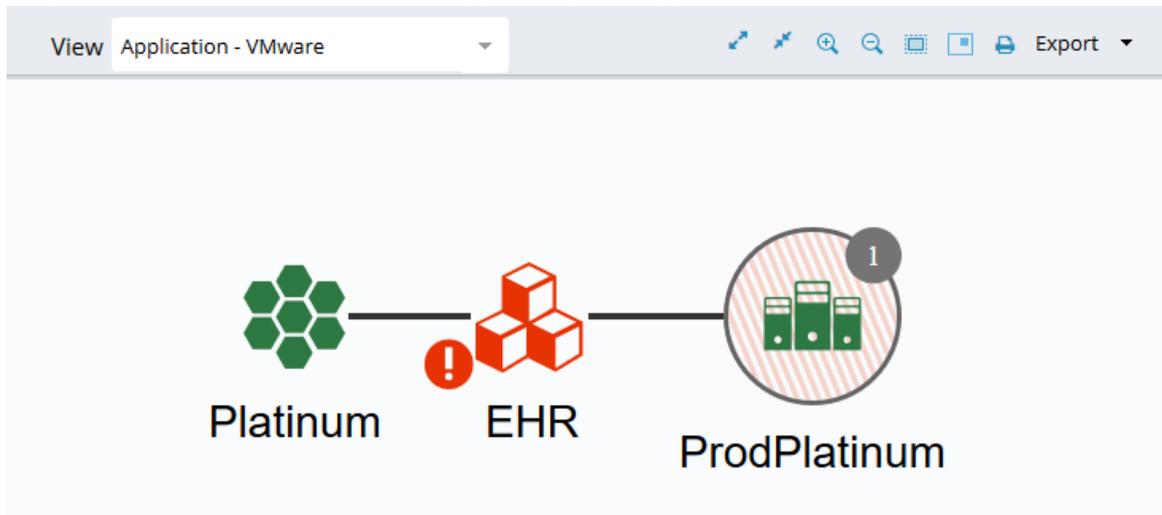


Selecting an alternate view

1. You can switch to a different view of the infrastructure by clicking on the down arrow in the View field.



2. The selected infrastructure will be drawn using data collected from the specified integration. In the example shown below, the same application entity is shown with only the VMware infrastructure supporting the application visible.



Saving Your Topology View

It can be useful to save your topology view to use in the future.

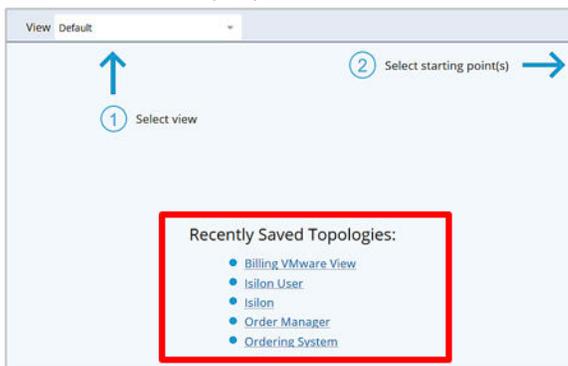
1. To save your view, select Save from the menu at the top right corner.



2. Enter a name and any other attributes you wish and click Save.

 A screenshot of the 'Topology Attributes' dialog box. It has a light blue background. At the top, it says 'Topology Attributes'. There are four input fields: 'Name *' (containing 'Billing VMware View'), 'Description', 'Tags', and 'Access Level' (a dropdown menu). At the bottom, there are two buttons: 'Save' (highlighted in yellow) and 'Cancel'.

3. Your view is displayed in the middle of the topology pane.



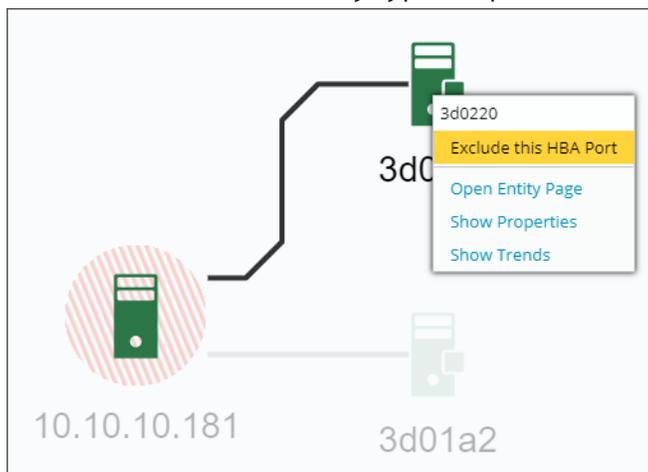
Excluding (Filtering) Nodes

VirtualWisdom provides the ability to further refine the topology you are viewing after selecting one or more starting entities. This allows you to more effectively manage the scope of the topology that you are viewing.

You can further restrict what you are visualizing by using the “exclude” feature to filter entities that you want to temporarily remove from the topology view. When you filter out entities, you also temporarily remove the downstream datapath of the entity.

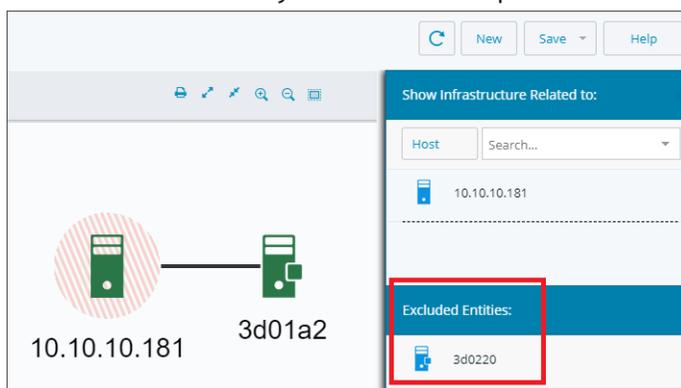
Steps

1. In the topology map, click the entity you wish to filter.
A popup menu appears, and some related entities in the map might be deemphasized, appearing grayed out.
2. Select “Exclude this <entity type>” option.



The entity is added to the Excluded Entities list and is removed from the topology map.

When filtering out an entity, all of the entities in the downstream datapath are also excluded, unless they are in the datapath of another included entity.



3. Repeat as desired to filter out additional entities.
4. To remove the filter:
 - a. Hover over the entity name in the Excluded Entities list.
 - b. Click the x icon associated with the name.

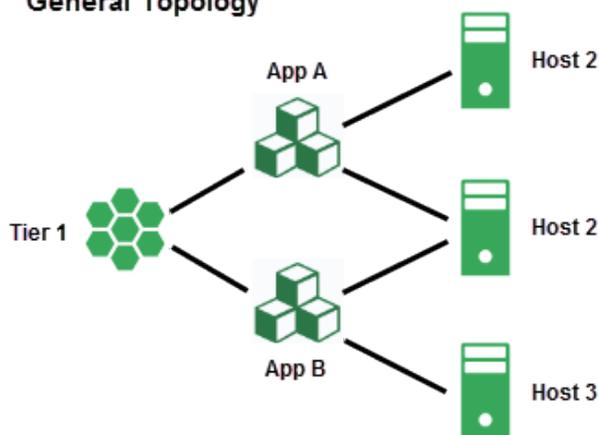
The excluded entity is removed from the list and the topology map reapplies the entity and all its paths.

Example

Assume you have the following environment:

- 3 Hosts: Host 1, Host 2, and Host 3
- 2 Applications: Application A and Application B
- 1 Tier: Tier 1
- Application A and Application B are members of Tier 1
- Host 1 and Host 2 are members of Application A
- Host 2 and Host 3 are members of Application B

General Topology

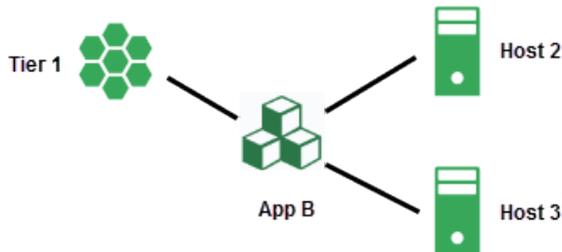


Scenario 1

You select Fibre Channel in the topology View field. Under “Show Infrastructure Related to” you select Tier as the entity type and Tier 1 as the starting entity instance. A topology displays (see General Topology above) with Tier 1, Applications A and B, and Hosts 1, 2, and 3.

If you filter out Application A, you see Tier 1, Application B, Hosts 2 and 3. You still see Host 2 because it is also a member of Application B, however you won't see Host 1 since that was in the downstream datapath from Application A, which was filtered out.

Topology: Application perspective, excluding App A

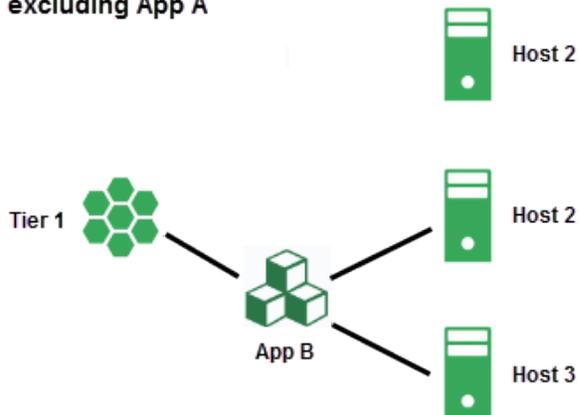


Scenario 2:

Again, you select Fibre Channel in the topology View field, but you select Host as the entity type and Hosts 1, 2, and 3 as the starting entities. The initial general topology looks similar to the one above, but because the entity type is Host, the perspective is from the hosts.

If you then filter out Application A, you would see Host 1 without a connection to Tier 1. The Tier 1 connection is excluded because it comes after Application A from the perspective of Host 1.

Topology: Host perspective, excluding App A



Entity Grouping

Topologies with a large number of entities can be difficult to view. To simplify the topologies, entities which are physically or logically contained within another entity are collapsed into a single entity group by default.

- Single entities are represented by a single icon.

For example, the following represents a cluster:



- Grouped entities are represented by an icon, surrounded by a circular halo.

For example, the following represents a cluster with grouped entities:

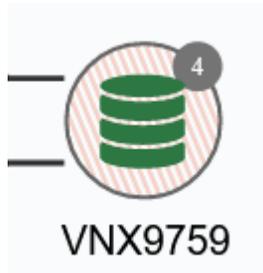


Hovering over a grouped entity shows you the number of child entities contained within that particular group.

Expanding an Entity Grouping

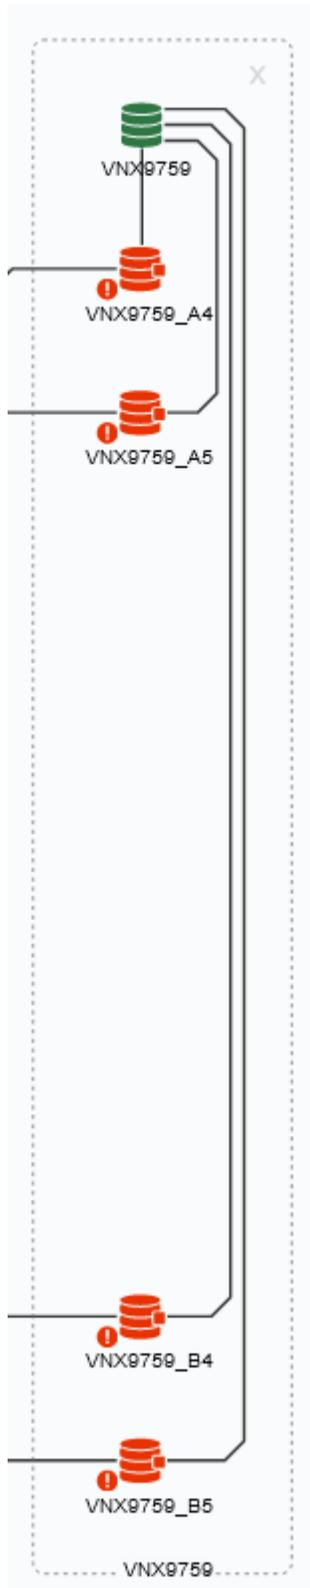
Grouped entities can be expanded in order to view the components of the group and the way they relate to each other. It is possible to have multiple levels of nested groups.

1. Identify the group to expand.

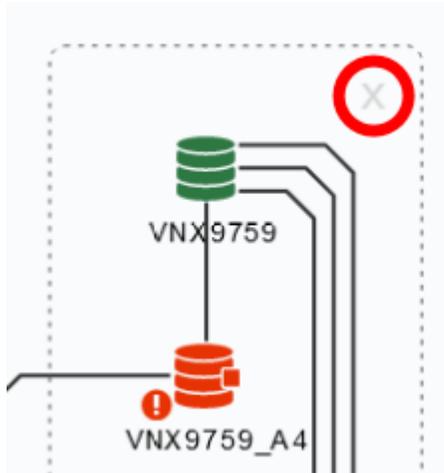


Any entity with grouped child entities has a halo around it and displays a number in the upper right if you hover over the icon.

2. Expand the group by double-clicking on the group icon.
The entity and its grouped child entities display, surrounded by a dotted-line expansion box.
If any of the child entities are grouped, you can also expand them.

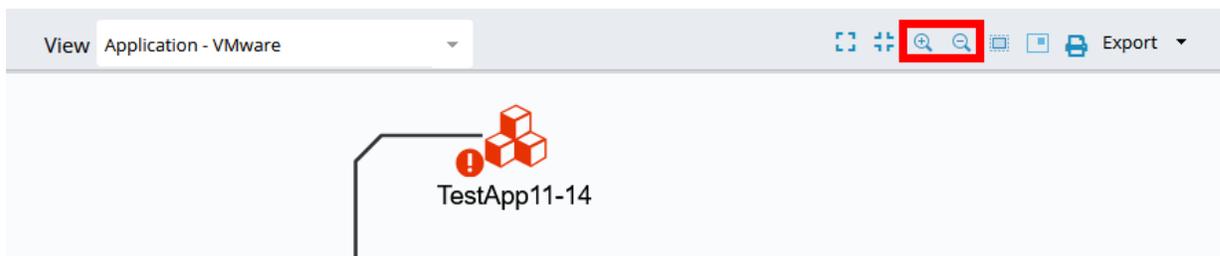


3. Click the X inside the expansion box to close the expanded grouping.



Using the Topology Map Controls

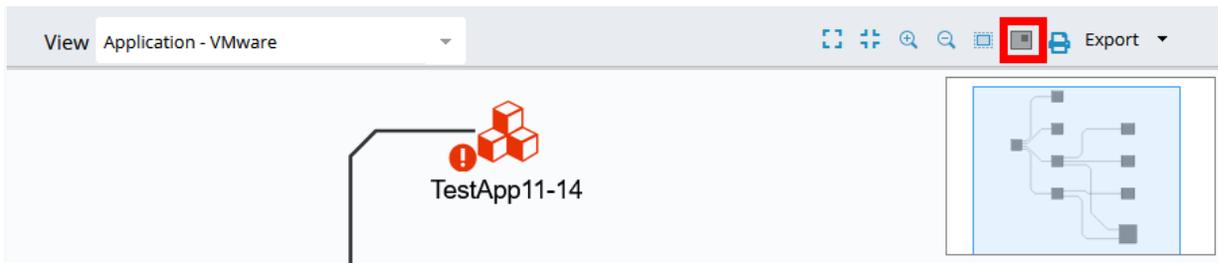
You can zoom in or out on topology by using the magnifying glass icons on the upper right.



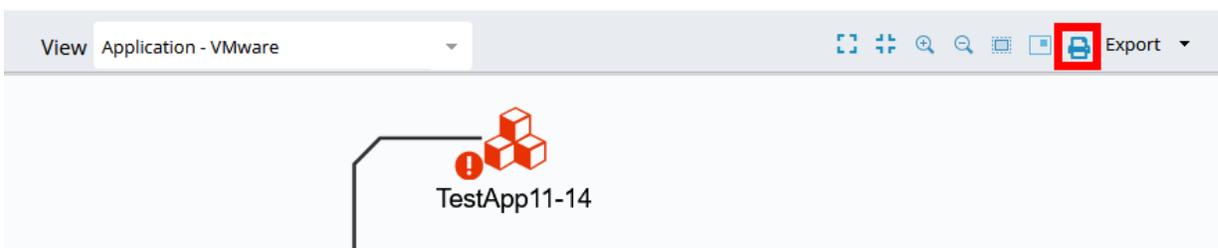
There is also a "fit to screen" option that shows the entire topology in the Topology map.



The "mini-map" icon lets you see and navigate the entire topology while displaying only a portion of it using a movable thumbnail.



You can display the entire topology view in a full screen browser window by using the Print icon.

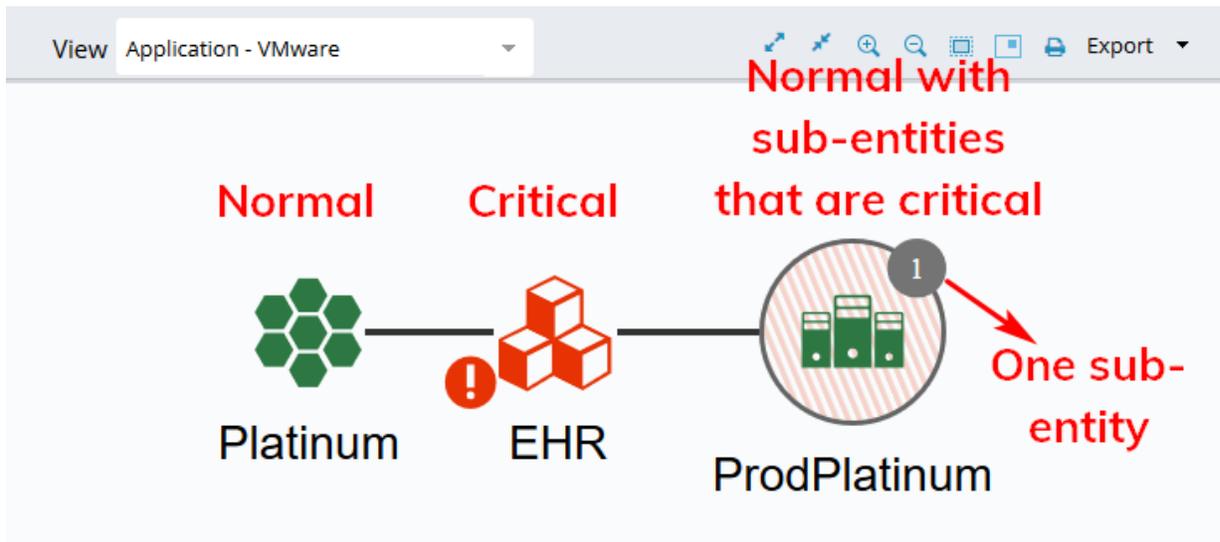


You can also export the topology view as a PNG, JPG, PDF, or SVG file.



Understanding Entity Status

The color of the entity's icon and its halo denotes its status and sub-entity relationships, just as they did in the Inventory view.



Entities with a halo have sub-entities. The number of sub-entities is shown in a circle next to the entity.

The color of the entity's icon indicates its status:

- Red indicates that the entity's status is critical
- Yellow indicates a warning
- Green indicates normal

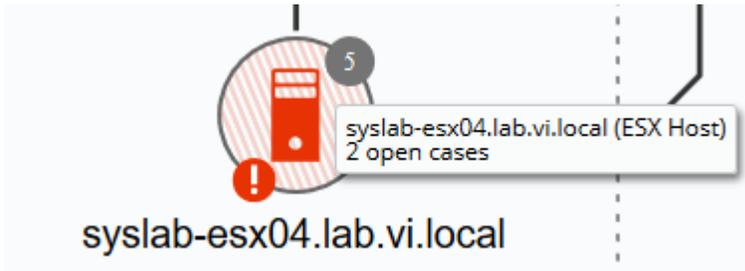
The color of the entity's halo indicates the status of its status:

- Red indicates that the entity has one or more sub-entities that are critical
- Yellow indicates that the entity has one or more sub-entities with warnings
- Green indicates that the entity's sub-entities are normal

Viewing Entity Data

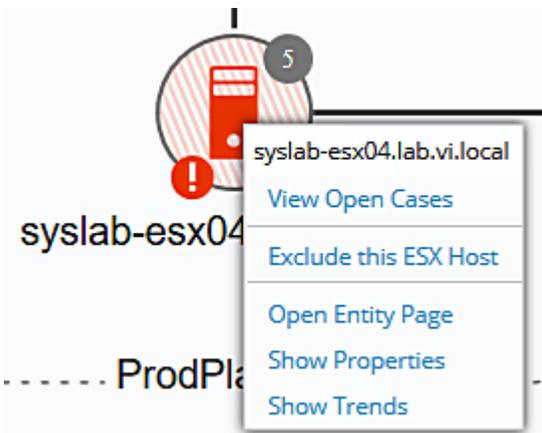
Hovering over an entity reveals the following data:

- The entity's type.
- The number of cases currently open for the entity.

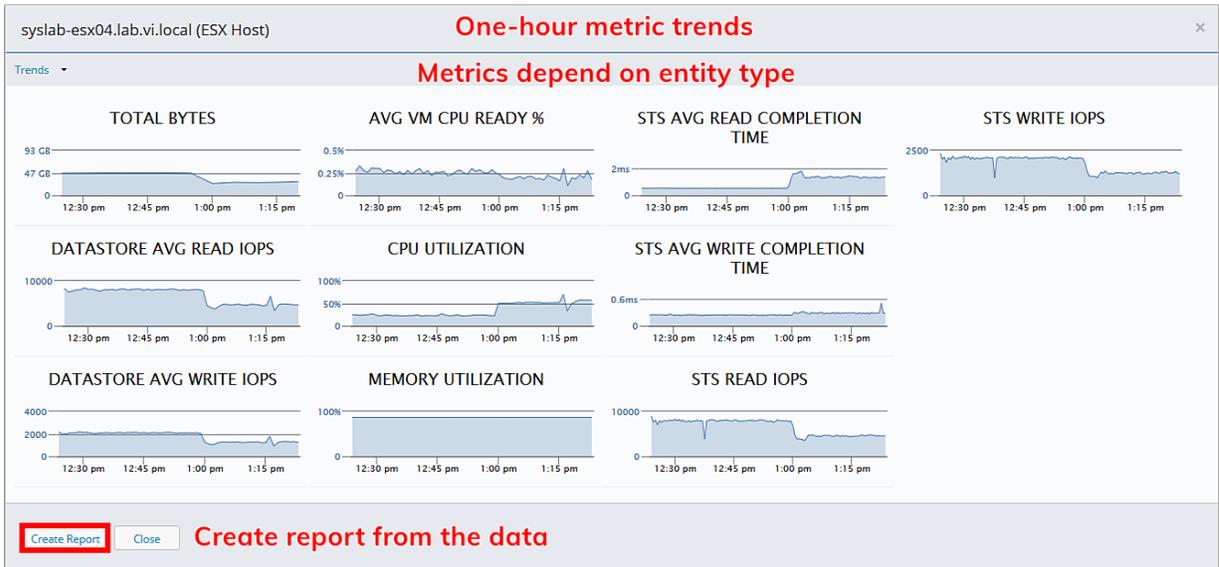


Clicking on an entity displays a menu of actions you can perform on the entity:

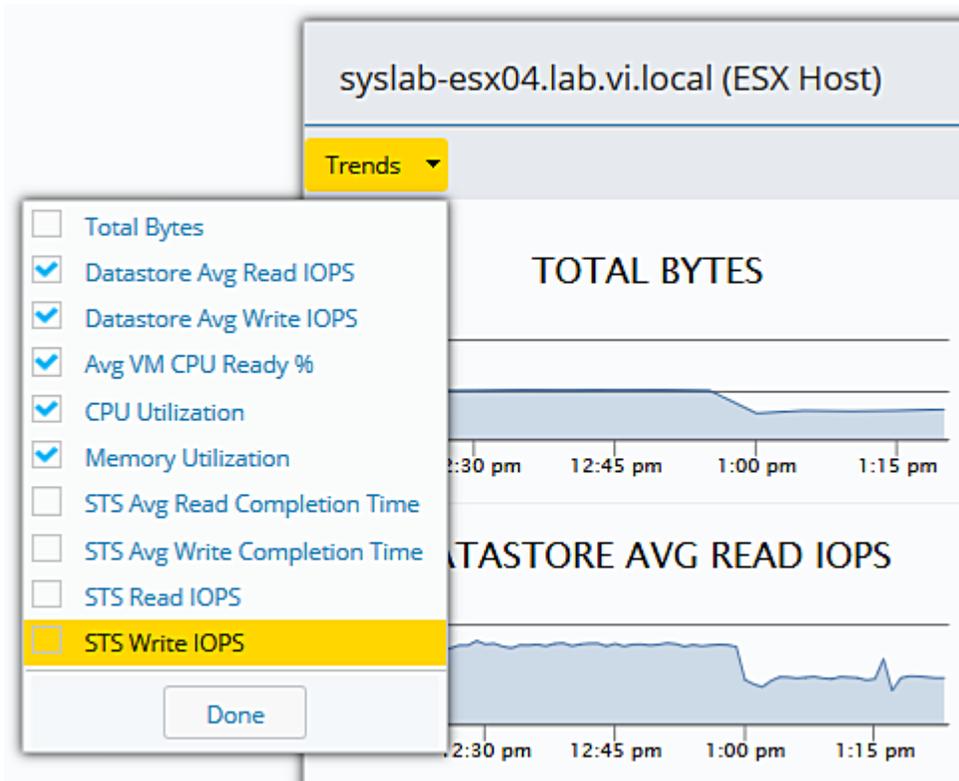
- View Open Cases displays a list of current open case on the entity.
- Exclude this entity hides the entity from the topology view.
- Open Entity Page takes you to the entity's inventory page.
- Show Properties displays a pop-up window that shows system properties of the entity.
- Show Trends displays a pop-up window with one-hour metric trends for the entity.



The pop-up window shows one-hour metric trends for the entity. The available trends depend on the entity type. You can open multiple trend windows to facilitate trend comparison. Note that trend windows must be explicitly closed, otherwise they remain open while navigating to other areas of the application.



You can select which trend charts to display and create a report from the trend charts.

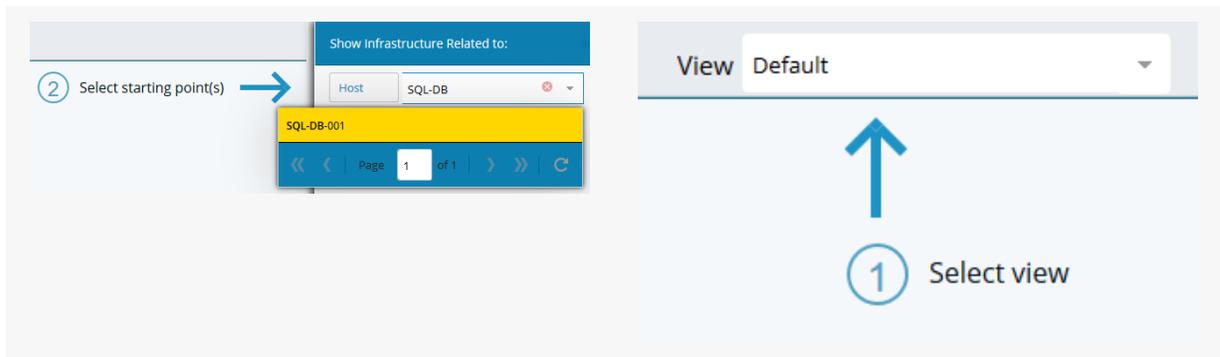


Topology Use Case

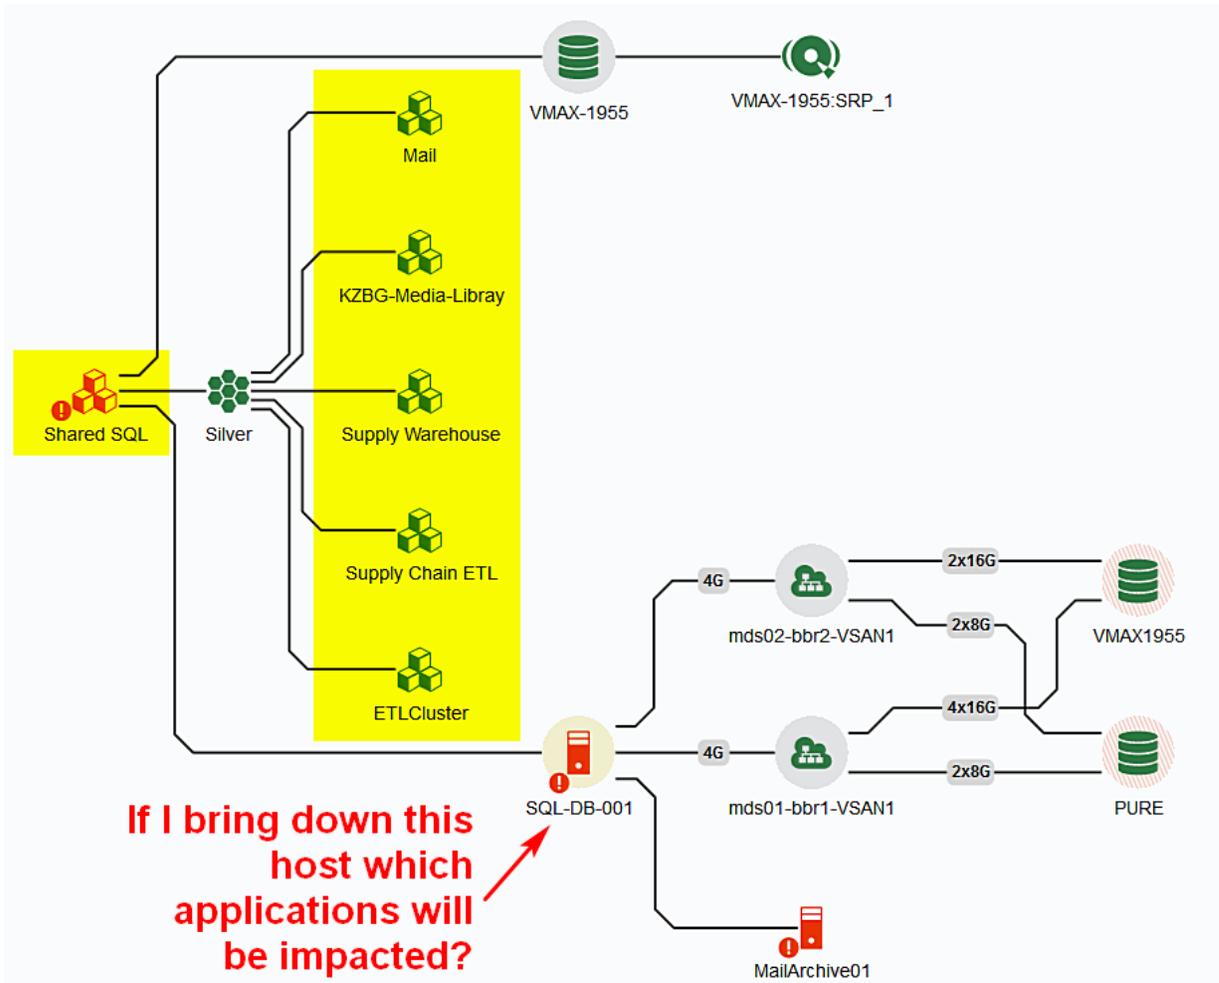
Let's look at a practical application of VirtualWisdom topology. In this use case, we'll answer the question: "If I bring down the SQL-DB-001 host for maintenance, which applications will be impacted?"

We'll need to show the topology for all infrastructure related to the host.

Select the host as our starting point for the topology. Choose the Default view to view all its related infrastructure across all integrations. The end-to-end topology is displayed, including all the applications that are related to the SQL-DB-001 host.

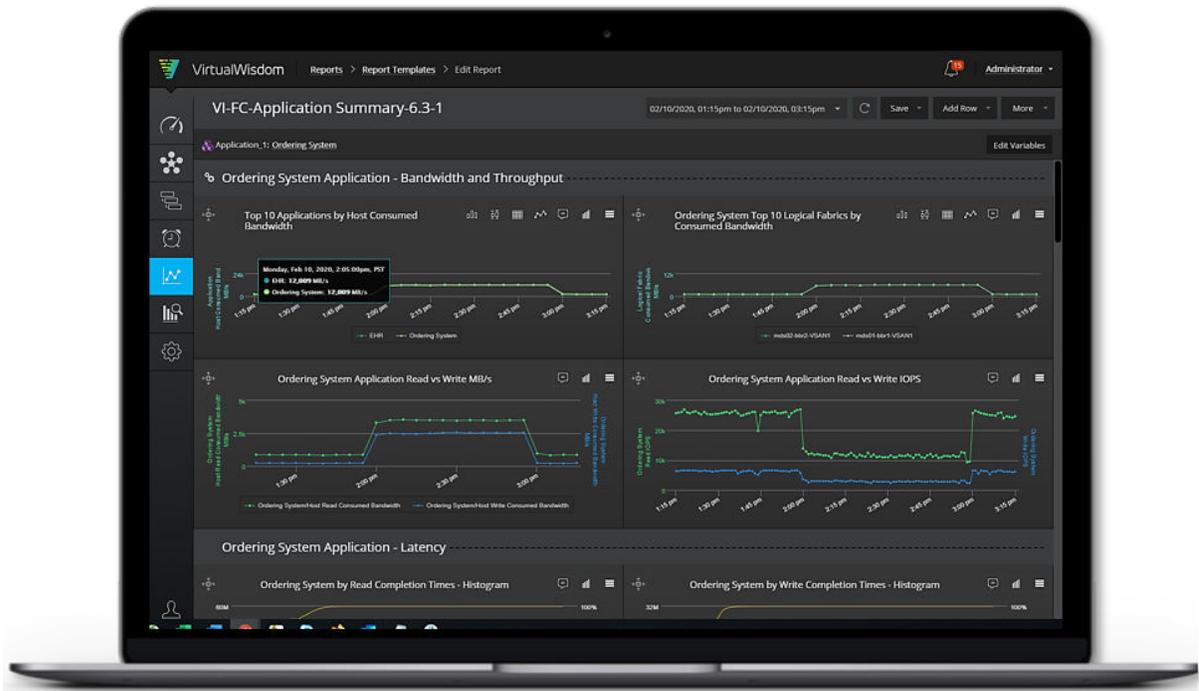


We observe that there are six applications that could be impacted if we bring down the host for maintenance.



Reports

VirtualWisdom provides live reporting capabilities that leverage the data collected by the VirtualWisdom probes and integrations.

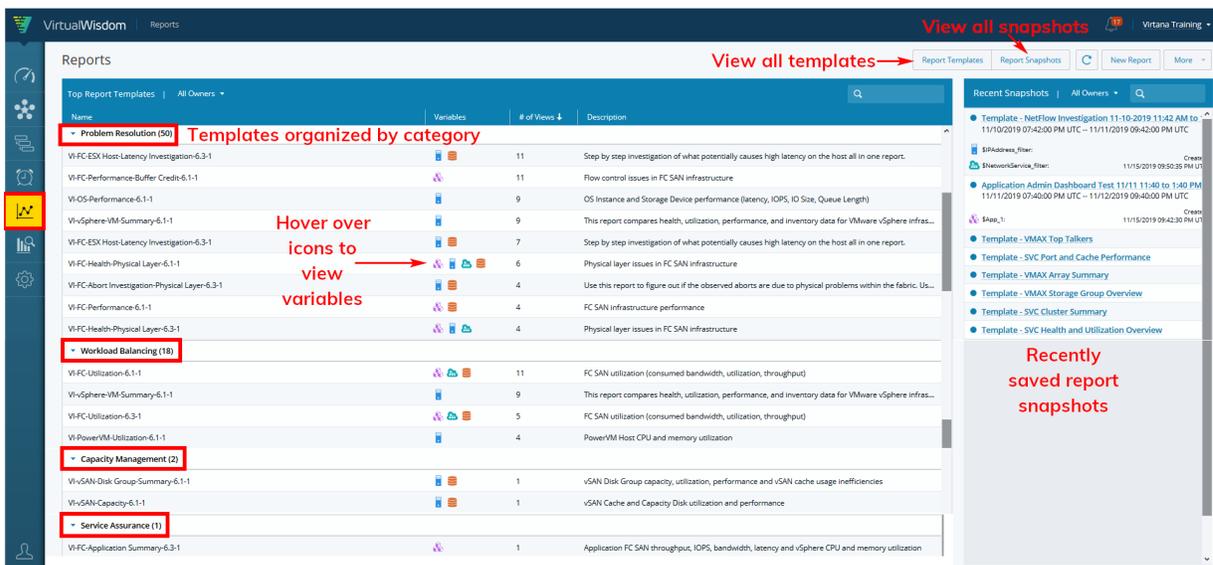


Reporting capabilities include multiple chart types designed to show and use VirtualWisdom collected data to perform analysis of application infrastructure health, utilization, and performance.

Our report framework lets the user organize their reports in an easy to find manner. Users can share reports and templates they have created with other users and groups. Reports can also be easily exported.

The most recently run reports are also easily accessible, letting users know what’s important now.

Reports Home Page



The Reports home page displays a list of report templates that you can use to start with. VirtualWisdom includes standard reports created by the Virtana Professional Services team, designed to help you resolve problems, balance workloads, manage capacity, and assure service levels. The report templates are organized into groups.

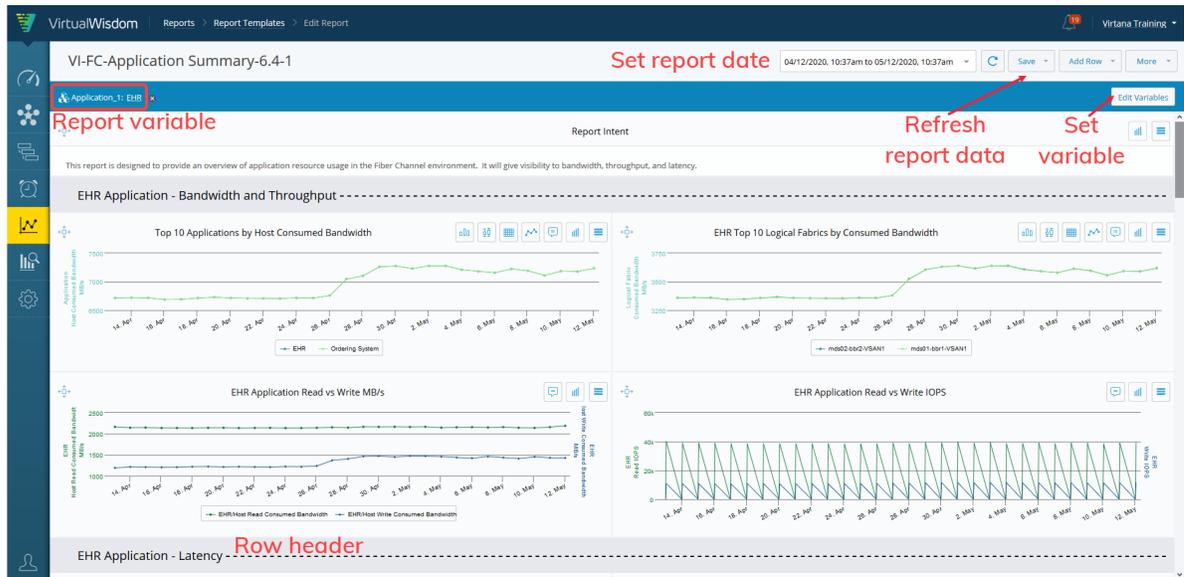
The list also shows you whether the report uses an entity variable and which type of entity is used for a variable. You can also see all report templates by selecting the Report Templates button at the top of the page.

To the right of the report templates list is a pane that displays recently saved report snapshots. A report snapshot captures the “point in time” output of a report. You can view all report snapshots by selecting the Report Snapshots button at the top of the page.

To view a report, click on its row in the list.

Report Page

Each report has a dedicated page that shows you its output. A report can include up to 15 rows.

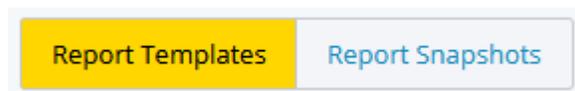


The report header includes a date field to set the date and time range for the report. You can select from a default range from 5 minutes up to 30 days or set a custom range. Use the circular arrow to refresh a report’s data any time you wish to view new data.

The report variable field lets you select a specific entity to filter the charts on the report. A chart may include a filter that is based on a report variable. This makes it easier to filter multiple charts by the same entity and change that entity easily.

The row headers are used to provide more information about the charts shown below them.

Report Templates



The Report Templates pages displays a list of all the saved report templates in your portal. You can use a report template as a starting point for VirtualWisdom reports.

Report templates are already populated with charts, entities, and metrics. Many include report variables that can be used to filter the report for a specific entity or set of entities.

Report Templates

Report Templates Report Snapshots New Report Edit More

(314 items) Searches Name, Owner, Tags fields Sort by these fields

Name	Owner	Access Level	Tags	Scheduled	Created On
<input type="checkbox"/> VI-PowerVM-Utilization-6.6-1	vi.training	Read, Write	Workload Balancing	--	10/02/2020 03:29:54 PM PDT
<input type="checkbox"/> Data Source Verify 6.3	vitoo.s	Read, Write		--	09/30/2020 03:24:46 PM PDT
<input type="checkbox"/> Application Rollup Investigation	administrator	Read, Write		--	09/28/2020 11:53:40 AM PDT
<input type="checkbox"/> Storage and Switch Inventory - For David D ...	marc.bachmeier	Read, Write		--	09/18/2020 10:48:16 AM PDT
<input type="checkbox"/> Template - NetApp Cluster Utilization and C...	administrator	Read, Write	Default,NetApp ONTAP	--	09/09/2020 08:08:26 AM PDT
<input type="checkbox"/> Specialty Dashboard - VMware Datastores	cdellaquila	Read, Write		--	09/03/2020 06:59:05 AM PDT
<input type="checkbox"/> Specialty Dashboard - VMware ESX VM Gue...	cdellaquila	Read, Write		--	09/03/2020 06:58:40 AM PDT
<input type="checkbox"/> Specialty Dashboard - VMware ESX Host	cdellaquila	Read, Write		--	09/03/2020 06:58:03 AM PDT
<input type="checkbox"/> Specialty Dashboard - VMware Cluster	cdellaquila	Read, Write		--	09/03/2020 06:57:32 AM PDT
<input type="checkbox"/> Error Investigation - Host Application ISCSI ...	cdellaquila	Read, Write		--	08/31/2020 08:39:07 AM PDT
<input type="checkbox"/> Error Investigation - Host Application HBA ...	cdellaquila	Read, Write		--	08/31/2020 08:38:21 AM PDT
<input type="checkbox"/> Executive Dashboard - ISCSI Application an...	cdellaquila	Read, Write	CDA	--	08/31/2020 07:02:54 AM PDT
<input type="checkbox"/> ISCSI SW Integration Metrics	cdellaquila	Read, Write	CDA	--	08/31/2020 06:49:50 AM PDT
<input type="checkbox"/> ISCSI Orphaned LUN Report	cdellaquila	Read, Write	CDA	--	08/31/2020 06:36:58 AM PDT

From the **Report Templates** page you can perform the following actions:

- Create new report templates
- Run a report template
- Edit a report template
- Save a copy of a report template
- Delete a report template
- Bulk edit report templates

Bulk Editing and Deleting Report Templates

You can add or remove tags from multiple report templates using the bulk edit feature.

1. Select the reports to bulk edit by using the check boxes next to their names in the list view.
2. Click **Edit**, then click **Add Tags** or **Remove Tags**.

Report Templates

Report Templates | Report Snapshots | Refresh | New Report | Edit | More

test (9 Items | 3 selected)

	Name	Owner	Access Level ↑	Tags	Scheduled	Created On
<input type="checkbox"/>	ProbeNAS 2x40G to 4x10G ...	vitoo.s	Read, Write	System Lab	--	10/02/2018 12:10:29 PM PDT
<input type="checkbox"/>	dashboard testing	administrator	Read, Write	System Lab	--	12/17/2019 07:49:00 PM PST
<input type="checkbox"/>	F5 Testing	administrator	Read, Write	System Lab	--	04/19/2019 04:28:56 PM PDT
<input type="checkbox"/>	StackPlay Testing	administrator	Read, Write	System Lab	--	01/02/2020 01:14:54 PM PST
<input type="checkbox"/>	NetFlow Test -LDX	vitoo.s	Read, Write		--	03/04/2020 09:36:50 PM PST
<input type="checkbox"/>	inventory chart test	david.ravlen	Read, Write		--	06/29/2020 02:00:38 AM PDT
<input checked="" type="checkbox"/>	Application SDS Performan...	administrator	Read, Write	Default, Test	--	12/05/2018 08:03:22 AM PST
<input checked="" type="checkbox"/>	Template - App Storage Pe...	administrator	Read, Write	Default, Test	--	11/26/2017 05:00:01 PM PST
<input checked="" type="checkbox"/>	Nutanix CVM - Host Integr...	administrator	Read, Write	Default, Test	--	10/07/2019 08:37:51 AM PDT

- You can also delete multiple report templates at once by selecting Delete. A template cannot be restored once it is deleted.

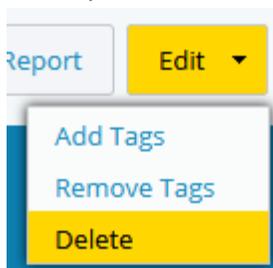
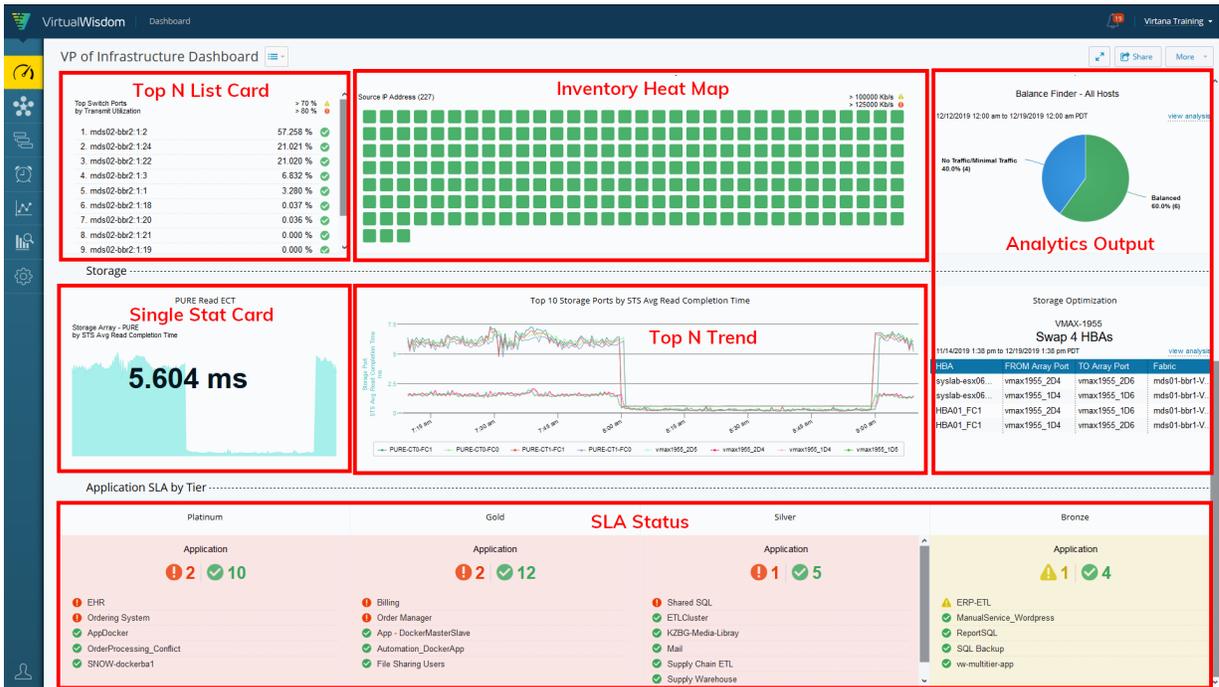


Chart Types

Each report is comprised of one or more rows, with one or more charts in a row. A single row can contain up to four charts.



Multiple rows are included in a report. A report can include up to 15 rows.

There are different chart types that you can use in a report. They are organized into categories:

Select Chart Type

<p>Basic</p> <ul style="list-style-type: none"> Single Stat Card Line Chart Free Form Inventory Bar Inventory Donut Inventory Pie 	<p>Top N</p> <ul style="list-style-type: none"> Bar Chart Box Plot Table List (Card) Trend 	<p>Comparison</p> <ul style="list-style-type: none"> Scatterplot Time Comparison 	<p>SLA Status</p> <ul style="list-style-type: none"> Histogram SLA Status 	<p>Special</p> <ul style="list-style-type: none"> Inventory Heatmap SCSI Status Open Case Summary Capacity Trend Analytics Output Topology
--	--	---	--	---

Chart Categories

Basic	Used to compare entity groupings and metrics, report on entity properties, and show basic text information in a chart.
Top N	Used to view a single metric trend for a group of entities.

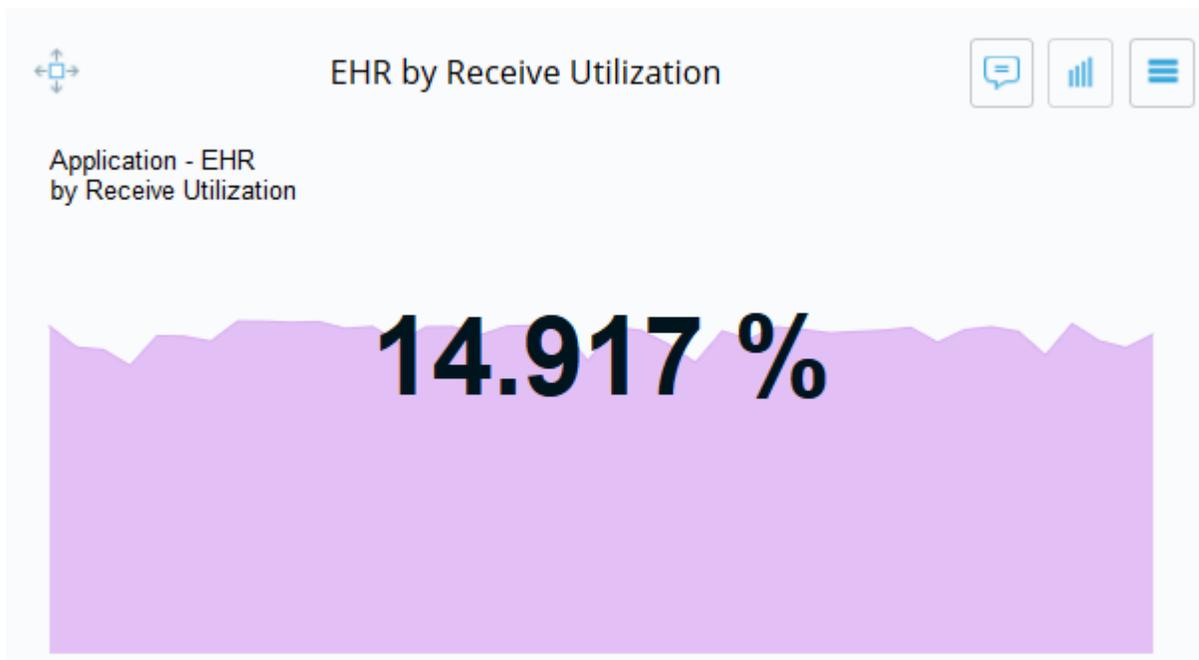
Comparison	Compares an entity-metric binding over time or compares two different metrics for a single entity.
SLA Status	Identifies problems using data that is averaged out. Useful for tracking SLA performance.
Special	Reports on specialized data such as analytics output, open cases, and SCSI status.

Basic Chart Types

The Basic Chart category contains six chart types:

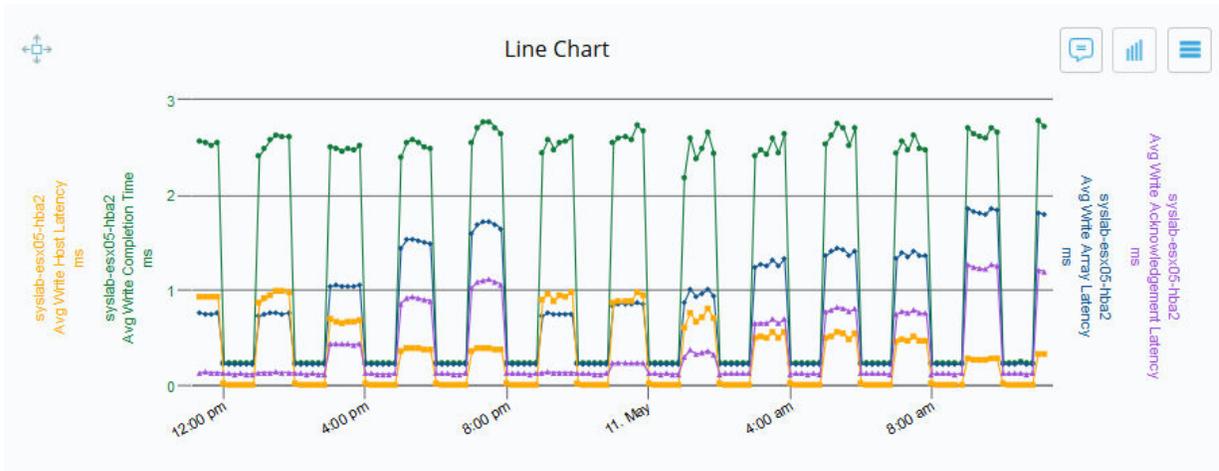
Single Stat Card

The Single Stat Card displays a single statistic that can be read from across the room.



Line Chart

The Line Chart compares different metrics and entity groups (hierarchies) over time to identify patterns that merit further investigation.



Free Form Chart

The Free Form Chart allows you to create a chart with free form content like comments, headers, images, bulleted or numbered lists, check boxes and format the content using different font styles and sizes.

The figure shows the "Free Form Chart" interface. At the top, it has a title "Free Form Chart" and two icons (a bar chart and a list icon). Below the title, there is a text block:

The free form chart type lets you create charts with "free form" content such as comments, headers, lists, tables, images and checkboxes.

 Below this text is an image of a laptop displaying a dashboard with multiple charts. At the bottom of the interface, there is another text block:

You can format this content using different font styles and colors.

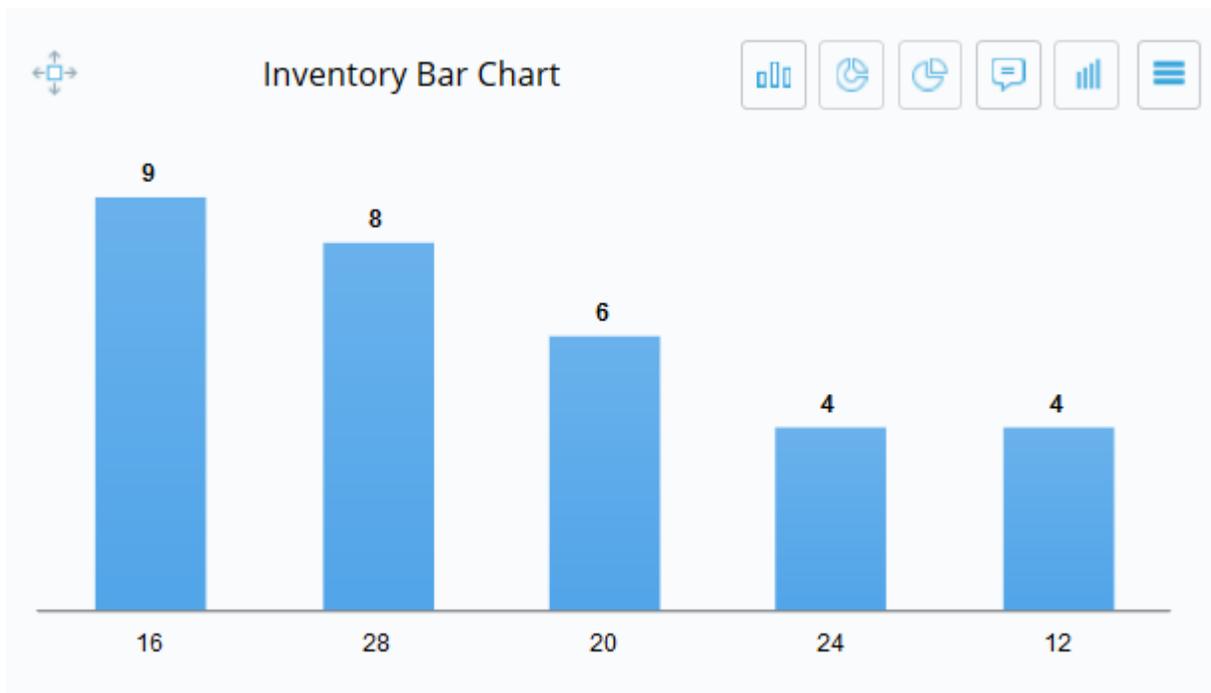
 The interface also includes a vertical scrollbar on the right side.

Inventory Chart Types

The Inventory chart types display information based on entity properties.

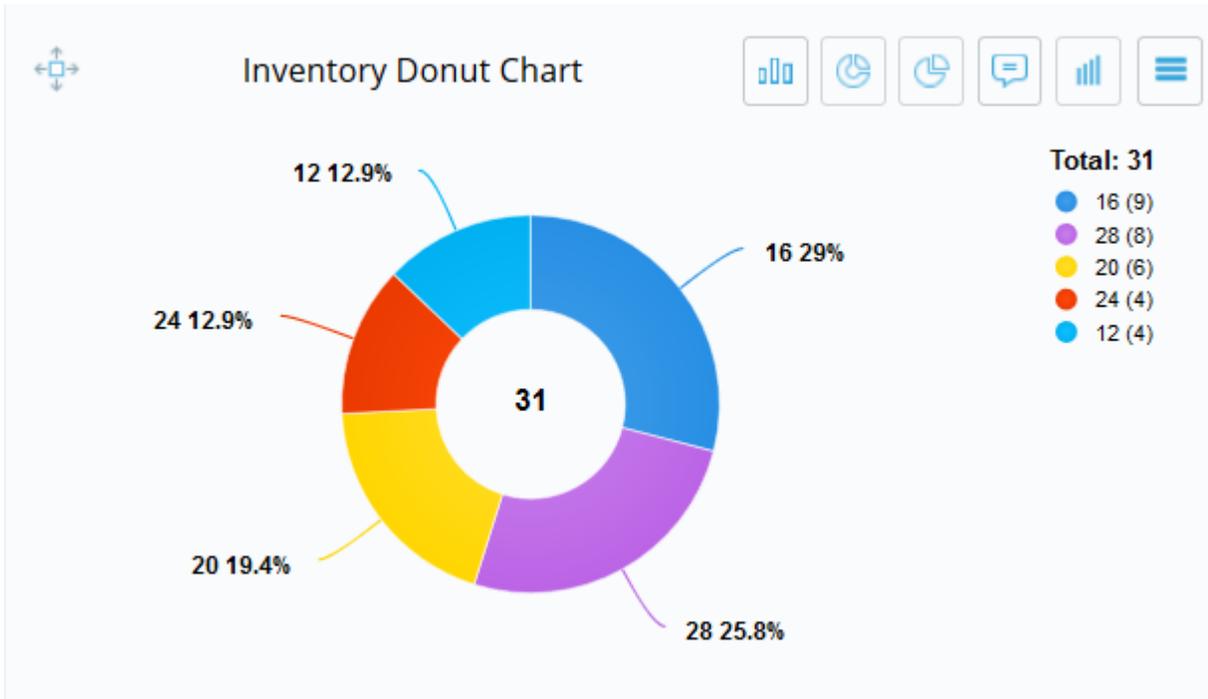
Inventory Bar Chart

The Inventory Bar Chart displays entity system or custom properties in a bar chart. Property types are grouped into bars, with the property value and the corresponding number displayed at each end of the bar.

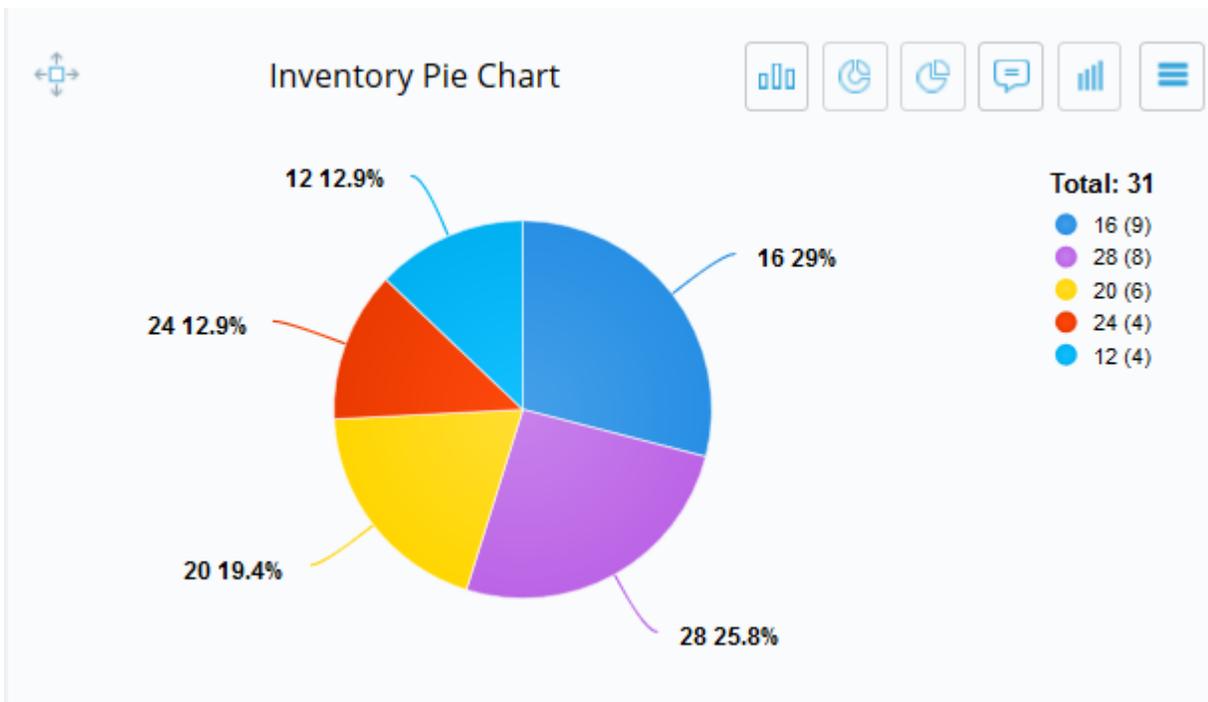


Inventory Donut Chart

The Inventory Donut Chart displays entity system or customer properties in a donut chart. The value displayed in the “donut hole” is the total number of entities examined, while the values shown on the “donut ring” show the property value and the corresponding percentage of entities that have that value.



Inventory Pie Chart



The Inventory Pie Chart is like the donut chart in that it shows the property value and the corresponding percentage of entities that have that value. However, the total number of entities is displayed to the right.

Inventory Table Chart

Use the Inventory Table Chart to view properties and metrics for a selected entity type in a table format.

ESX Host ↑	Number of CPU Packages	CPU Utilization (99th)
qe-appdisc-01.lab.vi.local	2	16.51 %
qe-appdisc-02.lab.vi.local	2	19.55 %
qe-appdisc-03.lab.vi.local	2	1 %
qe-appdisc-04.lab.vi.local	2	21.47 %
qe-appdisc-05.lab.vi.local	2	18.84 %
qe-esx4.lab.vi.local	2	33.64 %

Page 1 of 2 | Displaying Rows 1 - 50 of 52

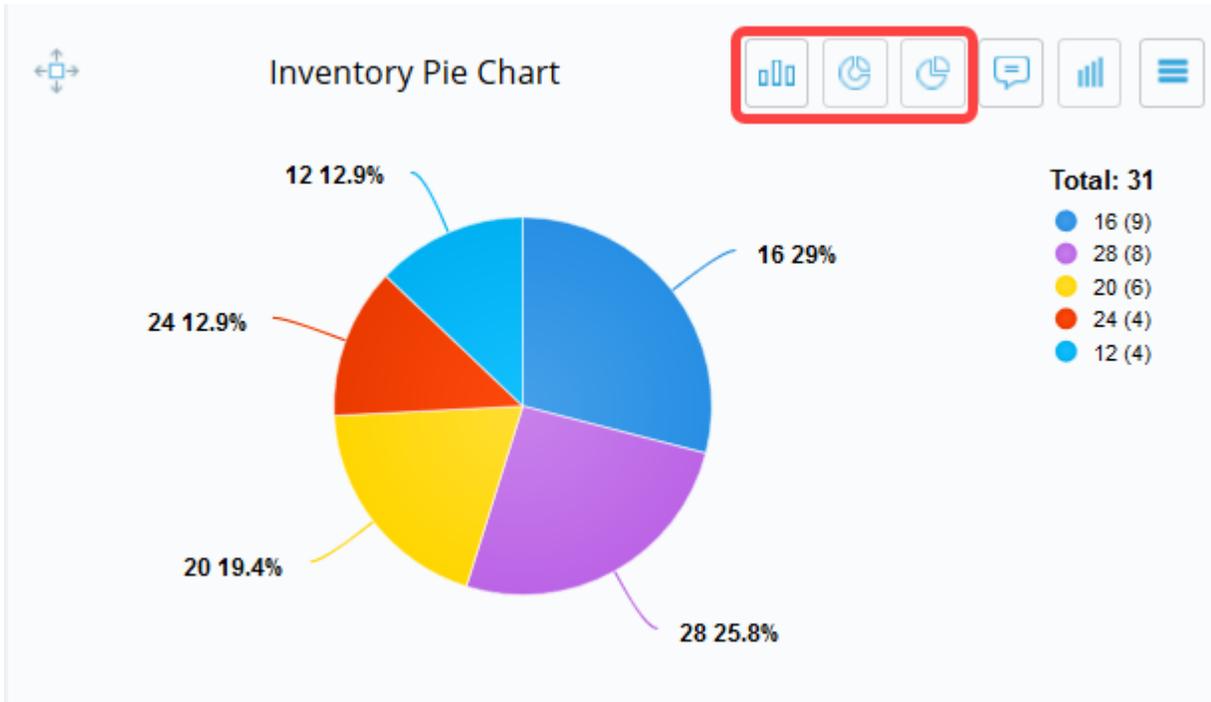
Scheduled reports and snapshots of the Inventory Table Chart are not currently supported. Export using CSV.

Threshold colors and icons are not included in data exports.

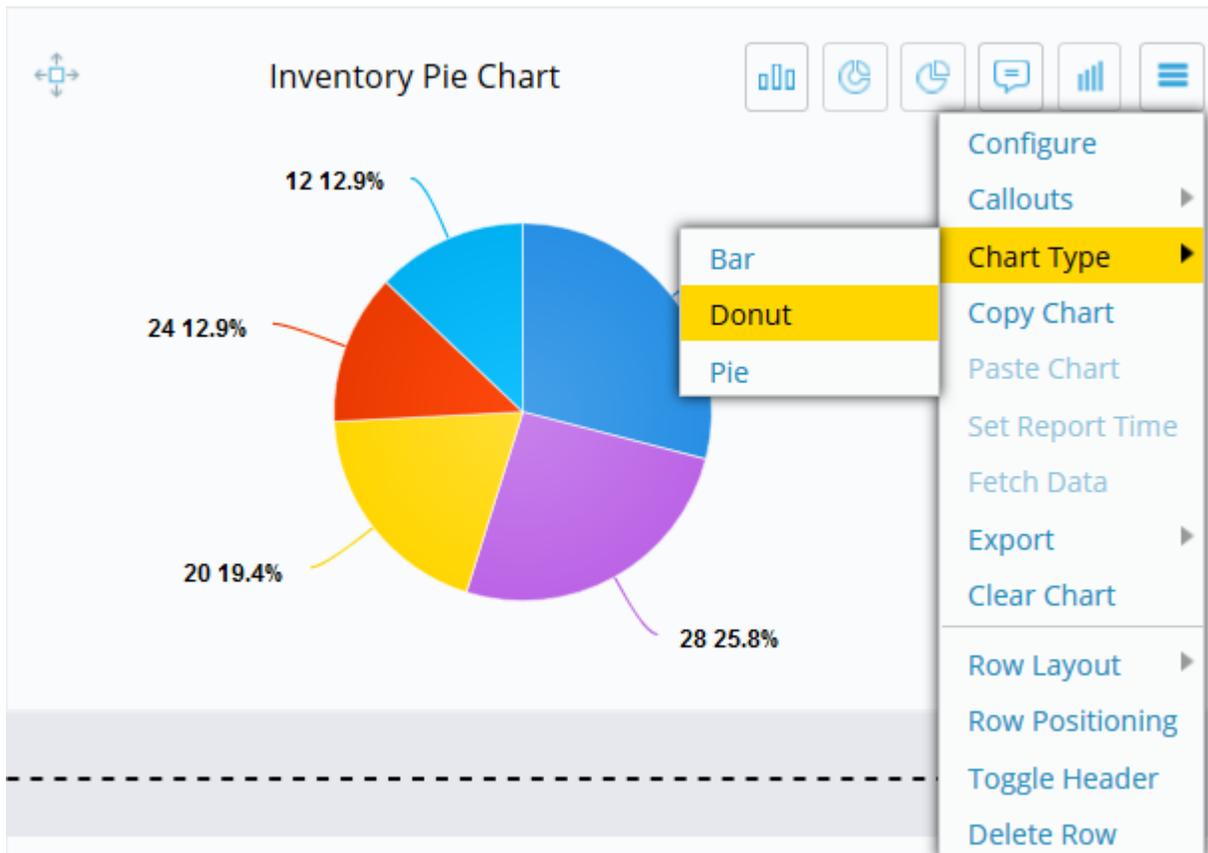
If you have configured the chart to display colors when approaching or crossing thresholds for multiple metrics, it is possible to have a yellow warning (approaching threshold) and a red warning (exceeding threshold) on different metrics for the same device. In such a case, the most serious warning color (red) will display.

Toggling Between Charts

You can toggle between the bar, donut, and pie charts by selecting an icon on the top right of the chart.



You can also use the chart menu to toggle between chart types.



Top N Charts

The Top N charts aggregate data from the top or bottom “n” events of configured metrics, where “n” is a variable that ranges in value from 1 to 50. Each Top N chart compares data for a specific metric and single entity group (hierarchy) and displays this data over time using different formats.

The Top N charts include five chart types.

Bar Chart

The Bar Chart is used to quickly compare the averaged or summed value of a single metric across multiple entities. It makes it easy to visualize how the average or sum for the entity compares to other metrics.

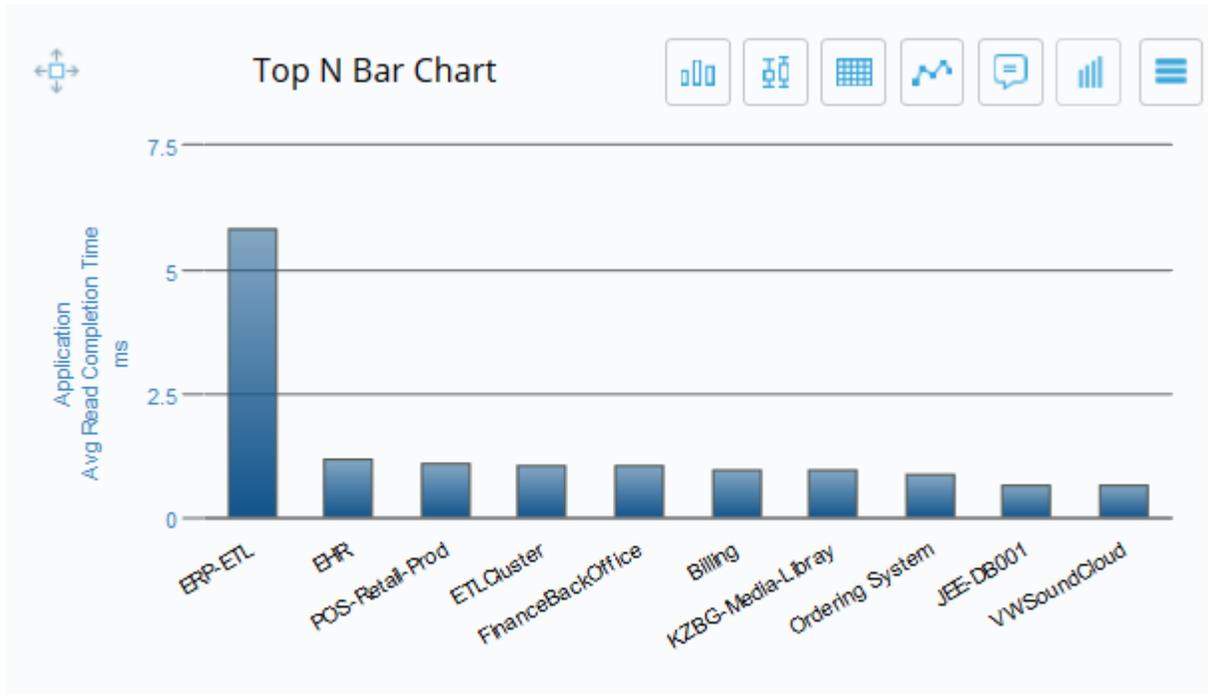


Table Chart

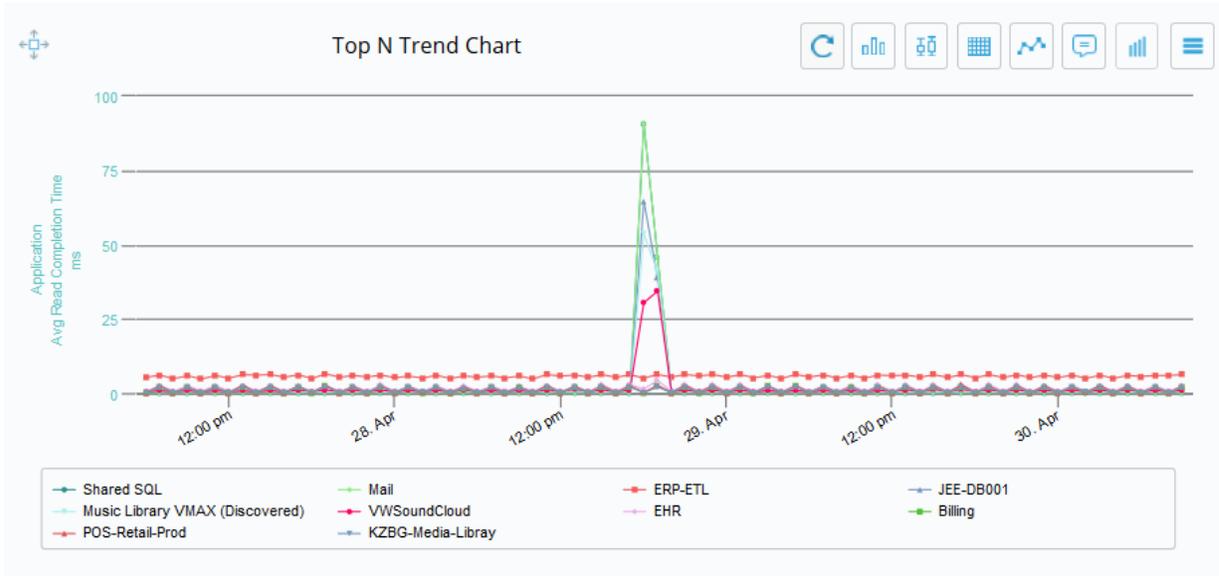
The Table chart displays the actual values of a metric in different categories (sum, min, max, percentiles) across multiple entities. We can view how the values compare for each category.

Top N Table Chart

Application	Avg Read Completion Time - ms							
	Su1	Mi1	5t1	25l	Me	75l	95l	Ma
ERP-ETL	5.8...	0.1...	0.1...	0.1...	5.4...	5.7...	6.3...	6.6...
EHR	1.2...	0.5...	0.5...	0.5...	0.67	2.3...	2.4...	2.8...
POS-Retail-Prod	1.1...	0.1...	0.1...	0.2...	0.7...	2.1...	3.7...	6.3...
FinanceBackOffice	1.0...	0.1...	0.2...	0.2...	0.6...	1.55	2.0...	2.6...
ETLCluster	1.0...	0.1...	0.2...	0.2...	0.6...	1.55	2.0...	2.6...
Billing	0.9...	0.1...	0.1...	0.24	0.8...	1.7...	2.7...	4.1...
KZBG-Media-Libray	0.9...	0.1...	0.1...	0.2...	0.8...	1.8...	2.7...	3.81
Ordering System	0.8...	0.5...	0.5...	0.5...	0.5...	1.5...	1.6...	1.9...
JEE-DB001	0.6...	0.1...	0.3...	0.4...	0.7...	1.0...	2.3...	6.1...
VWSoundCloud	0.6...	0.23	0.2...	0.3...	0.5...	0.9...	1.2...	1.4...

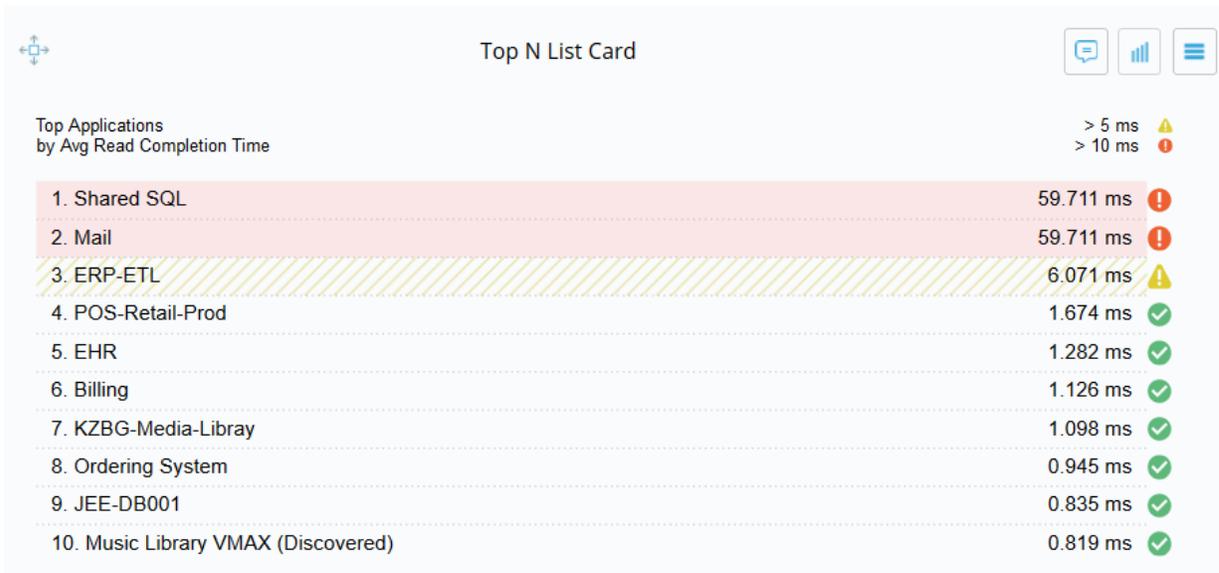
Trend Chart

The trend chart is basically a bar chart displayed over time. It compares data for a specific metric for a single entity group (hierarchy) but it also displays this data over a period of time.



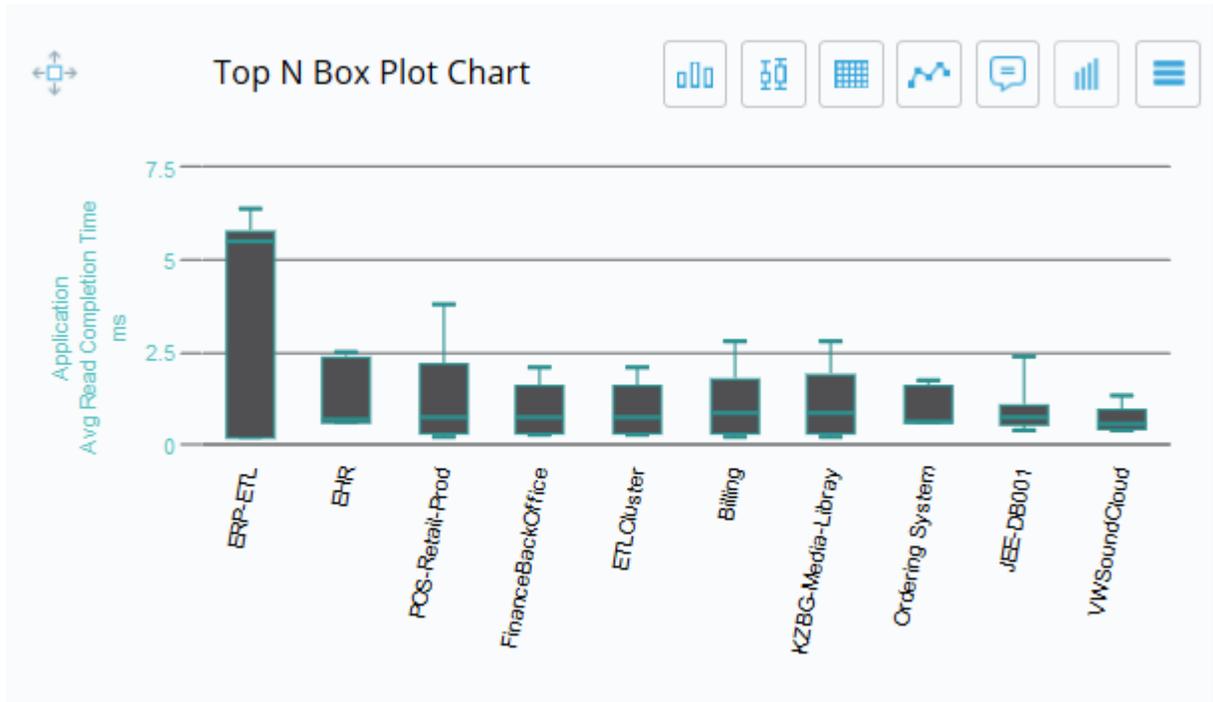
List Card

The List (Card) chart type displays Top N data in a card format.



Box Plot Chart

The Box Plot chart is used to compare the distribution of the metric data for each entity. Using a box chart helps us understand the variation of the data across entities.

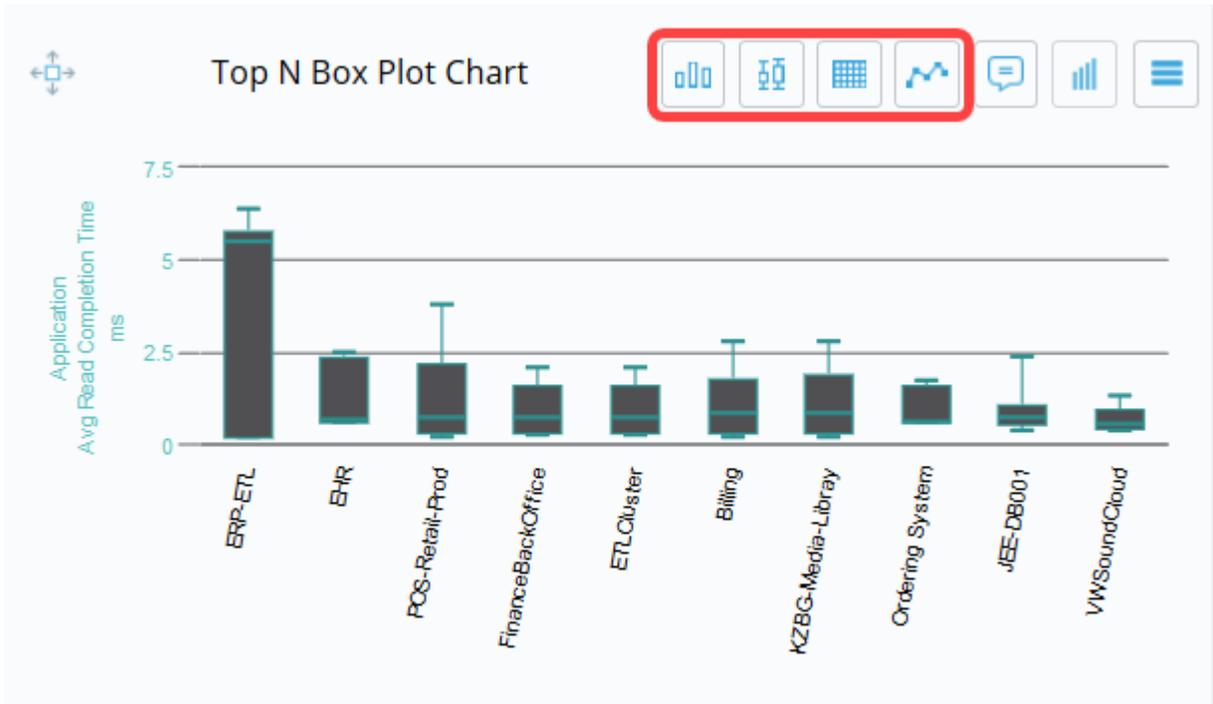


The top and bottom horizontal bars represent the lowest and highest value measured for the entity. The bottom edge of the box represents the lower quartile (25%) and the upper edge of the box represents the upper quartile (75%). The horizontal line in the box represents the median (50%) value.

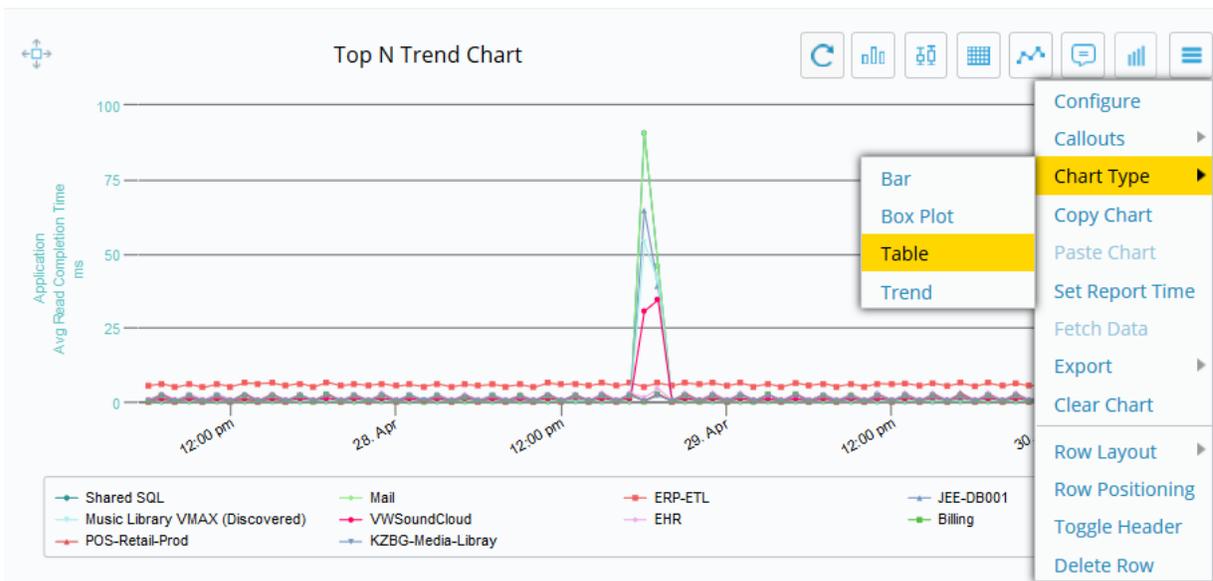
The spacing between the different parts of the box indicate the degree of dispersal or spread.

Toggle Between Top N Charts

You can toggle between bar, box plot, table, and trend charts by selecting an icon on the top right of the chart.



You can also use the chart menu to toggle between Top N charts.

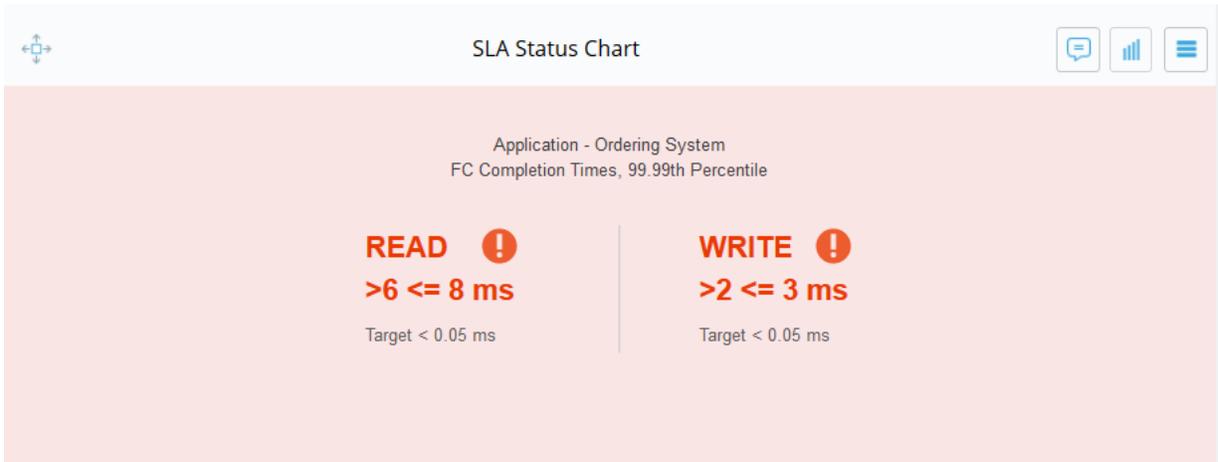


SLA Status Charts

The SLA Status charts are used to track and visualize Service Level Agreement adherence.

SLA Status Chart

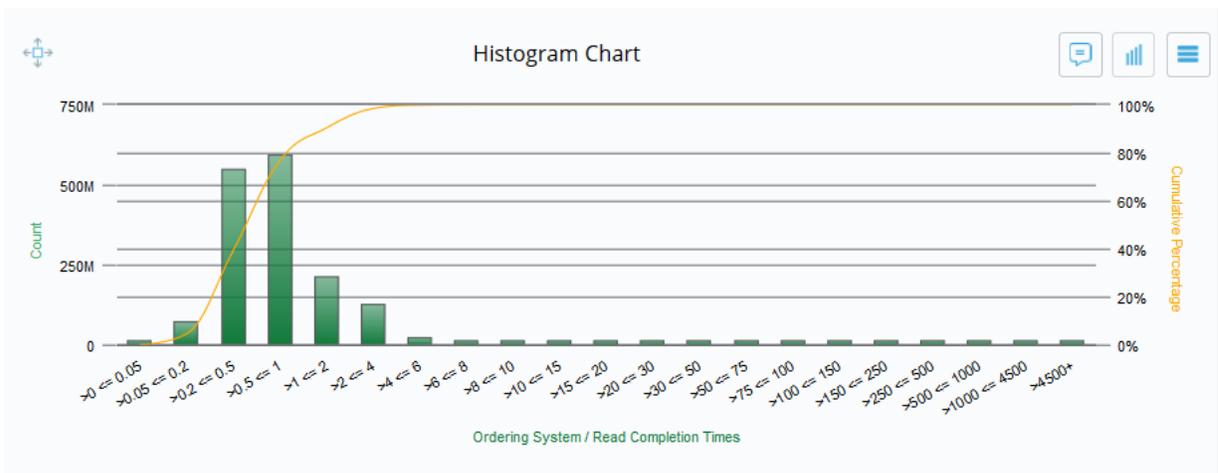
The SLA Status chart type is used to visualize Service Level Agreement adherence, e.g., How well are my storage subsystems are performing in terms of SLAs?



Histogram Chart

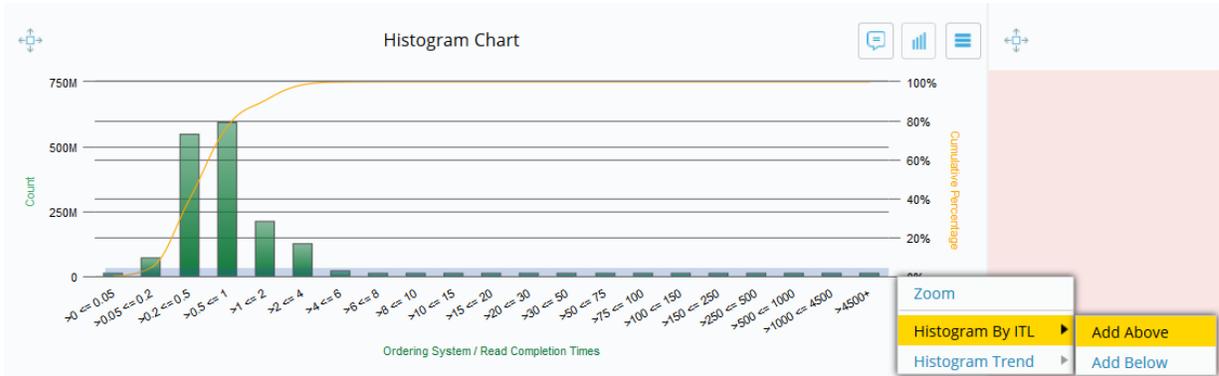
The Histogram chart displays the distribution of the data collected for a single metric over a reporting period. This provides a complete picture for transactions over any reporting period. Data is divided into ranges and the observed metric values are placed into their representative buckets. This allows us to more easily see the distribution of the metric data.

The chart represents the buckets using a bar chart, with the x-axis displaying the count and the y-axis displaying the range of buckets. A trend line that shows the cumulative percentage, from zero to 100%, is superimposed on the bar chart.



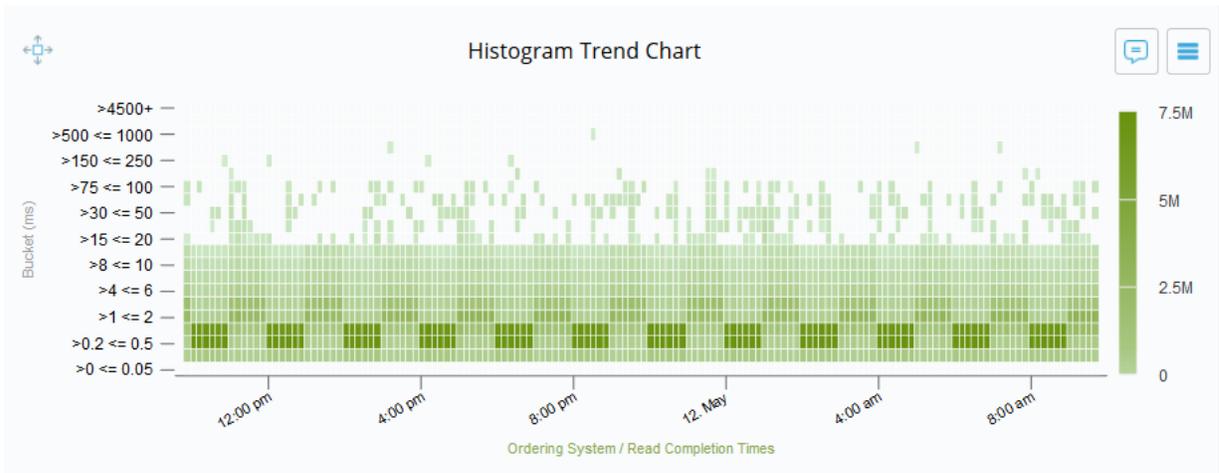
The Histogram chart type also includes two additional views of the data.

To view the base histogram chart using either of these views, drag the mouse across the chart then right-click and select either chart type and whether to add it above or below the existing chart.



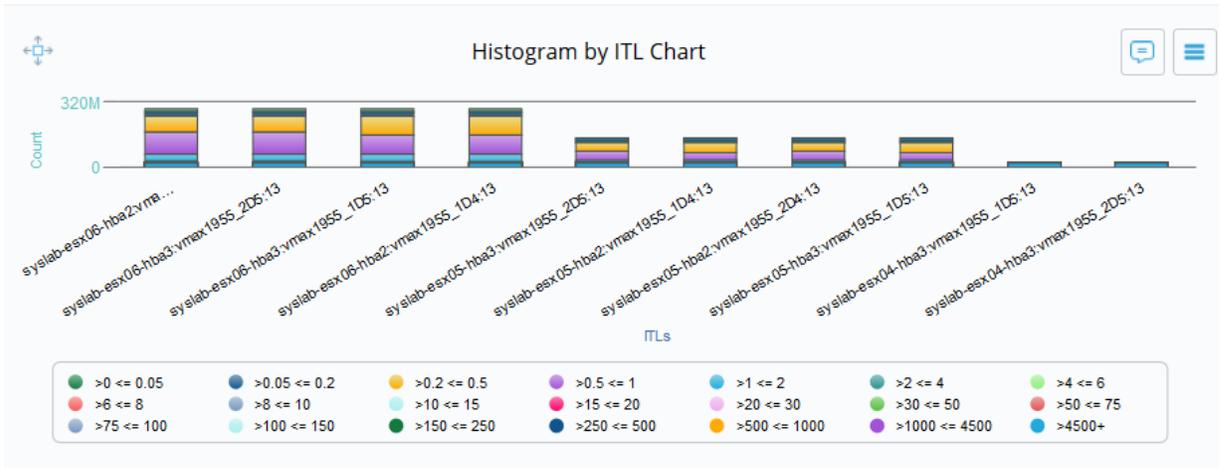
Histogram Trend

The Histogram Trend chart shows additional data for the selected bins. The timeline is displayed across the x-axis while the y-axis displays the bin values. The color density indicates the number of values in each bin. The darker the color, the more values were collected for the bin.



Histogram by ITL Chart

The Histogram by ITL chart shows the values collected for individual ITLs across the x-axis. The y-axis shows the total count of the values collected for each ITL. The colors indicate where the collected values for the ITL lie.

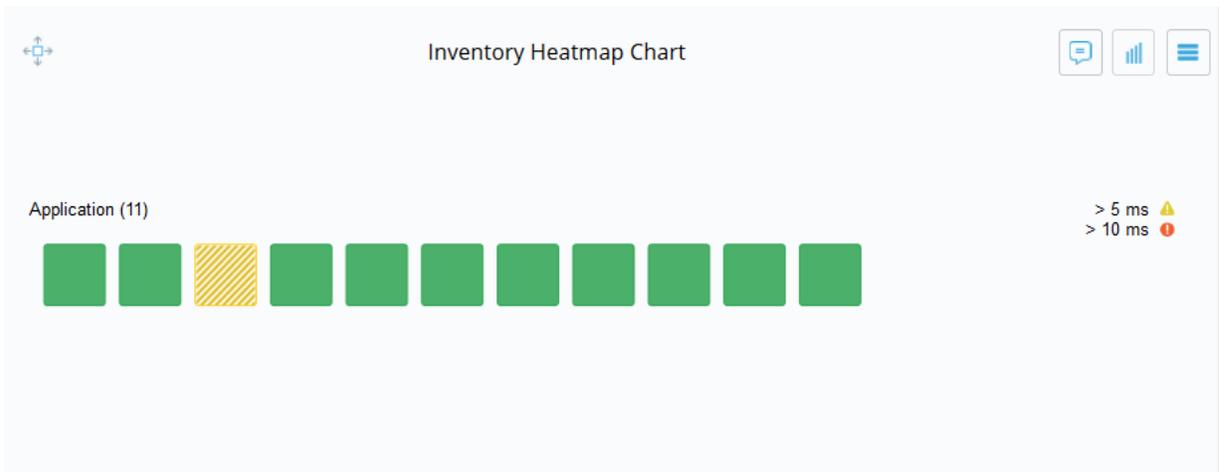


Special Charts

The special charts include charts that help you assess the health of the infrastructure at a glance, review capacity trends, view topology, and view the output of the VirtualWisdom analytics.

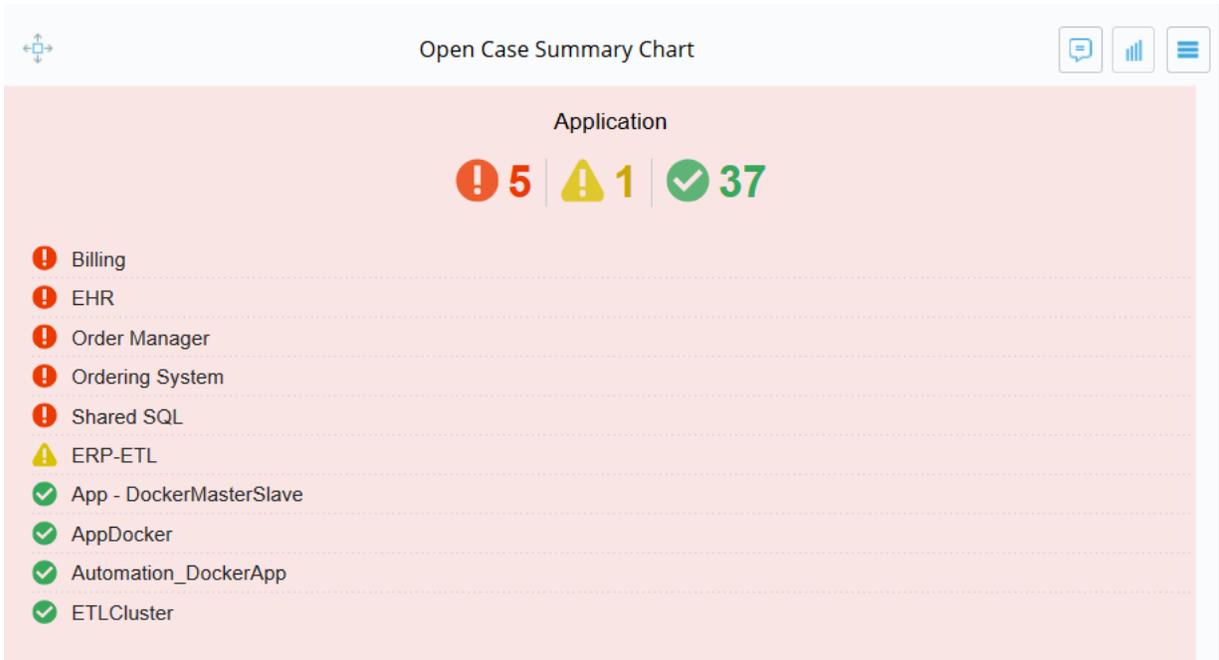
Inventory Heatmap Chart

The Inventory Heat Map chart type is a visual representation of the complexity of your environment. It is designed to quickly show you how many of the selected entity type makes up the filtered selection and how many have crossed a configurable threshold. You can click on an entity box to view open alarms, navigate to its entity page, or show its topology.



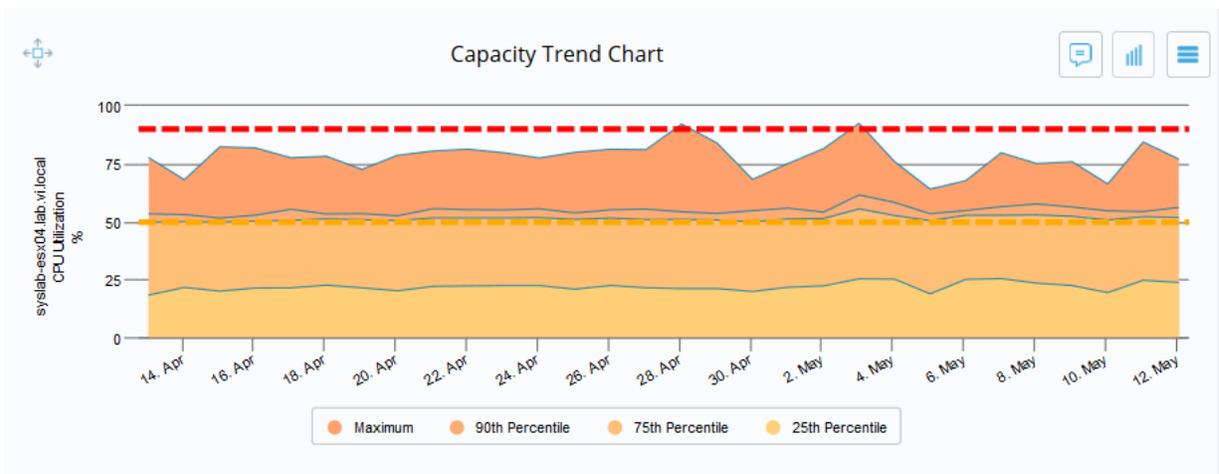
Open Case Summary Chart

The Open Case Summary chart shows you how many entities have open alarms and how serious they are.



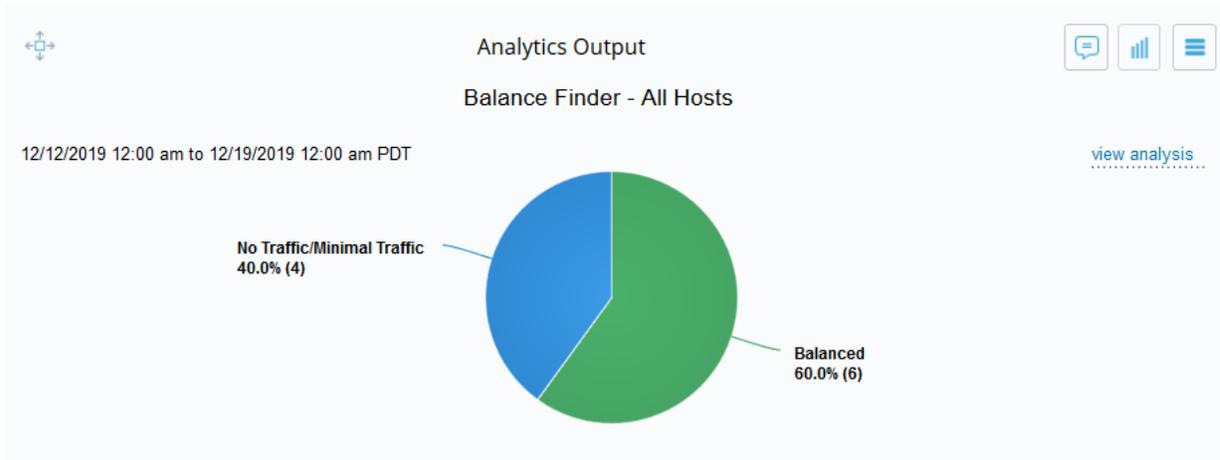
Capacity Trend Chart

The Capacity Trend chart provides a trend chart of 1-minute metric data saved using a statistical distribution across hourly and daily periods. This chart type provides support for analysis of metric data across longer durations.



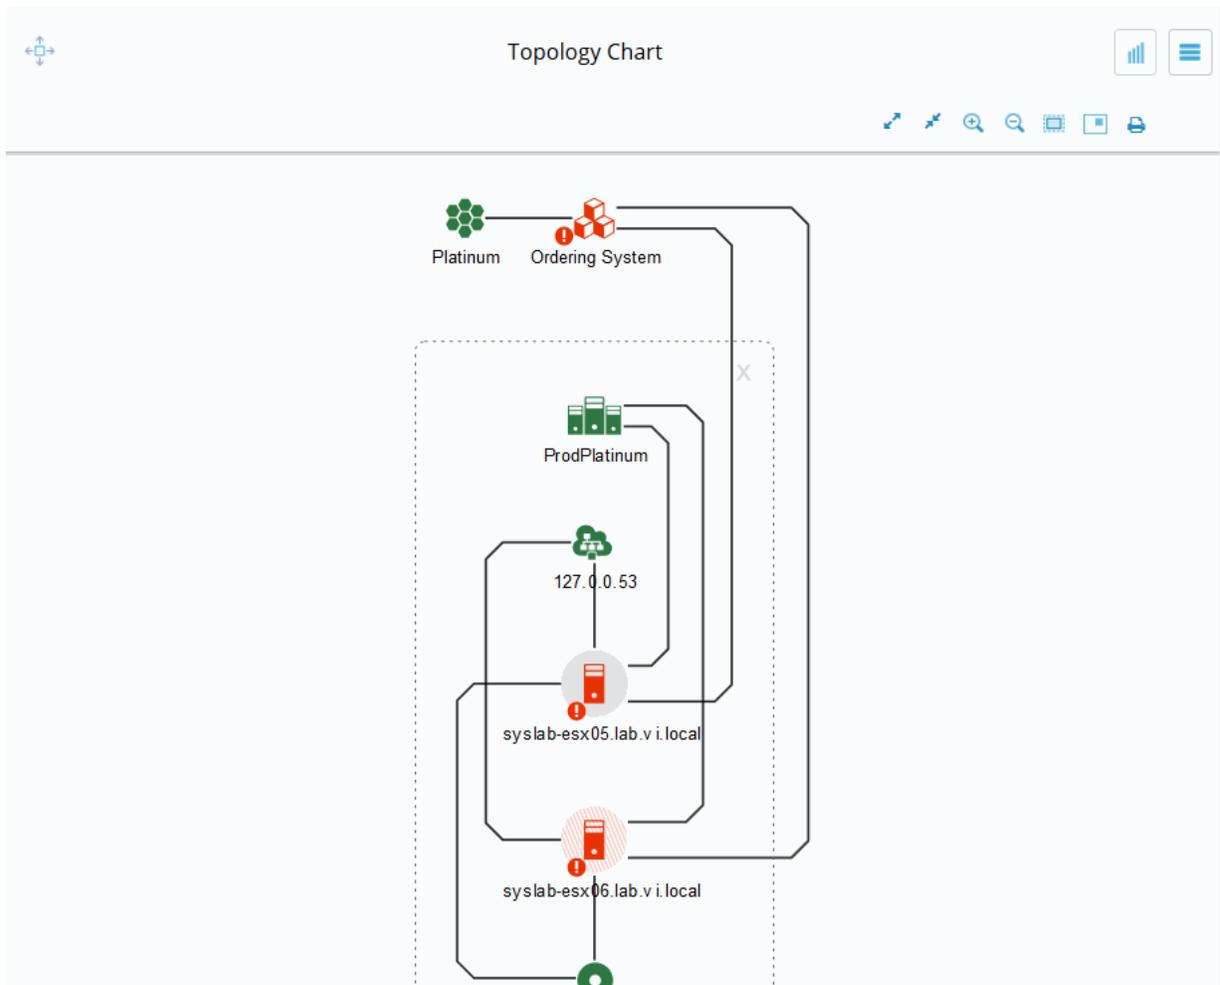
Analytics Output Chart

The Analytics Output chart displays VirtualWisdom View All Outputs in reports. Balance Finder, Storage Port Balancer, and VM Coordinator results can be displayed in a report.



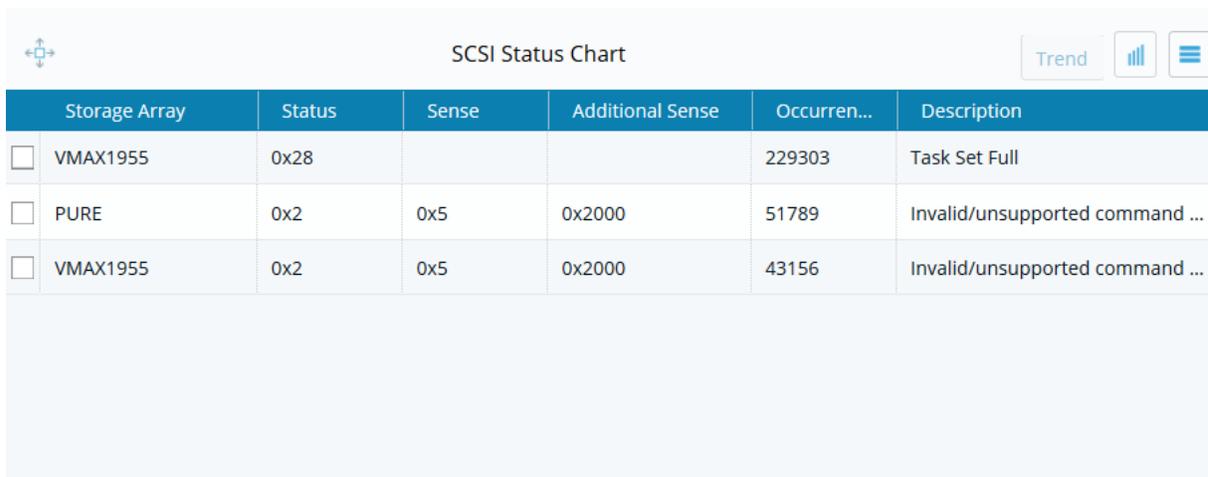
Topology Chart

The Topology chart type allows you to select a saved topology view to display as a chart in a report or dashboard. All the topology map controls (expand containers, zoom in/out, mini-map, etc.) are available and can be used in the chart. This chart type is very useful for dashboards.



SCSI Status Chart

The SCSI Status report displays a list of SCSI Status messages, for a selected entity, in a table format. This report is typically used for troubleshooting, for example, providing a vendor with SCSI status messages. This chart type is available only with the hardware probe.



	Storage Array	Status	Sense	Additional Sense	Occurren...	Description
<input type="checkbox"/>	VMAX1955	0x28			229303	Task Set Full
<input type="checkbox"/>	PURE	0x2	0x5	0x2000	51789	Invalid/unsupported command ...
<input type="checkbox"/>	VMAX1955	0x2	0x5	0x2000	43156	Invalid/unsupported command ...

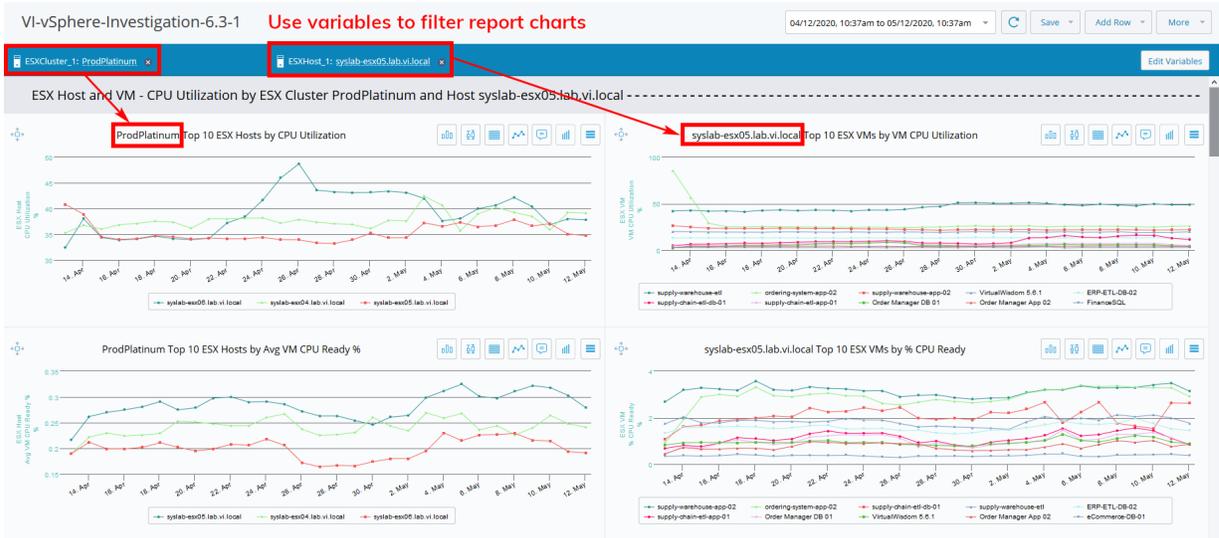
Using Reports

Using Report Variables

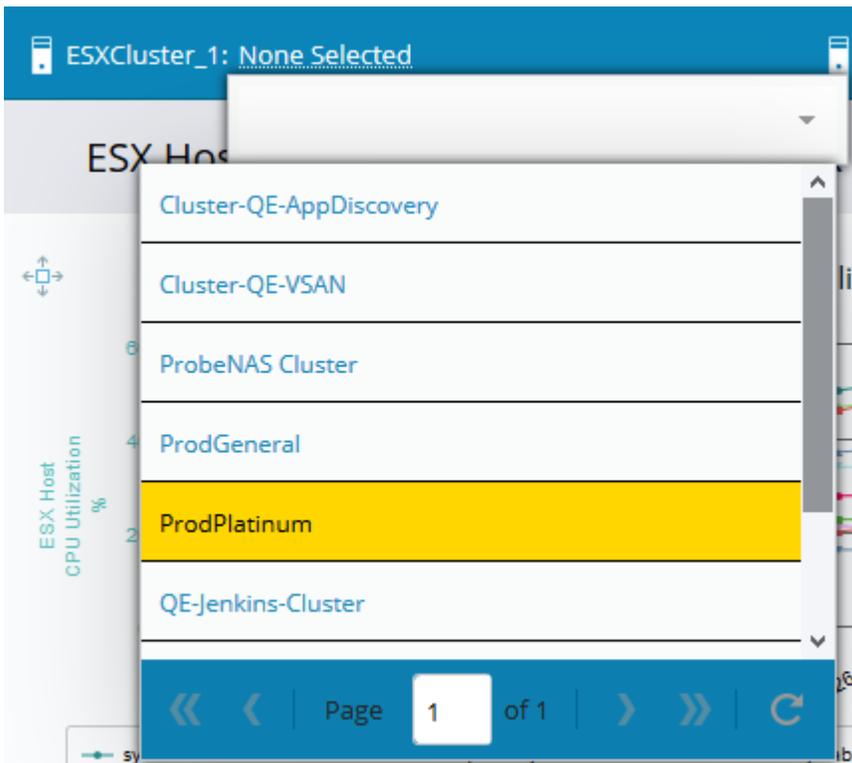
Report variables act as a short cut for using a report with different entities. They allow you to quickly and easily substitute a named entity without needing to re-select and filter every chart in the report, or create an entirely new report.

Report variables are defined at the report level but used to filter charts. Report variables can be created at any time to an existing report or added when creating a new report.

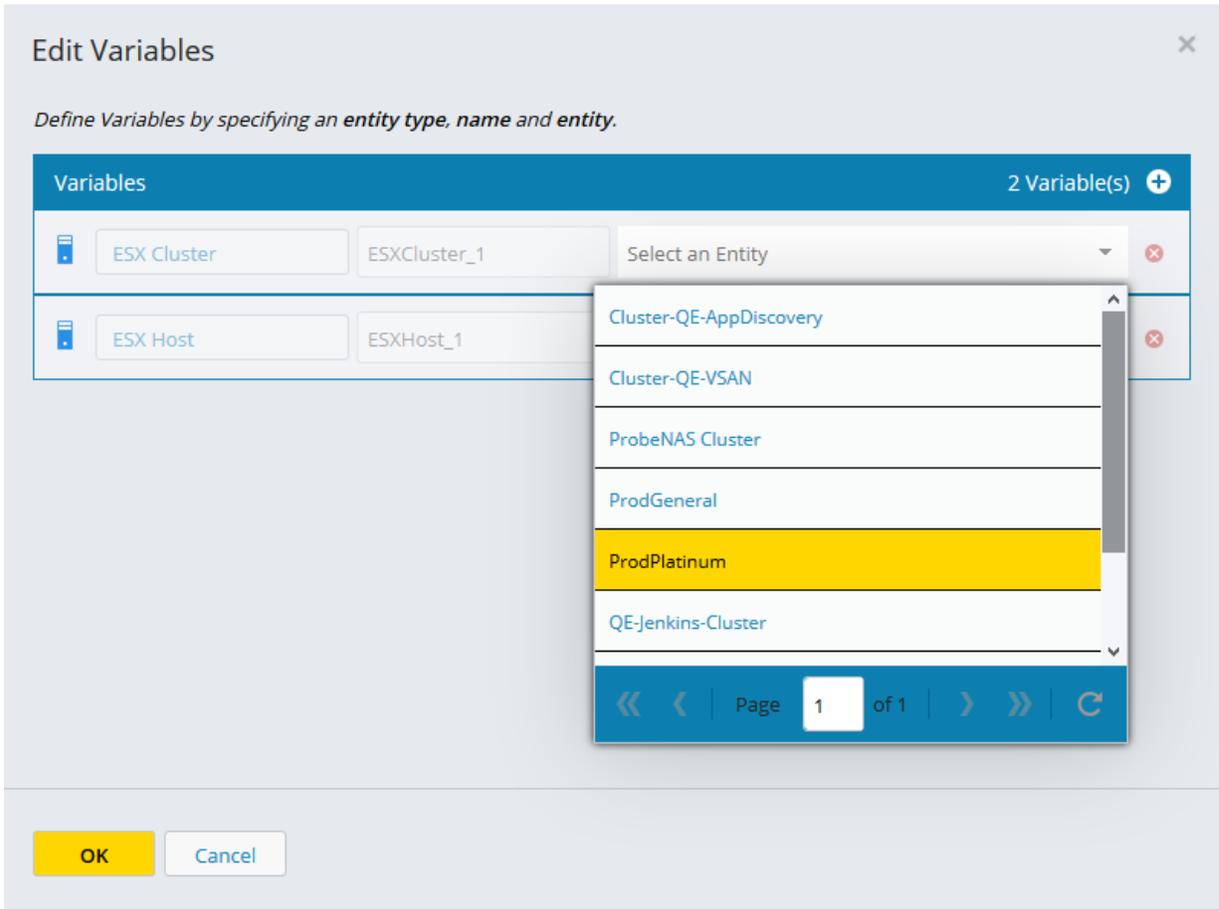
You can filter multiple charts using a single entity. Let's say that you have a report that has multiple charts that use ESX Cluster as the entity type. A global filter for ESX Cluster can be applied to the report by using a variable for ESX Cluster.



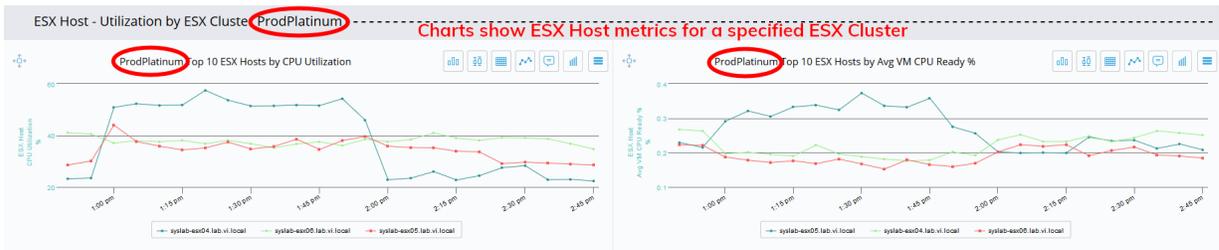
In the example below, we have selected a named ESX Cluster (ProdPlatinum) by using the ESX Cluster variable field.



Alternatively, you can also use the Edit Variables window to populate the variable field.



Once the variable is selected, every chart that uses the ESX Cluster variable in its filter selection is filtered for the ProdPlatinum cluster.

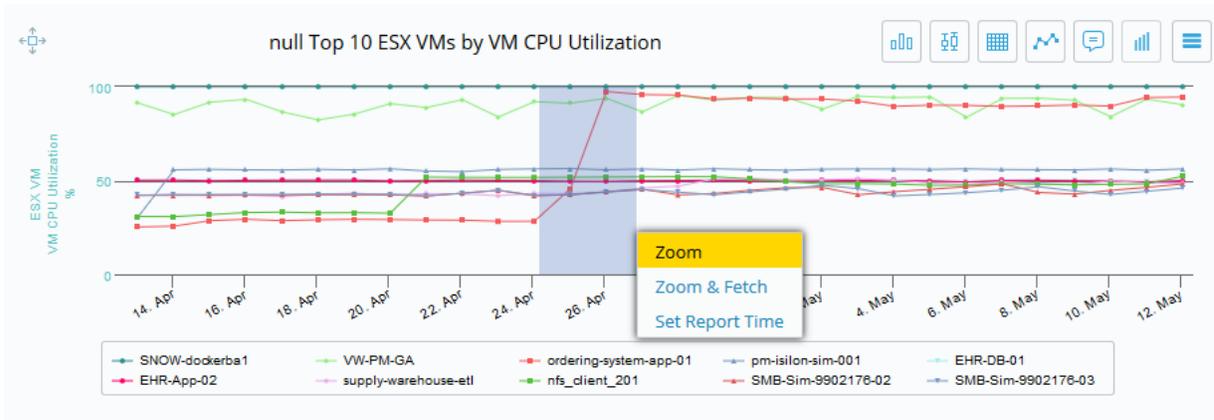


Zoom and Fetch and Set Report Time

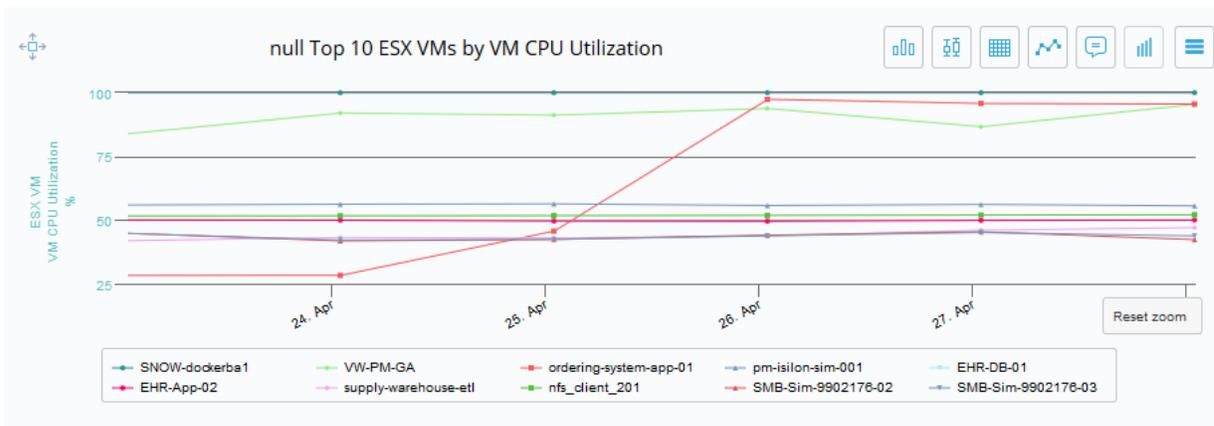
You can zoom in on a chart area to provide a larger view of the data in that time frame.

Zoom

Select a time frame on a chart by positioning your mouse to a starting point on the chart then click and hold the mouse button and drag it to the end point for the desired range, then release the button. This brings up a menu with three choices: Zoom, Zoom & Fetch, and Set Report Time. Clicking on **Zoom** displays the selected time frame across the full width of the chart without changing the data summary points.



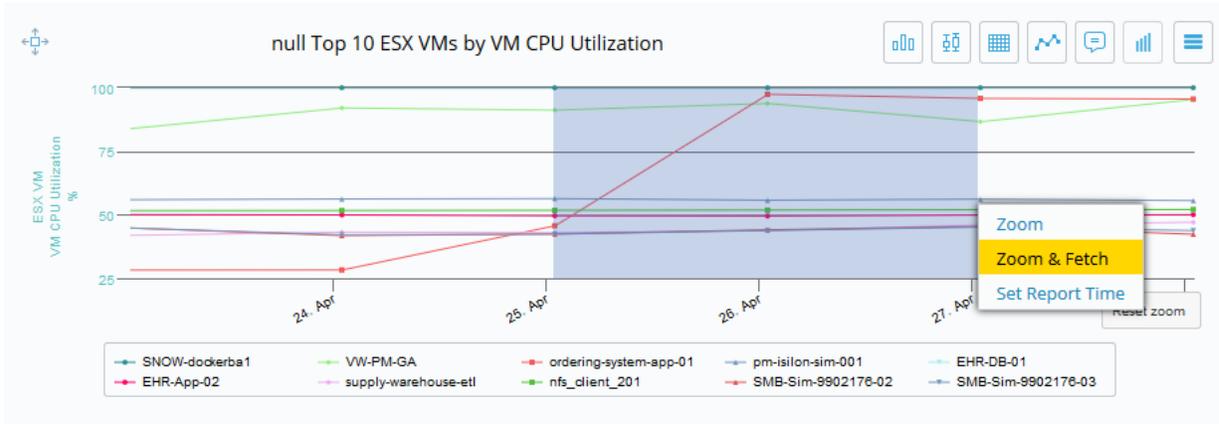
The chart is redrawn showing the data from the selected time frame.



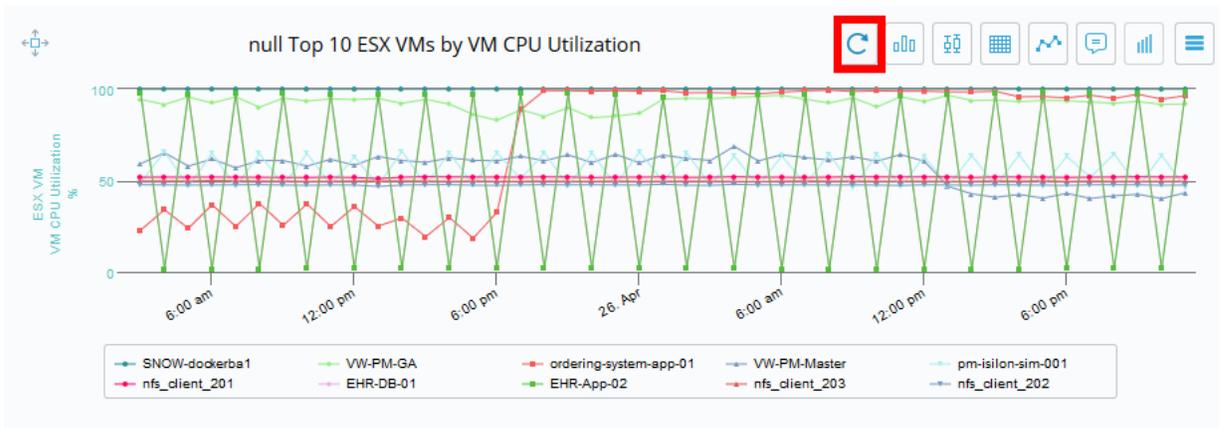
Zoom & Fetch

You can also zoom in on a chart area and load more data points for the selected period, if data is available.

Select the time frame then click on **Zoom & Fetch** to perform the zoom and load more data.

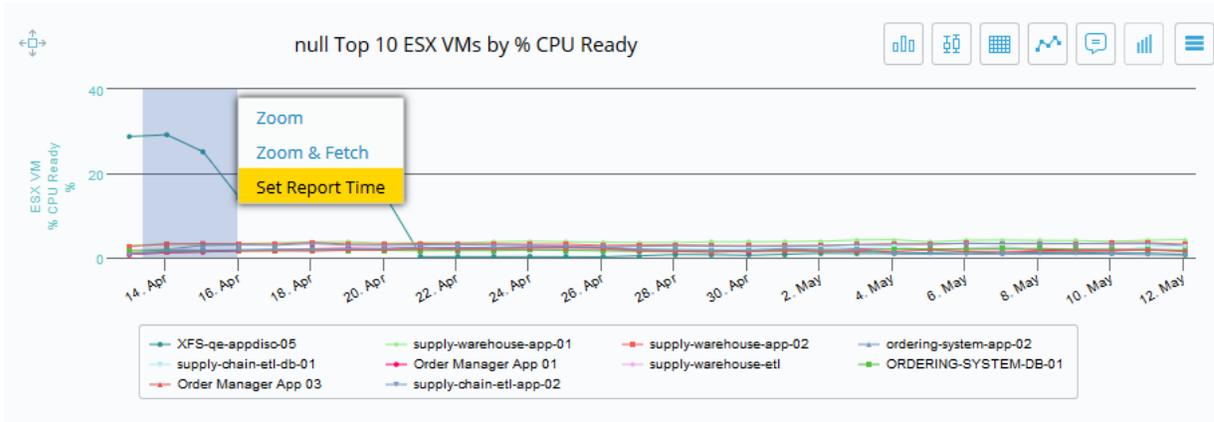


Zoom & Fetch displays the selected time frame across the full width of the chart and displays more data summary points, e.g., 10 min to 5 min. Click the Refresh button to return to the original view.



Set Report Time

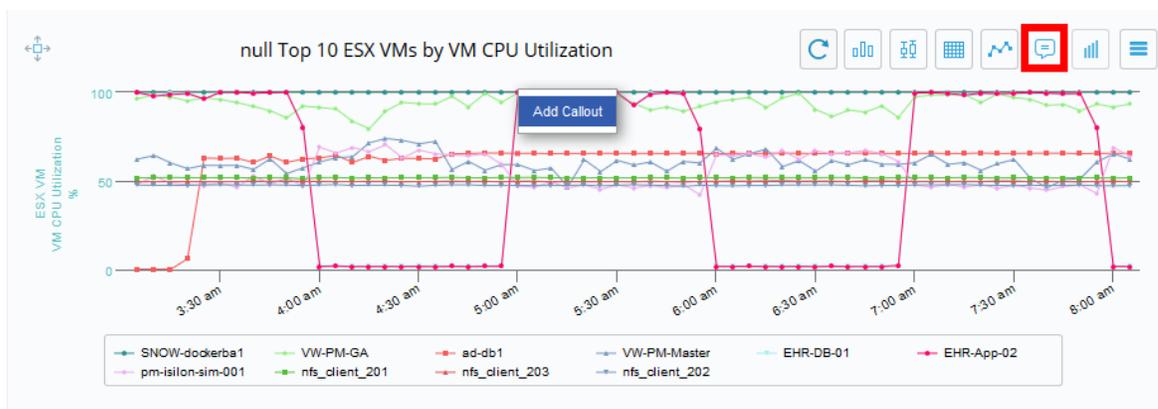
You can also set the time for the entire report by making a selection using your mouse then selecting Set Report Time. This changes the date range for the entire report.



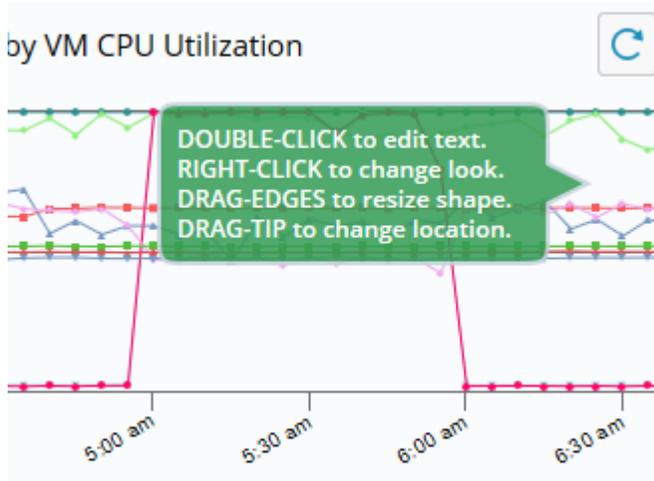
Report Callouts

Report callouts are information bubbles that can be added to a chart. Callouts can be added to all chart types except for Free Form, table charts (Top N & SCSI Status), and Topology.

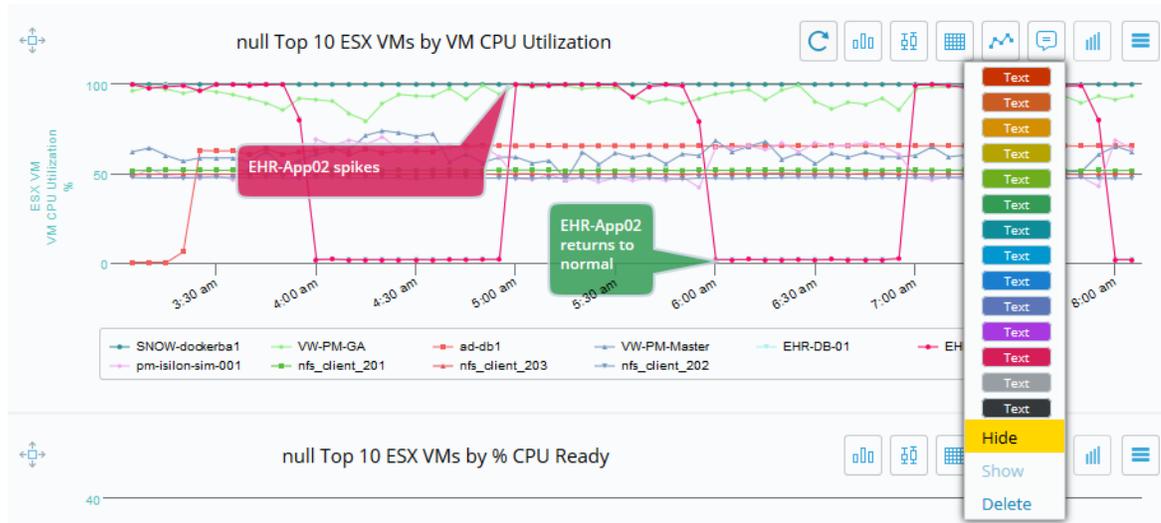
1. To add a callout, click on the callout button or right click within the chart itself and select Add Callout.



2. The callout's text can be edited, it can be resized, and you can change its look (color, transparency, or layer).



3. You can add more than one callout to a chart but you can only choose to show or hide all callouts.

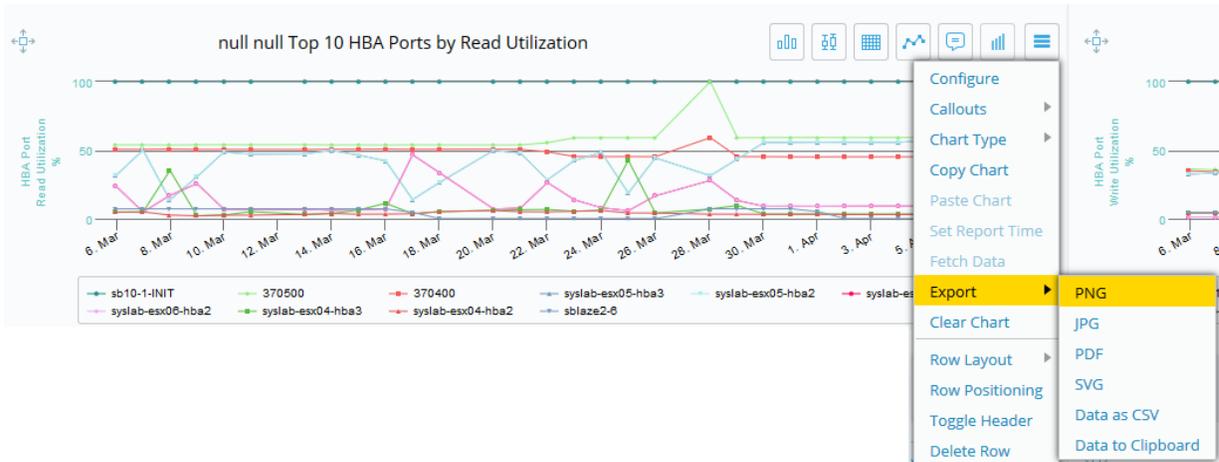


Callouts remain in the position where you created them. As the data changes, the callout remains in place.

Exporting a Chart or a Report

You can export a chart from VirtualWisdom as a PNG, JPG, PDF, or SVG file. This is useful for including in documents or in another report in VirtualWisdom.

To export a chart, select the hamburger icon then Export and choose the export file format.



You can also export the entire report as an image file or a PDF file. Click the More button then select Export and choose the export file format.



Sharing a Report Template

You can export and import reports from/to VirtualWisdom. This feature is used to import the standard Services reports into your VirtualWisdom platform. Reports are imported using the JSON file type.

Exporting a Report

Click on **More**, then select **Export > File for Import** to export a report using JSON format.

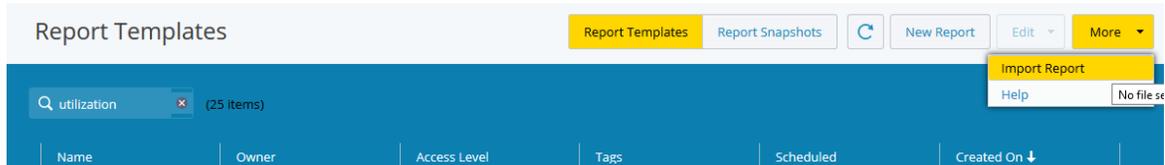


The file is downloaded to your local drive.



Importing a Report

1. From the Reports home page or the Report Templates page, click on the **More** button, and then select **Import Report**.
2. Select the report .json file to be uploaded.
Note that the file extension must be lowercase, e.g., .json (not .JSON).



The report is opened and displayed in VirtualWisdom. The report is automatically saved on import and will be visible in the list of saved reports on the Reports home page.

You may need to set the variable before data is displayed.

Bulk Report Import

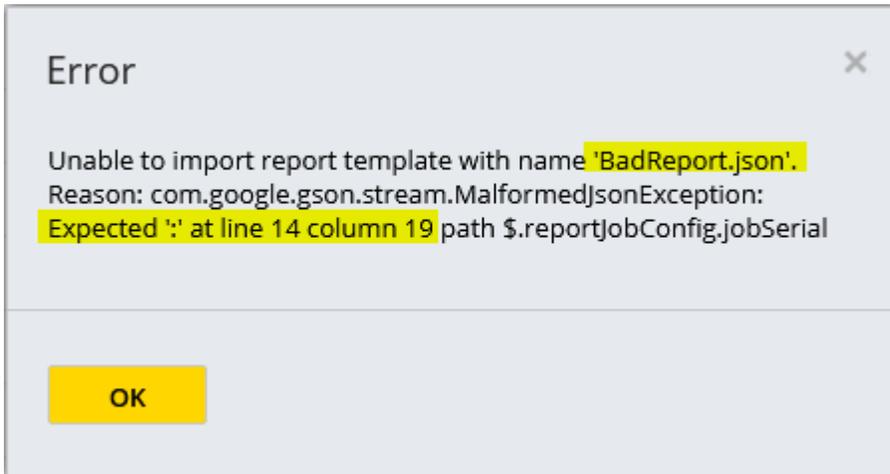
You can import multiple report templates by combining their .json files into a .zip file and uploading the .zip file using the **Import Report** feature.

Follow these tips to ensure a successful import:

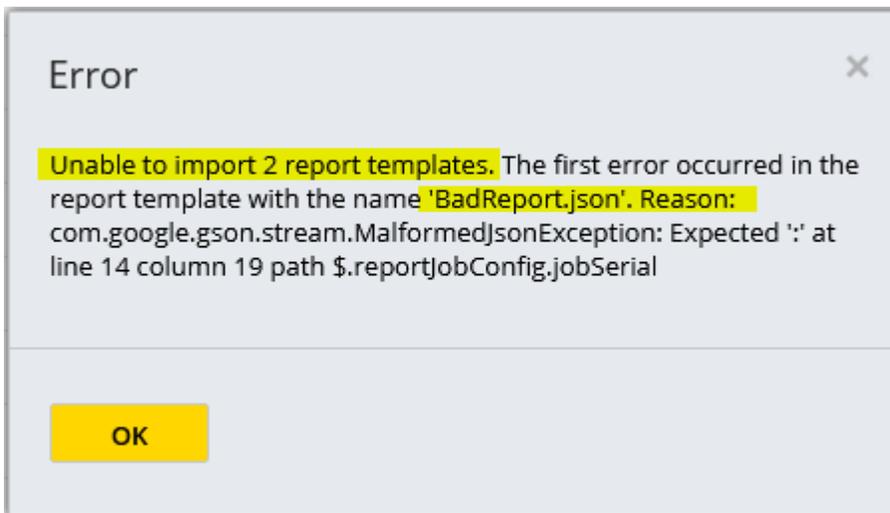
- All .json files in the .zip file must be valid VirtualWisdom report files
- The .zip file size should be less than 5 MB

- Do not import .json files which are not reports, or import empty .zip files

All .json files contained in the .zip file are validated prior to import. Any .json files failing validation will cause the import to fail with an error message containing the .json file name and the error encountered.



If multiple report template files fail validation, the number of files failing validation is shown, and the error encountered is shown for the first file that failed validation.



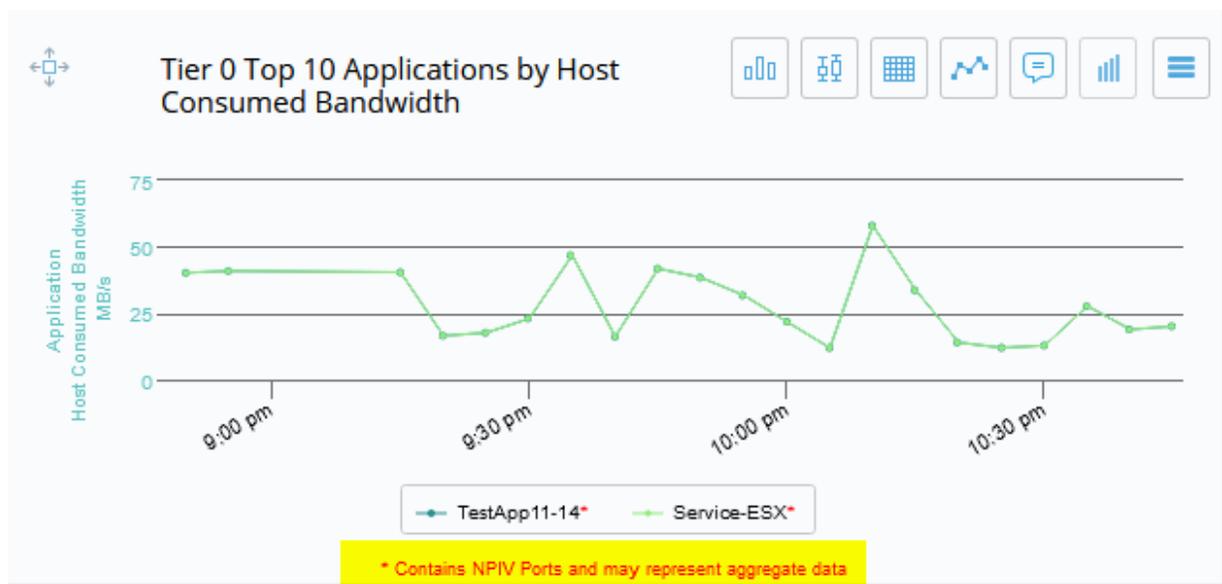


WARNING

Importing the same .zip file more than once will result in duplicate reports being created in VirtualWisdom.

NPIV Indicator Message

N_Port ID Virtualization (NPIV) has multiple hosts connected to a fabric through a single physical port. When VirtualWisdom displays data for a port connected by NPIV, it also displays an *NPIV indicator message*, warning that the data is an aggregate, and therefore not precise. VW shows precisely how much traffic is going through the port, but cannot tell how that traffic is distributed across the multiple hosts.



In the following examples, plotting a SAN Fabric metric, you see aggregated data (data that contains the object you want to see, potentially combined with data from multiple objects you do not want to see). In these cases, the *NPIV indicator* is displayed if:

- You plot out Cisco or Brocade SAN Integration metrics for an HBA attached to an NPIV port.
- You plot out Cisco or Brocade SAN Integration metrics for a Host that has an HBA attached to an NPIV port.
- Add a Host to an Application attached to an NPIV port.
- Add an I-T Conversation to an Application, and the Initiator is connected using NPIV.

- Add an I-T-L Conversation to an Application, and the Initiator is connected using NPIV.

Errors in Reports

When a report template is viewed, it is checked for errors. Any identified errors are displayed in the corresponding chart.

Errors can occur in Report Templates for a number of reasons:

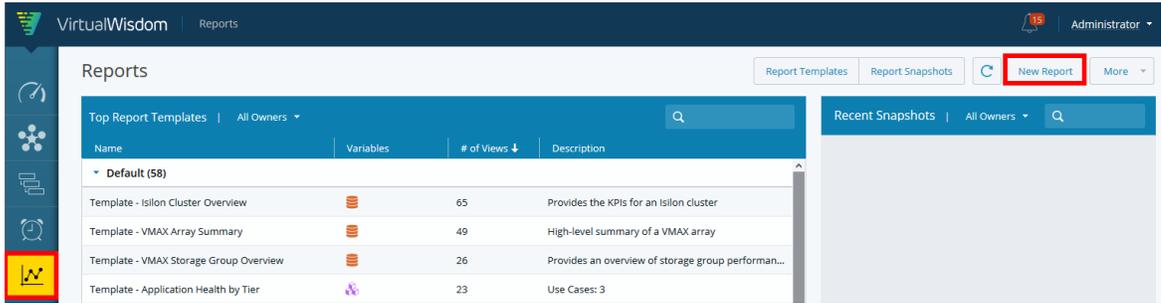
- A template was imported from another appliance and certain entities, entity types, metrics, or other details are not recognized by the new appliance.
- An entity was deleted from the appliance and is no longer recognized.
- A custom property was removed from all entities, and is no longer recognized
- No data is found, either due to no entities matching the query (filter is too tight) or because there is no data for that time frame.

Error Message	Description
Entity Type Not Found	The entity type is not recognized. Add/update an integration or update the Appliance, so that the entity type is defined, or select a different entity type.
Entity Not Found	The entity is not recognized. Select a different entity in the chart configuration.
Property Not Found	The property is not recognized. Select a different property.
Chart Type Not Found	The chart type is not recognized. Update the Appliance to a version that contains that chart type, or clear the chart.
Analytic Template Not Found	The analytic template is not recognized. Select a different analytic template. This might require going to the Analytic tab, editing the analytic of the type required by the chart, saving the configuration as a new analytic template, and selecting that template in the chart configuration for the report.
Analytic Not Found	The analytic type is not recognized. If the analytic is not found, you might need to upgrade the Appliance to a version that contains that analytic. You can also select a different analytic type or clear the chart.

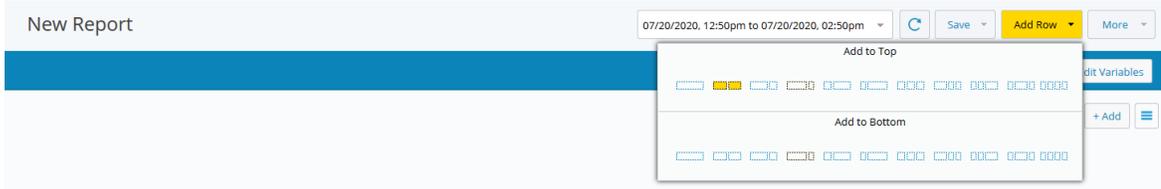
Error Message	Description
Variable Type Not Found	The variable name uses an unrecognized entity type. Add/update an integration or update the Appliance so that the entity type is defined, or select a different entity type.
Metric Not Found	Metrics in the metric list are not recognized. Add/update an integration or update the Appliance, so that the entity type is defined, or select a different entity type.
Topology Errors	Errors found in the selected topology's error list. Import a new saved topology.
SCSI Status Not Found	SCSI status code is not recognized. Open and save this chart configuration to clear the message
Empty Variable	Variable name in the chart is empty. Select a value for this variable. This message only displays if the variable is required to retrieve data for the chart. When variables are used in filters, they are ignored when empty, and that part of the filter is not applied.
Multiple Errors Occurred	A combination of the errors listed in the table occurred in one chart. The most common example of this error occurs when the report expects an integration that has not been defined or enabled on the Appliance. This could cause the both the entity type and metric to be unrecognized.
No Data	No data was found in the specified date/time range. If data should be available, try adjusting any filters used, or ask your administrator to check for data collection errors or notifications in VW Health.

Creating Reports

1. To create a new report, navigate to the Reports tab and click **New Report**.

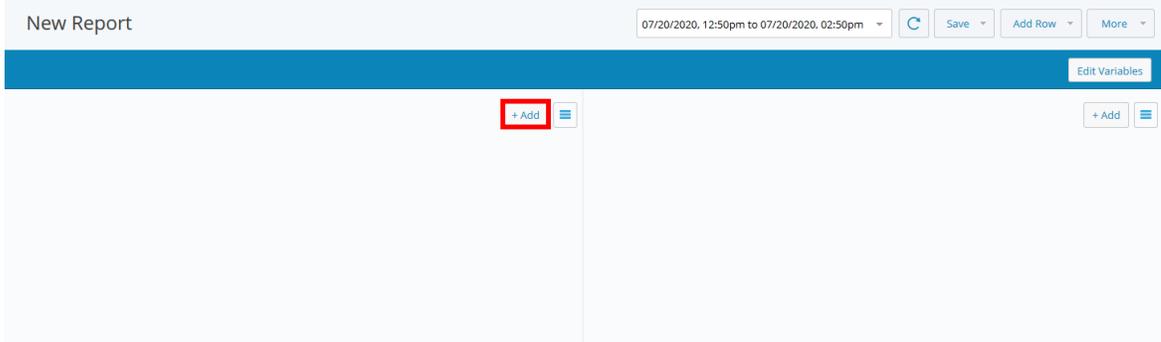


2. Click **Add Row**, then pick the row layout.



The chart areas are initially blank.

3. Click **Add** to choose a chart type to add to the report.



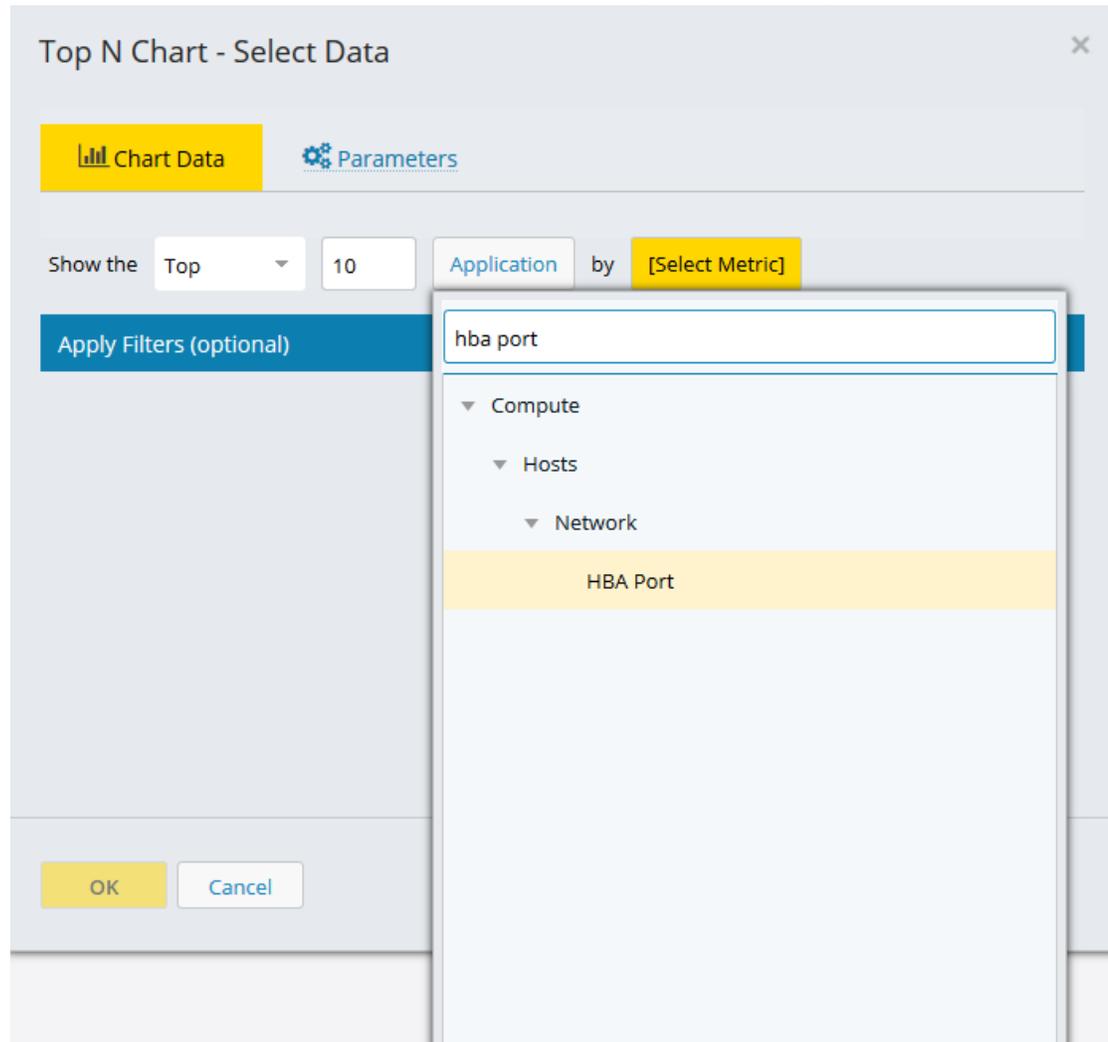
Select the chart type you wish to add to the report.



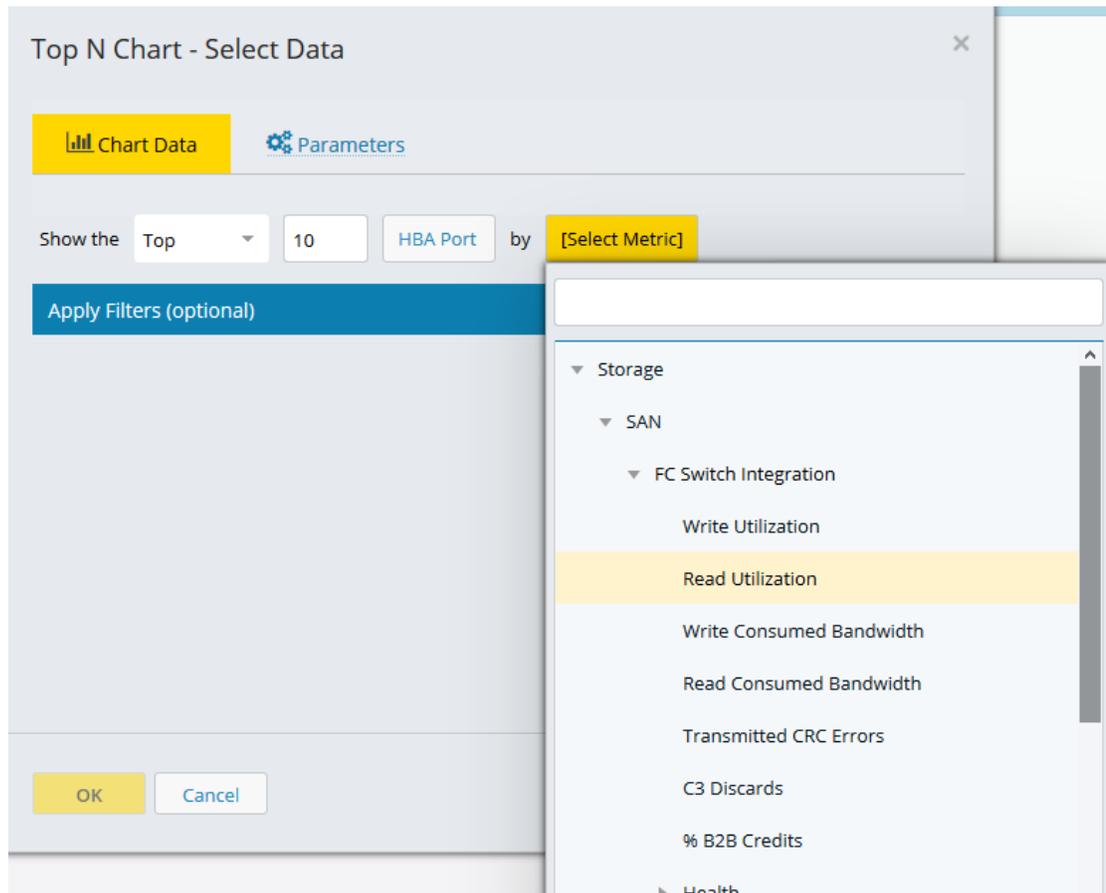
4. Select report data.
 - a. Depending on the chart type, you are prompted to select entities and metrics.

The availability of both depends on the configured integrations. If an integration is not configured, you will not be able to select that integration's entities and metrics.

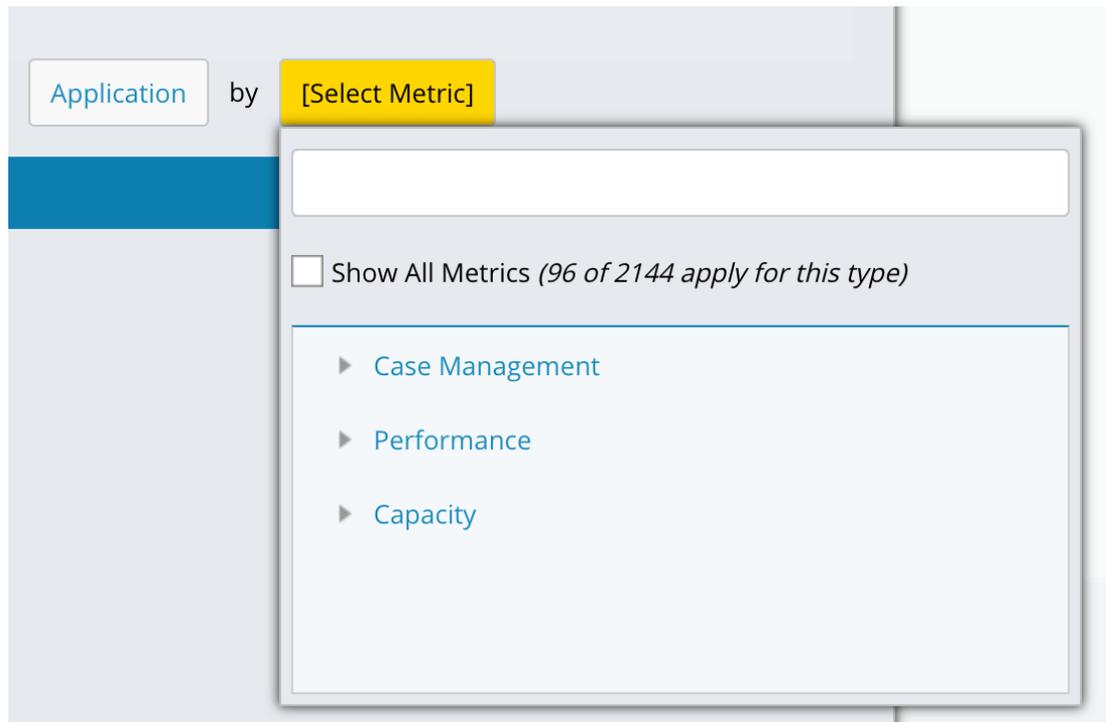
Select an entity type by clicking in the entity field and selecting from the drop-down list. You can use the search field to quickly find an entity type.



- b. Select a metric from the available metrics and click OK.



If you selected **Application** as the entity type, you can choose to show all metrics or to show only the metrics applicable to the Application entity type. This reduces the selection to only those entities that are currently assigned to at least one application and hides unassigned entities from the list. The number of applicable metrics is displayed in the selection modal.

**TIP**

If you plan to add entity types to an application at a later time and want to configure your report, dashboard, or alarm to include metrics from those entity types in advance, check the **Show All Metrics** box so they are available for selection.

5. The chart is populated with entity and metric data for the date range.

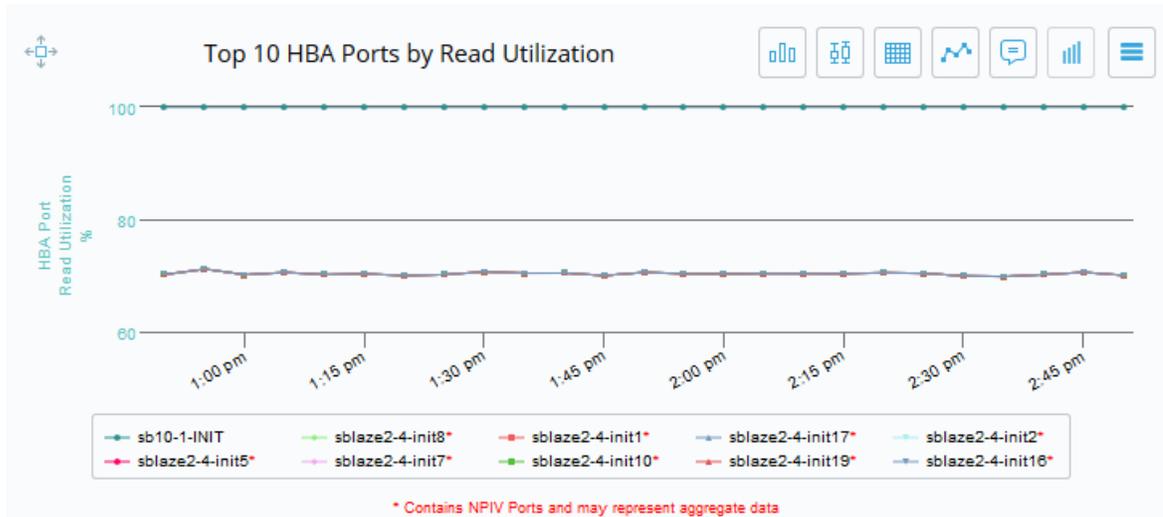


Chart Parameters

There are a number of parameters that can be configured for a chart. Every chart type provides a feature to auto-name the chart or specify a name of your choosing. In addition to this feature, you can also set thresholds, specify the number of entities to be included in the chart, e.g., show the top 10 entities, change the sort order, include a report variable in the chart name, and select where to view tooltips and legends.

To configure chart parameters, select the **Parameters** tab from the data selection modal. In the example shown below, the chart type allows you to specify a variable in the chart name, select the number of top entities to display, and change the short order.

Open Case Summary - Select Data

Chart Data Parameters

Name \$Tier_2 Applications by Open Case Criticality ↶

Top # * 10

Sort Descending ▾

OK Cancel

To reset the default chart name, click the back arrow next to the **Name** field.

Some chart types, e.g., Top N, Line, Inventory Heatmap, let you set thresholds that are displayed as yellow and red lines on the chart. In the example below, both red and yellow thresholds were set in order to easily see where the observed data exceeds a user-defined threshold.

Top N Chart - Select Data

Chart Data
Parameters

Name:

Tooltips:

Red Threshold:

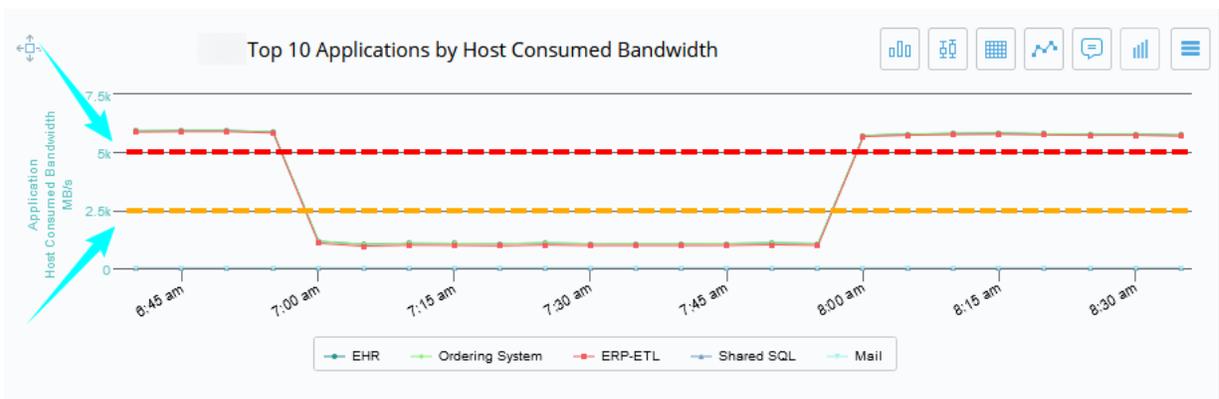
Yellow Threshold:

Markers:

Legend:

OK
Cancel

The thresholds are displayed on the chart as a yellow and red line, respectively.

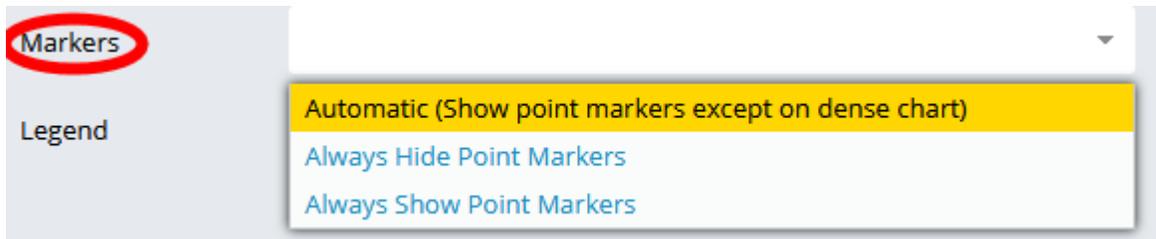




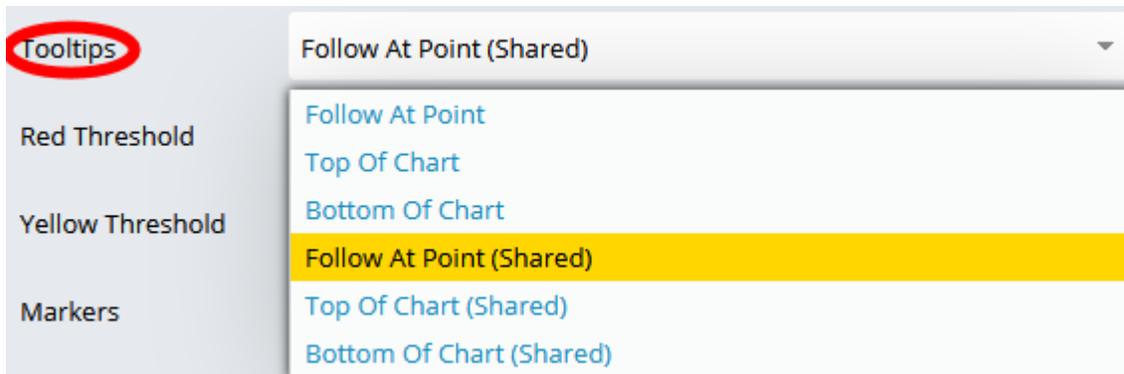
NOTE

A chart's y-axis might not start at zero; this is because VirtualWisdom auto-scales the y-axis min and max values based on the metrics values for the specified time range. When you specify a yellow and/or red threshold lines, charts are auto-scaled to include the horizontal threshold lines.

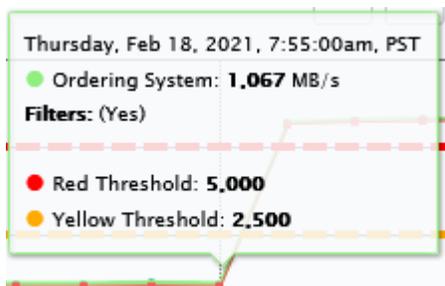
You can also specify where to show point markers automatically or always hide or show point markers.



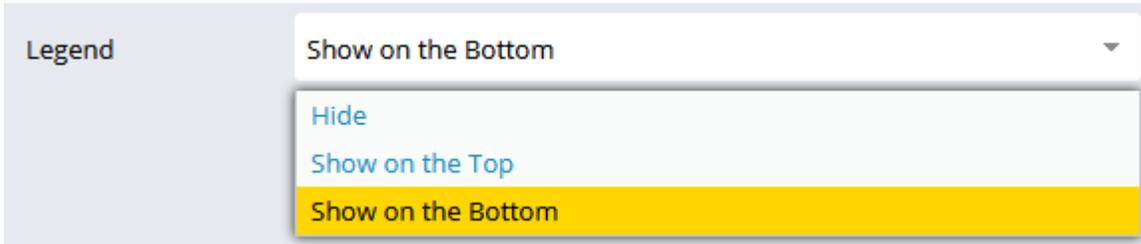
Specify where to view report tooltips on the chart.



Tooltips are displayed when you hover over a point in a chart and show you the value observed for the entity and metric at that point in time.

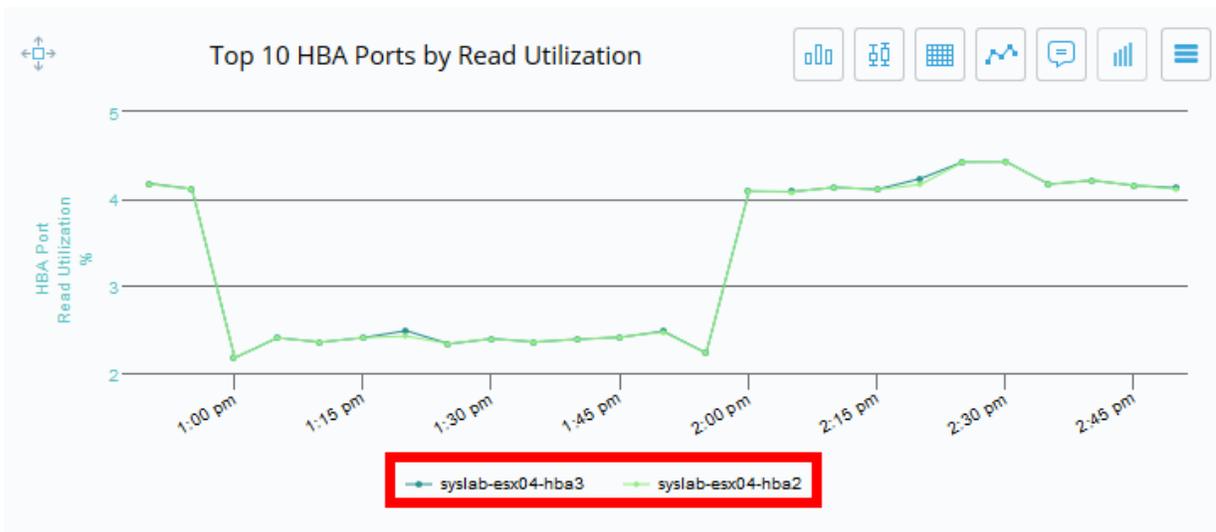


You can also choose where to show the chart legend.



Report Filtering

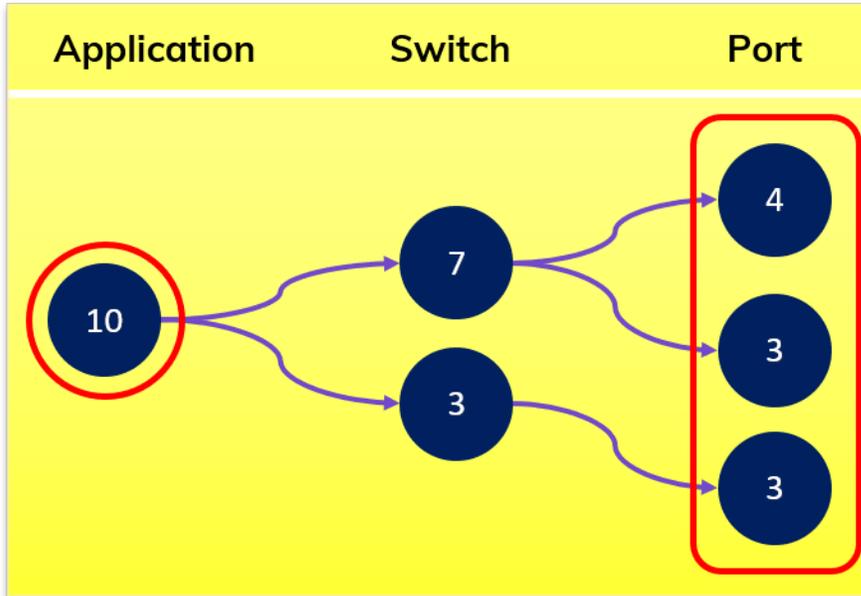
VirtualWisdom provides the ability to filter report data in multiple ways. For example, the chart shown below is filtered by only the selected entities outlined in red.



Entity Filtering

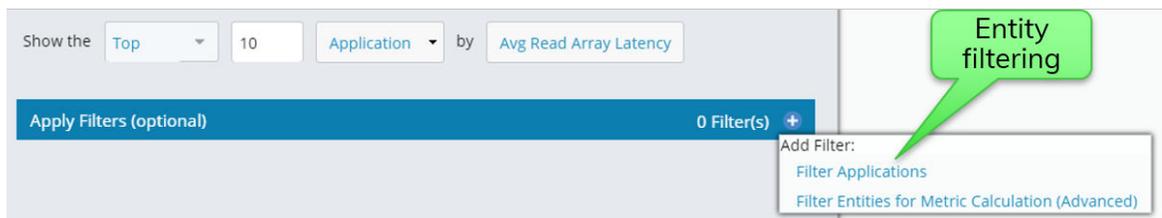
All sub-entities related to the entity are used in the aggregation. In the example shown below, the data is filtered for a single application, aggregating data collected from all related switch ports.

Entity Filtering



Entity filtering is only available on charts that display multiple entities, e.g., Top N charts.

1. To apply entity filtering, select the + sign, then choose Filter [Entity Type].



2. You can filter entities using four different options.

Top N Chart - Select Data

Chart Data Parameters

Show the Top 10 HBA Port by Read Utilization

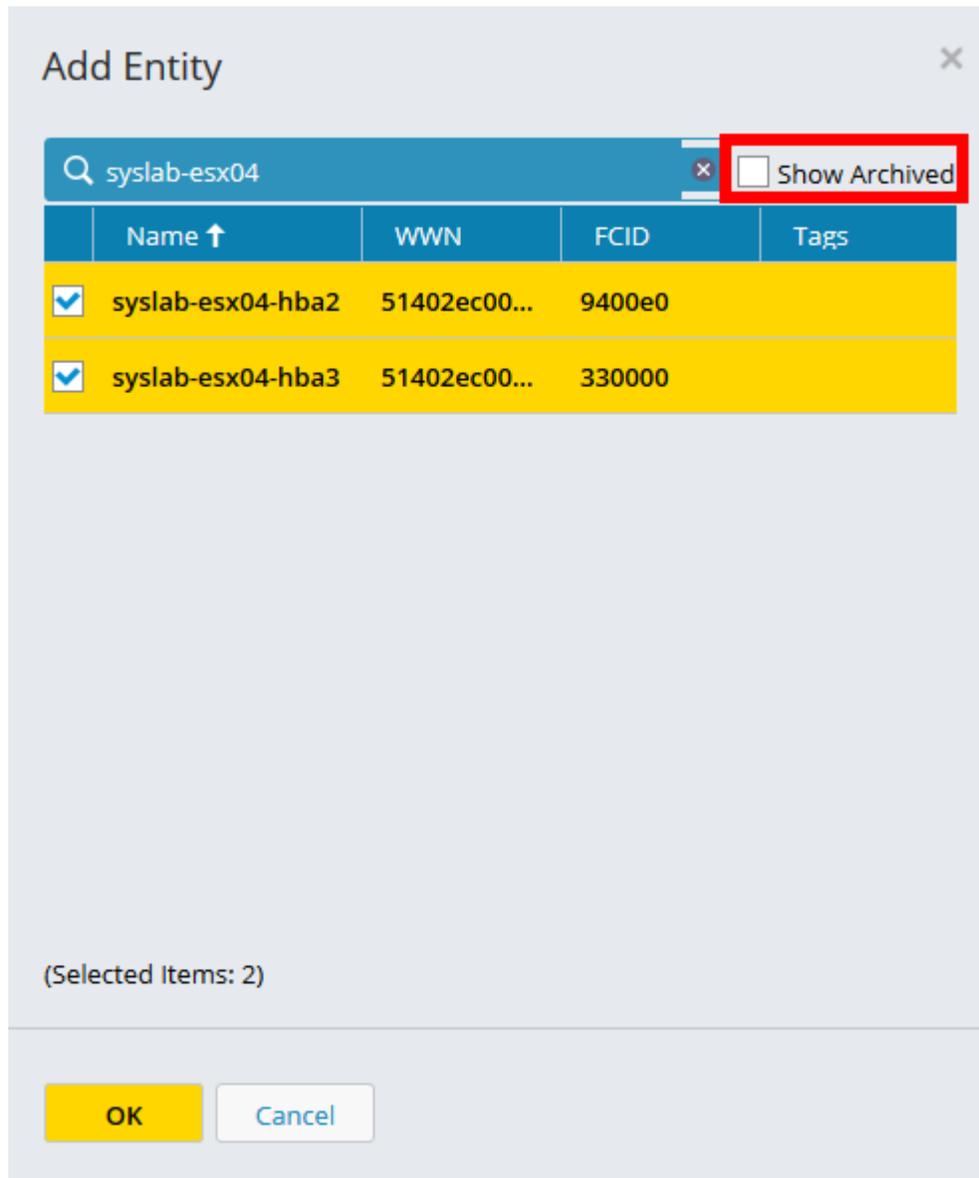
Apply Filters (optional) 1 Filter(s) +

Include HBA Ports

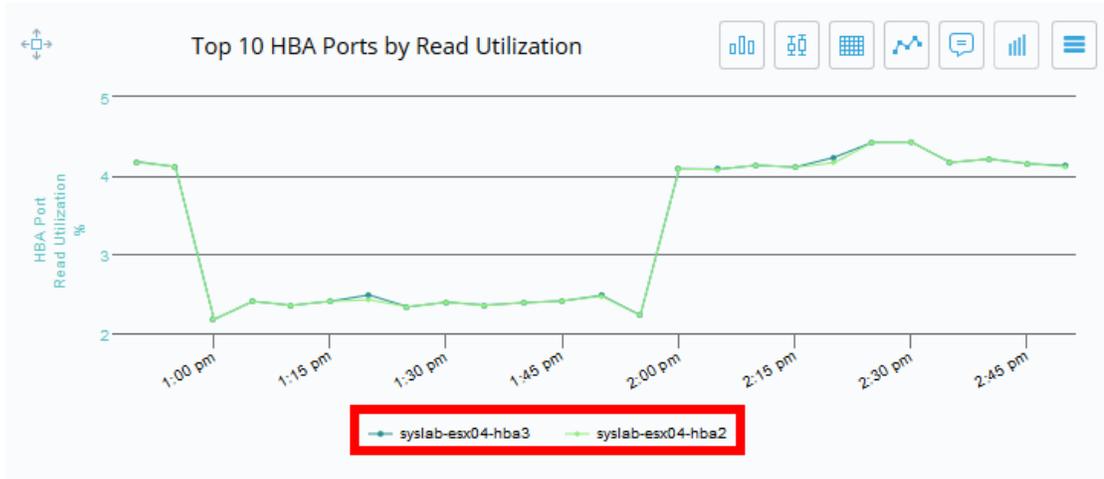
in this list Add/Remove 0 HBA Port(s)

in this list
not in this list
related to
with property

- a. **In this list:** The "in this list" option lets you choose specific named entities to filter by. Use the search box to search by entity name. Deselect the Show Archived box if you want to use only active entities. Check the entities you wish to filter by and click OK. Click OK in the Select Data box to apply your filter.



The chart shows data only for the selected entities.



- b. **Not in this list:** The "not in this list" is similar to the "in this list" filter except that it shows data only for entities that are not in the selected list.

Top N Chart - Select Data

Chart Data
 Parameters

Show the Top 10 HBA Port by Read Utilization

Apply Filters (optional) 1 Filter(s)

Include HBA Ports

not in this list Add/Remove 2 HBA Port(s) ✕

2100000e1ee5e4e1, 2100000e1ee5e630

OK
Cancel

- c. **Related to:** The "related to" filter lets you select entities that are related to the chart entity, to be used in the metric aggregation or calculation. The related to filter is useful when using report variables. While you can select a named entity to be used in the filter, you can also use a report variable, making it easy to filter the chart dynamically.

You can also limit the relation to a specified topology. This can speed up report rendering, allows you to be more selective about what is displayed in the chart, and lets you use topologies that are not included in filtering, like conversations.

The screenshot shows a dialog box titled "Top N Chart - Select Data". At the top, there are two tabs: "Chart Data" (highlighted in yellow) and "Parameters". Below the tabs, the configuration is as follows:

- Show the:** A dropdown menu set to "Top".
- Count:** A text input field containing "10".
- Entity:** A text input field containing "HBA Port".
- by:** A text input field containing "Read Utilization".

Below this is a section titled "Apply Filters (optional)" with a "1 Filter(s) +" indicator. Inside this section:

- Include HBA Ports:** A sub-section header.
- Filter:** A dropdown menu set to "related to", followed by a text input field containing "Application", a dropdown menu set to "in this list", and an "Add/Remove" button. To the right, it says "0 Application(s)".
- Limit to this topology:** A dropdown menu set to "Application - Fibre Channel".
- Instruction:** "Click the Add/Remove button to select Applications."

At the bottom of the dialog, there are two buttons: "OK" (highlighted in yellow) and "Cancel".

- d. **With property:** The "with property" filter lets you filter a chart based on a system property or tag. In this example, we're filtering HBA ports by their port speed, and using only those with 16 GB port speed in the chart aggregation.

Top N Chart - Select Data ✕

📊 Chart Data ⚙️ Parameters

Show the Top 10 HBA Port by Read Utilization

Apply Filters (optional) 1 Filter(s) +

Include HBA Ports

with property Port Speed ✕

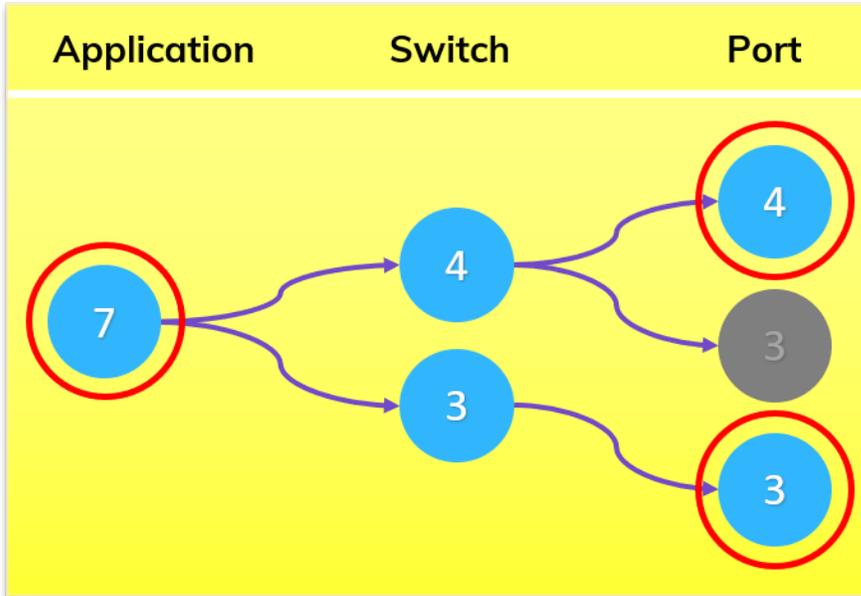
in this list SIXTEEN_GB ✕

OK Cancel

Data Filtering

Only the specified entity and metric are used in the aggregation. In the example shown below, the data is filtered for a single application, using data collected from only two of the related switch ports.

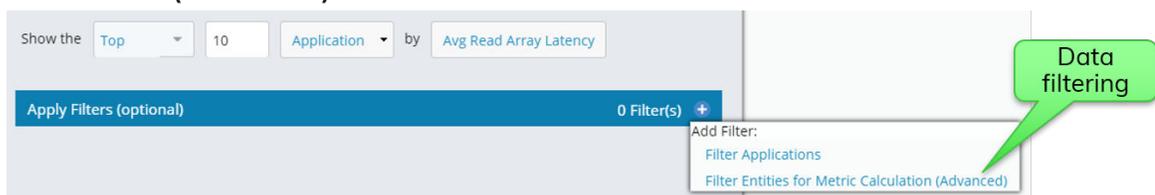
Data Filtering



Data filtering can be used on all charts but it's only available on Top N charts if bindings* are available for the selected metric. The Add filter button dynamically changes if bindings are available.

*Metric binding refers to an entity-metric combination, e.g., Switch Port-Consumed Bandwidth.

- To use data filtering in a chart, select the + sign, then choose **Filter Entities for Metric Calculation (Advanced)**.



There are three different selections that you can choose for a data filter.

- In this list:** "In this list" lets you select one or more named entities or a report variable.

Apply Filters (optional) 1 Filter(s) +

Calculate Read Utilization using data from

HBA Port in this list Add/Remove 0 HBA Port(s) x

Click the Add/Remove button to select HBA Ports.

in this list
not in this list
with property

- b. **Not in this list:** "Not in this list" reports on data collected from entities not in a specified list. You can also use a report variable with this selection.

Apply Filters (optional) 1 Filter(s) +

Calculate Read Utilization using data from

Application not in this list Add/Remove 1 Application(s) x

Application_1

- c. **With property:** The "with property" filter lets you filter a chart based on a system property or tag. In this example, we're filtering applications by their tier, and using only applications in the Platinum tier in the chart calculation.

Apply Filters (optional) 1 Filter(s) +

Calculate Read Utilization using data from

Application with property Tier x

in this list Platinum x

Creating Report Variables

Report variables are defined at the report level but used to filter charts. Report variables can be created at any time to an existing report or added when creating a new report.

Some chart types do not allow for use of variables with the primary data selection because you do not select specific entities. An example would be the Top N Chart. However, you can still use variables in metric filters for charts, to create relationships in any chart type that allows filtering.

If your report includes chart types that do not specify entities (such as Top N Chart), those charts are not affected by use of a variable on other charts in the report.

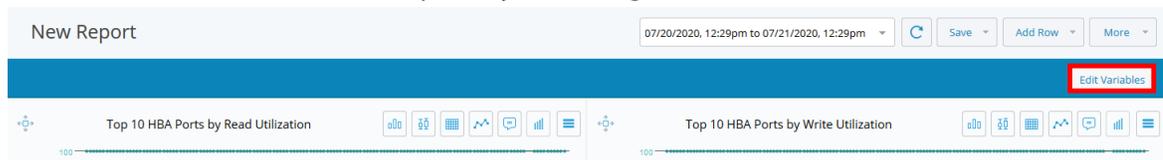
You can add multiple variables to a report and associate the variables with various entity types.



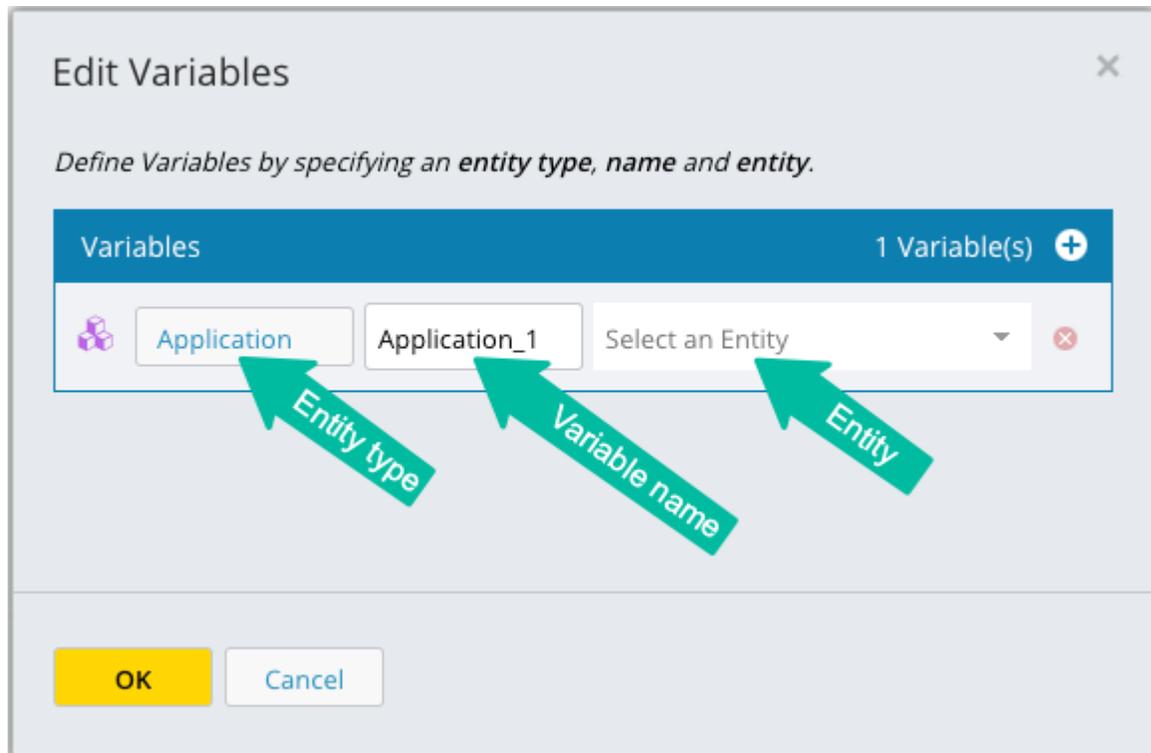
TIP

It saves some steps if you first create the variables you want for a report, then select the variables as you add the charts to the report.

1. Navigate to the **Reports** or **Reports Templates** page.
2. Click **New Report** or open an existing report.
3. Define new variables for the report by selecting **Edit Variables**.



4. Use the + sign to add a new variable.



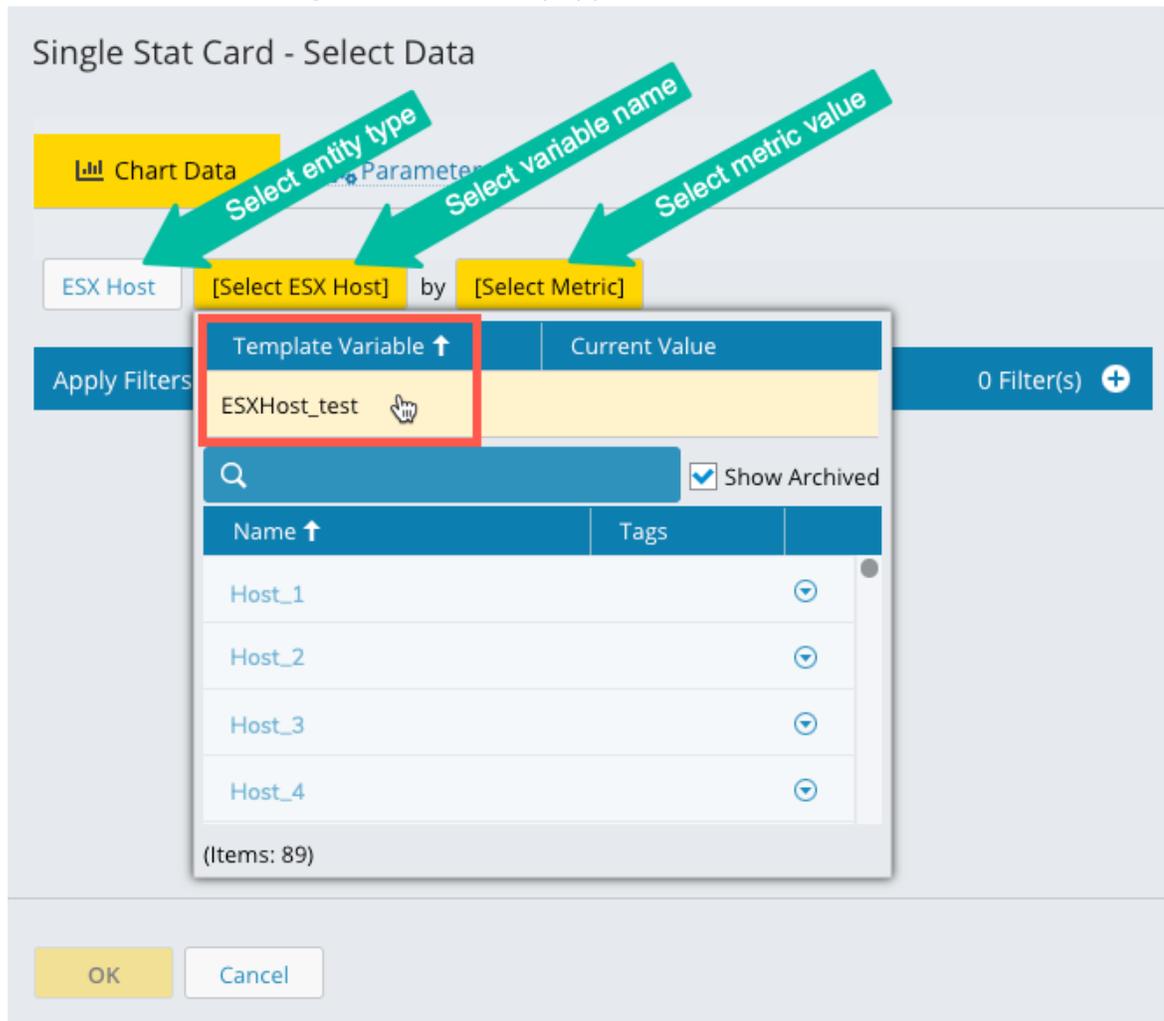
Select the entity type to be used in the variable.

The variable is named automatically but you can change the name using the second field. Only letters, numbers, underscores, and periods are allowed in the name.
Optional: Select a specific entity value in the third field.

- The variable is displayed at the top of the page.



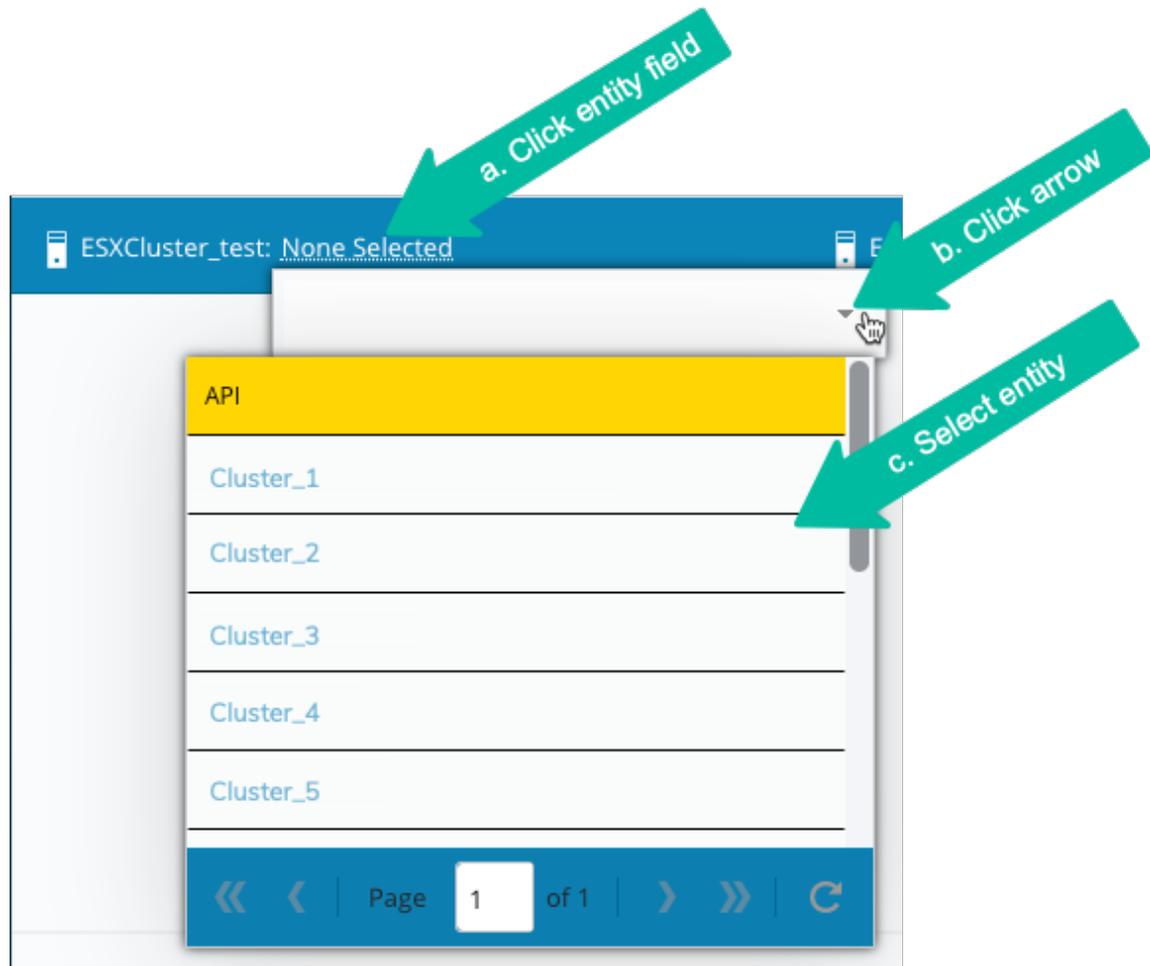
- The variable can now be used to filter charts in your report. To add a variable to a chart, click  to edit the chart.
- On the **Chart Data Page**, select an entity type in the first field.



In the second field, select from the list of report variables.

In the third field, select a metric value.

8. Click **OK**. The report page displays. You can now populate the report variable and the chart will be filtered for the value selected for the variable.



Changing Report Variables Using a URL

You can change the entity value for a report template variable either from the variable field on the report page, or by passing the entity name in a URL.

Prerequisites

- The report you want to work with must have been saved.
- The variable name must have been associated with the charts in the report.
- You need to know the names of the variable templates you want to use and the names of the entities for which you want to view data.

Steps

1. Navigate to the Reports or Report Templates page and select a saved report.
The browser URL field displays the static ID of the report in the format:
`https://appliance-id/#tab-name/page-name/report-uid`
The appliance name/ID + tab/page name + UID = the static URL of the report.
2. In the URL field, add the variable and entity name after the report ID in the format:
`/variable-name=entity-name`
Example:
`https://198.51.100.5/#reports/template/e9678-123-8910-b23eio/ESXCluster_test=Cluster_1`
198.51.100.5 = IP address (or name) of the VirtualWisdom appliance
#reports/template = The names of the tab and the page in the VirtualWisdom UI
e9678-123-8910-b23eio = The UID of the report
ESXCluster_test = The template variable name
Cluster_1 = The entity value
3. Click **Enter**.
The report page updates with the new data and the variable name and entity value display above the charts.

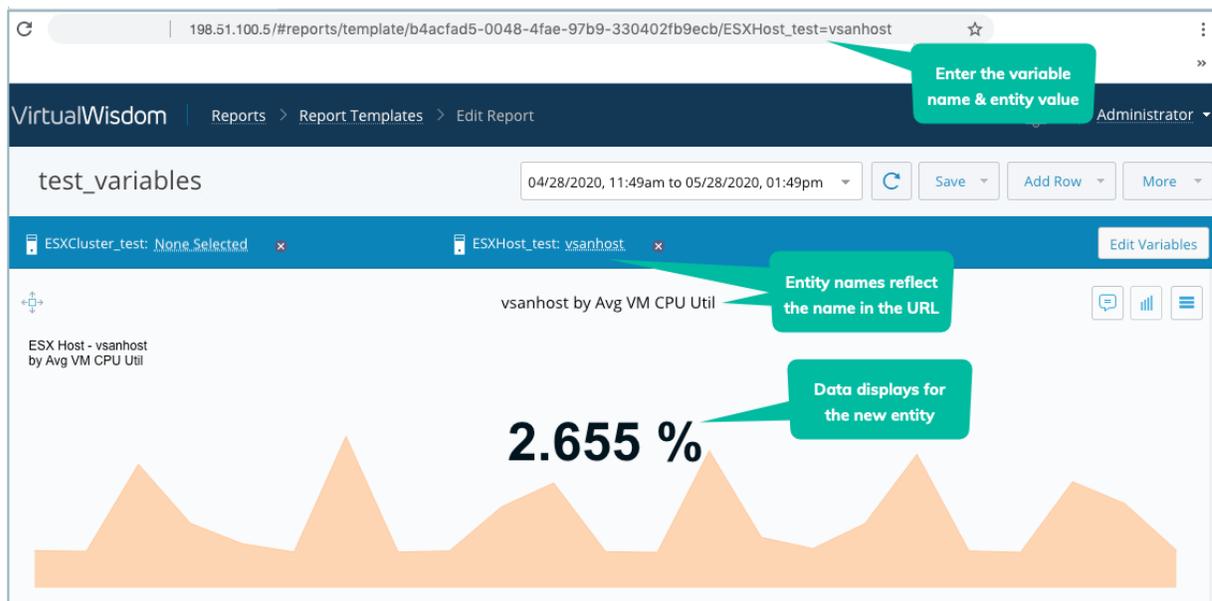
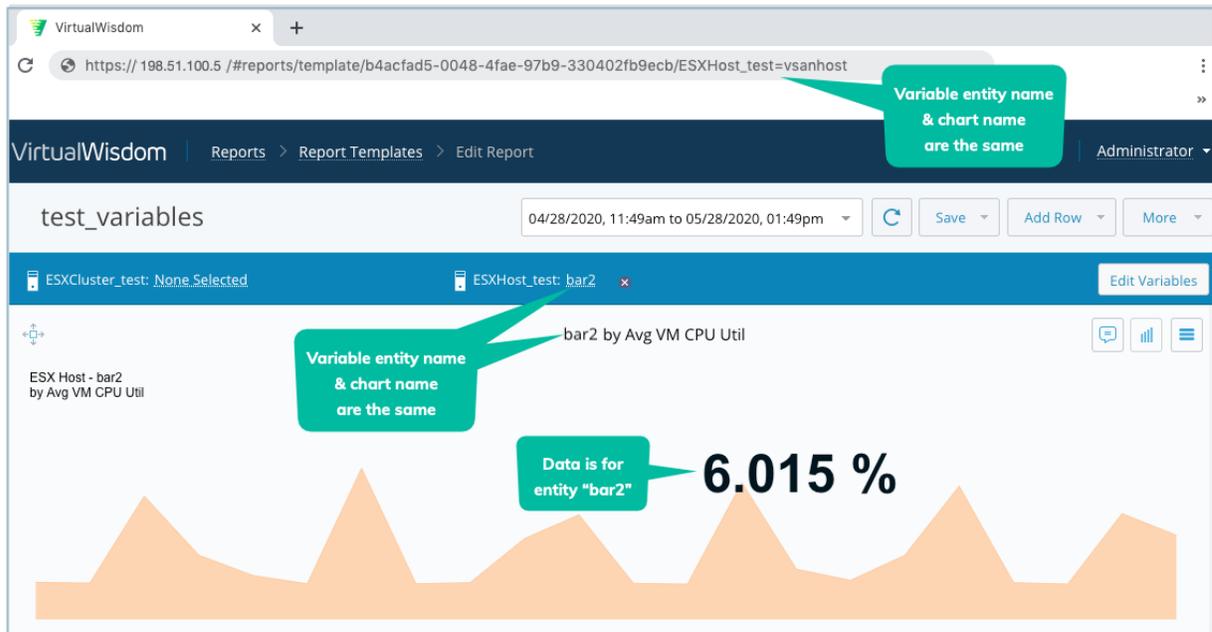


TIP

If you select a different entity using the field in the UI, it updates the data on the page for that entity, but the URL does not change to reflect the new entity name.

Example 1. Example: Using a URL to Pass Report Variables

The following images show how adding the template variable name and entity value to a report's static URL changes the data displayed in a report.



Report Snapshots



A snapshot is a point-in-time copy of a report. Snapshots allow you to keep an archived copy of report data. By design, Report Snapshots are read-only.

The *Report Snapshots* page shows the results of previous report executions that were saved as snapshots.

From the *Report Snapshots* page you can perform the following actions:

- Delete individual snapshots
- Configure automatic scheduled deletion of snapshots (Snapshots Cleanup)
- Add or remove tags on a snapshot
- Select a snapshot to view

If you select a snapshot in the *Report Snapshots* list, the content of the snapshot displays on a *View Results* page. From this page you can perform the following actions:

- Add callouts, if available (not available for some report types)
- Show, hide, or delete callouts
- View report attributes
- Export snapshot content
- Save changes to the snapshot or save the snapshot with a different name

Scheduling Snapshot Cleanup

Retaining a large number of report snapshots can make it difficult to manage the snapshots. Having a large number of snapshots also takes longer to display on the Snapshots page. You can delete snapshots by selecting individual snapshots and deleting them, or you can use Snapshots Cleanup to automatically delete snapshots older than a user-selected age.

About This Task

When you enable automatic deletion of report snapshots, all snapshots older than the time period you select will be automatically deleted and can only be retrieved by doing a restore from a backup. Be sure there are no snapshots you want to retain that would fall within the selected deletion timeframe.

Steps

1. Navigate to the **Reports** page and click **Report Snapshots**.
The *Report Snapshots* page displays.
2. Click **More > Snapshots Cleanup**.

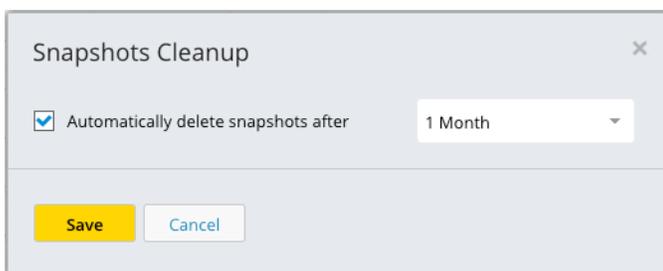


3. Enable automatic deletion of snapshots and select the time period after which snapshots will be deleted.



NOTE

All snapshots older than the selected time period will be deleted.



4. Click **Save**.
A warning displays, stating that snapshots will be immediately deleted.
5. Click **OK**.
A message displays stating that snapshots were deleted.
6. Click **OK**.
A message displays on the Report Snapshots page, informing you that snapshot cleanup is enabled.



TIP

You can disable Report Snapshots Cleanup by navigating to the Snapshots Cleanup window and clearing the checkbox.

Event Integration

Event integration creates an event framework from which VirtualWisdom can ingest events from probes and integrations. VirtualWisdom correlates these events with

monitored entity metrics. You can also overlay events onto a time-series report to provide additional context.

Using Reports and Charts, event integration gives you the ability to determine how abnormal events impact an entity's metrics during a specific period of time.

In VirtualWisdom 6.0, only AppDynamics's health violation events with critical severity can be pulled. These events are saved in a VirtualWisdom database. If you display a Line Chart of a specific host entity using metrics, VirtualWisdom pulls the entity-related events for that time period from the database and highlights them on the chart. You can turn the events overlay on or off for specific line charts.

AppDynamics's event integration uses the configuration settings associated with the AppDynamics APM Integration to access AppDynamics's controller to pull events. Currently, the VirtualWisdom UI does not allow you to specify preferred types of violation events to download and overlay on the chart. It also does not allow you to specify different severity levels, and so forth.

Event integration helps to establish correlation information about infrastructure and events. Using events, you can tell how application performance is impacted by infrastructure metrics. You can also tell why entity metrics charts have anomalies.

Enabling Event Polling

1. Click **Settings** in VW, *Integrations* in the *Probe Management* section, and then the **View** button for *AppDynamics APM*. The AppDynamics APM page displays.

AppDynamics APM

Authentication and Settings

Use SSL

Controller FQDN *
 e.g. myinstance.appdynamics.com

Port * C

Account Name *

Username *

Password *

Application Discovery Time and Frequency

Enable scheduled discovery

Frequency

Start Time

2. Ensure that the *Enable scheduled discovery* checkbox is checked. This is the default for the AppDynamics APM Integration.

Viewing Events

You can view events using Line Charts for AppDynamics applications.

When creating a chart, the default chart configuration is to “hide events.” To view events, select *Show related events* from the *Events* drop-down in the *Parameters* tab of the *Select Data* dialog when you configure your chart.

Line Chart - Select Data

Chart Data
 Parameters

Name:

Tooltips:

Axes Mode:

Red Threshold:

Yellow Threshold:

Markers:

Events:

Legend:

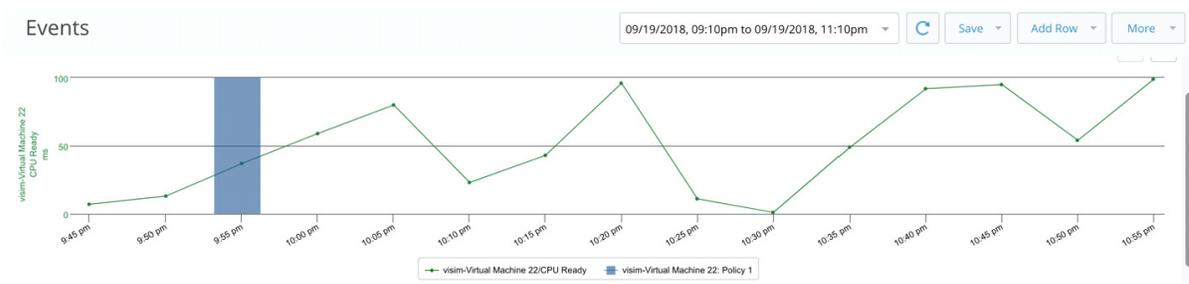
OK
Cancel

Event information is part of a scheduled report. Exported reports or charts of image type (PNG, JPG, and PDF) also capture events.

The following chart shows pulled events in a Line Chart.

Each event has an individual legend. You can use it to show or hide an individual event on a chart.

The thickness of the event line dictates how long the event was active.



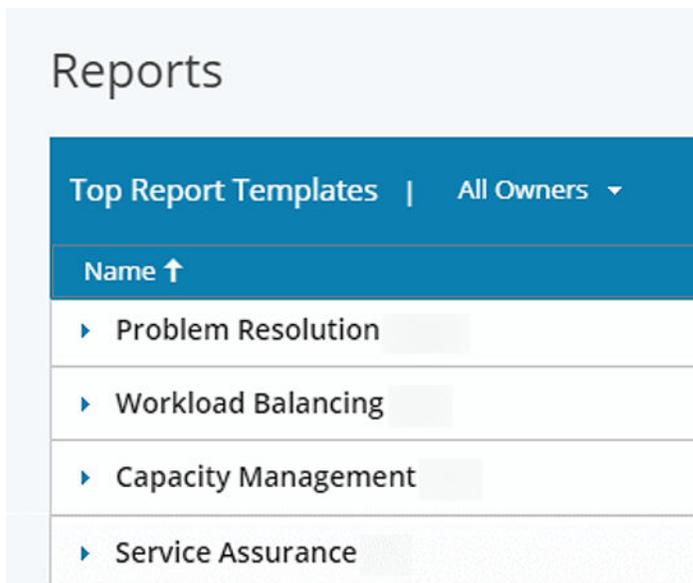
Disabling (Stopping) Event Polling

1. Unschedule AppDynamics application discovery.
2. Delete the specific AppDynamics APM Integration.

Standard Services Reports

The Virtana Services reports are a set of reports that can be installed in your VirtualWisdom portal to help you resolve issues and problems, analyze workload balance and utilization, manage capacity, and assure service levels.

The reports are organized into four categories:



Services reports are prefixed with "VI. The report title tells you the data source (probe/integration), report purpose, and version.



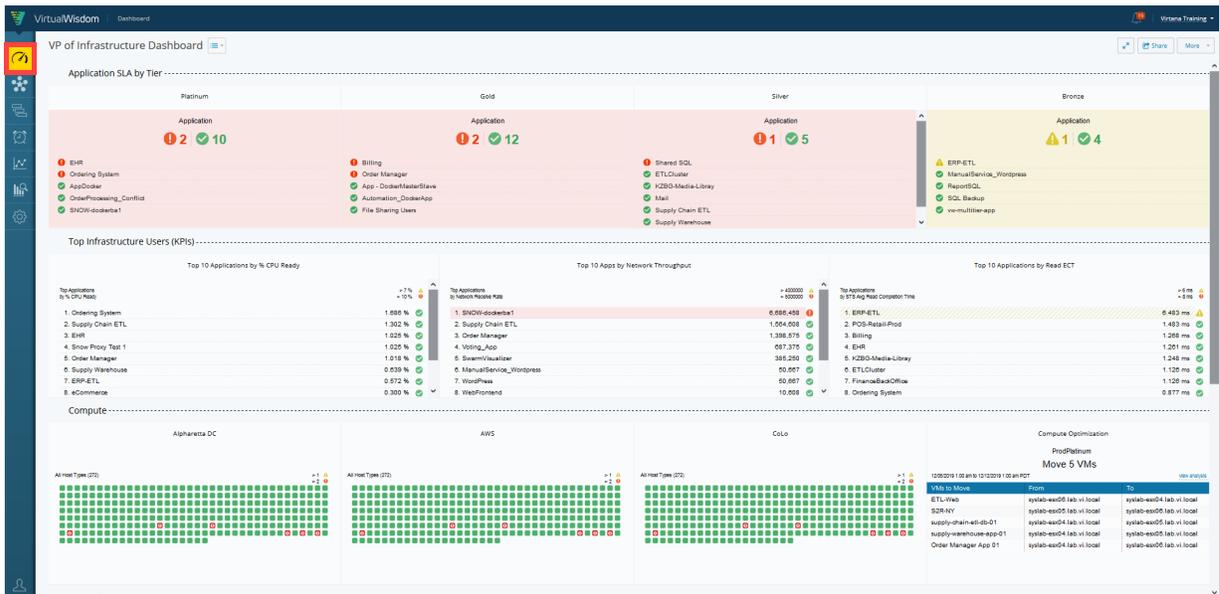


NOTE

Contact your Virtana Services team to have these reports installed in your portal.

Dashboards

Dashboards are a subset of reports and leverage some of the chart types we covered earlier.



The dashboard is a non-editable view of a report that can be used to provide an external view to other users, for example, for display in a NOC.

The main differences between a dashboard and a report are:

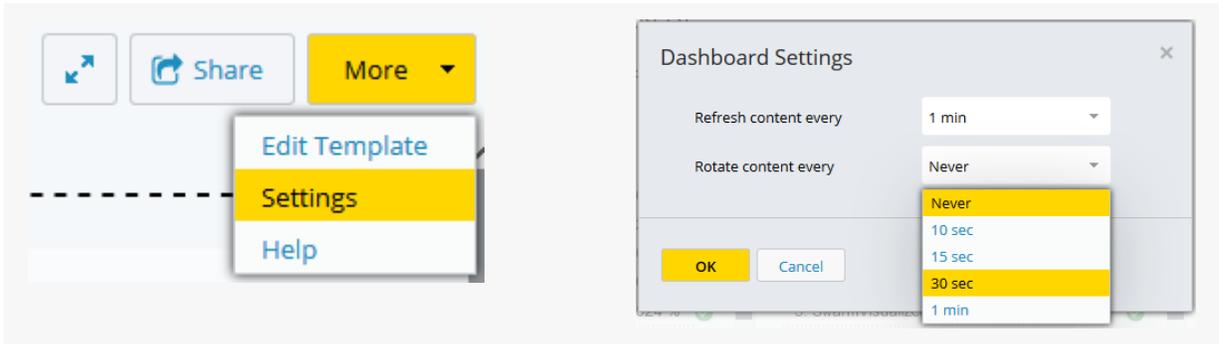
- The dashboard is displayed in the Dashboard module in the VirtualWisdom UI instead of the Reports module,
- The dashboard data continuously refreshes at a specific time interval while a report is run at a specific point of time, and
- The sections of the dashboard rotate from top to bottom at a specified interval while the sections of a report remain fixed.
- Dashboards do not use variables.

You can expand the dashboard to full screen by clicking on the two diagonal arrows in the header. This is also referred to as the External View of the dashboard and is useful for viewing the report in a monitoring center. Click the arrows again to minimize the view.

You can share the dashboard via email.



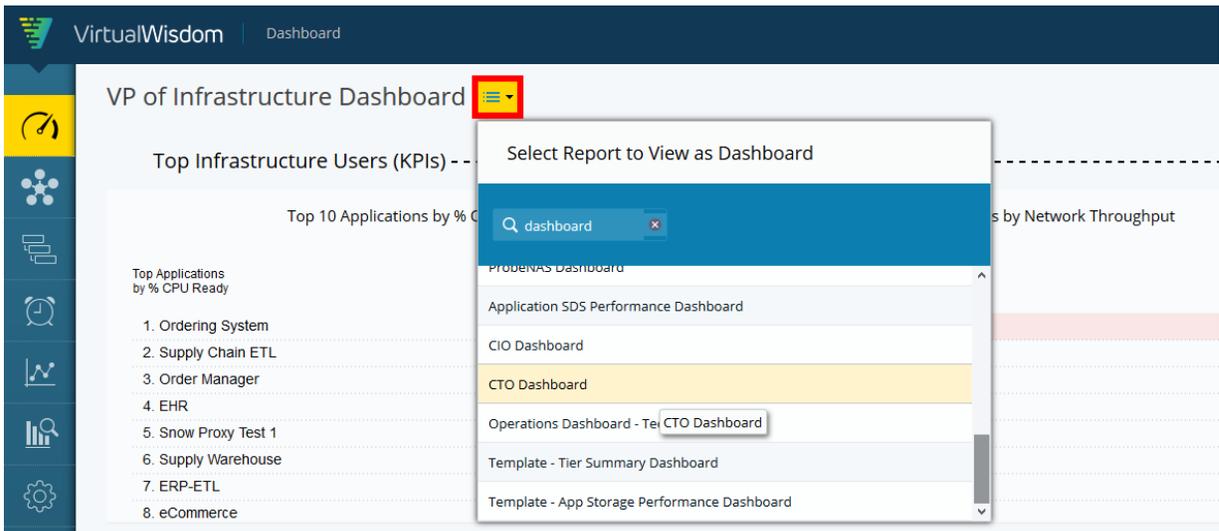
You can also change the interval at which the dashboard view refreshes its content and rotates through rows on the report. Select Settings to change the intervals.



There are standard dashboard templates available that you can use as a starting point for using dashboards. We suggest starting with these templates:

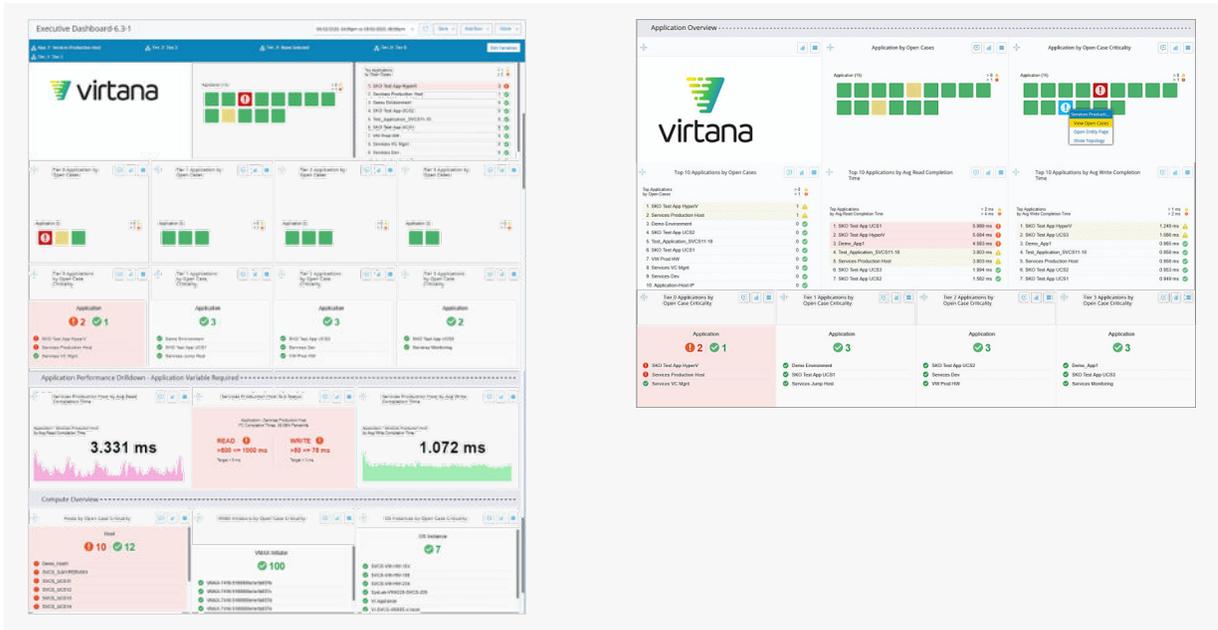
- Application Health by Tier: shows you which applications are experiencing issues and lets you drill down into open cases to investigate the issue and its possible causes.
- App Storage Performance Dashboard: presents a view into the performance of the storage supporting your application.
- Tier Summary Dashboard: shows you information for all the applications in a specific tier.

These templates can be found by selecting the hamburger icon on the Dashboard home page and selecting the template.



Visibility Dashboards

The Visibility Dashboards were designed to provide visibility to executives and application owners into how the infrastructure supporting their applications is performing.

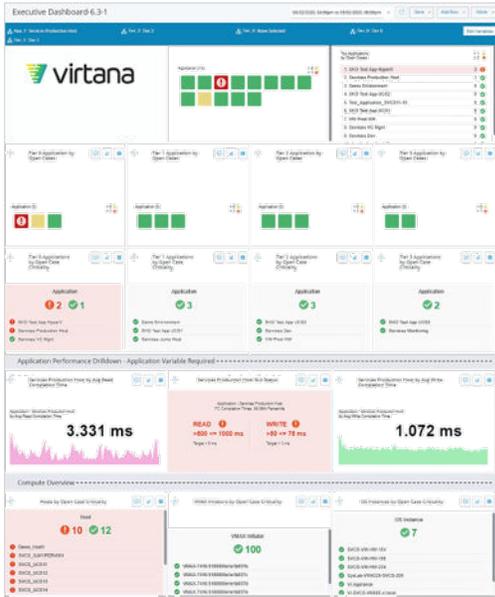


 **NOTE**

Contact your Virtana Services team to have these reports installed in your portal.

Executive Dashboard

The Executive Dashboard presents health and performance data for the infrastructure supporting the applications.



Report Variables

You can filter the report for up to four tiers. Setting the tier variables allows you to compare data for your VirtualWisdom tiers side-by-side.

The Application variable is used to view performance data for a single application.



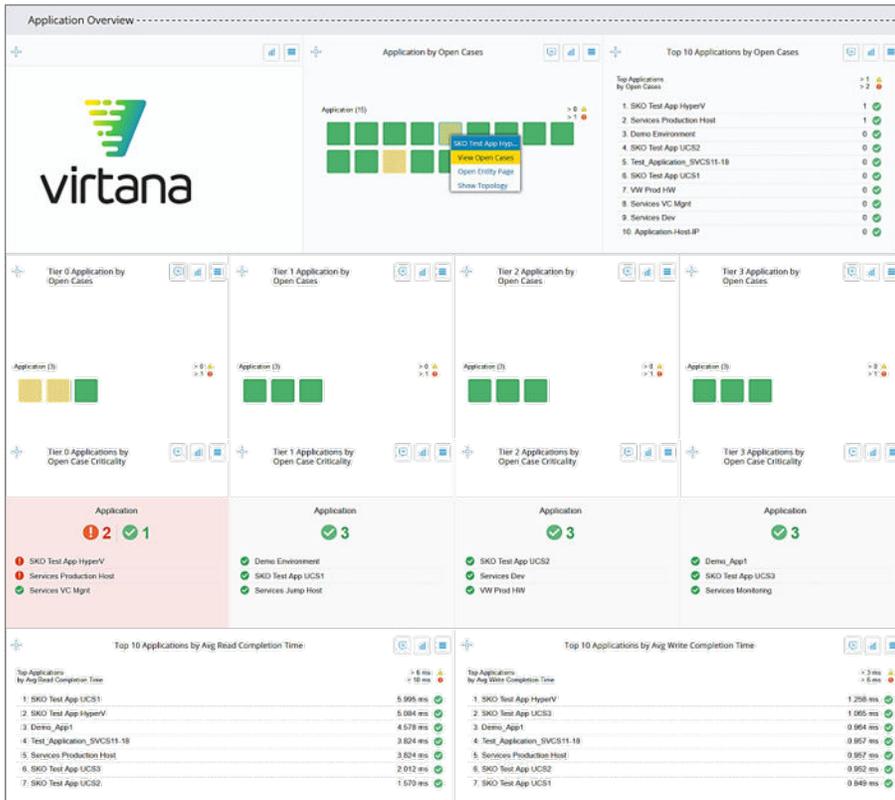
Report Sections

The report is divided into six sections:

- 1. Application Overview**

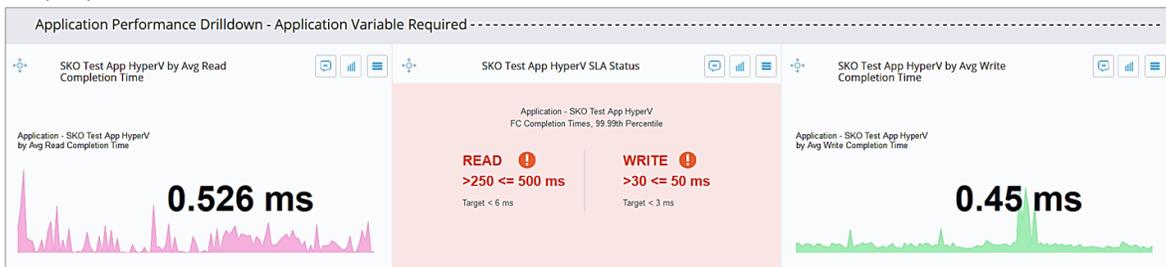
Use the Application Overview section to quickly observe where problems exist in the infrastructure supporting the applications.

Set the Tier variables to view summary health and performance data for applications. View open cases to use the investigations to troubleshoot and remediate issues. Note applications with issues to use to filter the report.



2. Application Performance Drill Down

Use the Application Overview section to identify an application you'd like to drill down on and select it in the App_1 report variable. Data for that application will be displayed.



Displays the application's average read completion time for the last 2 hours and its current value.

Displays the SLA status for FC completion times. In this example, the SLA threshold is set to 99.99%, with targets of < 6 ms for read completion and < 3 ms for write completion.

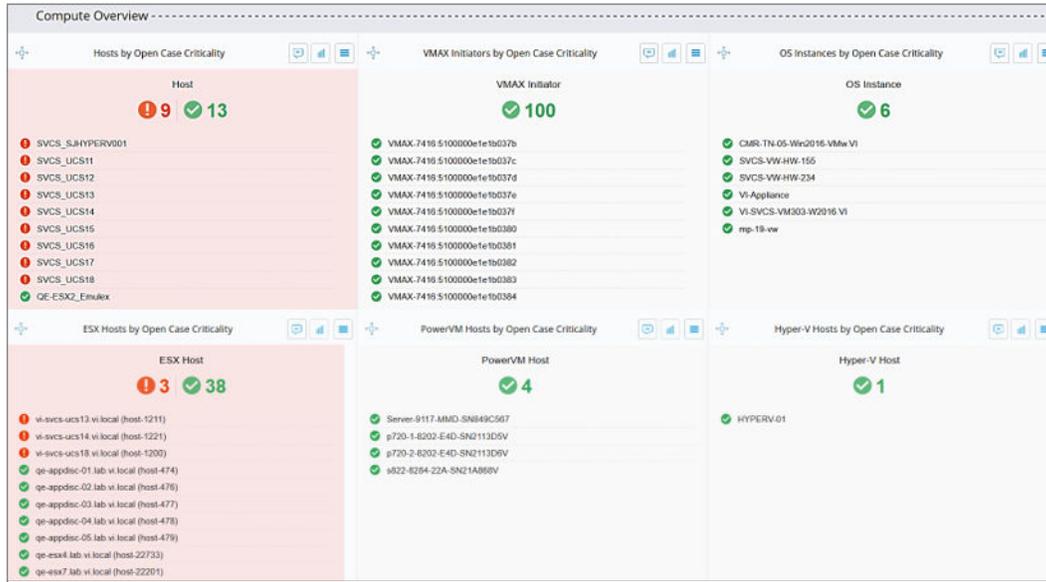
This chart shows the SLA status is red for this application.

Displays the application's average write completion time for the last 2 hours and its current value.

3. Compute Overview

This section shows you the overall health of your Compute environment by open case criticality. The default dashboard includes hosts, ESX Hosts, VMAX Initiators, OS Instances, PowerVM Hosts, and Hyper-V Hosts.

Drill down to view open cases, properties, and topology of any host.



4. Storage Overview

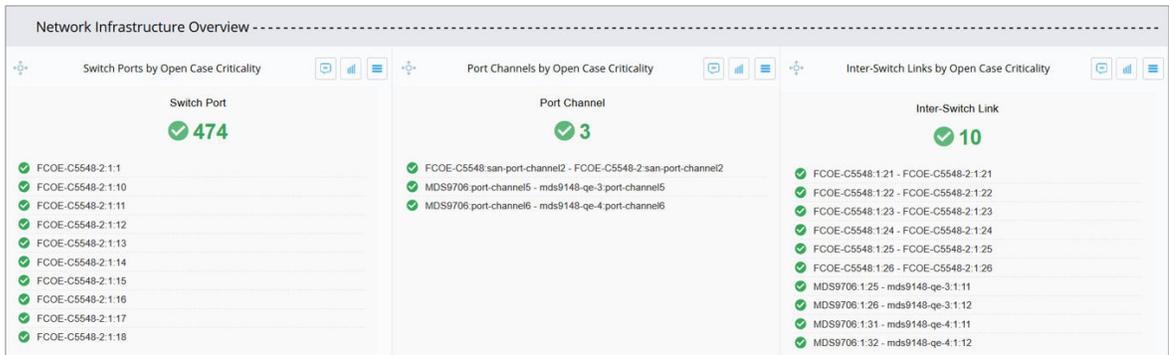
This section shows you the overall health of your Storage environment by open case criticality. The default dashboard includes SAN Storage Arrays, VMAX Storage Arrays, Isilon Clusters, SVC Clusters, VxFlex OS Systems, and NetApp Clusters.

Drill down to view open cases, properties, and topology of any storage component.



5. Network Infrastructure Overview

Use the Network Infrastructure Overview section to view the health of your SAN Fabric network infrastructure.



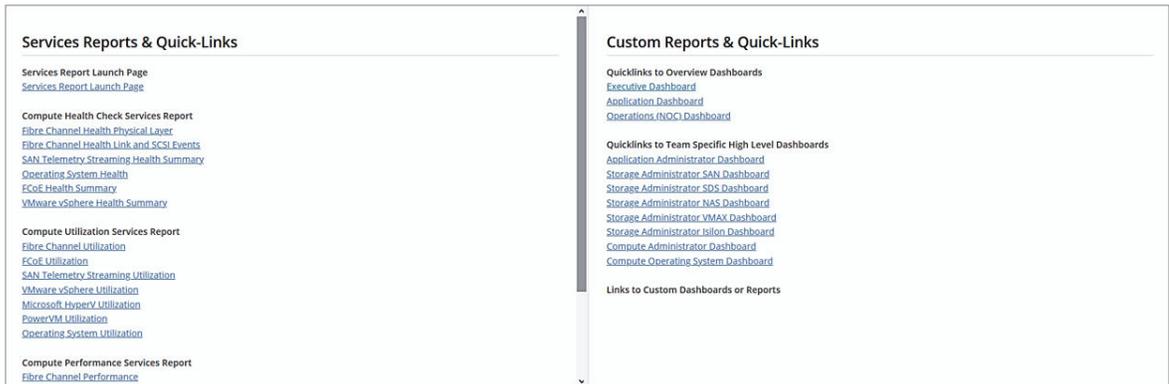
Displays the top switch ports by open case criticality.
Drill down to view open cases, properties, or topology.

Displays the top port channels by open case criticality.
Drill down to view open cases, properties, or topology.

Displays the top inter-switch links by open case criticality.
Drill down to view open cases, properties, or topology.

6. Report Quick Links

Use this section to view additional reports and dashboards.

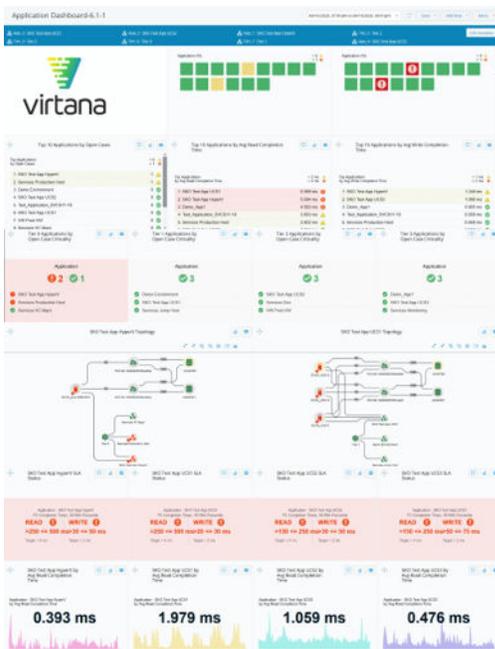


Displays links to the Services Reports that have been installed on your VirtualWisdom portal.

Displays links to the other standard dashboards and custom reports that you've created in your VirtualWisdom portal.

Application Dashboard

The Application Dashboard presents health and performance data for the applications.



Report Variables

You can filter the report for up to four applications and four tiers. Setting these variables allows you to compare data for your VirtualWisdom applications and tiers side-by-side.



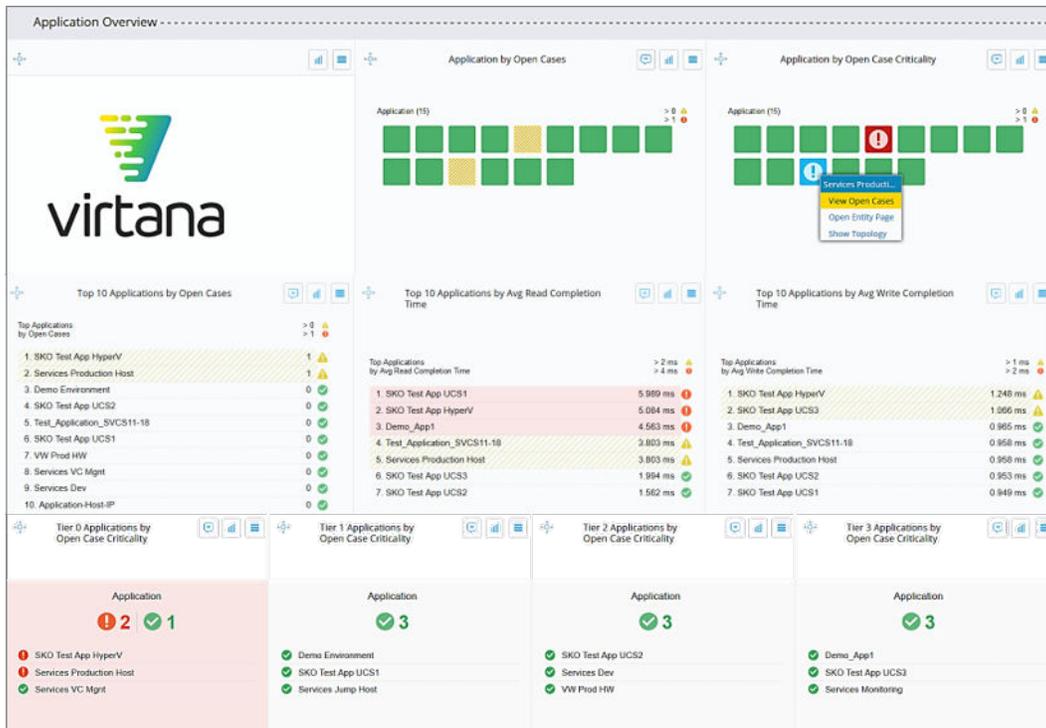
Report Sections

The report is divided into the following sections:

1. Application Overview

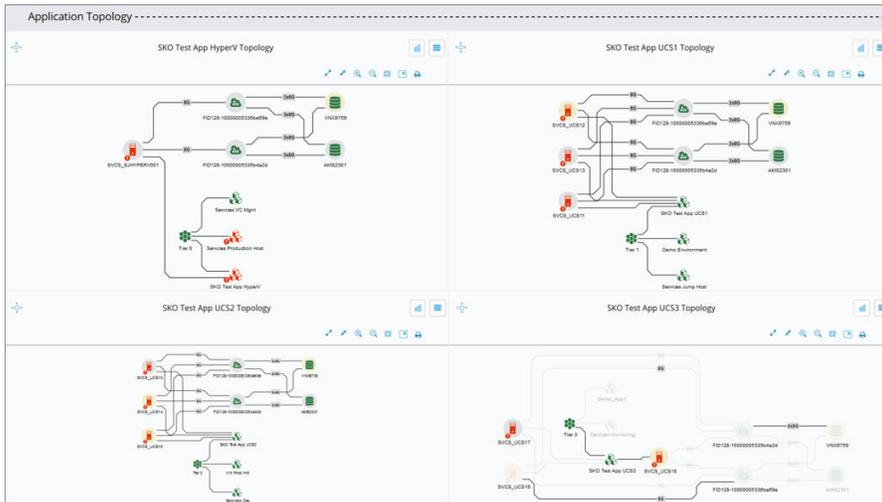
Use the Application Overview section to quickly observe which applications are experiencing issues.

Set the Tier variables to view summary health and performance data for applications. View open cases to use the investigations to troubleshoot and remediate issues. Note applications with issues to use to filter the report.



2. Application Topology

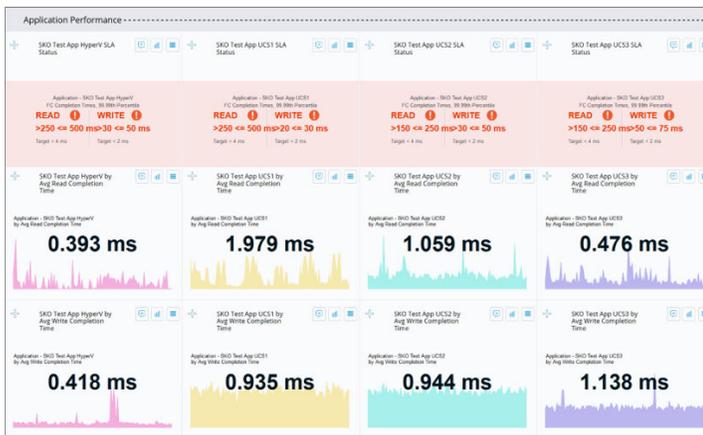
The Application Topology section shows the topology for up to four applications. Set the application variables to use this section.



The topology views for the applications filtered using the dashboard's variables are displayed in a grid.

3. Application Performance

This section reports on SLA status and performance for up to four applications. Set the application variables to use this section.



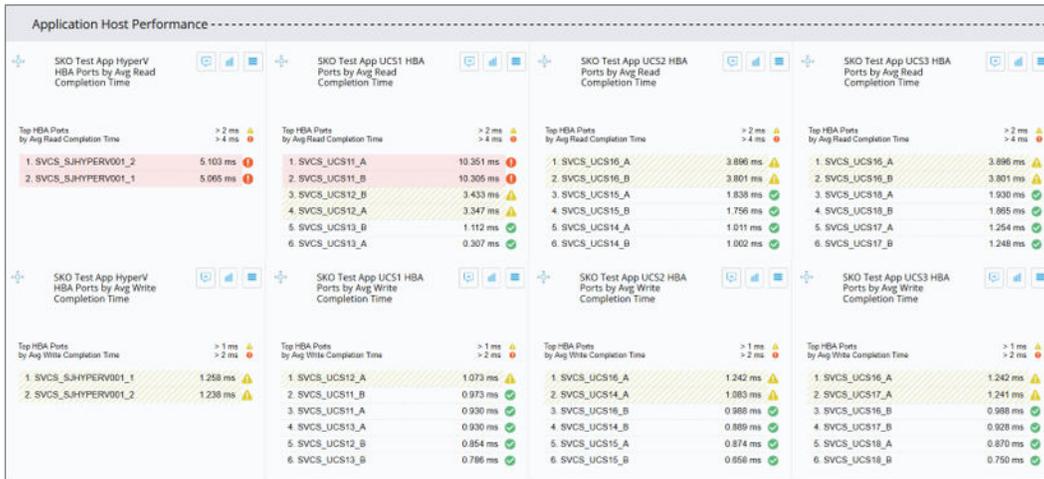
The first row shows you the SLA status for the four applications selected using the report filter.

The rows below show the average read and write completion times for the four applications.

Application 1 Application 2 Application 3 Application 4

4. Application Host Performance

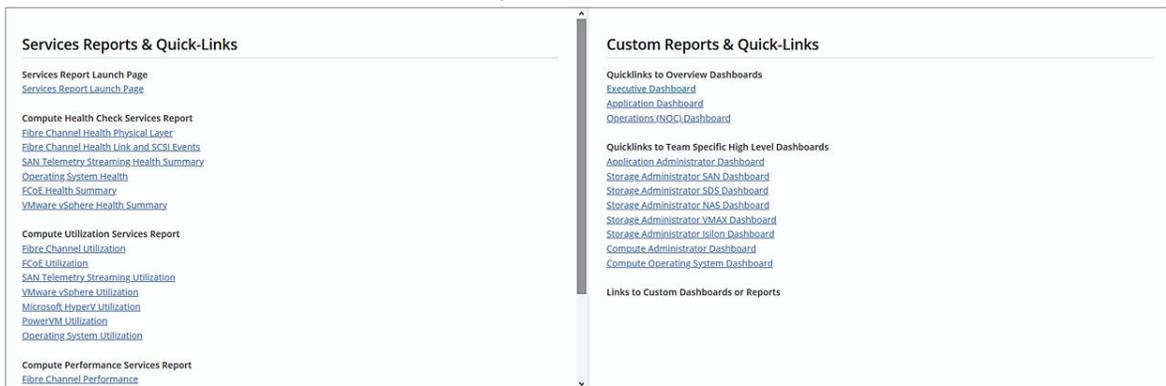
This section shows you the average read and write performance on the HBA ports of the hosts supporting the applications.



Application 1 Application 2 Application 3 Application 4

5. Report Quick Links

Use this section to view additional reports and dashboards.

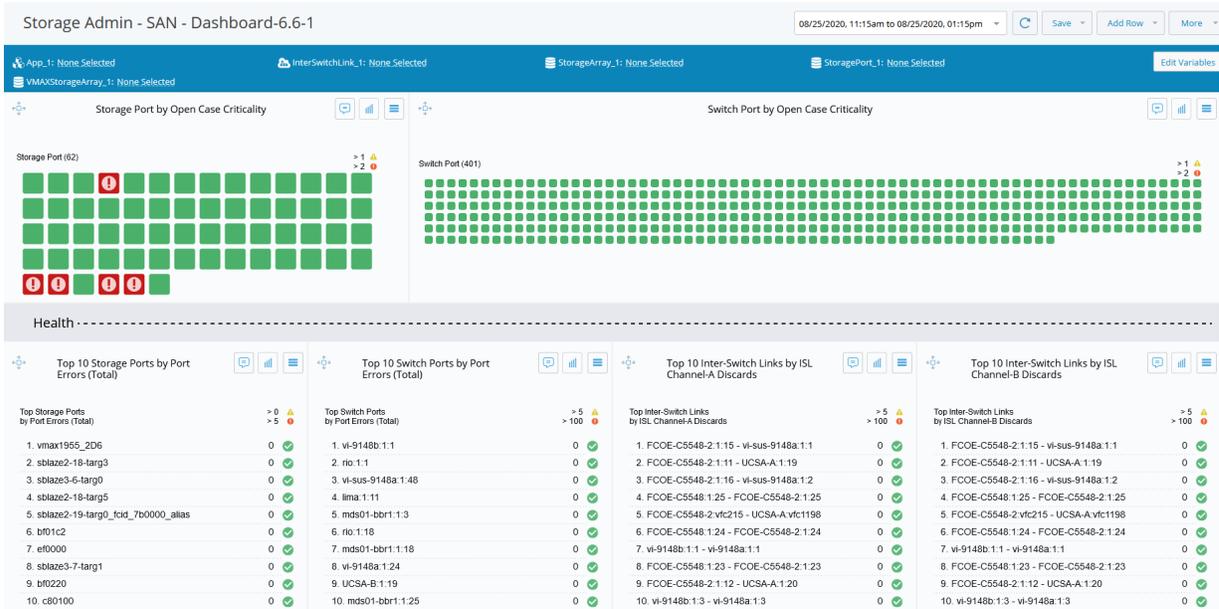


Displays links to the Services Reports that have been installed on your VirtualWisdom portal.

Displays links to the other standard dashboards and custom reports that you've created in your VirtualWisdom portal.

Admin Dashboards

Virtana Services has created a set of dashboards designed for use by infrastructure administrators. These admin dashboards can be used as a starting point for assessing your infrastructure, identifying issues, and performing troubleshooting exercises.

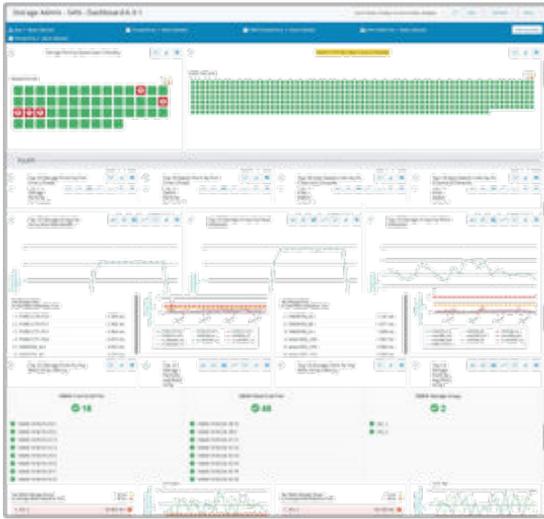


 **NOTE**
 Contact your Virtana Services team to have these dashboards installed in your portal.

Storage Admin Dashboards

The Storage Admin Dashboards present health, utilization, and performance data for storage infrastructure components.

The dashboard is available for **FC SAN**, **Isilon**, **NAS**, **SDS**, and **VMAX** integrations.



Report Variables

You can filter the report for an application, storage array, inter-switch link, or storage port. Use the filters to drill down on a specified infrastructure component.

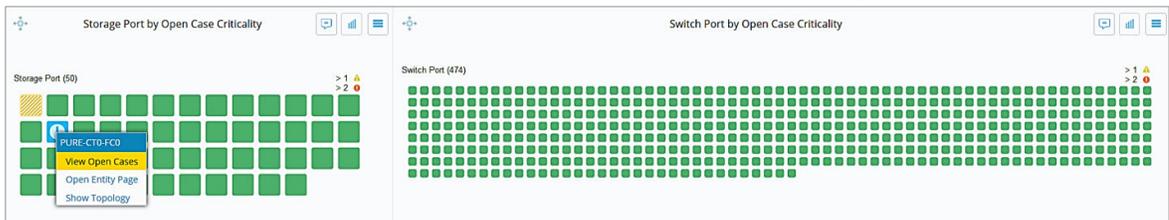


Report Sections

The report is divided into eight sections:

1. Storage Overview

Use the Storage Overview section to quickly observe where problems exist in the FC SAN infrastructure. View open cases to use the investigations to troubleshoot and remediate issues. Note entities with issues to use to filter the report.



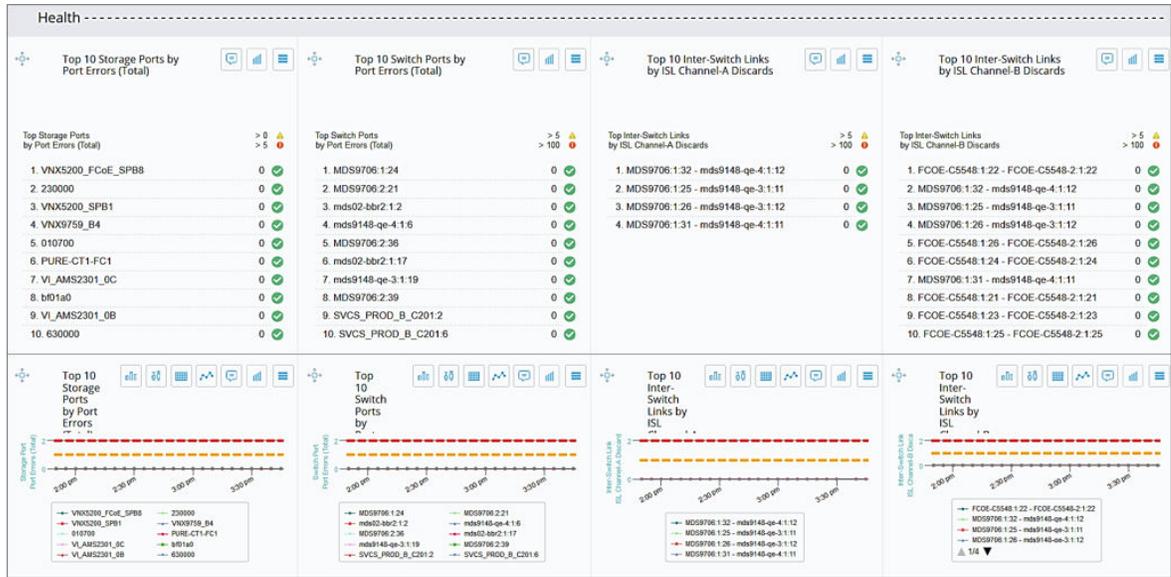
Displays the top storage ports by open case criticality.
 Drill down to view open cases, properties, or topology.

Displays the top switch ports by open case criticality.
 Drill down to view open cases, properties, or topology.

2. FC SAN Health

This section is filtered by application and/or storage port and shows you where port errors and Class 3 Discards are occurring for storage ports, switch ports, and inter-switch links.

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information and target these entities for physical layer issue investigation.

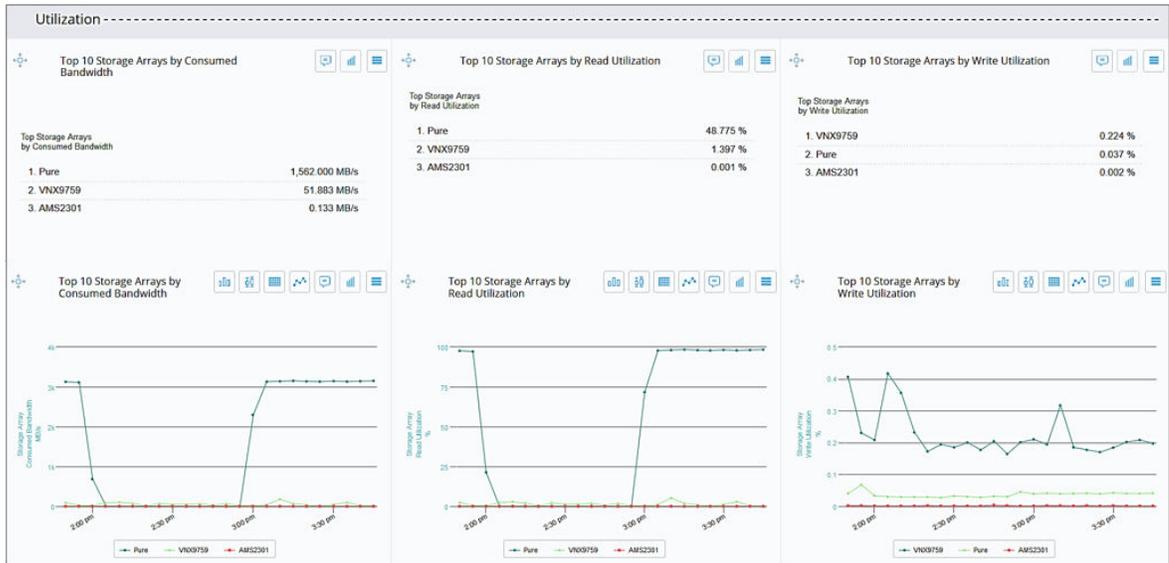


3. FC SAN Utilization

This section is filtered by application and shows you utilization data for storage arrays.

- Consumed bandwidth
- Read/write utilization

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information. Target these entities for workload balance and capacity investigation



4. FC SAN Array Latency

This section is filtered by application and storage port and shows you performance data for storage ports.

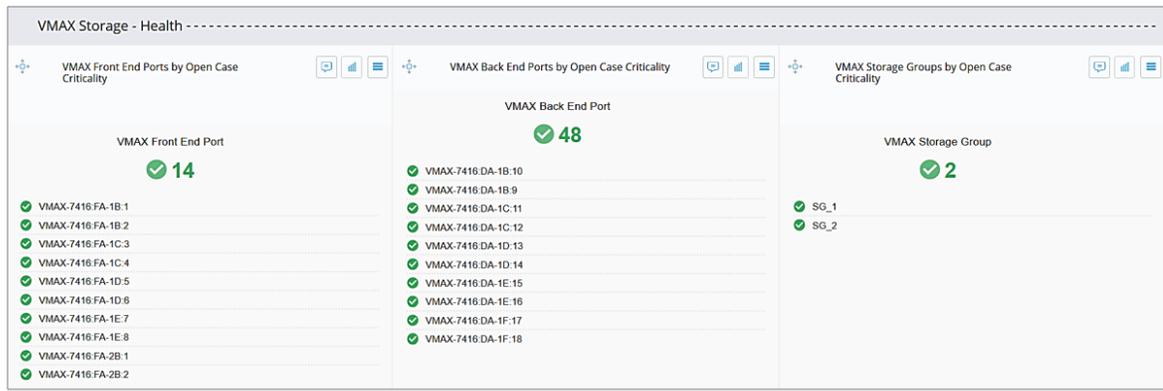
- Average read/write completion time
- Average read/write array latency

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information. Target these entities for flow control investigation.



5. VMAX Storage Health

Use the VMAX Storage Health section to quickly observe where problems exist in the VMAX infrastructure. View open cases to use the investigations to troubleshoot and remediate issues. Note entities with issues to use to filter the report.



Displays the top VMAX front end ports by open case criticality.

Drill down to view open cases, properties, or topology.

Displays the top VMAX back end ports by open case criticality.

Drill down to view open cases, properties, or topology.

Displays the top VMAX storage groups by open case criticality.

Drill down to view open cases, properties, or topology.

6. VMAX Storage Utilization

This section is filtered by VMAX storage array and shows you utilization data for VMAX front end and back end ports.

- Port utilization
- Write pending slot utilization
- RDF utilization

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information. Target these entities for workload balance and capacity investigation.

7. VMAX Performance/Latency

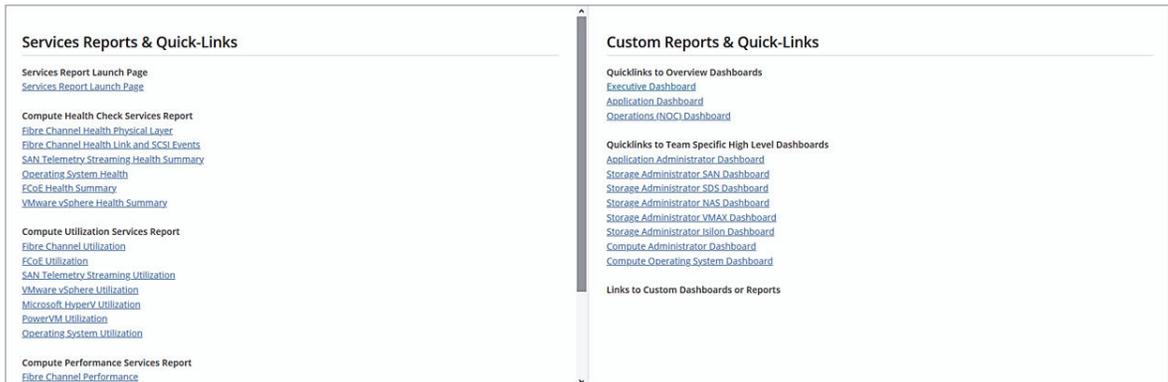
This section is filtered by VMAX storage port and shows you performance data for VMAX front end ports and storage groups

- Average read/write response time
- Average read/write array latency

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information. Target these entities for investigation.

8. Report Quick Links

Use this section to view additional reports and dashboards.



Displays links to the Services Reports that have been installed on your VirtualWisdom portal.

Displays links to the other standard dashboards and custom reports that you've created in your VirtualWisdom portal.

Compute Admin Dashboard

The Compute Admin Dashboard presents health, utilization, and performance data for compute infrastructure.

The dashboard is available for **FC SAN**, **VMware vSphere**, **Microsoft Hyper-V**, **IBM PowerVM**, and **Host OS** integrations.



Report Variables

You can filter the report for an **ESX host**, **host**, **VMAX host**, or **OS Instance**. Use the filters to drill down on a specified infrastructure component.

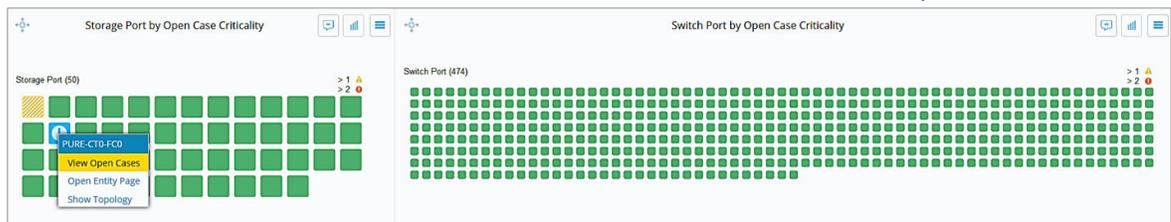


Report Sections

The report is divided into five sections:

1. Compute Overview

Use the Compute Overview section to quickly observe where problems exist in the Compute infrastructure. View open cases to use the investigations to troubleshoot and remediate issues. Note entities with issues to use to filter the report.



Displays the top storage ports by open case criticality.

Drill down to view open cases, properties, or topology.

Displays the top switch ports by open case criticality.

Drill down to view open cases, properties, or topology.

2. Compute Health

This section is filtered by host and ESX host and shows you where port errors, loss of sync, loss of signal, link errors, and discards are occurring.

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information. Target these entities for physical layer investigation.

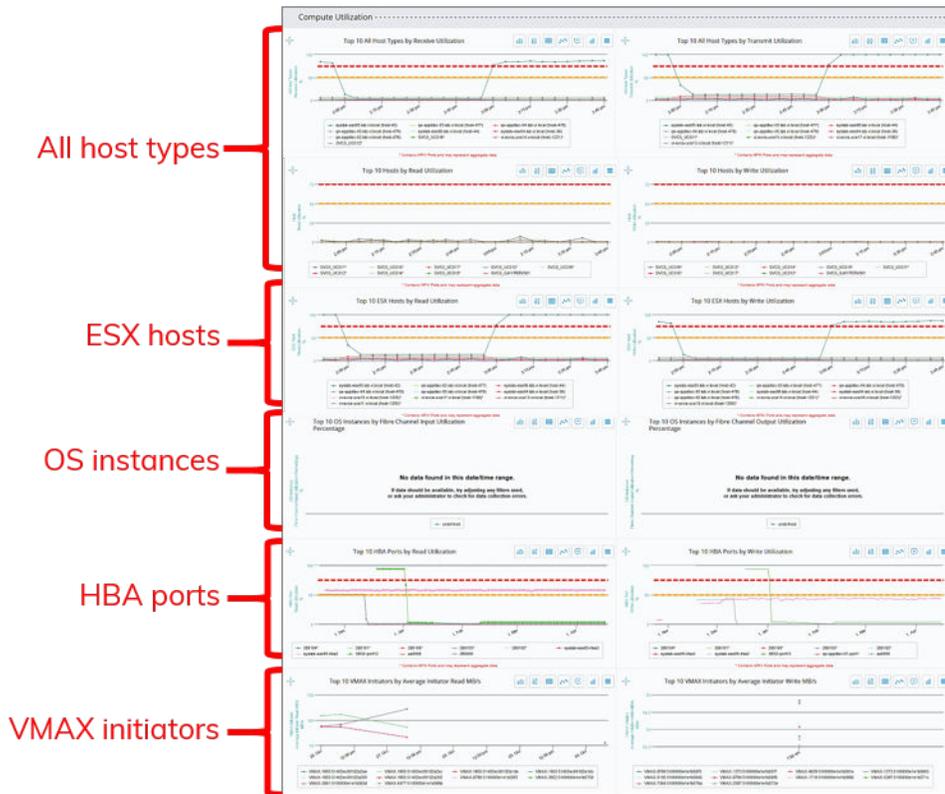


3. Compute Utilization

This section is filtered by host, ESX host, and OS instance and shows you utilization data for hosts, ESX hosts, Hyper-V hosts, PowerVM hosts, OS instances, HBA ports, and VMAX initiators.

- Receive/transmit utilization
- Read/write utilization
- FC input/output utilization
- Average initiator read/write MB/s

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information. Target these entities for more investigation.



4. Compute Performance

This section is filtered by application and storage port and shows you performance data for hosts.

- Average read/write completion time
- Average read/write response time

List card and trend charts are displayed for each entity type and metric. Drill down on entities in the list views for more information. Target these entities for more investigation.



5. Report Quick Links

Use this section to view additional reports and dashboards.

Services Reports & Quick-Links

[Services Report Launch Page](#)
[Services Report Launch Page](#)

Compute Health Check Services Report
[Fibre Channel Health Physical Layer](#)
[Fibre Channel Health Link and SCSI Events](#)
[SAN Telemetry Streaming Health Summary](#)
[Operating System Health](#)
[FCoE Health Summary](#)
[VMware vSphere Health Summary](#)

Compute Utilization Services Report
[Fibre Channel Utilization](#)
[FCoE Utilization](#)
[SAN Telemetry Streaming Utilization](#)
[VMware vSphere Utilization](#)
[Microsoft Hyper-V Utilization](#)
[PowerVM Utilization](#)
[Operating System Utilization](#)

Compute Performance Services Report
[Fibre Channel Performance](#)

Custom Reports & Quick-Links

Quicklinks to Overview Dashboards
[Executive Dashboard](#)
[Application Dashboard](#)
[Operations \(NOC\) Dashboard](#)

Quicklinks to Team Specific High Level Dashboards
[Application Administrator Dashboard](#)
[Storage Administrator SAN Dashboard](#)
[Storage Administrator SDS Dashboard](#)
[Storage Administrator NAS Dashboard](#)
[Storage Administrator VMAX Dashboard](#)
[Storage Administrator Isilon Dashboard](#)
[Compute Administrator Dashboard](#)
[Compute Operating System Dashboard](#)

[Links to Custom Dashboards or Reports](#)

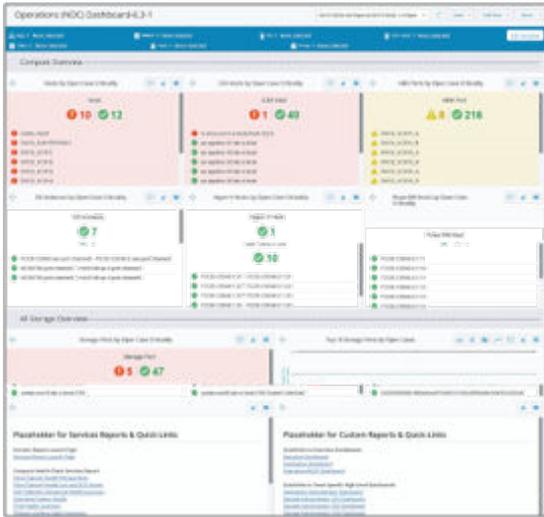
Displays links to the Services Reports that have been installed on your VirtualWisdom portal.

Displays links to the other standard dashboards and custom reports that you've created in your VirtualWisdom portal.

Operations (NOC) Dashboard

The Operations (NOC) Dashboard presents health data for applications, and compute, network, and storage infrastructure.

The dashboard Utilization is available for **FC SAN, VMware vSphere, Microsoft Hyper-V, IBM PowerVM, Host OS, VMAX, Isilon, vSAN, and VxFlex OS** integrations.



Report Variables

You can filter the report for an **application**, **ESX host**, **host**, **storage array**, **VMAX storage array**, **Isilon cluster**, or **OS Instance**. Use the filters to drill down on a specified infrastructure component.

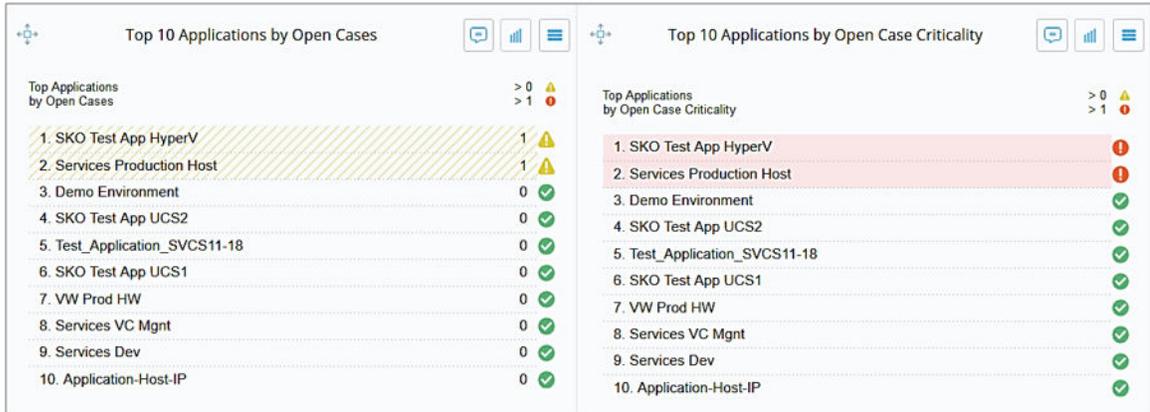


Report Sections

The report is divided into eight sections:

1. Application Overview

Use the Application Overview section to quickly observe where problems exist in the infrastructure supporting the applications. View open cases to use the investigations to troubleshoot and remediate issues. Note applications with issues to use to filter the report.



Displays the top applications by open cases.

Displays the top applications by open case criticality.

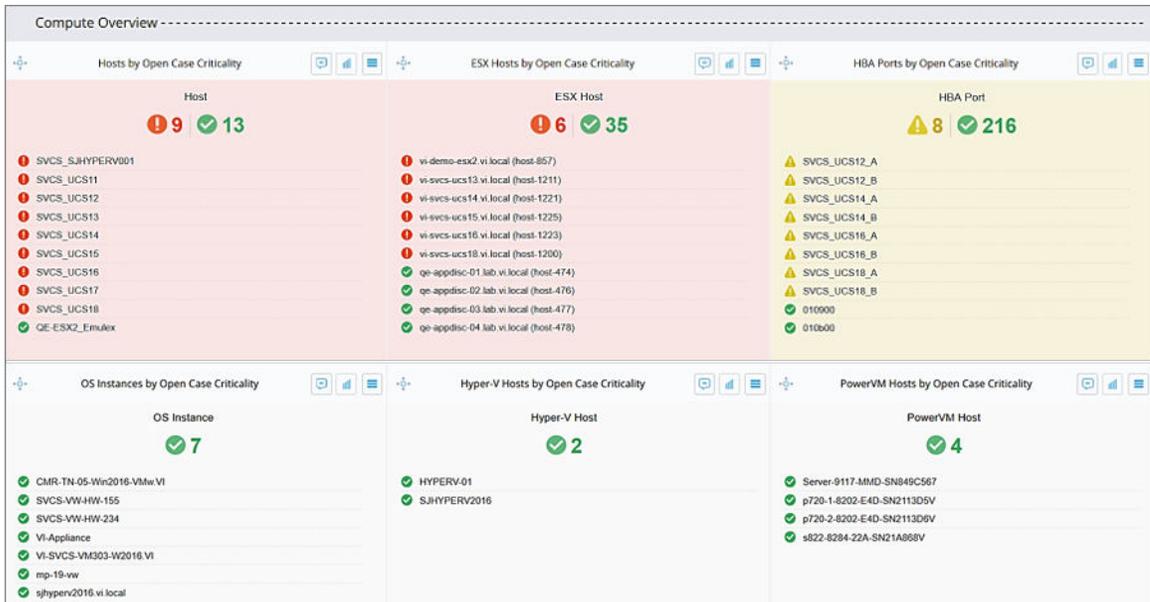
Drill down to view open cases, properties, or topology.

Drill down to view open cases, properties, or topology.

2. Compute Overview

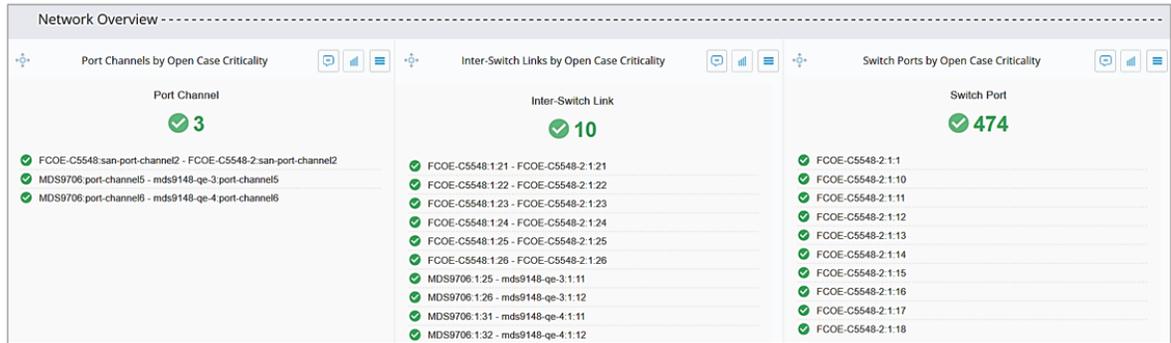
Use the Compute Overview section to quickly observe where problems exist in the Compute infrastructure. View open cases to use the investigations to troubleshoot and remediate issues.

This section is filtered by host, ESX host, and OS instance and shows you open case data for hosts, ESX hosts, Hyper-V hosts, PowerVM hosts, OS instances, and HBA ports.



3. Network Overview

Use the Network Overview section to quickly observe where problems exist in the Network infrastructure. View open cases to use the investigations to troubleshoot and remediate issues.



Displays the top port channels by open case criticality.

Drill down to view open cases, properties, or topology.

Displays the top inter-switch links by open case criticality.

Drill down to view open cases, properties, or topology.

Displays the top switch ports by open case criticality.

Drill down to view open cases, properties, or topology.

4. All Storage

Use the All Storage Overview section to quickly observe where problems exist in the Storage infrastructure. View open cases to use the investigations to troubleshoot and remediate issues.

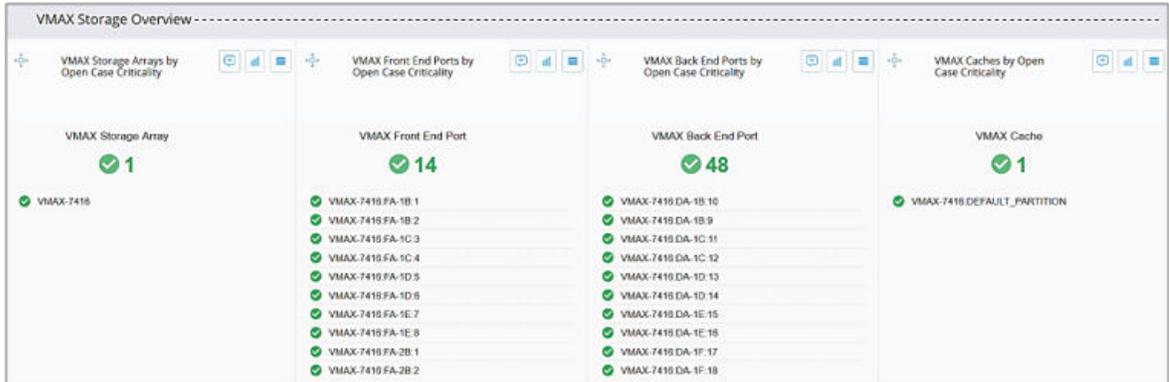
This section is filtered by storage array and shows you open case data for storage ports.



5. VMAX Storage Overview

Use the VMAX Storage Overview section to quickly observe where problems exist in the VMAX Storage infrastructure. View open cases to use the investigations to troubleshoot and remediate issues.

This section is filtered by VMAX storage array and shows you open case data for VMAX storage arrays, VMAX front end and back end ports, and VMAX caches.



6. Isilon Storage Overview

Use the Isilon Storage Overview section to quickly observe where problems exist in the Isilon Storage infrastructure. View open cases to use the investigations to troubleshoot and remediate issues.

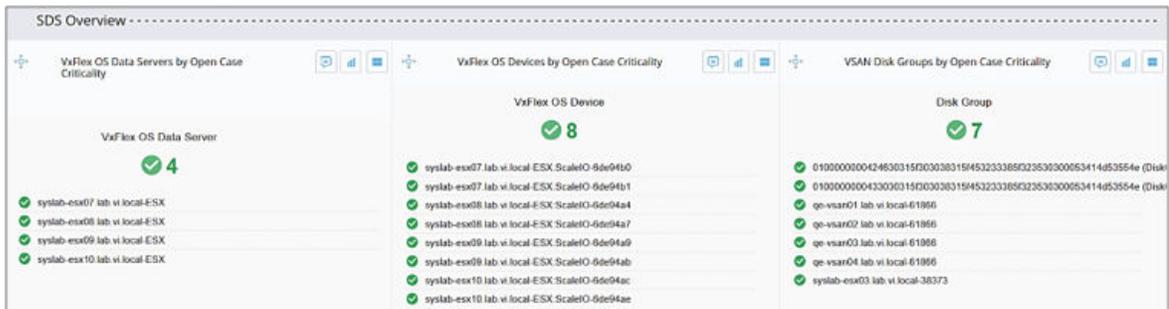
This section is filtered by Isilon Cluster and shows you open case data for Isilon Clusters, Nodes, and Node Ports.



7. SDS Storage Overview

Use the SDS Storage Overview section to quickly observe where problems exist in the vSAN or VxFlex OS storage infrastructure. View open cases to use the investigations to troubleshoot and remediate issues.

This section shows you open case data for VxFlex OS Data Servers and Devices, and vSAN Disk Groups.



8. Report Quick Links

Use this section to view additional reports and dashboards.

The screenshot shows two side-by-side panels. The left panel, titled 'Services Reports & Quick-Links', contains three main sections: 'Services Report Launch Page' with a link to 'Services Report Launch Page'; 'Compute Health Check Services Report' with links for 'Fibre Channel Health Physical Layer', 'Fibre Channel Health Link and SCSI Events', 'SAN Telemetry Streaming Health Summary', 'Operating System Health', 'FCoE Health Summary', and 'VMware vSphere Health Summary'; 'Compute Utilization Services Report' with links for 'Fibre Channel Utilization', 'FCoE Utilization', 'SAN Telemetry Streaming Utilization', 'VMware vSphere Utilization', 'Microsoft HyperV Utilization', 'PowerVM Utilization', and 'Operating System Utilization'; and 'Compute Performance Services Report' with a link for 'Fibre Channel Performance'. The right panel, titled 'Custom Reports & Quick-Links', contains three sections: 'Quicklinks to Overview Dashboards' with links for 'Executive Dashboard', 'Application Dashboard', and 'Operations (NOCI) Dashboard'; 'Quicklinks to Team Specific High Level Dashboards' with links for 'Application Administrator Dashboard', 'Storage Administrator SAN Dashboard', 'Storage Administrator SDS Dashboard', 'Storage Administrator NAS Dashboard', 'Storage Administrator VMAX Dashboard', 'Storage Administrator iSilon Dashboard', 'Compute Administrator Dashboard', and 'Compute Operating System Dashboard'; and 'Links to Custom Dashboards or Reports'.

Displays links to the Services Reports that have been installed on your VirtualWisdom portal.

Displays links to the other standard dashboards and custom reports that you've created in your VirtualWisdom portal.

Analytics

Analytics can help you to quickly identify and resolve problems with your integration.

VirtualWisdom analytics are designed to focus on our four key value areas.

1. Problem Resolution and Avoidance

Solve performance and uptime issues using a streamlined workflow powered by analytics. These analytics are designed to help you solve performance and available problems across your infrastructure.

- [Event Advisor \[259\]](#)
- [Trend Matcher \[261\]](#)

2. Workload Infrastructure Balancing

Proactively assure that your workloads and infrastructure are kept in optimal balance.

- [VM Coordinator \[268\]](#)
- [Workload Right Sizer \[278\]](#)
- [Storage Port Balancer \[274\]](#)
- [Queue Solver \[283\]](#)
- [Migration Analysis \[289\]](#)
- [Workload Analysis \[292\]](#)

3. Predictive Capacity Management

Forecast capacity needs across all your infrastructure services using the same solution that monitors your workloads.

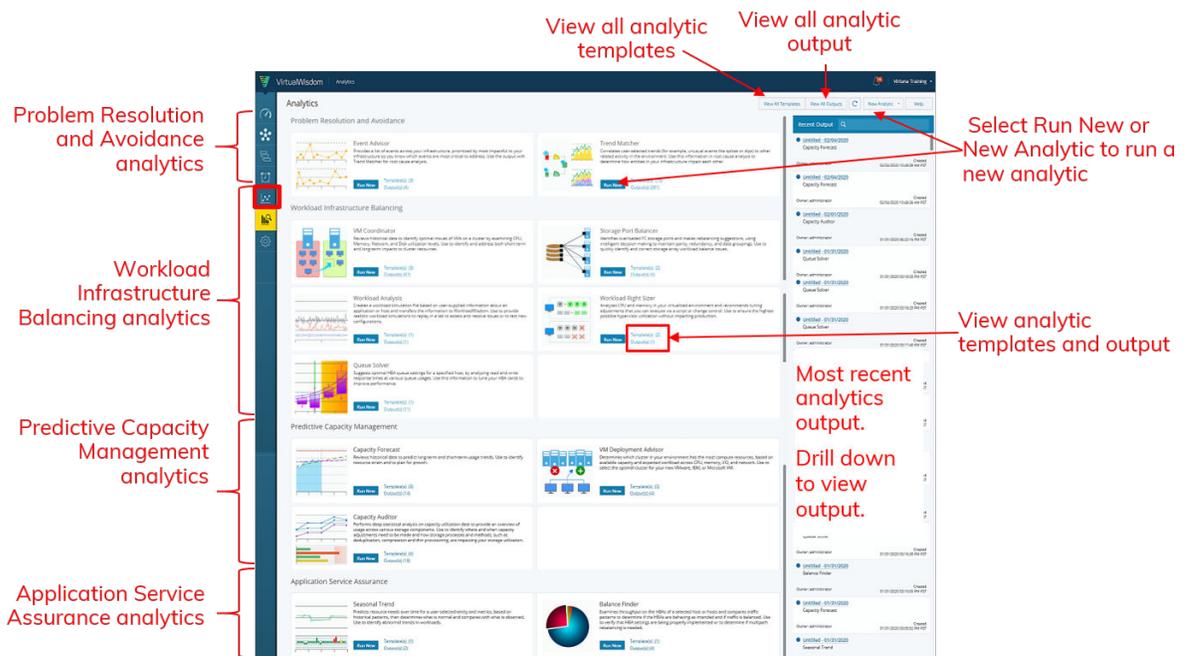
- [Capacity Forecast \[297\]](#)

- [VM Deployment Advisor \[301\]](#)
 - [Capacity Auditor \[305\]](#)
4. **Application Service Assurance**
- Ensure that your end-to-end infrastructure delivers your application services at the required service level. These analytics use predefined metrics and thresholds to perform detection of meaningful data patterns in a specified time period, identifying and isolate existing or potential problems in the infrastructure.
- [Balance Finder \[313\]](#)
 - [Seasonal Trend \[310\]](#)

Analytics Home Page

The VirtualWisdom Analytics are located under the Analytics icon on the left navigation bar. The Analytics are organized into four categories on the home page:

- Problem Resolution and Avoidance
- Workload Infrastructure Balancing
- Predictive Capacity Management
- Application Service Assurance



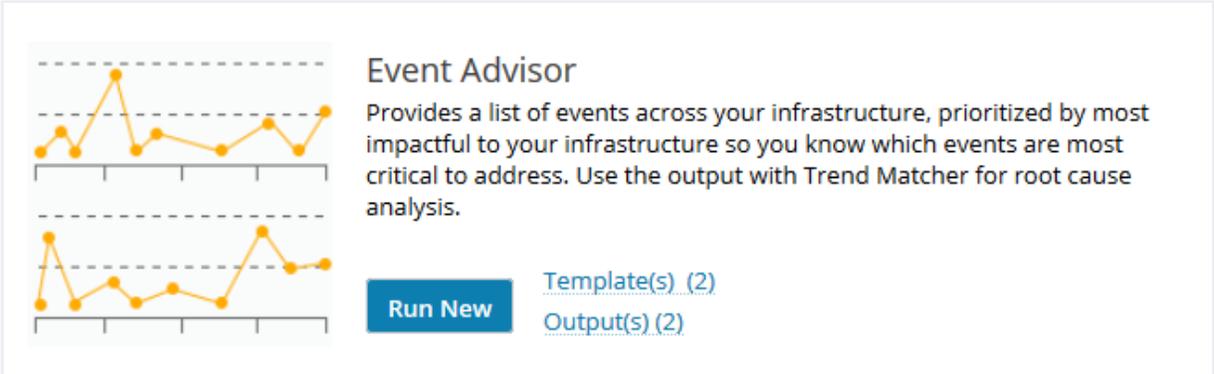
Each analytic provides a link to defined templates and saved outputs, if they are available.

A pane on the right side of the home page displays all recently saved outputs.

There are two buttons at the top of the page that allow you to view all saved analytics templates and all saved output.

Problem Resolution and Avoidance Analytics

The Problem Resolution and Avoidance analytics are designed to help you identify and resolve issues that occur in your infrastructure.

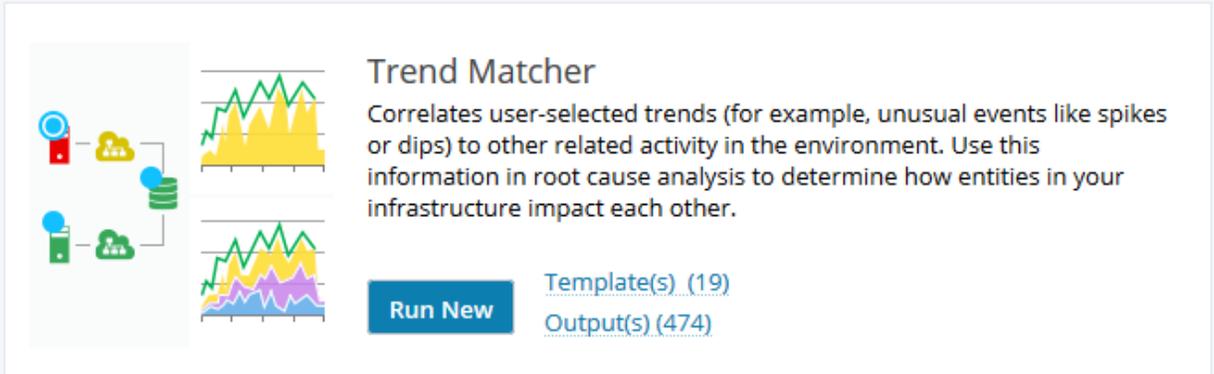


Event Advisor

Provides a list of events across your infrastructure, prioritized by most impactful to your infrastructure so you know which events are most critical to address. Use the output with Trend Matcher for root cause analysis.

[Run New](#) [Template\(s\) \(2\)](#)
[Output\(s\) \(2\)](#)

The **Event Advisor** analytic provides a prioritized list of potentially interesting events by finding spikes, relevant performance issues, and anomalous behavior, that require attention. The resulting list is ranked by magnitude and duration. Issue events can be transferred to Trend Matcher for further analysis. Event Advisor can be used as a starting point to find the root cause of an issue. The user can cast a wide net, looking for spikes or events that may shed light on where to start looking.



Trend Matcher

Correlates user-selected trends (for example, unusual events like spikes or dips) to other related activity in the environment. Use this information in root cause analysis to determine how entities in your infrastructure impact each other.

[Run New](#) [Template\(s\) \(19\)](#)
[Output\(s\) \(474\)](#)

Trend Matcher assists you in root cause analysis by providing a way to identify the impact that entities have on each other in your infrastructure. For example, application latency in

your data center can be due to any number of seemingly unrelated events. These can occur in silos (or domains) like servers, HBAs and NICs in servers, the SAN fabric, ports on your SAN attached or NAS storage array, etc. Using Trend Matcher, you can quickly and easily analyze your infrastructure to identify events or issues that correlate with a problem in your infrastructure.

Trend Matcher can be used in conjunction with Event Advisor or standalone.

Event Advisor

The **Event Advisor** analytic provides a prioritized list of potentially interesting events by finding spikes, relevant performance issues, and anomalous behavior, that require attention. The resulting list is ranked by magnitude and duration. Issue events can be transferred to [Trend Matcher \[261\]](#) for further analysis. Event Advisor can be used as a starting point to find the root cause of an issue. The user can cast a wide net, looking for spikes or events that may shed light on where to start looking.

In this example, we'll look for events on storage ports where buffer to buffer credits are high then use Trend Matcher to find correlating entities and events to help you troubleshoot the issue.

Running Event Advisor

1. Start by selecting the entity type (**Storage Port**) and metric (**% Time at Zero Transmit Credits**).

The screenshot displays the 'Event Advisor' interface. On the left, there is a table with columns for 'Type', 'Metrics', and 'Filter', and an 'Add' button. On the right, the 'Add Metrics' dialog box is open. It shows a list of 'Selected Entity Type' including Storage Port, Storage I/O Module, and various ViFlex OS components. The 'Storage Port' is highlighted. In the 'Available Metrics' section, '% Time at Zero Transmit Credits' is selected and highlighted. The 'Selected Item' section shows 'Storage Port / % Time at Zero Transmit Credits'. At the bottom, there are 'OK' and 'Cancel' buttons.

2. Select a date range to search for buffer credit events.

Last 2 Hours
 Last 6 Hours
 Last 24 Hours
Last 7 Days
 Last 30 Days
 Last 3 Months
 Last 6 Months
 Last 9 Months
 Last 12 Months
 Custom

Date Range -

Time Range -

< April 2020 >

S	M	T	W	T	F	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30		

< May 2020 >

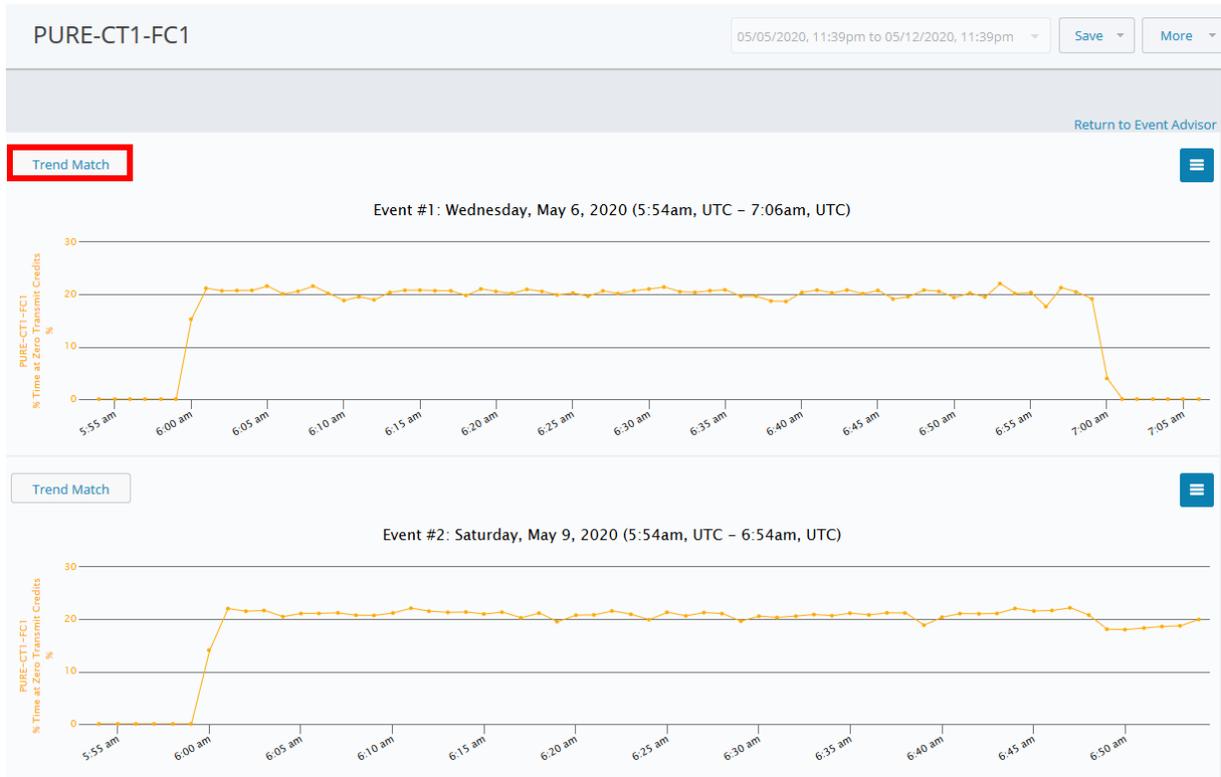
S	M	T	W	T	F	S
						1 2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
						31

Understanding Event Advisor Results

Event Advisor shows you a list of entities with interesting events, sorted by their severity, with 1 being the highest severity. Drill down on an entity to view the top ten events associated with the entity and metric.

Name	Number of Events	Severity	Total Length (min)	Severity of Worst Ev...	Max Length (min)	Tags
PURE-CT1-FC1	96	1	771	0.87	51	
PURE-CT0-FC1	97	0.99	772	0.86	51	
PURE-CT0-FC0	82	0.71	365	0.98	45	
PURE-CT1-FC0	74	0.68	335	1	44	

Select Trend Match to analyze the event in more detail using the Trend Matcher analytic.



Trend Matcher

Trend Matcher accepts a source trend identified by the user or by Event Advisor and a target metric. Trend Matcher uses intelligent trend matching to compare a base trend with other entities and metrics, for the timeframe specified, and provides a topology view showing the connected entities and trends that correlate with the base trend.

This helps you identify the source of a recognized issue. For example, if there is time spent at zero Buffer-to-Buffer Credits, the Trend Matcher could find an HBA that gets busy when the buffer credit problem starts and finishes its work when the buffer credit issue goes away.

Running Trend Matcher

1. Trend Matcher accepts a base trend identified by the user or by the Event Advisor and a target metric. Trend Matcher uses intelligent trend matching to compare a base trend with other entities and metrics, for the time frame specified, and provides a topology view showing the connected entities and trends that correlate with the base trend.

Match Trend for:

Storage Port PURE-CT1-FC1 by % Time at Zero Transmit Credits

Update Base Trend

Apply Filters (optional) 0 Filter(s) +

- The analysis mode offers a choice between Quick Mode and Robust Mode. Quick Mode lets you perform a quick search to find large correlations while Robust Mode provides a more comprehensive search but is slower.

Analysis Mode:

Quick Mode
A faster search looking for large scale features.

Robust Mode
A more comprehensive search and analysis that can be much slower.

Selected Base Trend:

The graph shows the percentage of time at zero transmit credits for PURE-CT1-FC1. The y-axis ranges from 0 to 30%. The x-axis shows time from 5:55 am to 7:05 am. The data shows a sharp increase starting at 6:00 am, peaking at approximately 22% around 6:05 am, and then gradually declining back to 0% by 7:00 am.

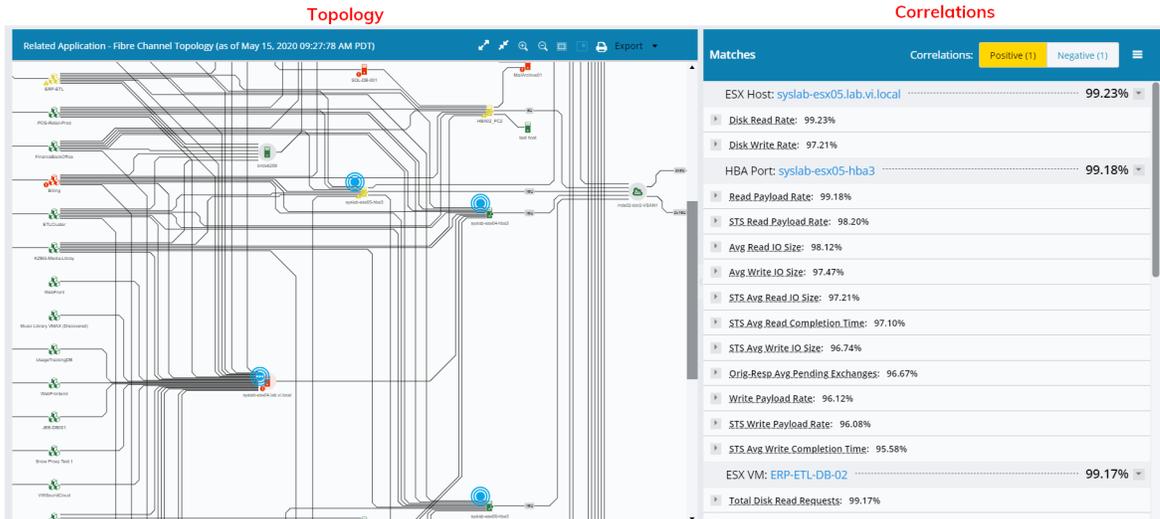
- There are multiple options for conducting the trend matching exercise. These options allow you to search for matches using a suggested or targeted situation, to limit the search to only connected entities, or search for matches across all entities.

Common Situations to Analyze: PURE-CT1-FC1

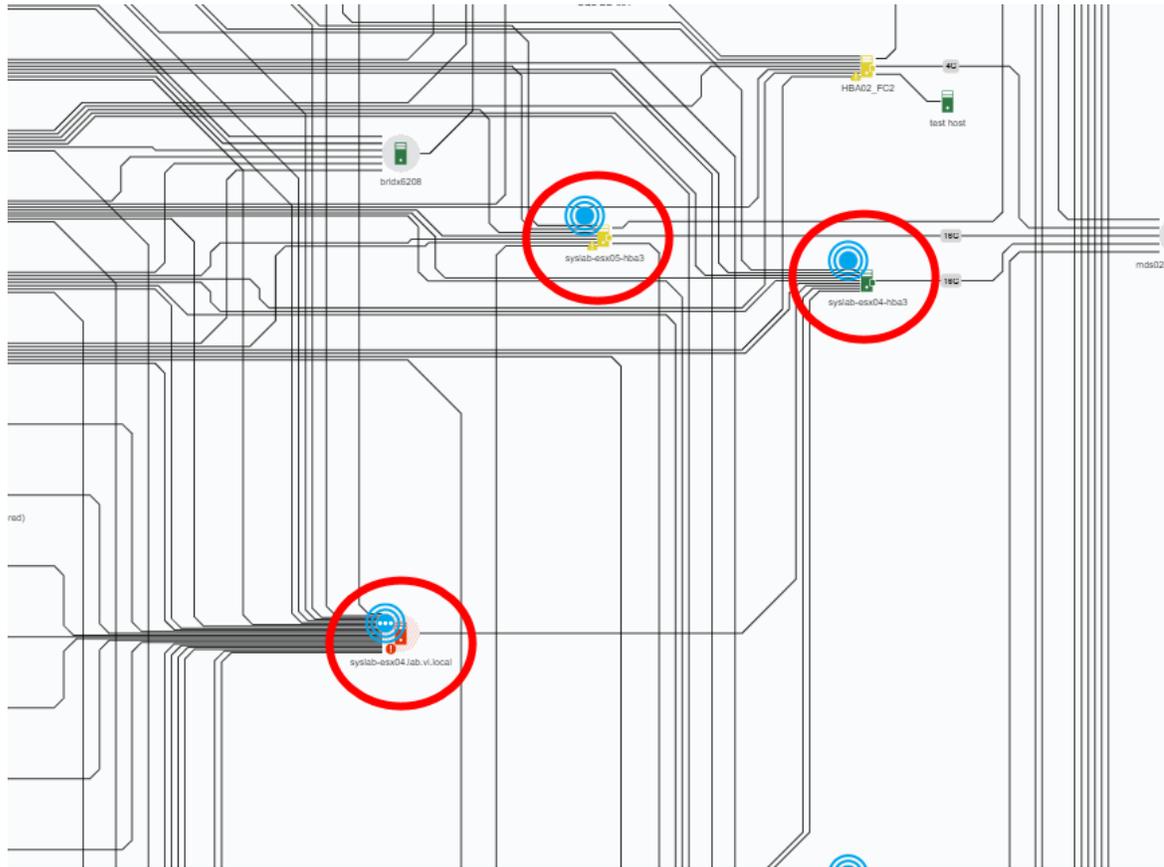
- Search All Connected Entities with Conversations
Starting with the base entity, this option crawls the entire topology (including all sub topologies relevant to the entity) outward and examines every entity encountered for significant correlation. Connected conversation entities will be examined. This can significantly increase processing time.
- Buffer Credit Starvation
This situation looks for devices that are the cause of buffer credit starvation.
- Search All Connected Entities without Conversations (Faster)
Starting with the base entity, this option crawls the entire topology (including all sub topologies relevant to the entity) outward and examines every entity encountered for significant correlation. No conversation entities will be examined.
- Targeted Situation
Analyze matches for selected entities
- Legacy Search
Not limited to any particular situation

Understanding Trend Matcher Results

1. The Trend Matcher results are divided into two sections: A Topology tab showing you where entities and metrics correlated with the base trend, and a list of correlations found that match the base trend.



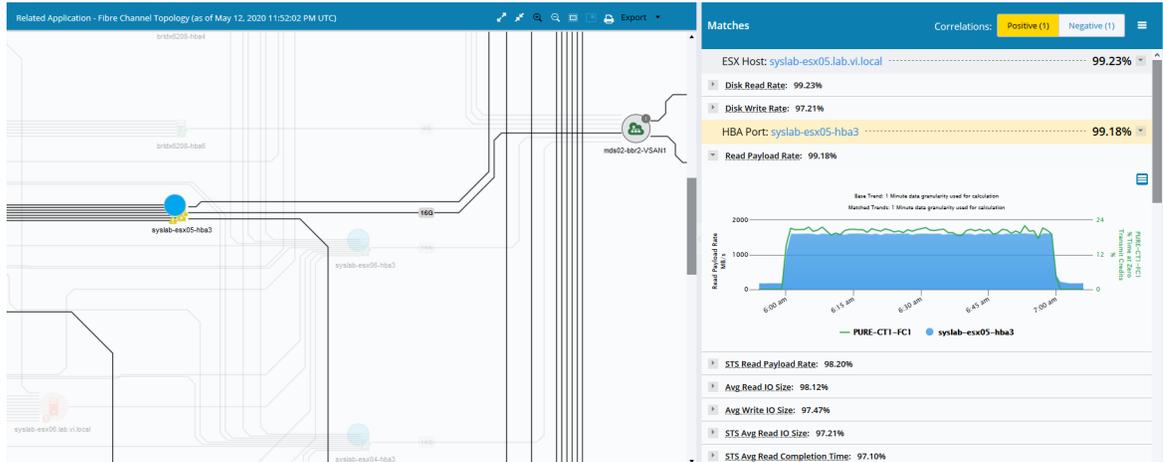
2. The solid blue circles indicate where there were matches. The circles around the solid circle indicate the degree of correlation: the more blue circles displayed around an entity, the higher the correlation. Entities displaying three dots at their center indicate a correlation at the sub-entity level – drill down to view more details.



3. To the right of the topology map is a list of all matching entities and events, ordered by percent correlation. The list displays both positive and negative correlations.

Matches		Correlations: Positive (1) Negative (1) ☰	
ESX Host: syslab-esx05.lab.vi.local	-----		99.23% ▾
▶ Disk Read Rate : 99.23%			
▶ Disk Write Rate : 97.21%			
HBA Port: syslab-esx05-hba3	-----		99.18% ▾
▶ Read Payload Rate : 99.18%			
▶ STS Read Payload Rate : 98.20%			
▶ Avg Read IO Size : 98.12%			
▶ Avg Write IO Size : 97.47%			
▶ STS Avg Read IO Size : 97.21%			
▶ STS Avg Read Completion Time : 97.10%			
▶ STS Avg Write IO Size : 96.74%			
▶ Orig-Resp Avg Pending Exchanges : 96.67%			
▶ Write Payload Rate : 96.12%			
▶ STS Write Payload Rate : 96.08%			
▶ STS Avg Write Completion Time : 95.58%			
ESX VM: ERP-ETL-DB-02	-----		99.17% ▾
▶ Total Disk Read Requests : 99.17%			
▶ Disk Read Rate : 99.17%			

- Expand the correlated entity and metric to see how it compares with the base trend. The correlated entity is highlighted on the topology map to the left. This gives you valuable information for analyzing the issue. For example, a flow control issue may be related to increased workload on an HBA port. Once the workload decreases, the buffer credit issue on the storage port also resolves.



You can also run Trend Matcher as a standalone analytic. You'll need to select a base trend to use for the analysis.

Workload Infrastructure Balancing Analytics

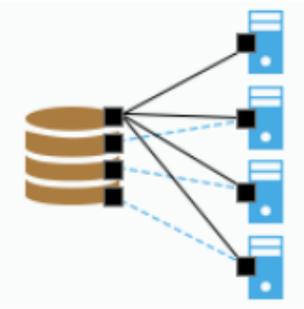
VirtualWisdom Workload Infrastructure Balancing analytics help you make decisions to balance your workload across the end-to-end infrastructure.

The screenshot shows the 'VM Coordinator' interface. On the left, a diagram illustrates two server racks (one red, one green) with VMs represented by blue icons. An arrow points from a VM in the red rack to a VM in the green rack, indicating a migration. The text on the right describes the tool's function:

VM Coordinator
Reviews historical data to identify optimal moves of VMs on a cluster by examining CPU, Memory, Network, and Disk utilization levels. Use to identify and address both short-term and long-term impacts to cluster resources.

Below the text is a blue 'Run New' button and two links: 'Template(s) (6)' and 'Output(s) (25)'.

VM Coordinator improves VM cluster performance by making recommendations to move VM resources in order to better balance the hosts. VM Coordinator provides optimal placement for VMs across a cluster using historical trends. VM Coordinator generates a vMotion script that can be downloaded and run to move the identified VMs.

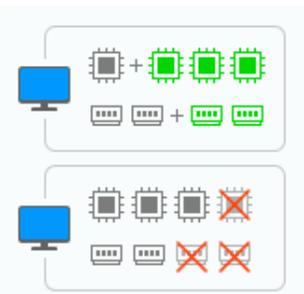


Storage Port Balancer

Identifies overloaded FC storage ports and makes rebalancing suggestions, using intelligent decision-making to maintain parity, redundancy, and data groupings. Use to quickly identify and correct storage array workload balance issues.

[Run New](#) [Template\(s\) \(2\)](#)
[Output\(s\) \(2\)](#)

Storage Port Balancer identifies overloaded storage ports and makes HBA/Host move suggestions that would rebalance the front-end ports of a Fibre Channel storage array.

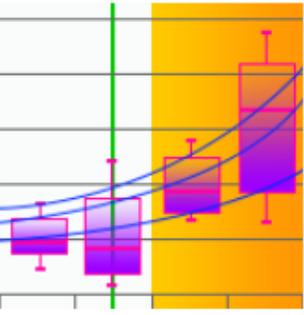


Workload Right Sizer

Analyzes CPU and memory in your virtualized environment and recommends tuning adjustments that you can execute via a script or change control. Use to ensure the highest possible hypervisor utilization without impacting production.

[Run New](#) [Template\(s\) \(3\)](#)
[Output\(s\) \(8\)](#)

Workload Right Sizer analyzes virtual CPU and memory and recommends tuning adjustments that can be executed using a script or via change control



Queue Solver

Suggests optimal HBA queue settings for a specified host, by analyzing read and write response times at various queue usages. Use this information to tune your HBA cards to improve performance.

[Run New](#) [Template\(s\) \(1\)](#)
[Output\(s\) \(14\)](#)

Queue Solver provides guidance in optimizing the settings for HBA queue depth. Queue Solver provides a visualization of the response time of that server as evidence for the recommendation.



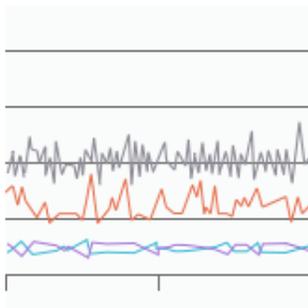
Migration Analysis

Prepares a detailed export for Virtana Migrate to enable application dependency mapping, suggest ideal resource groupings, and provide insights on right sizing and cost optimization across a multitude of cloud providers.

Run New

Output(s) (15)

Migration Analysis analyzes and exports data to enable application dependency, suggest ideal resource move groups, and provides insights for right sizing and optimizing cloud resources and costs, as part of a digital transformation migration project.



Workload Analysis

Creates a workload simulation file based on user-supplied information about an application or host and transfers the information to WorkloadWisdom. Use to provide realistic workload simulations to replay in a lab to assess and resolve issues or to test new configurations.

Run New

Output(s) (8)

Workload Analysis creates a workload simulation file that provides realistic workload simulations that can be replayed using WorkloadWisdom to assess and resolve issues or test new storage configurations.

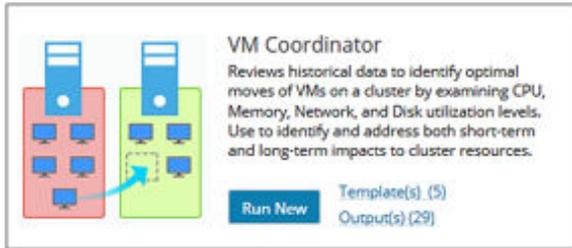
VM Coordinator

Use VM Coordinator to review ESX cluster balance and identify optimal moves for existing VMS to improve utilization and performance.

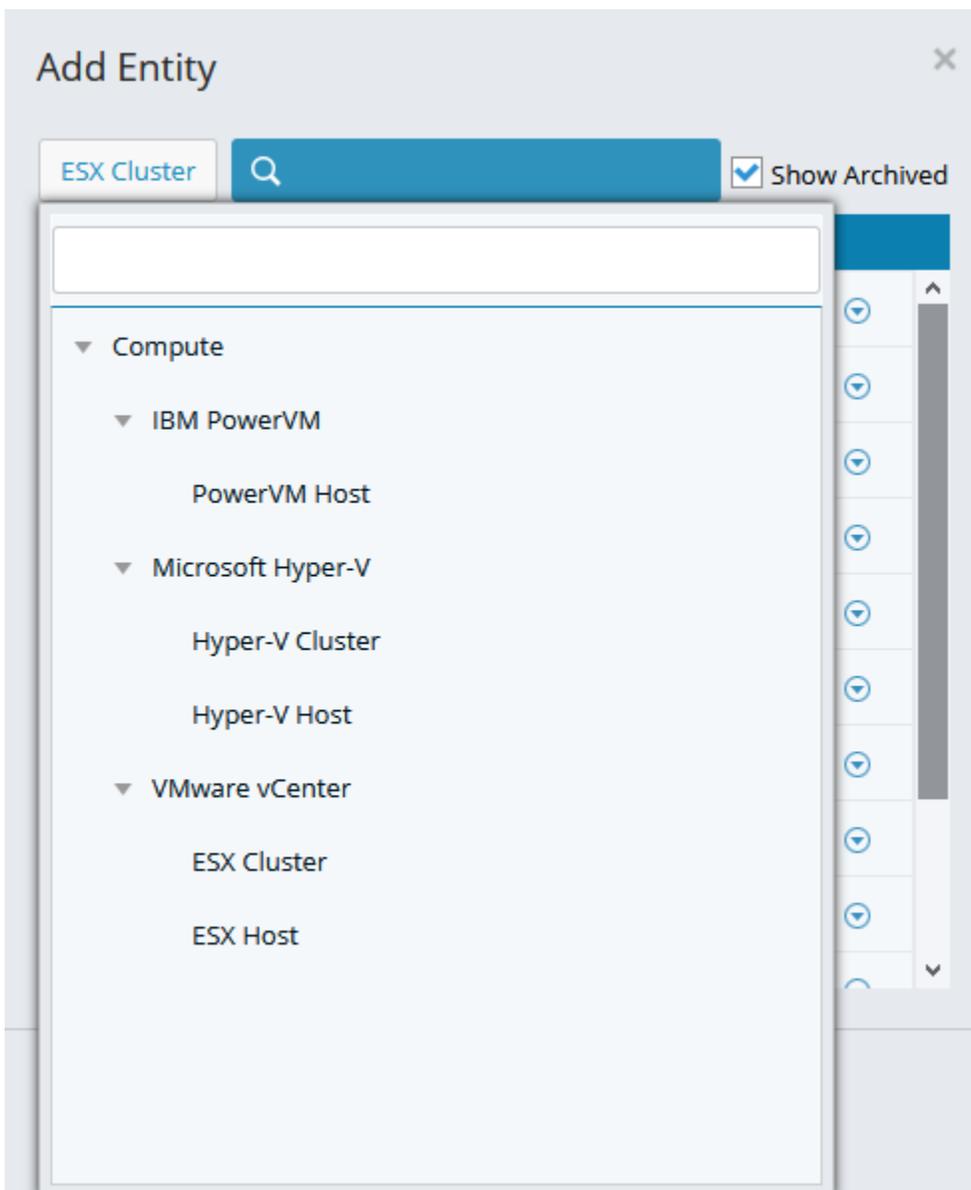
VM Coordinator reviews historical data to identify optimal moves of VMs on a cluster by examining CPU, memory, network, and disk utilization levels. Use VM Coordinator to identify and address both short-term and long term impacts to cluster resources

Running VM Coordinator

1. Start by running a new VM Coordinator from the Analytics page.



2. Select the entity type to analyze by clicking on the **Add** box and selecting from the following entity types: **ESX Cluster**, **ESX Host**, **Hyper-V Cluster**, **Hyper-V Host**, or **PowerVM Host**.



3. Select the named entity to analyze.

The screenshot shows the 'Add Entity' dialog box. At the top, there is a search bar containing 'ESX Cluster' and a magnifying glass icon. To the right of the search bar is a checkbox labeled 'Show Archived' which is checked. Below the search bar is a table with two columns: 'Name ↑' and 'Tags'. The table contains the following rows:

Name ↑	Tags
Cluster-QE-AppDiscovery	⌵
Cluster-QE-VSAN	⌵
ProbeNAS Cluster	⌵
ProdGeneral	⌵
ProdPlatinum	⌵
* Prod_Demo	⌵
QE-Jenkins-Cluster	⌵
ST_Cluster	⌵
* SVCS_Prod1	⌵

Below the table, it says '(Selected Items: 1)'. At the bottom of the dialog, there are two buttons: 'OK' (highlighted in yellow) and 'Cancel'.

4. **Set the Number of Swaps**
Set the parameter for the number of swaps. This limits the number of VMs that can be moved. The default number of swaps is five. The more moves you select, the longer it will take to run VM Coordinator.

Hosts

Name	
ProdPlatinum	

Number of VM Swaps: 0 5 10 15 20 25 30 35 Unlimited

Use Advanced Options

Run

5. Selecting Advanced Options

You can run the analytic using its default settings, or choose from its advanced settings.

- Check the **Use Advanced Options** box to display the advanced settings.
- The advanced options allow you to fine-tune the analytic to limit the network traffic and disk activity, the number of iterations allowed, and whether to optimize for CPU or memory or ignore DRS Affinity Rules (for ESX only).

VM Coordinator

07/21/2020, 09:46am to 08/25/2020, 09:46am

Hosts

Name	
ProdPlatinum	

Number of VM Swaps: 0 5 10 15 20 25 30 35 Unlimited

Use Advanced Options

Run

Advanced

Network Maximum * 9999 MB/s (per server)

Disk Maximum * 9999 MB/s (per server)

Run Time * 7920 Iterations

Optimization CPU 70% Memory 30%

Ignore DRS Affinity Rules (ESX Only)

6. Setting the Date Range

Run intervals may be based on cluster or ESX servers. The longer the date range and iterations, the more accurate the resource balancing will be.

The screenshot shows a configuration window for the VM Coordinator. On the left, there is a list of time range options: Last 2 Hours, Last 6 Hours, Last 24 Hours, Last 7 Days, Last 30 Days, Last 35 Days, Last 3 Months, Last 6 Months (circled in red), Last 9 Months, and Last 12 Months. Below these is a 'Custom' option. In the center, there are two calendar views for March 2020 and April 2020. The date range is set to 'Mar 16, 2020' to 'Apr 15, 2020'. The time range is set to '04:37 PM' to '04:37 PM'. At the bottom, there are 'Apply' and 'Cancel' buttons.

The first run of VM Coordinator should be a minimum of **30 days**, to catch monthly jobs on initial cluster rebalancing. VM Coordinator should be run weekly thereafter, to maintain cluster balancing.

7. Saving the Analytic as a Template

Save the analytic as a template by selecting the Save button at the top right and schedule it as a monthly job.

The screenshot shows the 'VM Coordinator' configuration window. At the top right, there is a 'Save' button. Below the title bar, there is a section for 'VM Coordinator Attributes'. The 'Save As' field is set to 'ProdPlatinum ESX Cluster VM Coordinator'. There is a 'Description' field. Below this is a 'Tags' field. Under the 'Schedule' section, the 'Run' frequency is set to 'Every Saturday', the 'Time of Day' is 'User: 12AM PDT / Appliance: 12AM PDT', and the 'Interval' is 'Last 7 Days'. There is a 'Pause Schedule' checkbox which is currently unchecked. At the bottom, there are 'Save' and 'Cancel' buttons.

Understanding VM Coordinator Results

The top section shows you how many VMs are recommended for moves, along with their names, and their suggested moves.

Recommendation for analyzed cluster ProdPlatinum

5

Based on 1861 iterations (28% completion) and the analysis of 3 ACTIVE servers and 88 ACTIVE VMs we recommend to reconfigure 5 VMs

VMs to be Moved	From ↓	To
OrderingReporting	syslab-esx06.lab.vi.local	syslab-esx05.lab.vi.local
ERP-ETL-App-01	syslab-esx05.lab.vi.local	syslab-esx06.lab.vi.local
Order Manager App 03	syslab-esx04.lab.vi.local	syslab-esx05.lab.vi.local
supply-warehouse-app-01	syslab-esx04.lab.vi.local	syslab-esx06.lab.vi.local
Order Manager App 01	syslab-esx04.lab.vi.local	syslab-esx06.lab.vi.local

Below is shown the projected impact on CPU and memory congestion, and disk and network usage.

Projected Group Changes for analyzed cluster ProdPlatinum

Group CPU Congestion: -6.3% Group Memory Congestion: -9.5% Group Disk Usage: 0.1% Group Network Usage: 1.8%

A list of the current measurement and projected impact for each host is shown below. Expand the details for a host by clicking the down arrow.

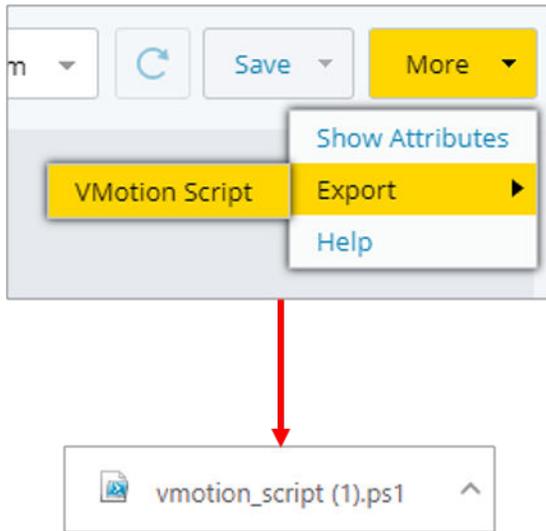


Download vMotion Script

Recommendations from VM Coordinator will require you to move several virtual machines using a vMotion script.

If you decide to use the recommendation, you must perform all the recommended moves. Performance may be negatively impacted if only some of the moves are made.

To download the script produced by VM Coordinator, select More, then Export.



Storage Port Balancer

Storage Port Balancer identifies overloaded storage ports and makes HBA/Host move suggestions that would rebalance the front-end ports of a fibre channel storage array.

Running Storage Port Balancer

1. Start by running a new Storage Portal Balancer from the Analytics home page by clicking **New Analytic** or **Run New** button.

A card titled 'Storage Port Balancer'. On the left is a diagram showing three storage arrays (represented by stacked disks) connected to four host icons. To the right of the diagram is the text: 'Storage Port Balancer Identifies overloaded FC storage ports and makes rebalancing suggestions, using intelligent decision-making to maintain parity, redundancy, and data groupings. Use to quickly identify and correct storage array workload balance issues.' Below this text is a blue 'Run New' button and two links: 'Template(s) (2)' and 'Output(s) (2)'.

2. Click the **Add** button to select a storage array.

 A screenshot of the 'Storage Port Balancer' configuration page. At the top, there is a title 'Storage Port Balancer' and a date range '09/09/2020, 01:58pm to 10/14/2020, 01:58pm' with 'Save' and 'More' buttons. Below this is a section titled 'Storage Array' with a search box labeled 'Name' and an 'Add' button.

3. Select a storage array from the list. You can use the search box to find a specific storage array. Click the **OK** button.

Add Entity ✕

Storage Array

 Show Archived

Name ↑	Tags	
HDS_2301	Entity Matching	⌵
NAS_SL		⌵
PURE		⌵
VMAX1955	Entity Matching	⌵
VNX9759	Entity Matching	⌵
syslab-c01		⌵

(Items: 6)

OK
Cancel

4. Use the slider to specify the **Number of HBA Swaps**. The number allows you to limit the number of HBAs that can be required to be moved as part of the recommendations. The default is eight.

Number of HBA Swaps

2

4

8

16

32

64

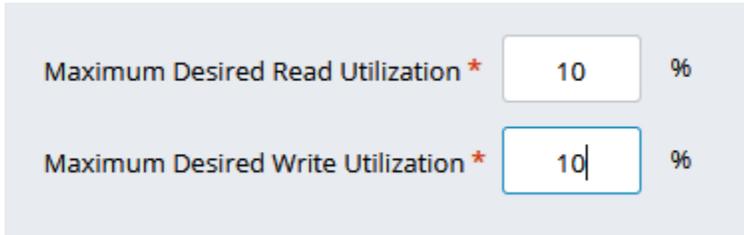
128

256

Unlimited

8

- Select the maximum desired read and write utilization post-move. The default value is 50%. Note that the desired utilization has been set low in this example to show results.

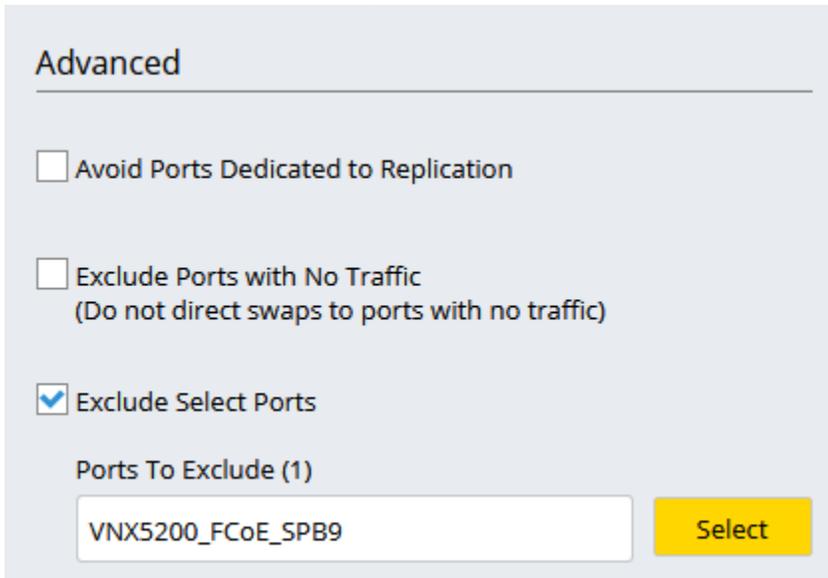


A screenshot of a configuration panel with a light gray background. It contains two rows of settings. The first row is labeled "Maximum Desired Read Utilization *" and has a text input field containing the number "10" followed by a percentage symbol "%". The second row is labeled "Maximum Desired Write Utilization *" and has a text input field containing the number "10" followed by a percentage symbol "%".

- If you click the *Use Advanced Options* check box (off by default), you can select additional check boxes (also off by default) to:

- Avoid Ports Dedicated to Replication
- Exclude Ports with No Traffic
- Exclude Select Ports

If you check **Exclude Select Ports**, the **Exclude Storage Ports** screen is displayed, and you can select one or more ports to exclude.

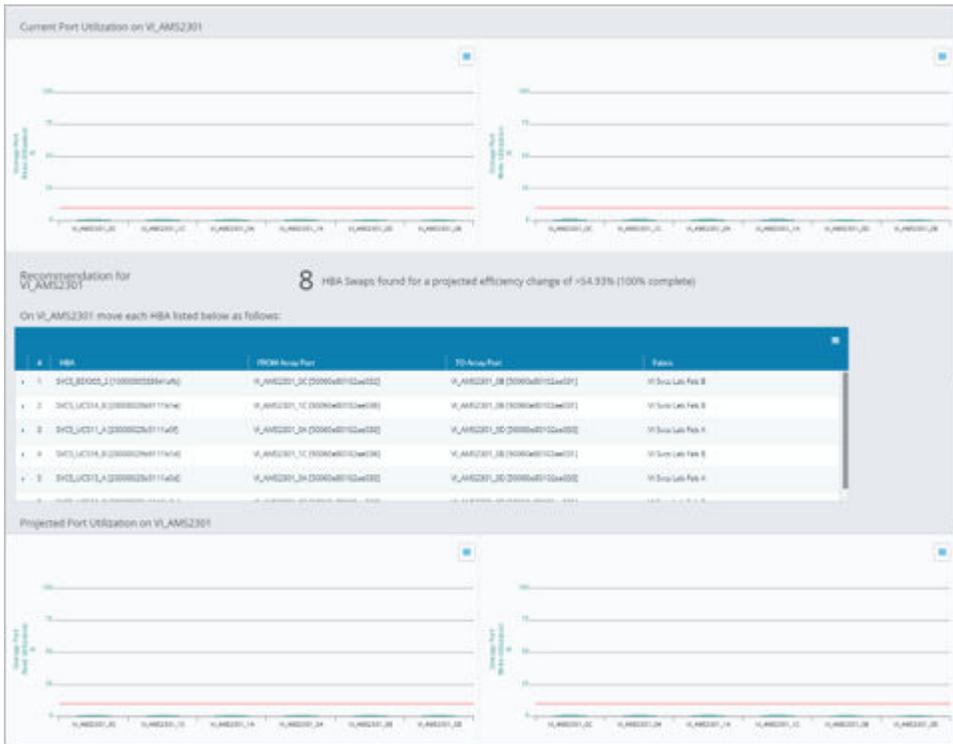


A screenshot of an "Advanced" configuration screen. At the top, the word "Advanced" is displayed in a bold, dark font. Below it, there are three check boxes: "Avoid Ports Dedicated to Replication" (unchecked), "Exclude Ports with No Traffic (Do not direct swaps to ports with no traffic)" (unchecked), and "Exclude Select Ports" (checked). Below the check boxes, the text "Ports To Exclude (1)" is shown. Underneath, there is a text input field containing the identifier "VNX5200_FCoE_SPB9" and a yellow "Select" button to its right.

- Click the **Run** button.

Understanding Storage Port Balancer Results

The recommendation for the selected array(s) is displayed, with a number of HBA swaps and a projected efficiency percentage change. If a more optimal configuration was not found based on the selected parameters, a dialog box is displayed suggesting parameter changes.



Click on the down arrow next to the HBA name to view more information about the swap.

HBA (selected)	FROM Array Port	TO Array Port	Fabric
Name: SVCS_ESK005_2 Host: v1-esk005-01.local WWN: 10000005336441afb Host Avg Block Size: 13.8 KB	Name: V1_AMS2301_DC WWN: 50060e80102ae032 Avg Block Size: 27.8 KB Port Speed: 8G	Name: V1_AMS2301_DB WWN: 50060e80102ae031 Avg Block Size: 50.4B Port Speed: 8G	V1 Sws Lvs Fab B

Impact: v1-esk005-01.local will be responsible for 14.28% of the traffic

You can save the move recommendation by clicking on the hamburger icon in the recommendations header, then exporting as a csv file or copying to the clipboard.

Press CTRL-C to copy to your clipboard

```
# HBA FROM Array Port TO Array Port Fabric
1 SVCS_ESK005_2 (10000005336441afb) V1_AMS2301_DC (50060e80102ae032) V1_AMS2301_DB (50060e80102ae031) V1 Sws Lvs Fab B
2 SVCS_UCS14_B (20000025b5111b1e1) V1_AMS2301_1C (50060e80102ae036) V1_AMS2301_JB (50060e80102ae037) V1 Sws Lvs Fab B
3 SVCS_UCS11_A (20000025b5111a0f) V1_AMS2301_0A (50060e80102ae030) V1_AMS2301_DC (50060e80102ae032) V1 Sws Lvs Fab A
4 SVCS_UCS16_B (20000025b5111b1d7) V1_AMS2301_1C (50060e80102ae036) V1_AMS2301_JB (50060e80102ae037) V1 Sws Lvs Fab B
5 SVCS_UCS13_A (20000025b5111a0e1) V1_AMS2301_0A (50060e80102ae030) V1_AMS2301_1C (50060e80102ae036) V1 Sws Lvs Fab A
6 SVCS_ESK004_1 (10000005334893c91) V1_AMS2301_1A (50060e80102ae034) V1_AMS2301_1C (50060e80102ae036) V1 Sws Lvs Fab A
7 SVCS_ESK004_2 (10000005334893c91) V1_AMS2301_1C (50060e80102ae036) V1_AMS2301_1A (50060e80102ae034) V1 Sws Lvs Fab A
8 SVCS_UCS13_B (20000025b5111b0e1) V1_AMS2301_DC (50060e80102ae032) V1_AMS2301_JB (50060e80102ae031) V1 Sws Lvs Fab B
```

OK

Workload Right Sizer

Workload Right Sizer analyzes CPU and memory in your virtualized environment and recommends tuning adjustments that you can execute via a script or change control.

Running Workload Right Sizer

1. Start by selecting the entity type to analyze: Application, ESX Cluster, ESX Host, or ESX VM, then select a named entity from the drop down list.

The image displays two screenshots of the Workload Right Sizer web interface. The left screenshot shows the 'Analyze' dropdown menu set to 'Application', with a list of applications below it. The 'Ordering System' application is highlighted in yellow. The right screenshot shows the 'Analyze' dropdown menu set to 'Ordering System', and the 'Use Advanced Options' checkbox is checked.

Name	Tags
Music Library VMAX (Discovered)	VMAX
Order Manager	RetailJSCSI
OrderProcessing	Retail
OrderProcessing_Conflict	SNOW
Ordering System	
POS-Retail-Prod	Retail
ReportSQL	Finance
SNOW-dockerba1	
SQL Backup	
SQL Cluster 1	Finance
SQL Cluster 2	

2. You can choose to use advanced options with Workload Right Sizer by checking the Use Advanced Options box. For example, you can change the CPU and memory over-subscription ratios, which are set to 2:1 by default. You can select whether a higher tier VM gets precedence for resources if there is a conflict. You can also set acceptable usage thresholds for vCPU and VM memory usage. This identifies candidates for right sizing.

Workload Right Sizer 04/08/2020, 08:04am to 05/13/2020, 08:04am ↻ Save More

Workload Right Sizer analyzes CPU and memory in your virtualized environment and recommends tuning adjustments that you can execute via a script or change control.

Analyze: Application Ordering System

Enforce Application Equilibrium
This option is not available as 'Ordering System' uses shared VMs

Use Advanced Options

Advanced

Target CPU Oversubscription Ratio 2 : 1 (Moderate)

Target Memory Oversubscription Ratio 2 : 1 (Conservative)

In case of a conflict, a higher tier VM gets the resources

Acceptable Usage Thresholds
Usage outside these ranges will identify workloads as candidates for right sizing.

Acceptable vCPU usage is above 20 % and below 80 %

Acceptable VM Memory usage is above 20 % and below 80 %

- You can also over-ride the thresholds for a specific tier. Say you wish to right size the Bronze tier differently than other tiers and select candidates based on different thresholds. You can do this by selecting “Override Thresholds by Tier”, selecting the tier and setting the usage thresholds to different values.

Override Thresholds by Tier + Add Tier

Tier Bronze

Acceptable vCPU usage is below 70 %

Acceptable VM Memory usage is below 70 %

- Set the date range for the analytic. When running Workload Right Sizer for the first time, choose a range where the entities have been in the same state. After the analytic is run for the first time, it's advisable to run it monthly.

04/08/2020, 08:04am to 05/13/2020, 08:04am

Last 2 Hours
 Last 6 Hours
 Last 24 Hours
 Last 7 Days
 Last 30 Days
 Last 35 Days
Last 3 Months
 Last 6 Months
 Last 9 Months
 Last 12 Months
 Custom

Date Range **Feb 13, 2020** - May 13, 2020

Time Range 08:13 AM - 08:13 AM

April 2020							May 2020							
S	M	T	W	T	F	S	S	M	T	W	T	F	S	
				1	2	3	4						1	2
5	6	7	8	9	10	11	3	4	5	6	7	8	9	
12	13	14	15	16	17	18	10	11	12	13	14	15	16	
19	20	21	22	23	24	25	17	18	19	20	21	22	23	
26	27	28	29	30			24	25	26	27	28	29	30	
							31							

Apply Cancel

5. You can save analytics in the same way you can save a report and schedule them to run periodically. Save the analytic as a template and schedule it as a monthly job.

Workload Right Sizer

Workload Right Sizer analyzes CPU and memory in your virtualized environment and recommends tuning adjustments that you can execute via a script or change control.

Advanced

Target CPU Oversubscription Ratio: 2 (1:1 (Default))

Save Cancel

Workload Right Sizer Attributes

Save As: Workload Right Sizer - Monthly - Ordering System Description: Monthly scheduled run of Workload Right Sizer for the Ordering System application

Tags: Application

Access Level

Schedule

Run: Every Sunday Time of Day: User: 12AM PDT / Appliance: 12AM PDT Interval: Last 30 Days Pause Schedule

+ Add Schedule

Save Cancel

6. You can run the analytic in the background and view its output later on the Outputs page.

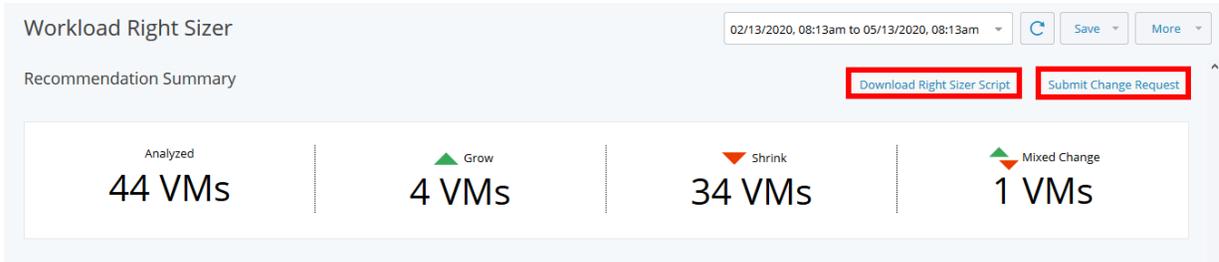
Analyzing possible solutions...

24% Completed

Notify me when completed

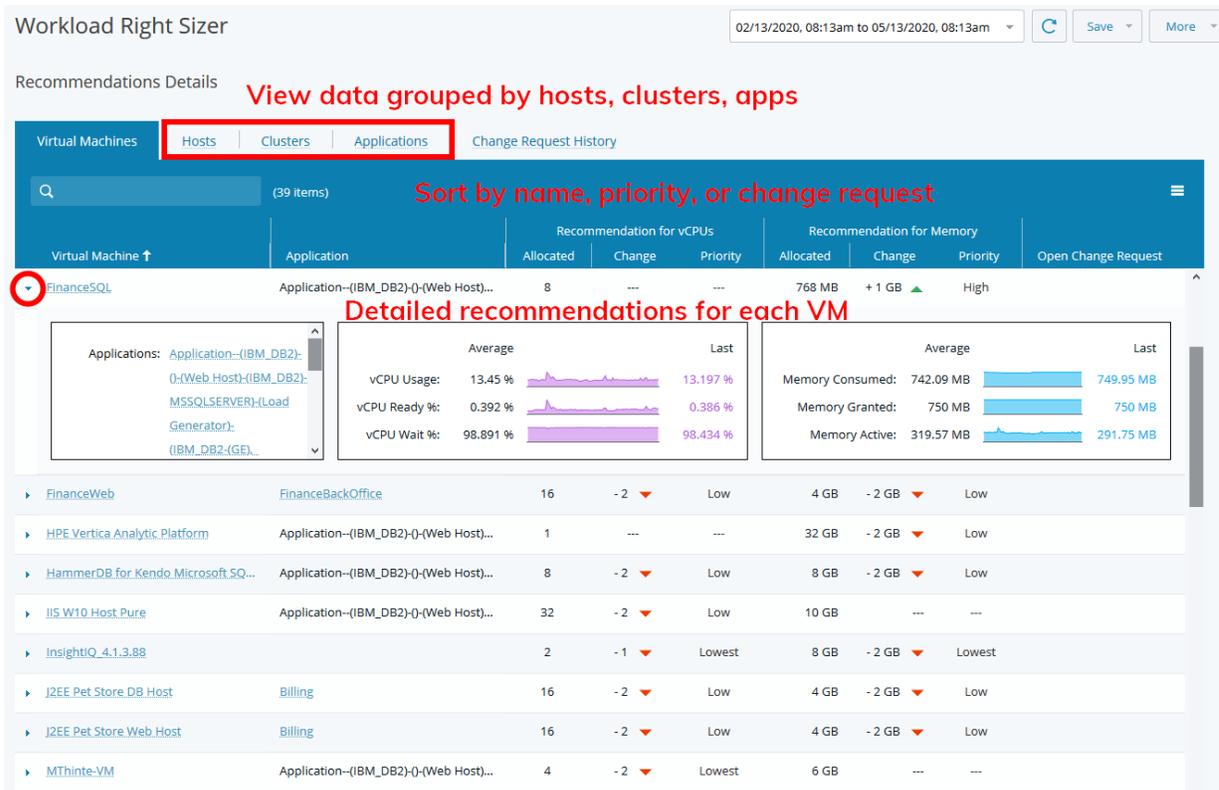
Understanding Workload Right Sizer Results

The top section shows you how many VMs were analyzed and a summary of the recommendations in terms of VM to grow, shrink, or a combination of both. Links to download the Right Sizer script or to submit a change request via the ServiceNow integration are also provided here.

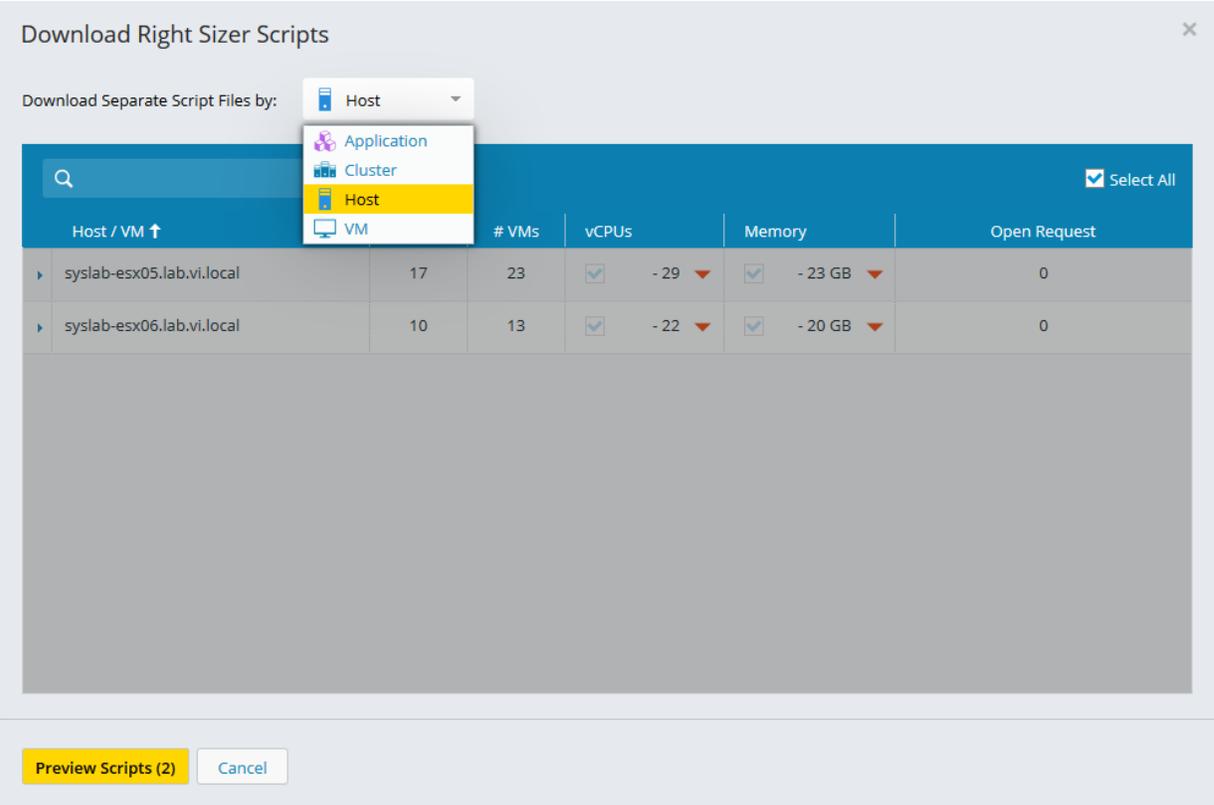


The Virtual Machines tab shows a list of the analyzed VMs and recommendations for vCPU and memory. You can expand the results for each VM by clicking the down arrow next to its name. The results can be sorted by name, priority, or open change request.

Three additional tabs share analysis and recommendations grouped by hosts, clusters, and applications.



You can download a script to make the recommended changes.



Download Right Sizer Scripts

Download Separate Script Files by: Host

Select All

Host / VM ↑	# VMs	vCPUs	Memory	Open Request	
syslab-esx05.lab.vi.local	17	23	- 29	- 23 GB	0
syslab-esx06.lab.vi.local	10	13	- 22	- 20 GB	0

Preview Scripts (2) Cancel

You can also initiate and track ServiceNow change requests (requires the ServiceNow integration).

Open Change Request in ServiceNow ✕

You are about to create **2 change requests** in ServiceNow.
Requests are grouped by **Host** and will affect **36 Virtual Machines**. Override file groupings and use 1 file

Target Maintenance Window:

Review request details:

Host / VM ↑	# Apps	# VMs	Requested Changes	
			vCPUs	Memory
▶ syslab-esx05.lab.vi.local	17	23	- 29 ▼	- 23 GB ▼
▶ syslab-esx06.lab.vi.local	10	13	- 22 ▼	- 20 GB ▼

Back
Submit Change Requests (2)
Cancel

Queue Solver

Queue Solver provides guidance in optimizing the settings for HBA queue depth. It enables you to maximize SAN traffic queue depth without hitting the avalanche point of backlogged requests that slows host response time.



NOTE

Queue Depth refers to submitted but unanswered read or write requests.

IF SET TOO HIGH

- Storage ports overrun or congested
- Poor application performance
- Potential data loss or corruption

IF SET TOO LOW

- May impair HBA performance
- Storage ports underutilized because network is underutilized
- Critical apps may wait while lower priority tasks proceed
- May cause erratic transaction latency

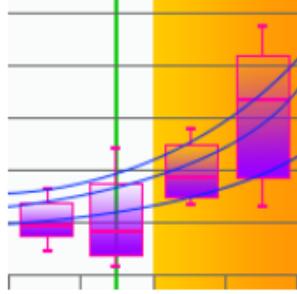
The analytic displays a graph of the inferred queue depth vs. read or write acknowledgement latency, and provides one of four recommendations about queue depth:

- No change to the queue depth limit proposed.
- It has been determined that the queue depth limit should be lowered. Lower it to {x}.
- Based on this data set, it was determined that the queue depth limit is too low. If you feel the chosen data range represents normal operation, your system might benefit from raising the queue depth limit. Proceed with caution in incremental amounts.
- Based on this chosen time range, your queue depth limit setting looks fine. You might see some gains if you lower the queue depth limit to {x}; however, the time gains might not be significant or noticeable.

The Queue Solver analytic examines network traffic for an HBA card, but does not directly read the current queue depth from the HBA BIOS.

To create a Queue Solver analytic:

1. Start by running a new Queue Solver from the Analytics home page by click the **New Analytic** or **Run New** button.

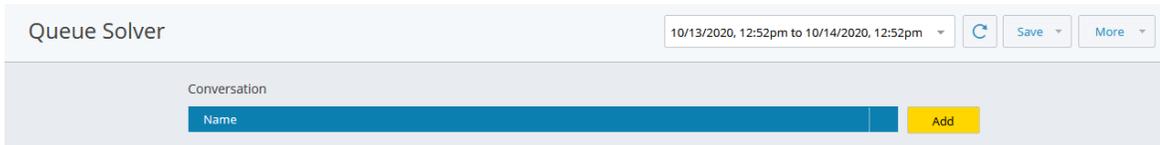


Queue Solver

Suggests optimal HBA queue settings for a specified host, by analyzing read and write response times at various queue usages. Use this information to tune your HBA cards to improve performance.

[Run New](#) [Template\(s\) \(1\)](#)
[Output\(s\) \(6\)](#)

2. Click the **Add** button to select a host or ESX host to analyze.



Queue Solver 10/13/2020, 12:52pm to 10/14/2020, 12:52pm C Save More

Conversation

Name

[Add](#)

3. Select a Host or ESX Host, then click the **OK** button. You can use the search box to filter the display of Hosts/ESX Hosts.

Add Entity ✕

Host Show Archived

Name ↑	Tags
syslab-esx04	⌵
syslab-esx05	⌵
syslab-esx06	⌵
syslab-ntx-n01	⌵
syslab-ntx-n02	⌵
syslab-ntx-n03	⌵

(Items: 6)

4. Run Queue Solver for at least two weeks, making sure to select a date range that is a good representation of your business cycle, e.g., includes month-end processing.

Last 2 Hours
 Last 6 Hours
 Last 24 Hours
 Last 7 Days
 Last 30 Days
 Last 3 Months
 Last 6 Months
 Last 9 Months
 Last 12 Months

Date Range: Sep 30, 2020 - Oct 14, 2020
 Time Range: 12:52 PM - 12:52 PM

September 2020							October 2020						
S	M	T	W	T	F	S	S	M	T	W	T	F	S
		1	2	3	4	5			1	2	3		
6	7	8	9	10	11	12	4	5	6	7	8	9	10
13	14	15	16	17	18	19	11	12	13	14	15	16	17
20	21	22	23	24	25	26	18	19	20	21	22	23	24
27	28	29	30				25	26	27	28	29	30	31

Subsequent runs should be made after queue depth settings have been changed.

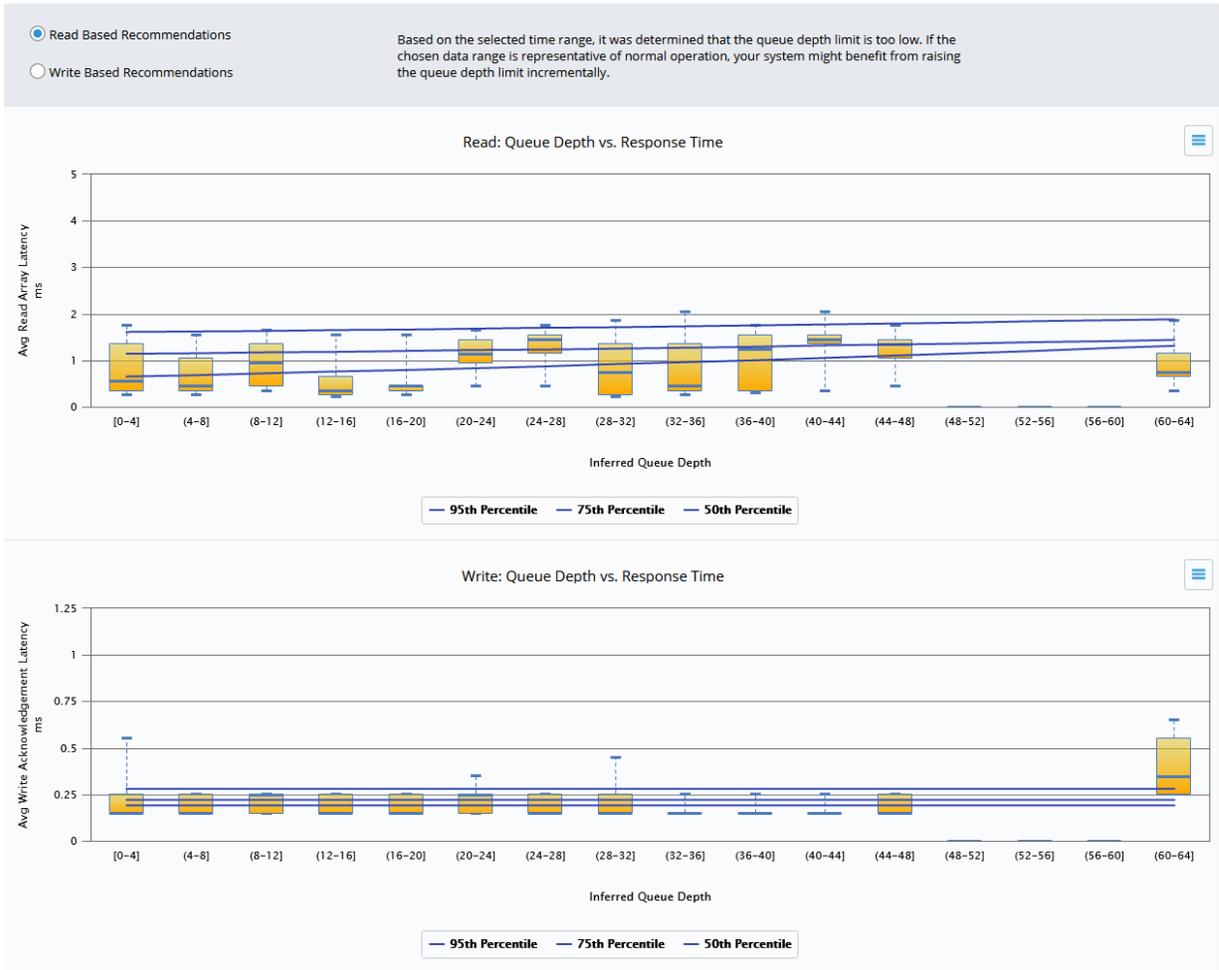
- Click the **Run** button.
- If the host does significantly more writes than reads, click the **Write-Based Recommendations** radio button.
The queue solver analytic screen is redisplayed, showing the write-based recommended action to be taken on the queue depth setting. In this case, the read- and write-based recommendations are different.
- You can click the percentile legend selections to show only some or all of the information.
Only the selected percentile curve is displayed.
- Hover over the graph lines to display detail information.
- To export a copy of the displayed graph, click the options menu (horizontal bars), and select the **Export** format from the displayed menu.

Understanding Queue Depth Results

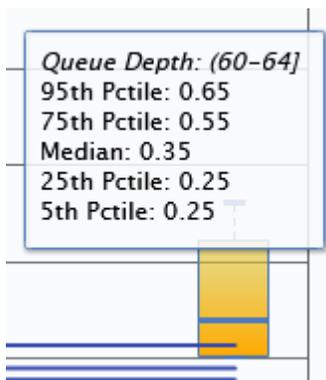
Two graphs are displayed at the bottom of the page for **Read** and **Write**-based recommendations. You can toggle between the recommendations using the radio buttons.

A summary of the recommendation is displayed at the top.

The charts present the **Avg Read Latency** or **Avg Write Latency** along the y-axis, and the inferred queue depth along the x-axis. As expected, latency increases with queue depth as port utilization increases.



Hover over a box to display percentile values for the box charts.



! IMPORTANT

When you implement Queue Solver recommendations, examine the potential impact of the change on other devices sharing the same target storage controller, especially different HBA queue depth settings to the same controller port.

A single HBA with a significantly lower queue depth setting than the other HBAs sharing the controller can result in proportionally less work done, with a potential degradation of throughput for that host.

Therefore, if a large queue depth reduction is recommended for a single HBA, carefully consider reducing the queue depth setting for all HBAs sharing the storage port.

If a host has multiple HBA cards, the same queue depth should be specified for all the HBAs.

Migration Analysis

The Migration Analysis analytic is used to prepare data collected by VirtualWisdom for use in [Virtana Migrate](#). Application and device mappings, and inventory, network, workload and other data is collected and organized into a zip file that can be downloaded from VirtualWisdom. If you have configured a connection to the Virtana Platform, then the collected and organized data can be sent directly to Virtana Migrate. Migration Analysis makes it possible for you to use the most realistic data possible, collected from your production infrastructure, to complete cloud migration readiness assessments and other digital transformation projects using Virtana Platform.

Running Migration Analysis

1. Start by running a new Migration Analysis by clicking the New Analytic or Run New button.



Migration Analysis
Prepares a detailed export for Virtana Migrate to enable application dependency mapping, suggest ideal resource groupings, and provide insights on right sizing and cost optimization across a multitude of cloud providers.

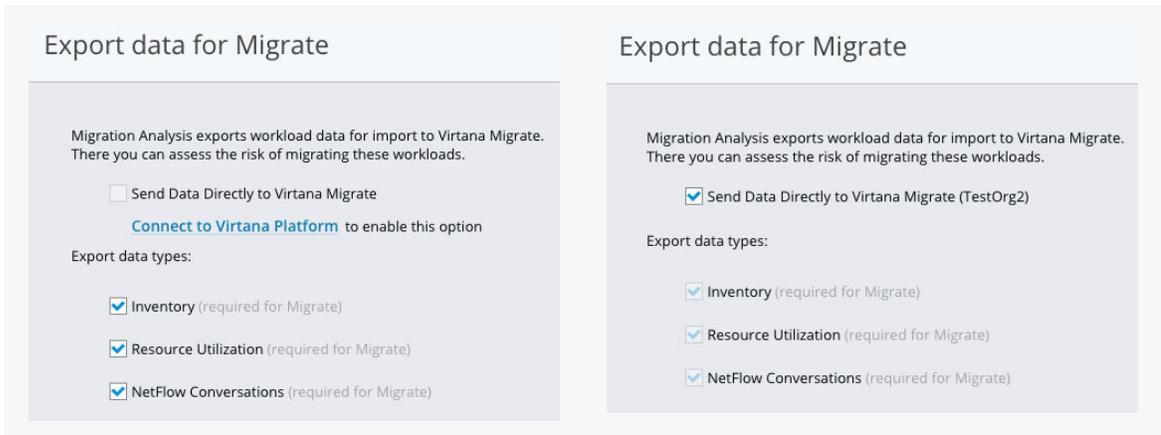
[Run New](#) [Output\(s\) \(15\)](#)

2. Select the types of data you wish to export. There are two options. First, if connectivity to Virtana Platform has not yet been configured, as shown below in the left-hand image, you select which data you want exported: **Inventory**, **Resource Utilization**, or **NetFlow Conversation** data. If connectivity to Virtana Platform has been configured, as shown below in the right-hand image, the form will already be filled out according to the Virtana Platform connection settings. NetFlow Conversations include both NetFlow and netstat conversations.

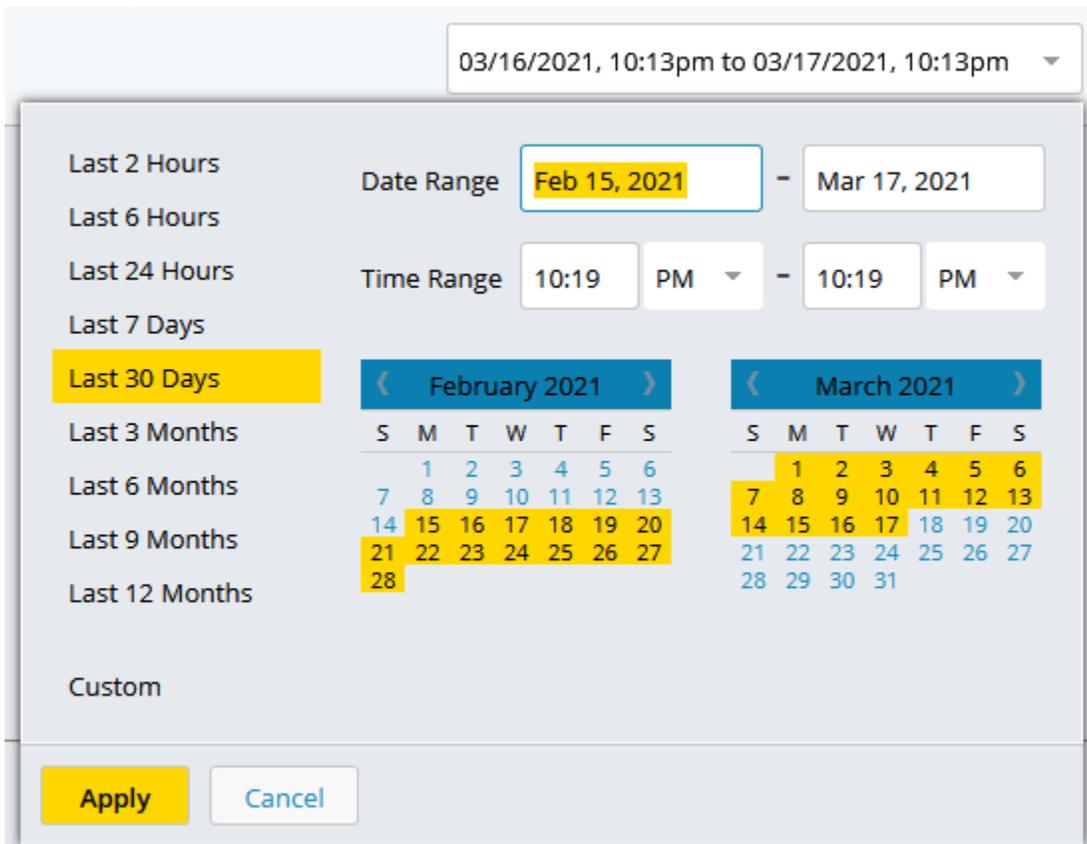


TIP

If you plan to import the data into Virtana Migrate, you must check all three data types. This will be done by default as shown below if you have already configured a connection to Virtana Platform. Uncheck the **Send Data Directly to Virtana Migrate (Org)** checkbox to be able to select the desired data types.



3. Select a time range for the export by clicking on the date field. The ideal date range for extracting the most granular data (5-minute) is 30 days. The granularity of data extracted for ranges over 30 days will be based on VirtualWisdom's [data persistence \[97\]](#) policy.



4. Click Prepare Export to start the export process if Virtana Platform connectivity has not yet been configured. Otherwise, click the **Send to Virtana Migrate** button. The

process runs in the background. You can choose to be notified when it is complete by checking the **Notify me when completed** box.



5. A message is displayed when the analysis is complete. Two possibilities exist as shown below. The left-hand picture depicts what is displayed when Virtana Platform connectivity is not configured. In this case, click the **Download Results** button to download the file. The right-hand picture depicts what is displayed when Virtana Platform connectivity is configured. For this case, downloading is optional, as the data has already been pushed to Virtana Migrate.



TIP

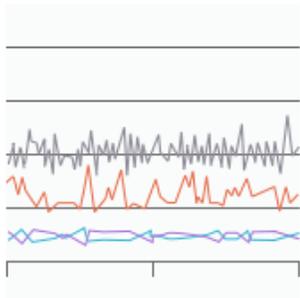
After the results of Migration Analysis have been sent to Virtana Migrate, VirtualWisdom will retain the results until Migration Analysis is run again. Once it's run again, the previously stored results data will be overwritten with new data. If you wish to retain a copy of the results data, you should download it for safe keeping.

Workload Analysis

The Workload Analysis analytic is used to prepare and transfer workload data collected by VirtualWisdom probes into WorkloadWisdom. Workload data is analyzed and grouped by VirtualWisdom and a file is prepared for download. Workload Analysis makes it possible for you to use the most realistic data possible, collected from your production infrastructure to perform validation and testing of storage using WorkloadWisdom.

Running Workload Analysis

1. Start by running a new Workload Analysis by clicking the **New Analytic** or **Run New** button.



Workload Analysis

Creates a workload simulation file based on user-supplied information about an application or host and transfers the information to WorkloadWisdom. Use to provide realistic workload simulations to replay in a lab to assess and resolve issues or to test new configurations.

[Run New](#) [Template\(s\) \(1\)](#)
[Output\(s\) \(1\)](#)

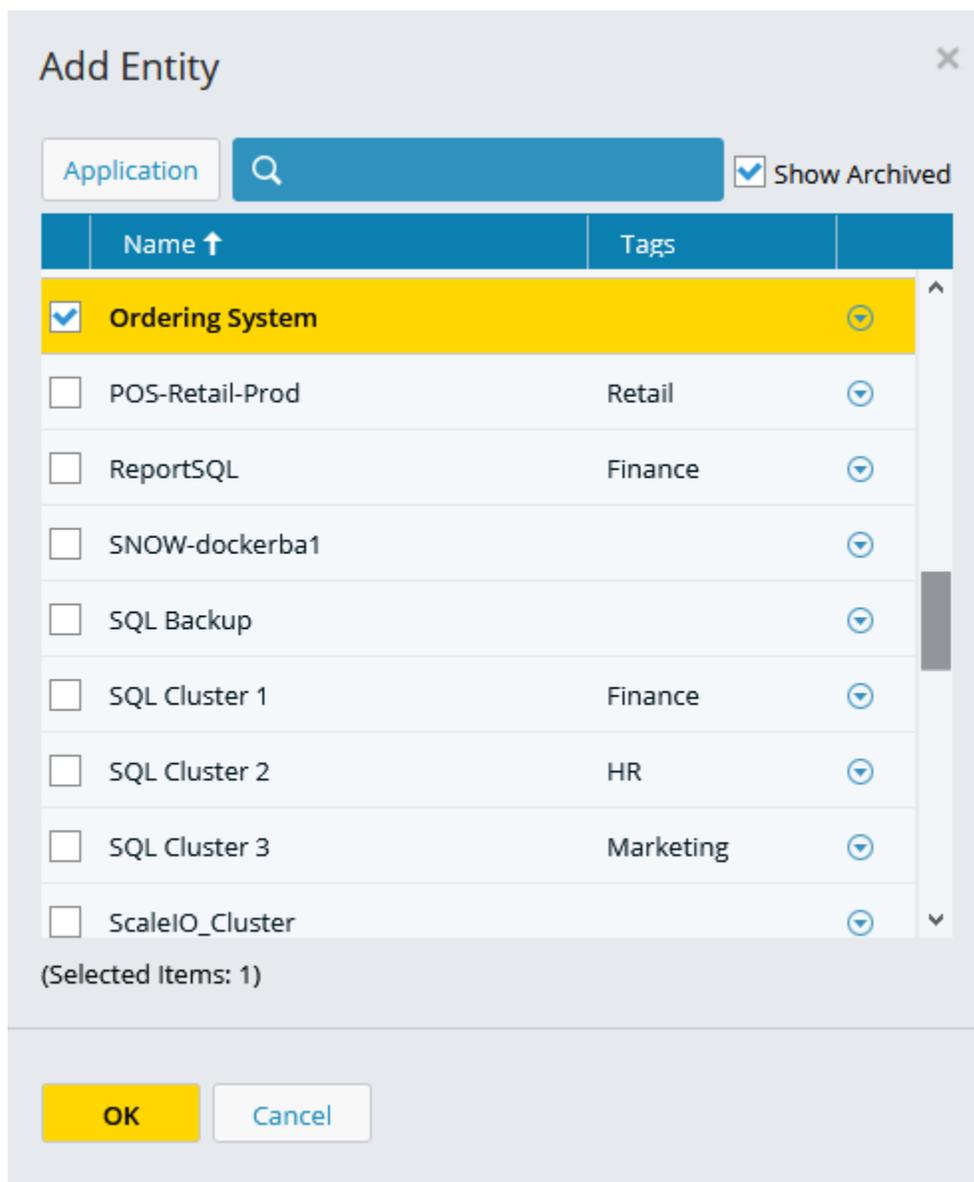
2. Click **Add** to select the entity to be analyzed.

Export data for Workload Analysis

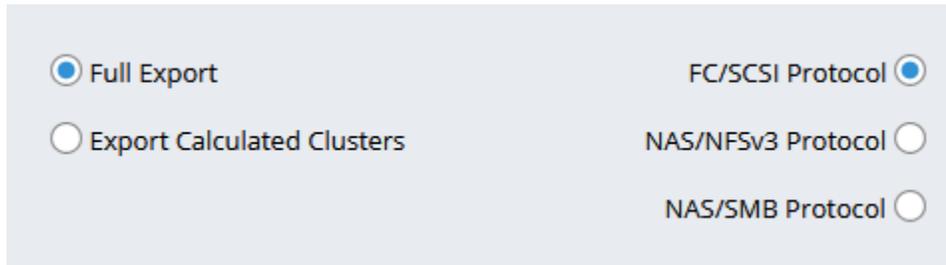
Entities

Name		Add
------	--	-----

3. Select an entity type and entity.



- Specify a full export or an export with calculated clusters, groups of similar data profiles.
The full export generates a large raw file, and the clustered option produces a smaller file, with the *Advanced Options* of either auto-detecting or limiting the maximum number of clusters.
Select the protocol using the radio buttons.



Form showing export options:

- Full Export
- Export Calculated Clusters
- FC/SCSI Protocol
- NAS/NFSv3 Protocol
- NAS/SMB Protocol

5. If exporting calculated clusters, check the **Use Advanced Options** and choose to auto-detect or limit the maximum number of clusters to a specified value. The default setting is to limit the maximum number of clusters to eight.

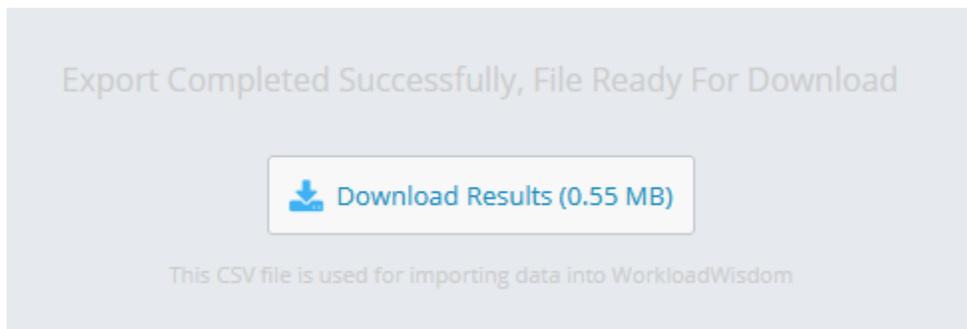


Advanced options form:

Advanced

- Auto-detect number of clusters
- Limit maximum number of clusters to

6. Click the **Prepare Export** button. The file is generated.
7. A message is displayed when the analysis is complete. Click the **Download Results** button to download the file.



Export Completed Successfully, File Ready For Download

[Download Results \(0.55 MB\)](#)

This CSV file is used for importing data into WorkloadWisdom

Predictive Capacity Management Analytics

Predictive Capacity Management analytics forecast your capacity needs using the same solution that monitors your workloads.



Capacity Forecast

Reviews historical data to predict long-term and short-term usage trends. Use to identify resource strain and to plan for growth.

[Run New](#) [Template\(s\) \(5\)](#)
[Output\(s\) \(10\)](#)

The **Capacity Forecast** analytic reviews historical resource usage data to predict short and long-term usage trends. Capacity Forecast can be used to identify resource strain before it leads to performance and availability issues, and to plan for future infrastructure growth to support critical business applications.

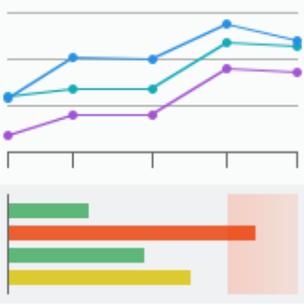


VM Deployment Advisor

Determines which cluster in your environment has the most compute resources, based on available capacity and expected workload across CPU, memory, I/O, and network. Use to select the optimal cluster for your new VMware, IBM, or Microsoft VM.

[Run New](#) [Template\(s\) \(4\)](#)
[Output\(s\) \(8\)](#)

The **VM Deployment Advisor** analytic optimizes the deployment of new Virtual Machines by examining historical usage and identifying which cluster and VM to deploy workload on for optimal performance and balance. The analytic uses historical usage to determine how a new VM will have to fit in terms of CPU, Memory, Network and Disk Usage.



Capacity Auditor
Performs deep statistical analysis on capacity utilization data to provide an overview of usage across various storage components. Use to identify where and when capacity adjustments need to be made and how storage processes and methods, such as deduplication, compression and thin provisioning, are impacting your storage utilization.

[Run New](#) [Template\(s\) \(4\)](#) [Output\(s\) \(6\)](#)

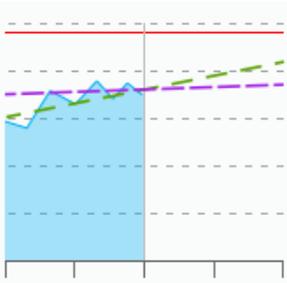
Capacity Auditor performs deep statistical analysis on capacity utilization data to provide an overview of usage across various storage components. Use Capacity Auditor to identify where and when capacity adjustments need to be made and how storage processes and methods, such as deduplication, compression and thin provisioning, are impacting your storage utilization.

Capacity Forecast

Capacity Forecast reviews historical usage data to predict short and long-term usage trends. Use Capacity Forecast to identify resource strain before it leads to performance and availability issues, and to plan for the growth required to support critical business applications.

Running Capacity Forecast

1. Start by running a new Capacity Forecast from the Analytics home page.



Capacity Forecast
Reviews historical data to predict long-term and short-term usage trends. Use to identify resource strain and to plan for growth.

[Run New](#) [Template\(s\) \(5\)](#) [Output\(s\) \(10\)](#)

2. Select an entity type to analyze from among Compute, Network, and Storage entities.

Capacity Forecast

Capacity Forecast
See how fast capacity usage is growing and view predictions for when capacity will be fully used

Show Capacity Forecast for ESX Host [Select ESX Host]

Use Advanced Options

Show Archived

Name ↑	Tags
syslab-esx01.lab.vi.local	⌵
syslab-esx02.lab.vi.local	⌵
syslab-esx03.lab.vi.local	⌵
syslab-esx04.lab.vi.local	⌵
syslab-esx05.lab.vi.local	⌵
syslab-esx06.lab.vi.local	⌵
syslab-esx07.lab.vi.local	⌵
syslab-esx08.lab.vi.local	⌵
syslab-esx09.lab.vi.local	⌵
syslab-esx10.lab.vi.local	⌵

(Items: 10)

3. Check the box to use advanced options. You can adjust the forecast thresholds, select the metrics you wish to analyze, and change the percentile value to base the forecast on.

Capacity Forecast 04/08/2020, 09:44am to 05/13/2020, 09:44am ▾

Capacity Forecast
See how fast capacity usage is growing and view predictions for when capacity will be fully used

Show Capacity Forecast for

Use Advanced Options

Advanced

Forecast Thresholds
Forecasts estimate when usage is likely to reach full capacity.

! Show Critical if less than ▾

! Show Warning if less than ▾

Capacity Metrics

CPU Usage

Memory Usage

Network (IP) Usage (Receive)

Network (IP) Usage (Transmit)

Disk Space Usage

Usage Percentile
Calculate forecast based on the ▾ of capacity usage.

Run

4. Run Capacity Forecast for a minimum of 6 months. Using a longer date range is recommended. You can run the analytic in the background and review its output later the Outputs page.

Last 2 Hours
 Last 6 Hours
 Last 24 Hours
 Last 7 Days
 Last 30 Days
 Last 35 Days
 Last 3 Months
 Last 6 Months
 Last 9 Months
 Last 12 Months
 Custom

Date Range: -

Time Range: -

April 2020						
S	M	T	W	T	F	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30		

May 2020						
S	M	T	W	T	F	S
					1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
						31

Understanding Capacity Forecast Results

The results are displayed in panes below the settings pane. The topmost panel shows a summary of the results, letting you quickly pinpoint where there may be an issue.

Capacity Forecast 11/13/2019, 09:49am to 05/13/2020, 09:49am

Capacity Forecast - syslab-esx04.lab.vi.local

CPU Usage
Usage is Stable.

Memory Usage
May reach capacity in **49 days**

Network (IP) Usage (Receive)
Daily Rate of Change: 0.029 MB/s

Network (IP) Usage (Transmit)
Daily Rate of Change: 0.654 MB/s

Disk Space Usage
10% of ESX Datastores at or near capacity.

The summary pane shows potential issues with disk space usage on this host. The details are shown below on the Disk Space Usage Details pane. 5 datastores may reach the limit of their usage within 180 days.

Disk Space Usage Details

ESX Datastore Capacity Summary

2 ESX Datastores
May reach capacity within 180 days

18 ESX Datastores
Usage is Stable.

ESX Datastore	Estimated Capacity Reaching Date	Latest 95th Percentile Usage	Capacity
Pure_Vol01	07/13/2020 61 days	1672.596 GB	1999.756 GB
vMax-Datastore01-2TB	10/24/2020 164 days	1593.843 GB	1999.756 GB

The data shown will vary based on the entity type analyzed. Shown above is the data analyzed for a storage array.

Capacity Forecast

11/13/2019, 09:49am to 05/13/2020, 09:49am

Capacity Forecast - PURE

Capacity Summary

Port Usage (Transmit)
100% of Storage Ports at or near capacity.

Port Usage (Receive)
Usage is Stable.

Port Usage (Transmit) Details

Storage Port Capacity Summary

2 Storage Ports
May reach capacity within 30 days

2 Storage Ports
May reach capacity within 180 days

Storage Port	Estimated Capacity Reaching Date	Latest 95th Percentile Usage	Capacity
PURE-CT0-FC0	06/18/2020 36 days	786.361 MB/s	800.000 MB/s
PURE-CT0-FC1	06/09/2020 27 days	786.243 MB/s	800.000 MB/s
PURE-CT1-FC0	06/18/2020 36 days	786.433 MB/s	800.000 MB/s
PURE-CT1-FC1	06/09/2020 27 days	786.301 MB/s	800.000 MB/s

Port Usage (Receive) Details

Storage Port Capacity Summary

4 Storage Ports
Usage is Stable.

VM Deployment Advisor

VM Deployment Advisor identifies the optimal cluster and host to which to deploy a VM, based on available capacity and expected VM workload across CPU, memory, I/O, and network.

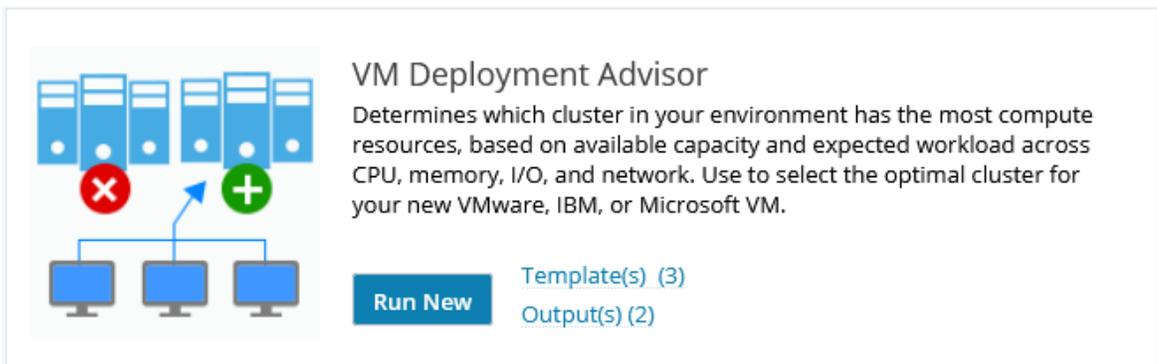
A VM Deployment Advisor analytic tells you how many more VMs fit in each host, and in turn, each cluster. You can sort by Cluster Name, or by availability. The clusters can be

color coded for vacancy amount. The size of the VM and the maximum fill point are hidden by default but an advanced user can adjust these parameters. Each Cluster has a link to VM Coordinator such that it can be launched to reorganize that cluster. Each cluster has a dropdown to reveal the host occupancy.

This analytic is used with the [VM Coordinator \[268\]](#) analytic.

Running VM Deployment Advisor

1. Start by running a new VM Deployment Advisor by clicking the **New Analytic** or **Run New** button.

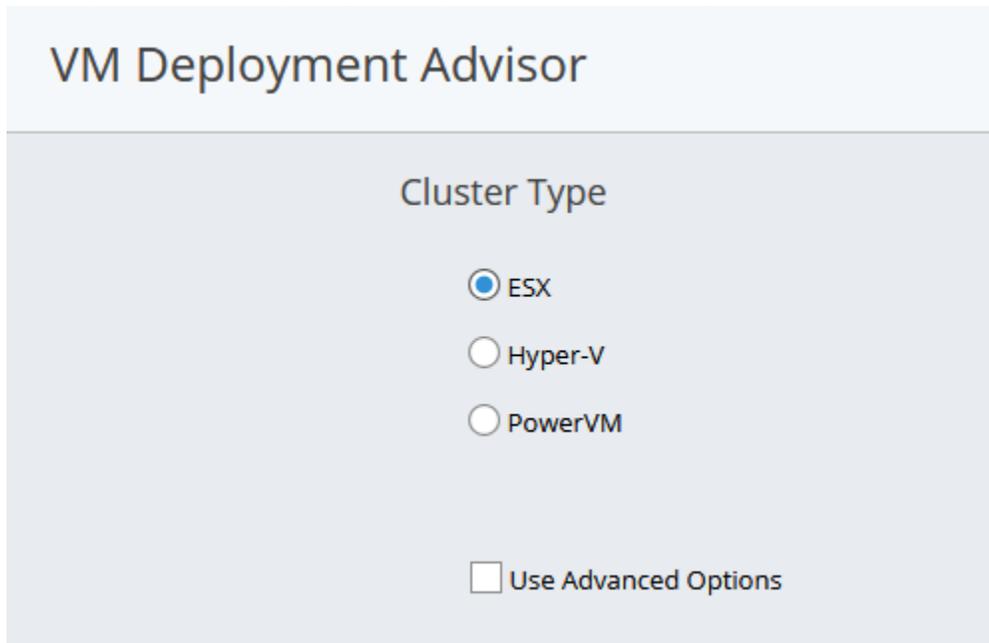


VM Deployment Advisor

Determines which cluster in your environment has the most compute resources, based on available capacity and expected workload across CPU, memory, I/O, and network. Use to select the optimal cluster for your new VMware, IBM, or Microsoft VM.

[Run New](#) [Template\(s\) \(3\)](#)
[Output\(s\) \(2\)](#)

2. Specify the Cluster Type (ESX, Hyper-V, or PowerVM).



VM Deployment Advisor

Cluster Type

ESX

Hyper-V

PowerVM

Use Advanced Options

3. If you check the **Use Advanced Options** box, you can specify desired resource limits for CPU, memory, disk and network throughput.

Cluster Type

ESX
 Hyper-V
 PowerVM

Advanced

Use auto calculated VM Size
 Use manually specified VM Size

Resource Limits

Max CPU * %
 Typical CPU * %
 Max Memory * %
 Typical Memory * %
 Typical Disk Throughput * MB/s
 Typical Net Throughput * MB/s

Use Advanced Options

[Restore Default](#)

Run

- Run the analytic for the past seven days or longer.

Last 2 Hours
 Last 6 Hours
 Last 24 Hours
Last 7 Days
 Last 30 Days
 Last 35 Days
 Last 3 Months
 Last 6 Months
 Last 9 Months
 Last 12 Months
 Custom

Date Range -

Time Range -

September 2020

S	M	T	W	T	F	S
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30			

October 2020

S	M	T	W	T	F	S
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

Apply [Cancel](#)

- Click **Run**.

Understanding VM Deployment Advisor Results

1. VM Deployment Advisor displays its results in a table presented below the settings pane. Included are the names of the clusters and the total number of VM slots available.

Clusters with Available Slots

Name ↑	Total VM Slots Available ↑
Cluster-QE-AppDiscovery	11
Cluster-QE-VSAN	5
ProdGeneral	6
ProdPlatinum	0
QE-Jenkins-Cluster	3
ScaleIO	16
ST_Cluster	10

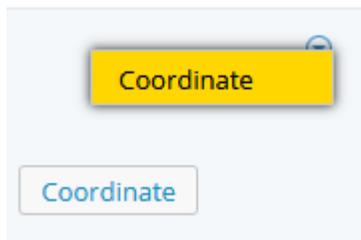
Host Name ↓	Max CPU	Typical CPU	Max Memory	Typical Memory	Average Disk Throughput	Average Net Throughput	Available VM Slots
qe-jenkins-esx01.lab.vi.local	8.23%	6.14%	38.04%	36.32%	2.36 MB/s	142.97 KB/s	2 Coordinate
qe-essystest11.lab.vi.local	52.72%	51.37%	93.8%	93.72%	187 B/s	13.04 MB/s	Over
qe-essystest10.lab.vi.local	4.86%	4.33%	54.07%	54.07%	1.42 MB/s	82.45 KB/s	1

2. An exclamation point next to the cluster name indicates that the host is already over-provisioned. Click the down arrow on the left side to expand the information displayed for the cluster.

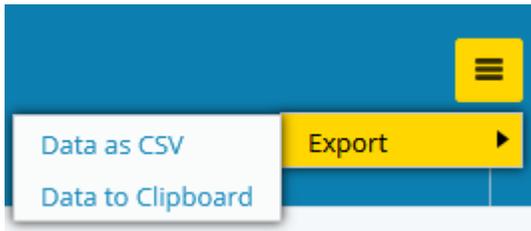
! QE-Jenkins-Cluster 3

Host Name ↓	Max CPU	Typical CPU	Max Memory	Typical Memory	Average Disk Throughput	Average Net Throughput	Available VM Slots
qe-jenkins-esx01.lab.vi.local	8.23%	6.14%	38.04%	36.32%	2.36 MB/s	142.97 KB/s	2 Coordinate
qe-essystest11.lab.vi.local	52.72%	51.37%	93.8%	93.72%	187 B/s	13.04 MB/s	Over
qe-essystest10.lab.vi.local	4.86%	4.33%	54.07%	54.07%	1.42 MB/s	82.45 KB/s	1

The **Available VM Slots** field shows you which of the hosts in the cluster have available slots and how many are available. You may see the work **Over** displayed next to the hosts that VM Deployment Advisor has determined are over-provisioned. Click on **Coordinate** to run [VM Coordinator \[268\]](#) to address the over-provisioning situation. You can also access VM Coordinator from the drop-down menu associated with the down arrow on the right.



3. To save the move recommendations, click on the hamburger icon in the header and export them to a csv file, or copy to the clipboard.



Capacity Auditor

Capacity Auditor considers how storage processes such as deduplication, compression, and thin provisioning are impacting your storage utilization and identifies where and when capacity adjustments need to be made to storage.

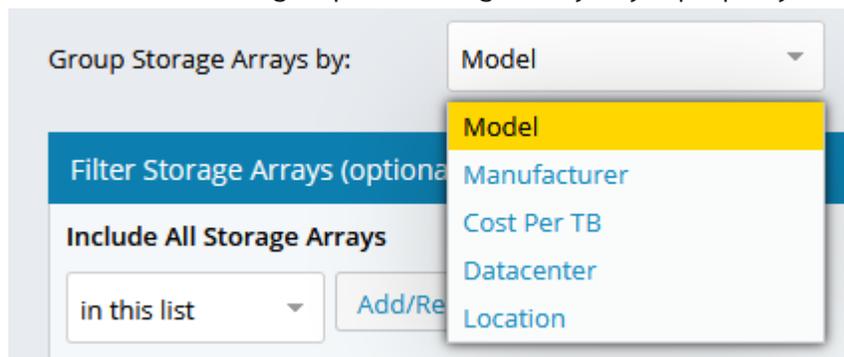
Capacity Auditor shows you capacity data and trends for your **VMAX** and **Isilon** storage arrays.

Running Capacity Auditor

1. Select **Run New** Capacity Auditor from the Analytics home page.

A screenshot of the Capacity Auditor interface. It features a line chart with three data series (blue, green, purple) and a bar chart with four bars (green, orange, green, yellow). To the right of the charts is the text: "Capacity Auditor Performs deep statistical analysis on capacity utilization data to provide an overview of usage across various storage components. Use to identify where and when capacity adjustments need to be made and how storage processes and methods, such as deduplication, compression and thin provisioning, are impacting your storage utilization." Below this text are two buttons: "Run New" (yellow) and "Template(s) (4) Output(s) (6)" (blue).

2. Set Capacity Auditor parameters.
 - a. Choose whether to group the storage arrays by a property.

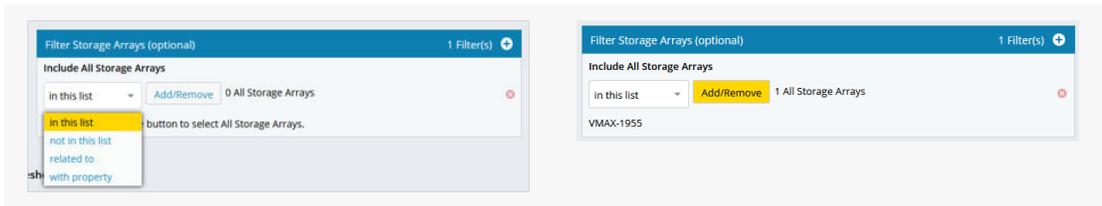




TIP

You can define [custom properties \[134\]](#) for storage arrays and use them in the analysis.

- b. To analyze a subset of arrays, use the filter to select one or more named storage arrays. You can filter for storage arrays in/not in a specified list, related to other entities, or with a specified property.



- c. Select threshold values for red/yellow severity warning levels.

Threshold Values

	Capacity Usage	Overprovisioning
● Red Threshold:	85 %	150 %
● Yellow Threshold:	70 %	100 %

- 3. Choose a date range. You should run the analytic for the past twelve months, if data is available, then monthly thereafter.

Last 2 Hours
 Last 6 Hours
 Last 24 Hours
 Last 7 Days
 Last 30 Days
 Last 3 Months
 Last 6 Months
 Last 9 Months
Last 12 Months

Date Range: **Aug 25, 2019** - Aug 25, 2020
 Time Range: 12:59 PM - 12:59 PM

< August 2019 >

S	M	T	W	T	F	S
				1	2	3
4	5	6	7	8	9	10
11	12	13	14	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

< August 2020 >

S	M	T	W	T	F	S
						1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31					

Custom

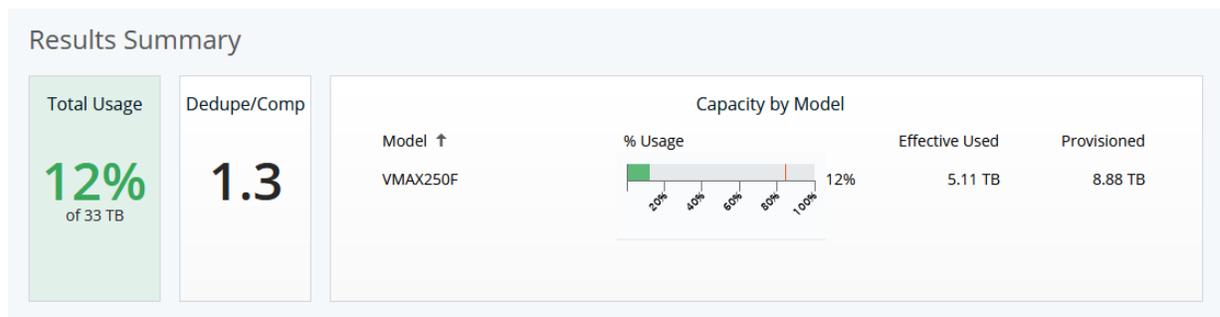
Apply Cancel

4. Click **Run** to run the analytic.

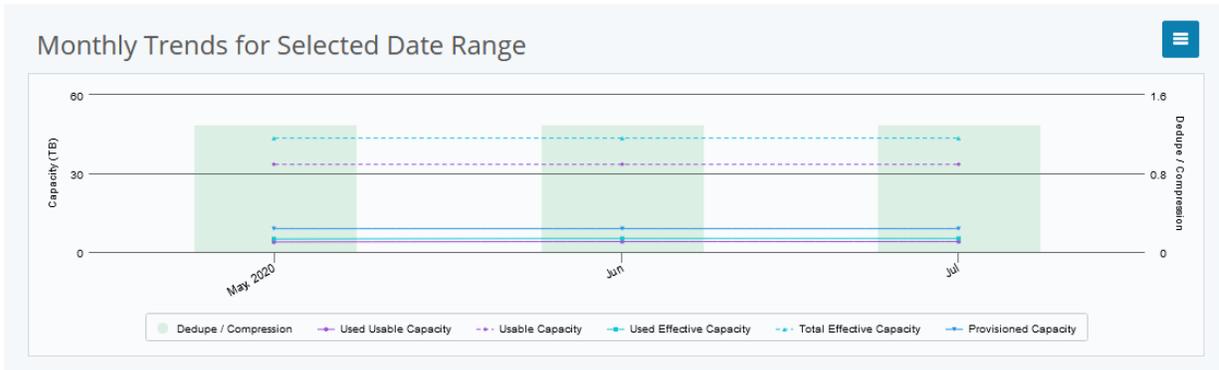
Understanding Capacity Auditor Results

Capacity Auditor results are displayed in three sections on a single page.

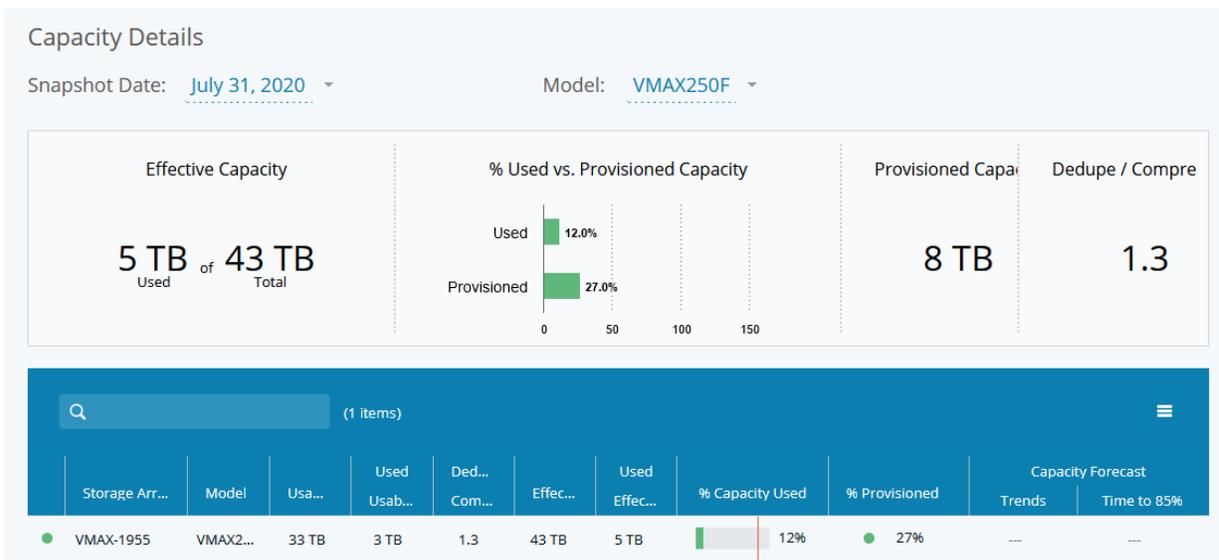
The first section shows you usage statistics for the storage array(s).



The middle section displays a trend chart showing monthly trends for capacity for the storage array(s).

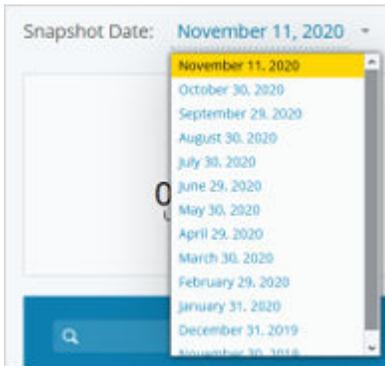


The bottom section displays more details about the selected storage array grouping.

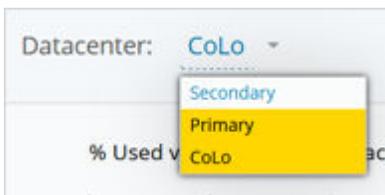


Rows display details for each storage array, including the time to reach the "red" capacity threshold chosen for the analysis. Changing this threshold changes the forecast.

To view details for a previous analysis, click the date and select a previous snapshot.



To view results for a different property value, click the value and select from the list.



Application Service Assurance Analytics

Our Application Service Assurance analytics help you assure the performance and health of your application infrastructure.



Balance Finder

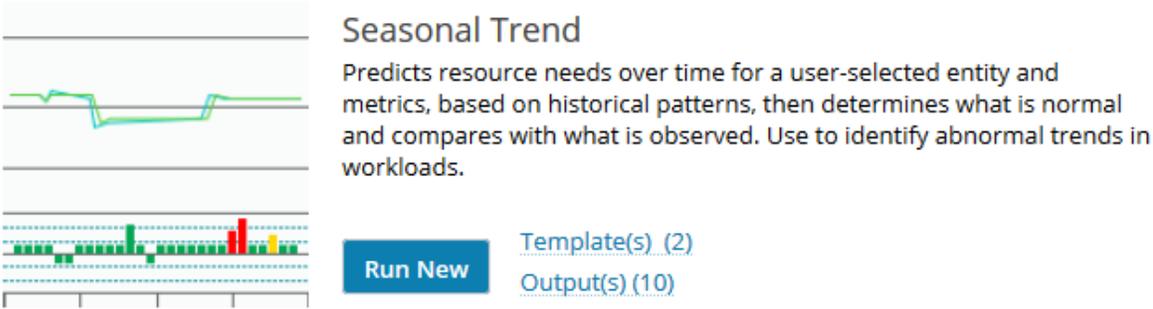
Examines throughput on the HBAs of a selected host or hosts and compares traffic patterns to determine if the HBAs are behaving as intended and if traffic is balanced. Use to verify that HBA settings are being properly implemented or to determine if multipath rebalancing is needed.

[Template\(s\) \(2\)](#)

[Output\(s\) \(13\)](#)

[Run New](#)

Balance Finder examines traffic patterns and determines if they are balanced accurately across the hosts. In large organizations the goal is to make sure that the host workload and the associated traffic is appropriately balanced across the environment. Balance Finder validates that servers in the environment have functioning path failover, ensuring availability to and from storage in event of a failure. It also validates whether there is available capacity in the event of a path or component failure in the environment.



Seasonal Trend

Predicts resource needs over time for a user-selected entity and metrics, based on historical patterns, then determines what is normal and compares with what is observed. Use to identify abnormal trends in workloads.

[Run New](#) [Template\(s\) \(2\)](#) [Output\(s\) \(10\)](#)

Seasonal Trend predicts resource needs over time for a user-selected entity and metrics, based on historical patterns, then determines what is normal and compares with what is observed. Use Seasonal Trend to identify abnormal trends in workloads and tune settings for the Seasonal Trend alarm to consider cyclical workload patterns.

Seasonal Trend

Seasonal Trend predicts resource needs over time for a user-selected entity and metrics, based on historical patterns, then determines what is normal and compares with what is observed. Use Seasonal Trend to identify abnormal trends in workloads and tune settings for the Seasonal Trend alarm to consider cyclical workload patterns.

At least 2 hours of data is required for a result, modeling seasonality requires:

- 1 hour of hourly data for daily/weekly
- 4 hours of 4-hour data for weekly/monthly
- 1 day of daily data for monthly/quarterly/yearly

Running the Seasonal Trend Analytic

1. Start by running a new Seasonal Trend from the **Analytics** home page.



Seasonal Trend

Predicts resource needs over time for a user-selected entity and metrics, based on historical patterns, then determines what is normal and compares with what is observed. Use to identify abnormal trends in workloads.

[Run New](#) [Template\(s\) \(2\)](#) [Output\(s\) \(10\)](#)

- Click the **Add** button to specify which entity and metric to run the analytic for.

Seasonal Trend 09/09/2020, 12:15pm to 10/14/2020, 12:15pm ↻ Save More

Entities		Add
Name	Metrics	

- Choose an entity and metric and click **Select**. Click **OK** to apply the selection.

Add Metrics ✕

Application 🔍 Show Archived

Name ↑	Tags	
Ordering System		⊕
POS-Retail-Prod	Retail	⊕
ReportSQL	Finance	⊕
SNOW-dockerba1		⊕
SQL Backup		⊕
SQL Cluster 1	Finance	⊕
SQL Cluster 2	HR	⊕
SQL Cluster 3	Marketing	⊕
ScaleIO_Cluster		⊕
Shared SQL		⊕

Available Metrics

- Storage
 - SAN
 - FC Switch Integration
 - FC-SCSI
 - Avg Read Completion Time
 - Read IOPS
 - Avg Read IO Size
 - Read Payload Rate
 - Avg Write Completion Time
 - Write IOPS

Select

Selected Item

Ordering System / Read IOPS ✕

Ordering System / Read IOPS

OK Cancel

- Specify the date range, and click the **Apply** button.



NOTE

It is recommended that you run the analytic for at least the last 30 days in order to see seasonal patterns.

Last 2 Hours
 Last 6 Hours
 Last 24 Hours
 Last 7 Days
Last 30 Days
 Last 35 Days
 Last 3 Months
 Last 6 Months
 Last 9 Months
 Last 12 Months
 Custom

Date Range: **Sep 14, 2020** - Oct 14, 2020
 Time Range: 12:23 PM - 12:23 PM

September 2020							October 2020							
S	M	T	W	T	F	S	S	M	T	W	T	F	S	
			1	2	3	4	5					1	2	3
6	7	8	9	10	11	12	4	5	6	7	8	9	10	
13	14	15	16	17	18	19	11	12	13	14	15	16	17	
20	21	22	23	24	25	26	18	19	20	21	22	23	24	
27	28	29	30				25	26	27	28	29	30	31	

Apply Cancel

5. Click the **Run** button.

Understanding Seasonal Trend Results

The results are displayed as a trend chart (above) and a bar chart (below).

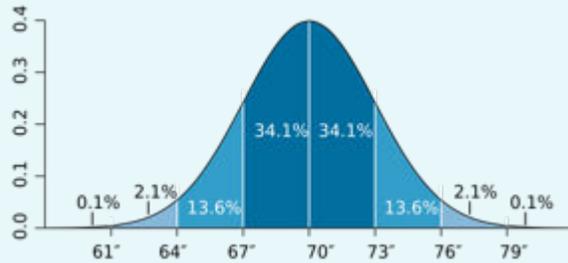
- The trend chart shows the **expected** pattern of behavior for the entity and metric, based on historical data, overlaid by the **actual** behavior.
- The bar chart displays the standard deviation between expected behavior vs. actual behavior, with the color of the bar indicating the standard deviation value.

	Remainder less than 1 S.D.
	Remainder less than 3 S.D.s
	Remainder over 3 S.D.s

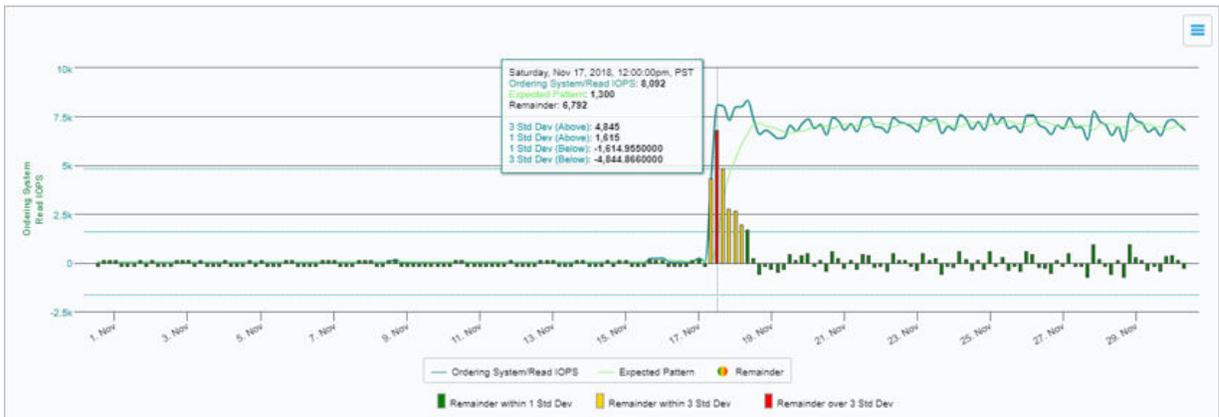


NOTE

The standard deviation measures how different the numbers in a range are from each other.



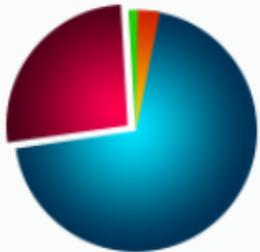
In the example below, we can see where the actual behavior differed from the expected behavior by over 3 standard deviations.



This information can be used to set up a Seasonal Trend alarm on the application workload to alert when seasonal patterns are not followed.

Balance Finder

Balance Finder examines traffic patterns and determines if they are balanced accurately across the hosts. In large organizations the goal is to make sure that the host workload and the associated traffic is appropriately balanced across the environment. Balance Finder validates that servers in the environment have functioning path failover, ensuring availability to and from storage in event of a failure. It also validates whether there is available capacity in the event of a path or component failure in the environment.

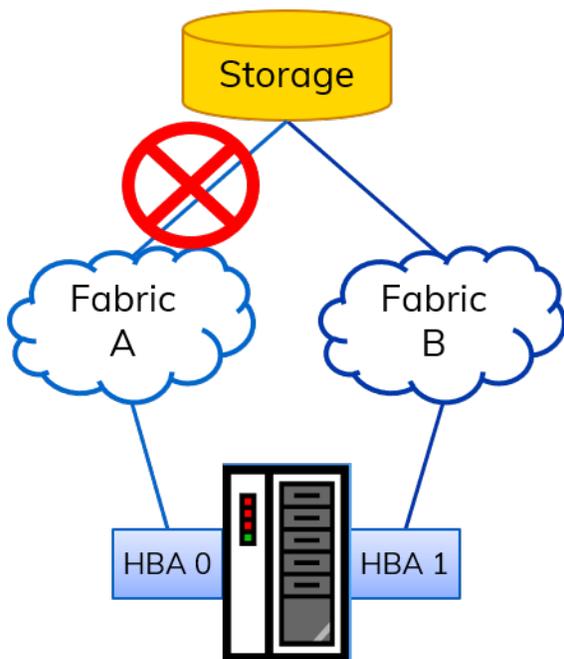


Balance Finder

Examines throughput on the HBAs of a selected host or hosts and compares traffic patterns to determine if the HBAs are behaving as intended and if traffic is balanced. Use to verify that HBA settings are being properly implemented or to determine if multipath rebalancing is needed.

Run New
Template(s) (2)
Output(s) (13)

Balance Finder validates that the hosts supporting your applications have functioning path failover, ensuring availability to and from storage in the event of a failure.



Balance Finder shows you where your multi-pathing configuration may pose a risk in the event of a failure in the path from your hosts to storage.

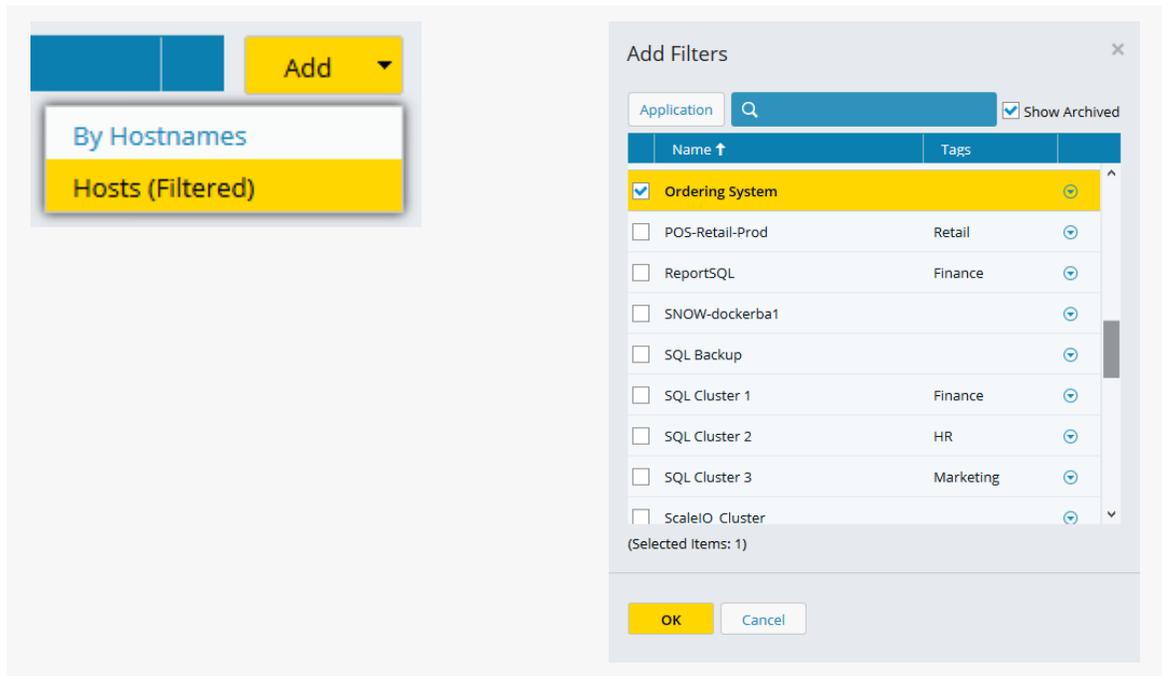
Running Balance Finder

1. The first step in using Balance Finder is to specify a host or a set of hosts to analyze.

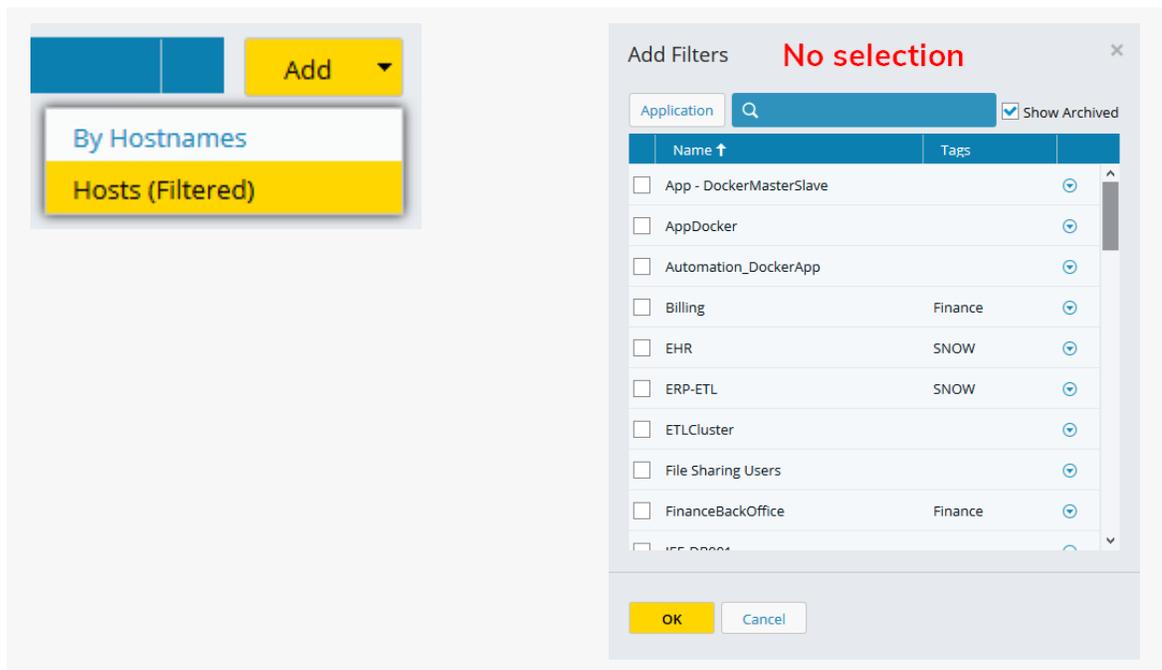
You can run Balance Finder against one or more named hosts. Select **By Hostnames**, then check the named hosts to run Balance Finder against.

Name	Tags
<input checked="" type="checkbox"/> DataMover02	
<input checked="" type="checkbox"/> DataMover03	
<input type="checkbox"/> Host1	
<input type="checkbox"/> Host10	
<input type="checkbox"/> Host100	
<input type="checkbox"/> Host101	
<input type="checkbox"/> Host102	
<input type="checkbox"/> Host103	
<input type="checkbox"/> Host104	

Alternately, you can run Balance Finder against hosts that are related to an entity. Select **Hosts (Filtered)**, then select the entity type and choose the entities to use in the "related to" filter.



Finally, you can run Balance Finder against all hosts in your environment. Select **Hosts (Filtered)** then click OK from the Add Filters dialog to run Balance Finder against all hosts.



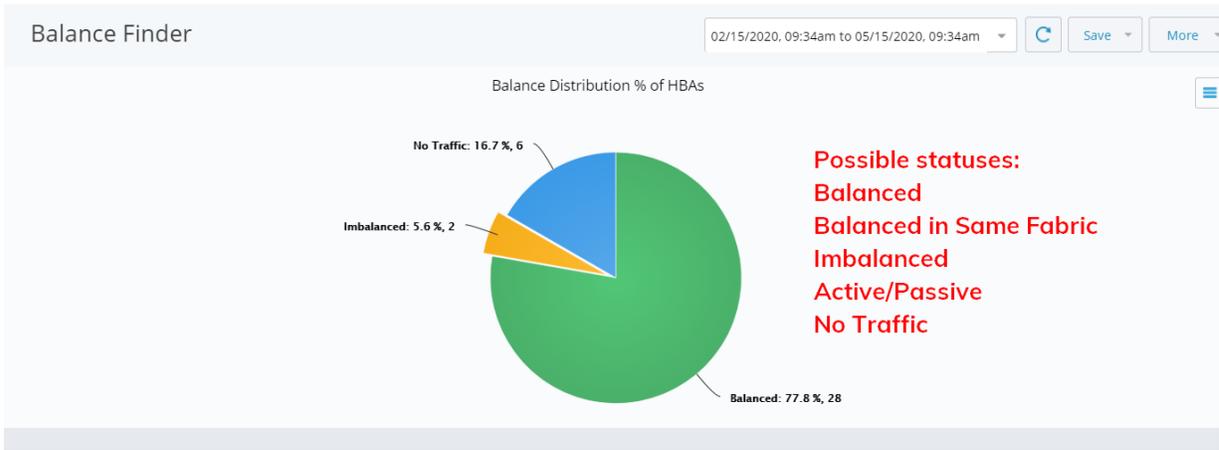
2. Select the date range.

Run Balance Finder for the last 7 days then weekly thereafter. You can choose to run the report in the background and receive a notification when it is completed. Note that the results will be automatically saved and will be available under the View All Outputs tab on the Analytics home page.

The screenshot shows a configuration panel for the Balance Finder tool. On the left, there is a list of pre-defined date ranges: Last 2 Hours, Last 6 Hours, Last 24 Hours, Last 7 Days (highlighted with a red box), Last 30 Days, Last 3 Months, Last 6 Months, Last 9 Months, Last 12 Months, and Custom. To the right, the Date Range is set to May 6, 2020 - May 13, 2020, and the Time Range is 10:23 AM - 10:23 AM. Below these fields are two calendar views for April 2020 and May 2020. The dates May 6, 7, 8, and 9 are highlighted in yellow in the May 2020 calendar. At the bottom, there are 'Apply' and 'Cancel' buttons.

Understanding Balance Finder Results

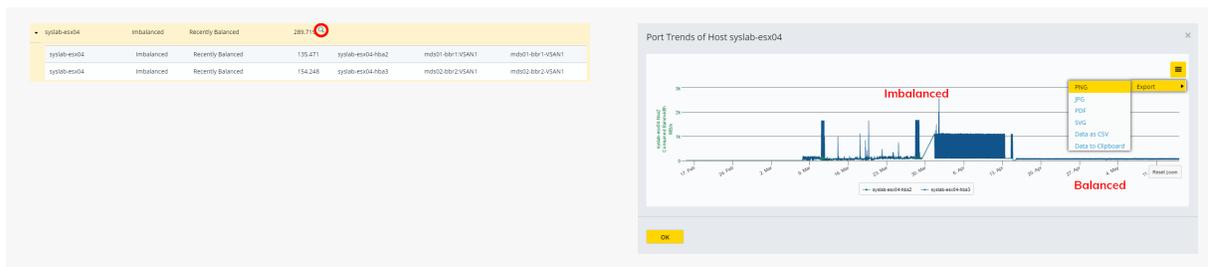
Balance Finder displays the results of its analysis in a pie chart at the top of the results page. The possible multipath statuses include balanced, balanced in the same fabric, imbalanced, active/passive, and no traffic.



The classification status, sub-status, and MB/s for each host are shown on the results grid. The Status column shows an aggregate of the status for all ports included with the server. Click on the down arrow to expand the results grid, show results by HBA, and display additional information (host port, switch, and fabric).

Hostname	Status	Sub-status	MB/s ↓	Host Port	Switch	Fabric
▾ sbлаze2-4	Balanced	Balanced In Same Fabric	1193.278*			
▾ syslab-esx06	Balanced		987.073			
⊕ syslab-esx04	Imbalanced	Recently Balanced	289.719			
syslab-esx04	Imbalanced	Recently Balanced	135.471	syslab-esx04-hba2	mds01-bbr1:VSAN1	mds01-bbr1:VSAN1
syslab-esx04	Imbalanced	Recently Balanced	154.248	syslab-esx04-hba3	mds02-bbr2:VSAN1	mds02-bbr2:VSAN1
▾ SQL-DB-001	Balanced		0.622			
▾ MailArchive01	Balanced		0.622			
▾ test host	Balanced		0.622			
▾ brldx6208	No Traffic		0.000			

You can view the port trends for a host by clicking on the magnifying glass in the host's row. The workload trends for the host's ports are displayed in a pop up window. Like all charts, the chart can be exported by clicking on the hamburger icon then selecting Export.



Saved Analytics



The *Saved Analytics* screen shows the results of previous analytics executions.

It lists analytics you have already run and can use as templates to clone new analytics. Click the *Templates* field next to the **Run New** button of an analytic to display the *Saved Analytics* screen, filtered to show only the selected type. The number in parentheses shows how many of these templates exist.

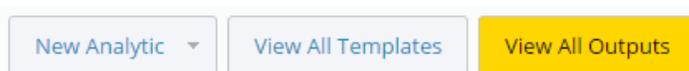
You can select a *Saved Analytics* row to be displayed, and then click the **Run** button to execute the analytic.

If you click the **Save** button and select the *Save As* option, you can modify and save a copy of the output.

The options menu (horizontal bars) enables you to export any or all of the displayed *Saved Analytics* as either a CSV file or data on the clipboard.

To display a saved analytic, click the row. To delete a saved analytic, place the pointer in its row and click the V displayed for the analytic to be deleted.

Analytics Output



The *Analytics Output* screen shows the results of previous analytics executions.

You can select an analytics output result row to be displayed.

If you click the **Save** button and select the Save As option, you can modify and save a copy of the output.

The options menu (horizontal bars) enables you to export any or all of the displayed output as either a CSV file or data on the clipboard.

To display an analytic, click the row. To delete an analytic, place the pointer in its row and click the V displayed for the analytic to be deleted.

Deleting Analytics Output

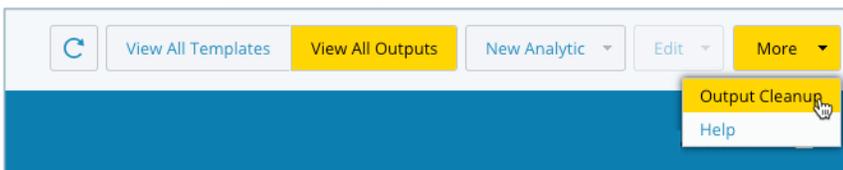
Retaining a large number of saved analytics output files can make it difficult to manage the output list. A large number of saved files also takes longer to display on the Analytics Output page. Using Output Cleanup, you can automatically delete saved output files older than a user-selected age..

About This Task

When you enable automatic deletion of analytic output, all output files older than the time period you select will be automatically deleted and can only be retrieved by doing a restore from a backup. Be sure there are no output files you want to retain that would fall within the selected deletion timeframe.

Steps

1. Navigate to the **Analytics** page and click **View All Outputs**.
The Analytics page displays.
2. Click **More > Output Cleanup**.

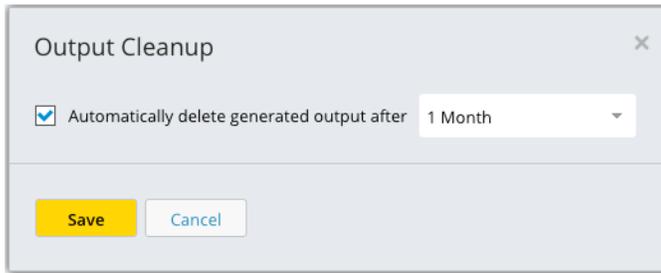


3. Enable automatic deletion of output files and select the time period after which the files will be deleted.



NOTE

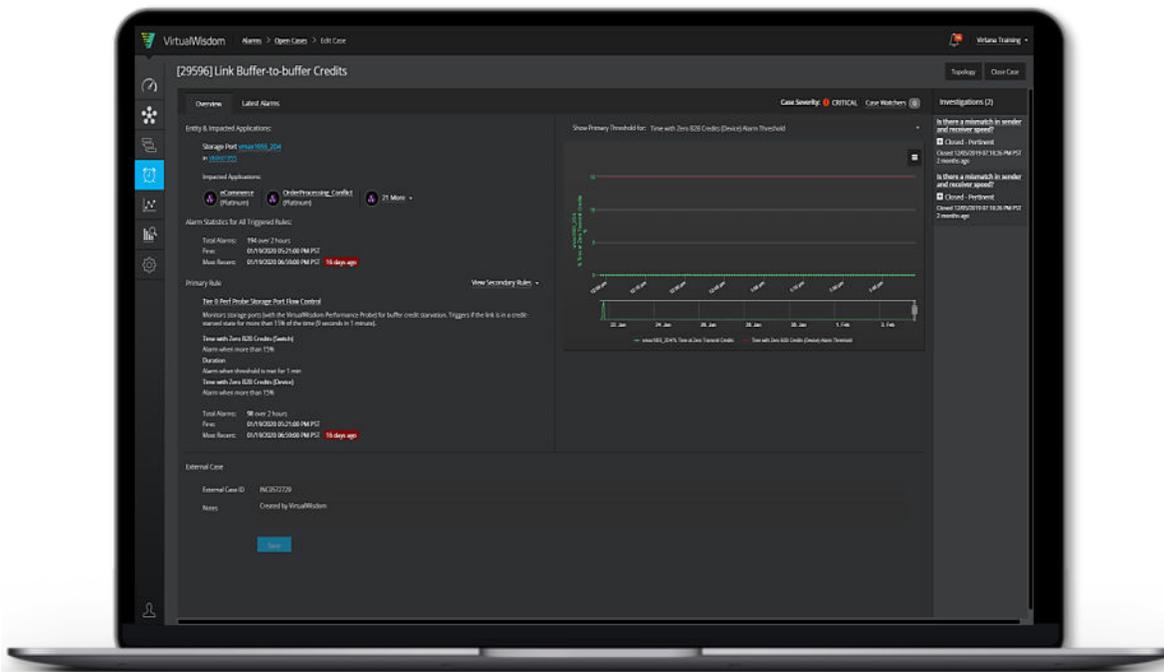
All generated output older than the selected time period will be deleted.



4. Click **Save**.
A warning displays, stating that output files will be immediately deleted.
 5. Click **OK**.
A message displays stating that output files were deleted.
 6. Click **OK**.
A message displays on the Analytics Output page, informing you that output cleanup is enabled.
- Tip:** You can disable Analytics Output Cleanup by navigating to the Output Cleanup window and clearing the checkbox.

Alarms and Cases

Quickly see where there are issues and use investigative workflows to solve them



VirtualWisdom's case-based alarms compare data collected through VirtualWisdom monitoring using configured thresholds and act when the defined conditions are met. The VirtualWisdom application opens or updates cases when an alarm is triggered. All subsequent alarms related to the initial event populate that case, allowing the user to quickly understand trends, and determine correlation and severity on a case-by-case basis.

VirtualWisdom is populated with alarm rule templates based on best practices collected from hundreds of customers, to drive a shift from reactive alarm management to proactive infrastructure performance management. These templates remove the confusion over what alarms should be created and instead are built to define what matters to the customer.

Our alarm rule templates are associated with tiers, enabling the user to manage alarms for critical applications differently than less critical applications.

Alarms Home Page

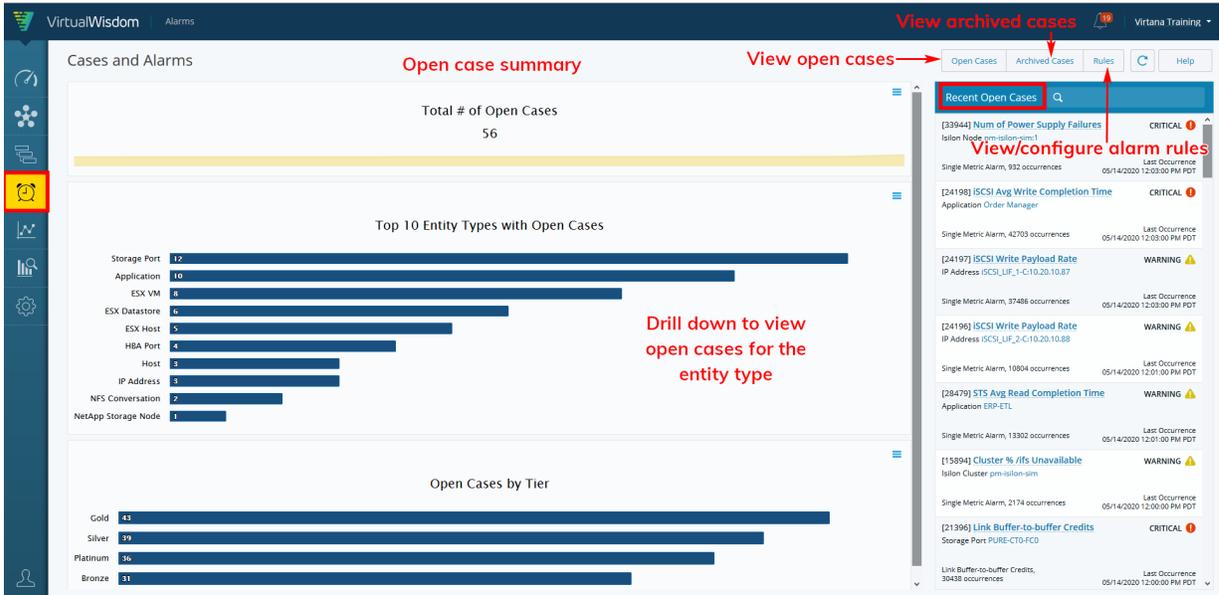
The Alarms module is located on the alarm clock icon in the VirtualWisdom Management software's user interface. The purpose of the module is to compare data collected through VirtualWisdom monitoring with configured thresholds act if conditions are met. Included in the module is a Case Management feature which opens or updates cases upon triggering of an alarm, and an Investigations features to assist users in troubleshooting and resolving the issue that caused the alarm.



NOTE

Clicking the **Open Cases by Tier** bar on the Alarms landing page does not filter open cases by tier.

The default view presents a graphical representation of the open cases on the left. You can drill down on the entity type to view all open cases on that type.



On the right is a pane showing recent open cases, ordered by date and time descending (most recent at the top). Drill down on a case to view more data.

At the top of the page are two buttons to view all open cases and all archived cases.



An archived case is a case that was closed by the user or archived by the VirtualWisdom platform.



There is also a button to view and configure alarm rule templates. VirtualWisdom is pre-populated with these standard Alarm Rule Templates.

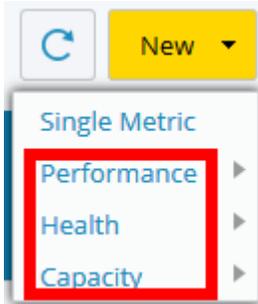


Alarm Rule Templates

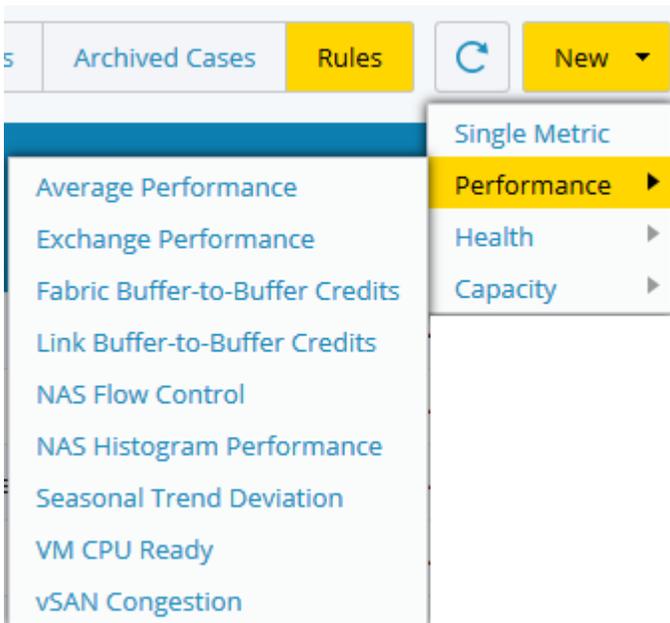
Rules are the criteria for triggering alarms when specific error thresholds are exceeded. The Alarms module contains a set of predefined rule templates that can be used to create

a new rule. These alarm rule templates are designed to reduce the complexity of rule set up and allow the administrator to quickly identify and create rules based on common errors. Rule templates are available for different error types. VirtualWisdom is pre-populated with these standard alarm rule templates.

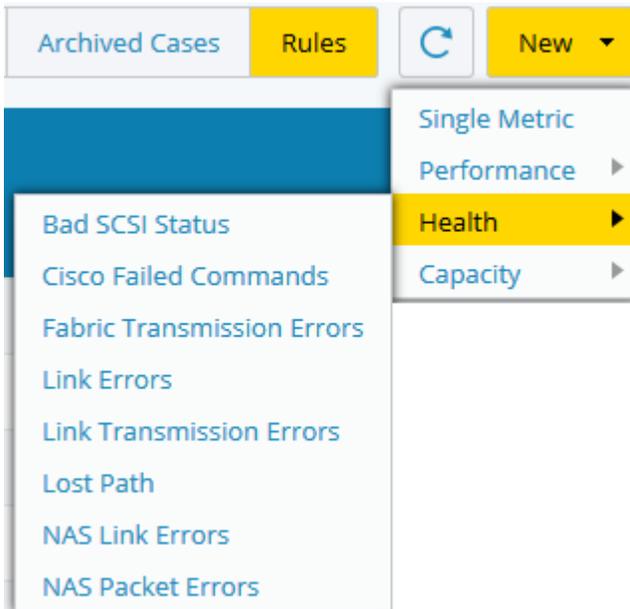
There are three categories of alarm rule templates: **Performance**, **Health**, and **Capacity**.



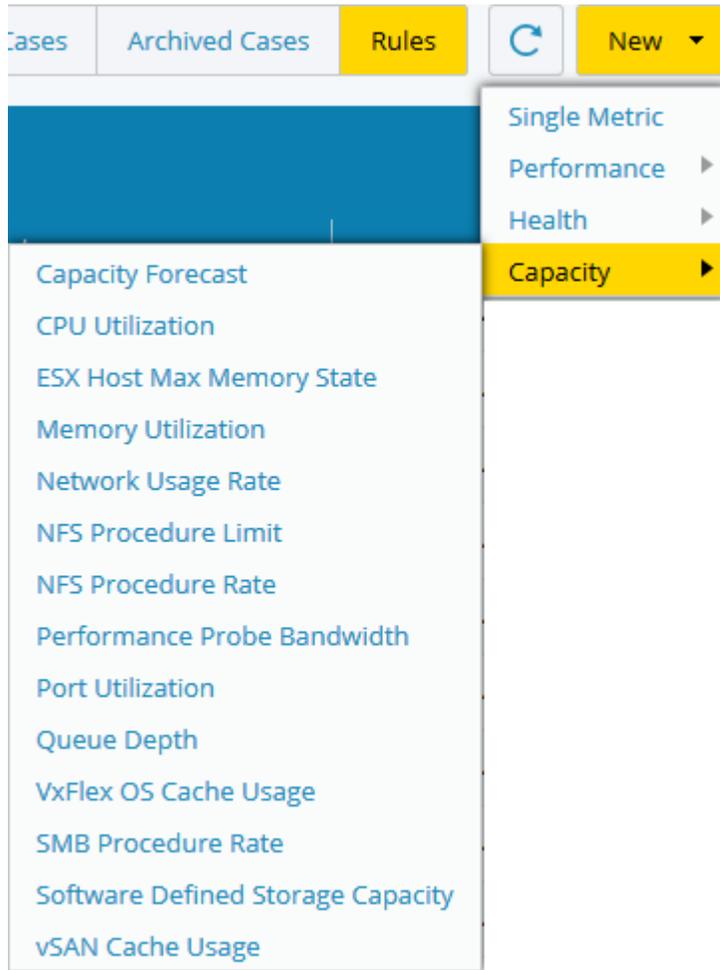
Performance alarms capture issues in the infrastructure that can impact the performance of applications and workloads, such as latency, flow control, CPU contention, and software-defined storage congestion.



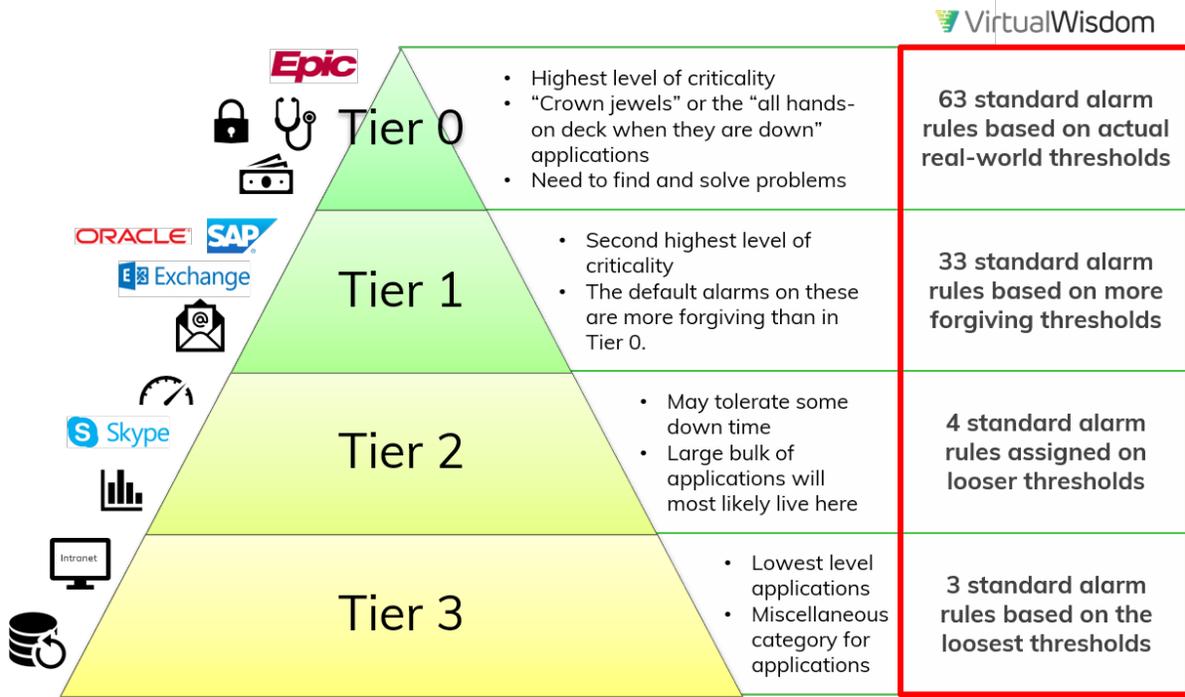
Health rules monitor and alert on common health issues in fibre channel and NAS infrastructure: physical layer errors, bad SCSI status, and communication errors.



Capacity rules capture problems involving utilization and capacity on multiple infrastructure types (Compute, Network, Storage). Capacity rules alert you when issues like high CPU utilization, port utilization, cache usage, or memory utilization occur.



Standard Rule Templates for Application Tiers



VirtualWisdom tiered alarms are designed to operate with your applications. Tier 0 alarms are configured to work with your most critical applications to find and solve problems. The lower tiers utilize more forgiving thresholds and are designed for your less critical applications.

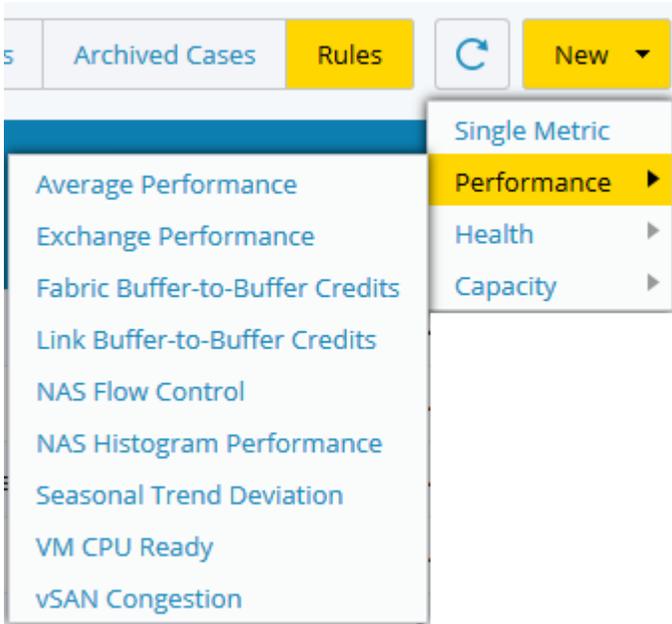
Each of these tiers and their associated alarms can be tailored for your specific requirements. Your Virtana Services team can work with you to determine which applications should be placed into which tiers, and to tailor and configure your applications tiers and their associated rules.

Standard Alarm Rule Templates

VirtualWisdom includes standard templates for configuring alarm rules. These rules are triggered when certain conditions are met. You can use the standard rule templates to quickly identify and set up alarms based on common infrastructure error conditions.

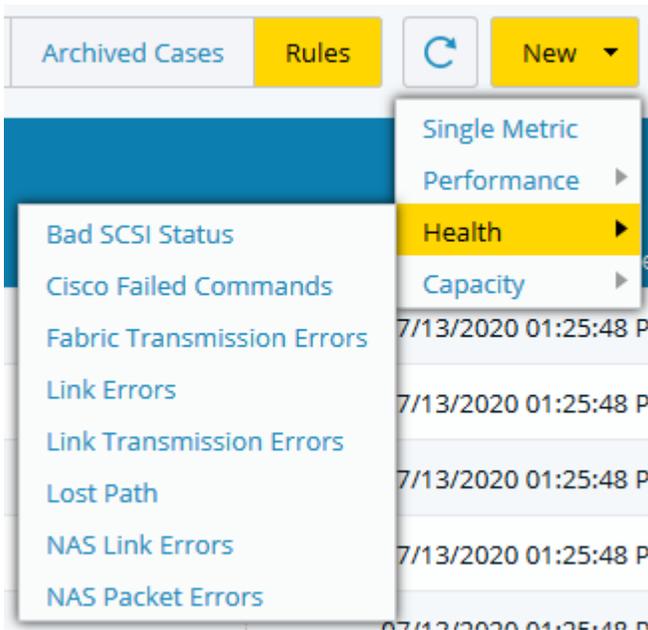
There are four types of alarm rule templates to choose from.

Performance Alarm Rule Templates



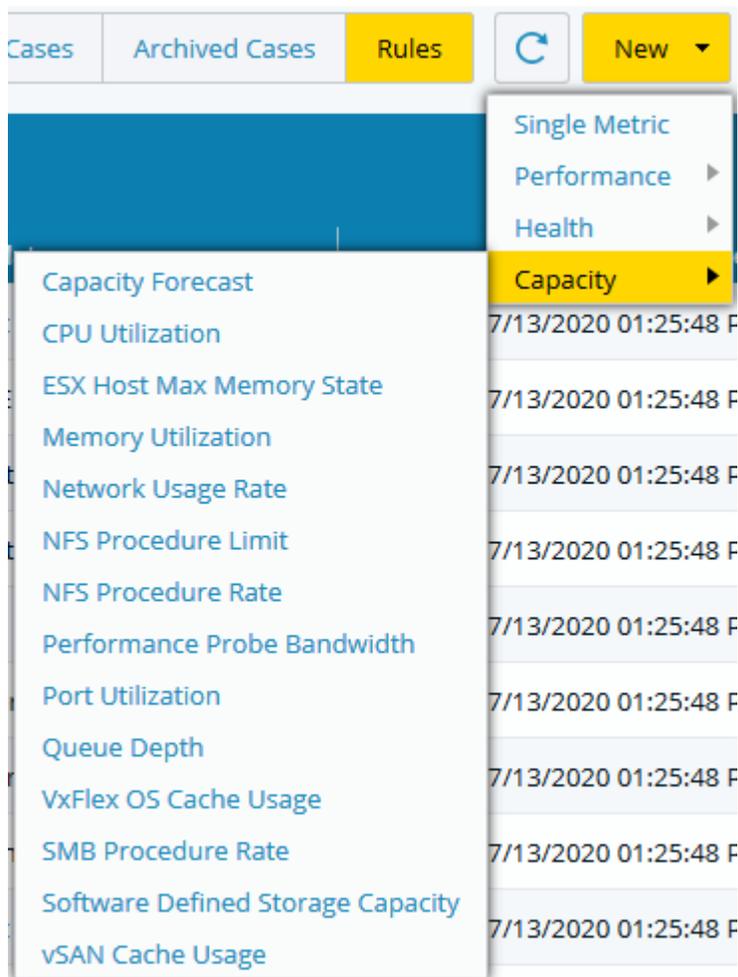
Performance alarm rules monitor the performance of your infrastructure and alert you when issues like flow control, high read response time, latency, and vSAN congestion occur.

Health Alarm Rule Templates



Health alarm rules monitor the health of your infrastructure and alert you when issues like bad SCSI statuses, transmission errors, lost path errors, failed commands, and packet errors occur.

Capacity Alarm Rule Templates



Capacity alarm rules monitor the capacity of your infrastructure and alert you when issues like high CPU utilization, high or low port utilization, cache usage, or high memory utilization occur.

Single Metric Alarms

VirtualWisdom also offers single metric alarms, which are alarms that are based on many of the single entity types and metrics available in your portal.

Configuring a Single Metric Alarm

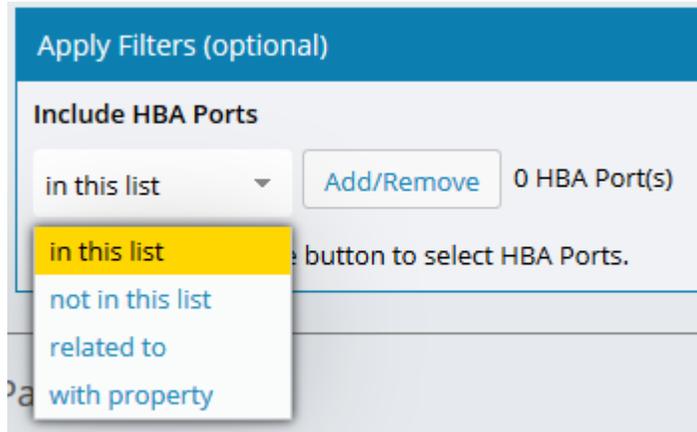
1. Enter a **Name** and **Description** for your alarm.
2. Select an entity type to alarm on. You can choose to alarm on all entities or only the entities assigned to a tier.

3. Apply a filter.
You can use entity filtering or data filtering. Refer to [Report Filtering \[208\]](#) for more information on entity and data filtering. You can apply a filter to the entity selection if you wish to alarm only on specific entities. Alarm filtering is similar to report entity filtering.

- a. To filter for specific entities, select the plus sign to add a filter, then Filter [Entity Type].



- b. You can filter for entities in/not in a list, related to an entity, or with a specified property.



4. Select a metric.

The screenshot shows the configuration interface for an alarm rule template. It is divided into three main sections: Entities, Parameters, and Notifications. In the Entities section, the 'Alarm on' dropdown is set to 'All', and there is an 'Apply Filters (optional)' button. The Parameters section includes 'Alarm when' with a '[Select Metric]' dropdown, 'Duration' set to '1 min', and 'Severity' set to 'Critical'. A modal window is open over the '[Select Metric]' dropdown, showing a search bar and a list of metrics under the 'Storage' category. The metrics listed are: SAN, FC Switch Integration, FC-SCSI, iSCSI, FCoE-SCSI, Cisco STS, NAS, Link Layer, and Case Management. At the bottom of the interface, there are checkboxes for 'Enable Rule' (checked) and 'Enable SNMP' (unchecked).

If you selected **Application** as the entity type, you can choose to show all metrics or to show only the metrics applicable to the Application entity type. This reduces the selection to only those entities that are currently assigned to at least one application and hides unassigned entities from the list. The number of applicable metrics is displayed in the selection modal.

Application by [Select Metric]

Show All Metrics (96 of 2144 apply for this type)

- ▶ Case Management
- ▶ Performance
- ▶ Capacity



TIP

If you plan to add entity types to an application at a later time and want to configure your report, dashboard, or alarm to include metrics from those entity types in advance, check the **Show All Metrics** box so they are available for selection.

5. Configure the alarm parameters: threshold, **Duration**, and **Severity**.

Parameters

Alarm when is

Duration

Severity

6. Configure notifications.

You can set up a notification plan for the alarm rule to notify users when the alarm is triggered.

To add users to the notification plan, click the Add box and check the users to be added. You can also use an email distribution list for the notification plan.



NOTE

Available users are determined by the users listed in the LDAP and User Management sections of the Settings tab.

Notifications

Name	Email ↑	
vi.training (VI Training)	training@virtualinstruments.com	

Also Notify
john.smith@virtana.com

Enable Rule Enable SNMP

The **Also Notify** field provides the ability to email alarm notifications to people who are not registered users of VirtualWisdom. Unregistered users added to this field can only receive initial notifications about new cases, but do not receive any succeeding notifications, such as investigation updates.

Check the **Enable SNMP** box to send alarm details to a configured SNMP trap. You can use SNMP traps to automatically trigger actions based on your specific requirements. For more information on configuring SNMP traps, see the [SNMP Traps](#) topic in the Administering Your VirtualWisdom Portal section of the VirtualWisdom Administrator Guide.

Viewing Your Configured Alarms

You can view your configured alarm rules by selecting Rules from the top of any Alarms page.

Rules Open Cases Archived Cases **Rules** C New Help

Sort list by any field

Rule Name ↑	Description	Enabled	Tier	Severity	From Template	Last Modified
Tier 0 Compute ESX Host CPU Utiliza...	Monitors ESX Host CPU Utilization. Triggers if the CPU exc...				CPU Utilization	04/15/2020 06:30:40 PM PDT
Tier 0 Compute ESX Host Memory St...	Monitors the memory management conditions of an ESX...	Yes	Platinum	Critical	ESX Host Max Memory State	04/15/2020 06:30:41 PM PDT
Tier 0 Compute ESX Host Memory UT...	Monitors ESX Host Memory utilization. Triggers if the hos...	No	Platinum	Critical	Memory Utilization	04/15/2020 06:30:40 PM PDT
Tier 0 Compute ESX VM CPU Ready	Monitors ESX Virtual Machine CPU Ready. Triggers if the ...	Yes	Platinum	Critical	VM CPU Ready	04/15/2020 06:30:41 PM PDT
Tier 0 Compute ESX VM CPU Utilizati...	Monitors ESX Virtual Machine (VM) CPU Utilization. Trigge...	Yes	Platinum	Critical	CPU Utilization	04/15/2020 06:30:40 PM PDT
Tier 0 Compute HBA Flow Control	Monitors metrics gathered directly from Fibre Channel s...	Yes	Platinum	Critical	Fabric Buffer-to-Buffer Credits	04/15/2020 06:30:42 PM PDT
Tier 0 Compute HBA Utilization	Monitors HBA links for utilization by monitoring the amo...	Yes	Platinum	Warning	Port Utilization	04/15/2020 06:30:41 PM PDT
Tier 0 Compute HBAs Link Errors	Monitors HBA port links for link-level errors. Triggers if m...	Yes	Platinum	Critical	Link Errors	04/15/2020 06:30:40 PM PDT
Tier 0 Compute HBAs Transmission ...	Monitors HBA port links for fabric transmission errors. Tr...	Yes	Platinum	Critical	Fabric Transmission Errors	04/15/2020 06:30:41 PM PDT
Tier 0 Compute Transmission Errors	Monitors links for abort sequence frames. Triggers if a ho...	Yes	Platinum	Critical	Link Transmission Errors	04/15/2020 06:30:41 PM PDT
Tier 0 Destination IP Address Write P...	Monitors an iSCSI target for excessive write payload (thro...	Yes	Platinum	Warning	Single Metric	04/15/2020 06:30:40 PM PDT
Tier 0 Host Excessive Queue Depth	Detects hosts where maximum queue depth usage exce...	Yes	Platinum	Critical	Queue Depth	04/15/2020 06:30:41 PM PDT
Tier 0 Host Lost Path	Monitors host links: when a link has at least 5 KB/s traffic ...	No	Platinum	Critical	Lost Path	04/15/2020 06:30:41 PM PDT
Tier 0 Host Port Link Errors	Monitors host port links for link-level errors. Triggers if a...	Yes	Platinum	Critical	Link Errors	04/15/2020 06:30:40 PM PDT
Tier 0 Host Read Response Times	Monitors a host's read response time (the aggregated me...	Yes	Platinum	Critical	Exchange Performance	04/15/2020 06:30:41 PM PDT
Tier 0 Host Write Response Times	Monitors a host's write response time (the aggregated m...	No	Platinum	Critical	Exchange Performance	04/15/2020 06:30:41 PM PDT
Tier 0 Isilon Cluster Disk Unavailable...	Dell EMC recommends to not let a single Isilon Cluster co...	Yes		Warning	Single Metric	04/15/2020 06:30:41 PM PDT
Tier 0 Isilon Cluster System CPU Usa...	Depending on your sales cycle, a cluster that is operating ...	Yes	Platinum	Critical	Single Metric	04/15/2020 06:30:41 PM PDT
Tier 0 Isilon Disk Usage exceeds 85%	A general rule of thumb is to keep disk utilization levels b...	Yes	Platinum	Critical	Single Metric	04/15/2020 06:30:41 PM PDT

Drill down to view configured rule

A list of configured rules is displayed, sorted by Last Modified Date. You can change the sort field by clicking on any of the field headers.

Drill down on a row to view the configured alarm.

Tier 0 Compute ESX Host CPU Utilization Help

Rule Name * Description

Entities

Alarm on Platinum ESX Host Tier and entity being monitored

Apply Filters (optional) 0 Filter(s)

Parameters

Alarm when

CPU Utilization exceeds 70 % Threshold, duration, and alarm severity parameters

Duration 5 mins

Severity Critical

Notifications

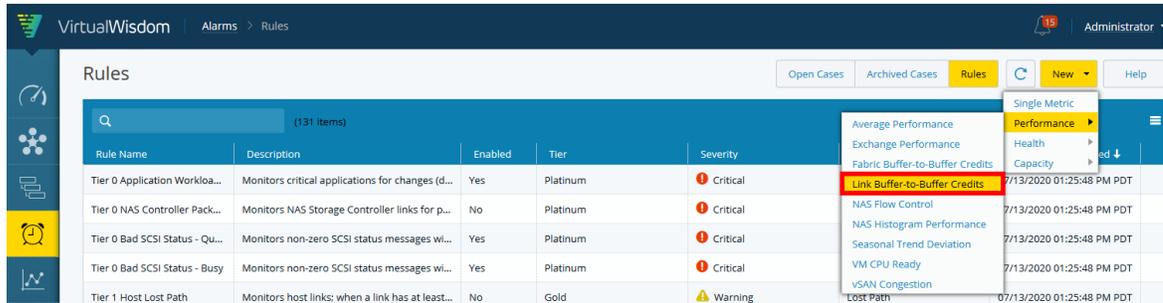
Enable Rule Enable SNMP Rule status and notification plan

Save Close Delete

The configured rule page shows you the tier and entity being monitored for an alarm situation, the threshold, duration and severity of the alarm, and the status (enabled/not enabled) and the notification plan set up for the alarm.

Configuring an Alarm Rule Template

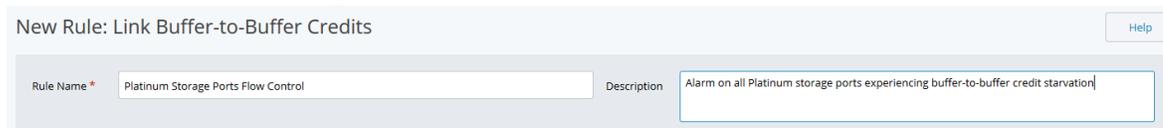
1. Select an alarm rule from the list of standard alarm rule templates.



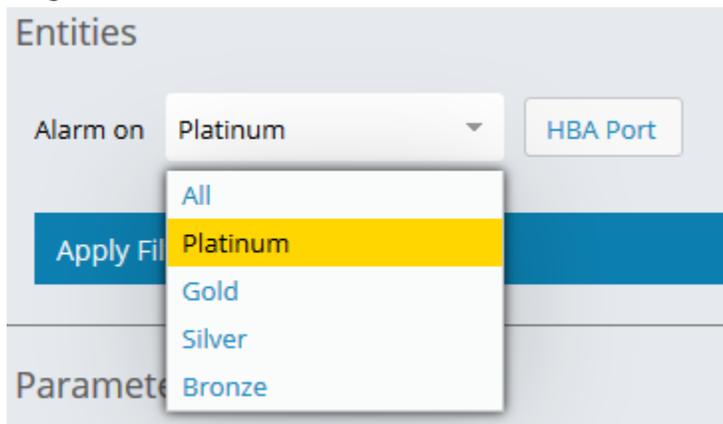
2. The **New Rule** screen is displayed. All rule screens have the following common areas:

- Name / Description
- Entities
- Parameters
- Notifications / Enable

Enter a name and description for the new alarm.



3. If tiers have been configured, you can choose whether to alarm on all tiers, or just a single tier.



**NOTE**

Alarms default to all tiers. You must select a tier to override this.

4. Choose the entity type to be observed by the alarm rule.

Entities

Alarm on All ▼ **HBA Port**

You can use the search field to find a specific entity type.

HBA Port

storage port

- ▼ Storage
 - ▼ SAN
 - Storage Port

**NOTE**

Only the entity types appropriate to the alarm rule are displayed.

5. **Applying Filters**

You can use entity filtering or data filtering. Refer to [Report Filtering \[208\]](#) for more information on entity and data filtering. You can apply a filter to the entity selection if you wish to alarm only on specific entities. Alarm filtering is similar to report entity filtering.

- a. To filter for specific entities, select the plus sign to add a filter, then Filter [Entity Type].

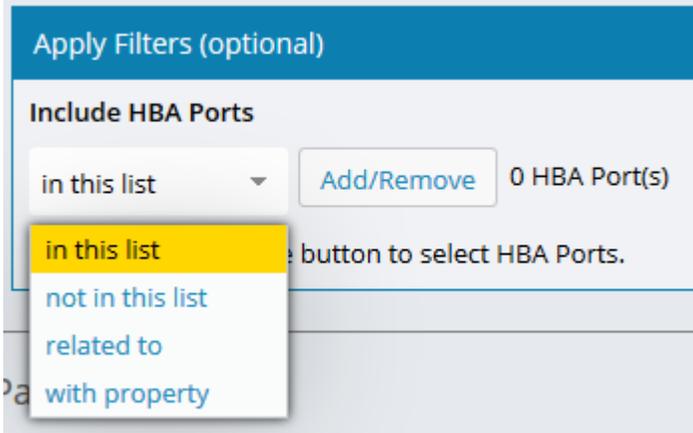
Apply Filters (optional) 0 Filter(s) +

Parameters

Add Filter:

- Filter HBA Ports
- Filter Entities for Metric Calculation (Advanced)

- b. You can filter for entities in/not in a list, related to an entity, or with a specified property.



6. Configure the alarm rule parameters. In the parameter section, set the threshold(s), duration, and severity. Thresholds may differ for each alarm rule but there is always a duration and severity level for each alarm rule template.

Parameters

Alarm when

Switch Zero Buffer-to-buffer Credits exceeds  50 %

Device Zero Buffer-to-buffer Credits exceeds  50 %

Duration

Severity  Critical



NOTE

The **Parameters** area on each page is unique to the rule template.

7. Configure notifications. You can set up a notification plan for the alarm rule to notify users when the alarm is triggered.

To add users to the notification plan, click the Add box and check the users to be added. You can also use an email distribution list for the notification plan.



NOTE

Available users are determined by the users listed in the LDAP and User Management sections of the Settings tab.

Name	Email ↑	Add
vi.training (VI Training)	training@virtualinstruments.com	

Also Notify
john.smith@virtana.com

Enable Rule Enable SNMP

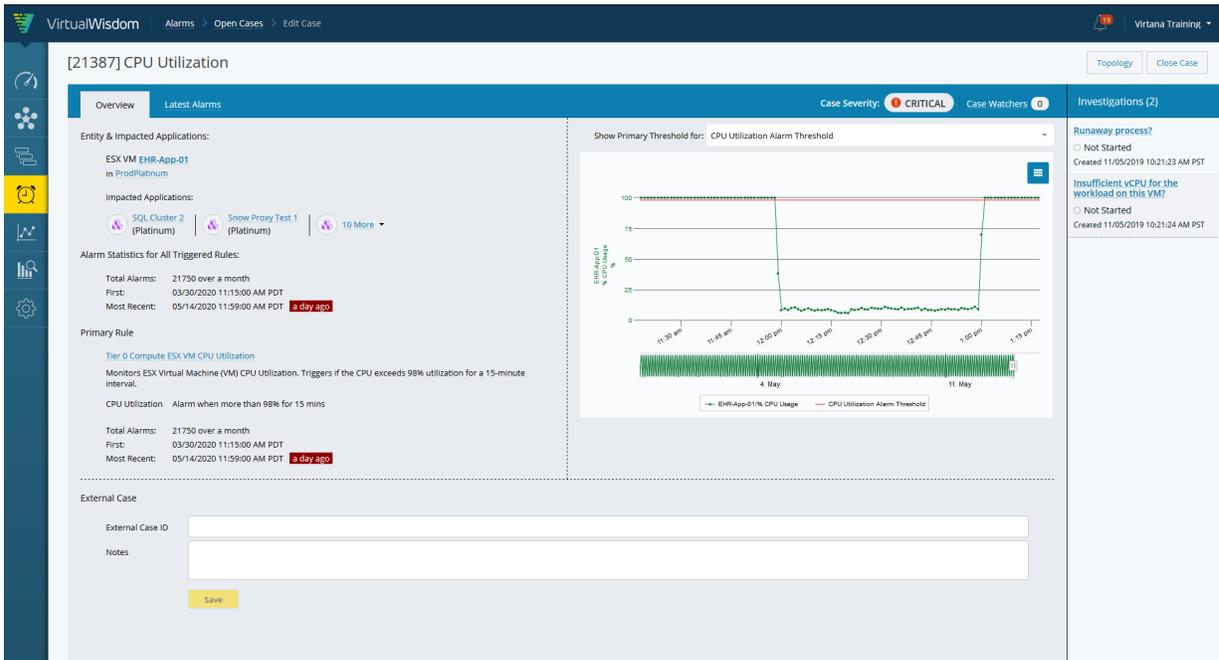
Save Cancel

The **Also Notify** field provides the ability to email alarm notifications to people who are not registered users of VirtualWisdom. Unregistered users added to this field can only receive initial notifications about new cases, but do not receive any succeeding notifications, such as investigation updates.

8. Check the **Enable SNMP** box to send alarm details to a configured SNMP trap. You can use SNMP traps to automatically trigger actions based on your specific requirements. For more information on configuring SNMP traps, see the [SNMP Traps](#) topic in the Administering Your VirtualWisdom Portal section of the VirtualWisdom Administrator Guide.

Cases

Cases are used to group alarm events.



Cases let you handle issues that arise and are identified through alarms that are triggered based on the defined alarm rules. Cases are opened when an alarm is triggered and are comprised of groupings of alarms to minimize the noise-to-signal ratio.



View open cases by selecting the Open Cases button. A list of all open cases is displayed, sorted by the Most Recent Occurrence date.

Table 7. Open Case List View Fields

Column ID	Definition
Severity	Alarm Severity
Case ID	VirtualWisdom Case ID
Case Type	VirtualWisdom Case Type
Entity	Entity affected by the case
Entity Type	Entity type affected by the case
External Case ID	Text entry box for tracking a ticket number from an external ticketing system
Open On	Date and time case opened

Column ID	Definition
Most Recent Occurrence	Last date and time case updated

You can sort the list view by these fields: Severity, Case ID, Case Type, Open On, and Most Recent Occurrence. Click the down arrow to expand the row and view description and number of occurrences. Drill down on a row to view the open case.

The screenshot shows the 'Open Cases' interface in VirtualWisdom. At the top, there are navigation buttons: 'Open Cases' (highlighted in yellow), 'Archived Cases', 'Rules', 'Close Cases', and 'Help'. Below the navigation is a search bar and a 'Sort by these fields' instruction. The main table lists cases with columns for Severity, Case ID, Case Type, Entity, Entity Type, External Case Id, Open On, and Most Recent Occurrence. A red box highlights the 'Most Recent Occurrence' column header. A red arrow points to the first row (Case ID 21387), with the text 'Drill down to view open case' next to it. Below this row, a sub-section titled 'Case Type: CPU Utilization' is expanded, showing a description and the number of occurrences (21750). A red arrow points to the first row of this expanded view, with the text 'Expand to see description and # of occurrences' next to it. The table contains various cases with different severities (Critical, Warning) and entities (EHR-App-01, EHR-DB-01, PURE-CT0-FC0, PURE-CT1-FC0).

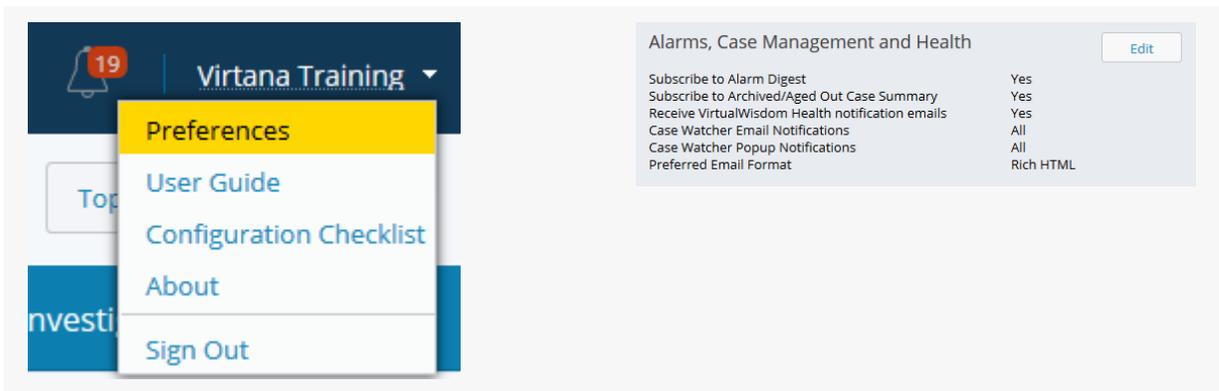
The Latest Alarms tab shows you a list of the most recent alarm events. It includes the alarm rule that triggered the alarm, the metric being monitored, the threshold, and the metric value that triggered the alarm.

[21387] CPU Utilization

Timestamp ↓	Rule triggering the alarm	Metric	Alarm threshold value	Measured value
05/14/2020 11:59:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:58:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:57:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:56:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:55:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:54:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:53:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:52:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:51:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:50:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:49:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:48:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:47:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:46:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:45:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:44:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:43:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:42:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:41:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:40:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100
05/14/2020 11:39:00 AM PDT	Tier 0 Compute ESX VM CPU Utilization	OS Instance CPU Utilization	98	100

Daily email summaries of open cases (the Alarm Digest) are sent to individuals with the vw-admin role every day at 5:00 AM. The email contains summary details of open cases and occurrence details for the last 24 hours. Closed cases are filtered from the view.

Email preferences for the Alarm Digest can be set in the user profile. Navigate to Preferences > Alarms, Case Management and Health.



You can set your alarm notification preferences in Account Information.

Open Case Page

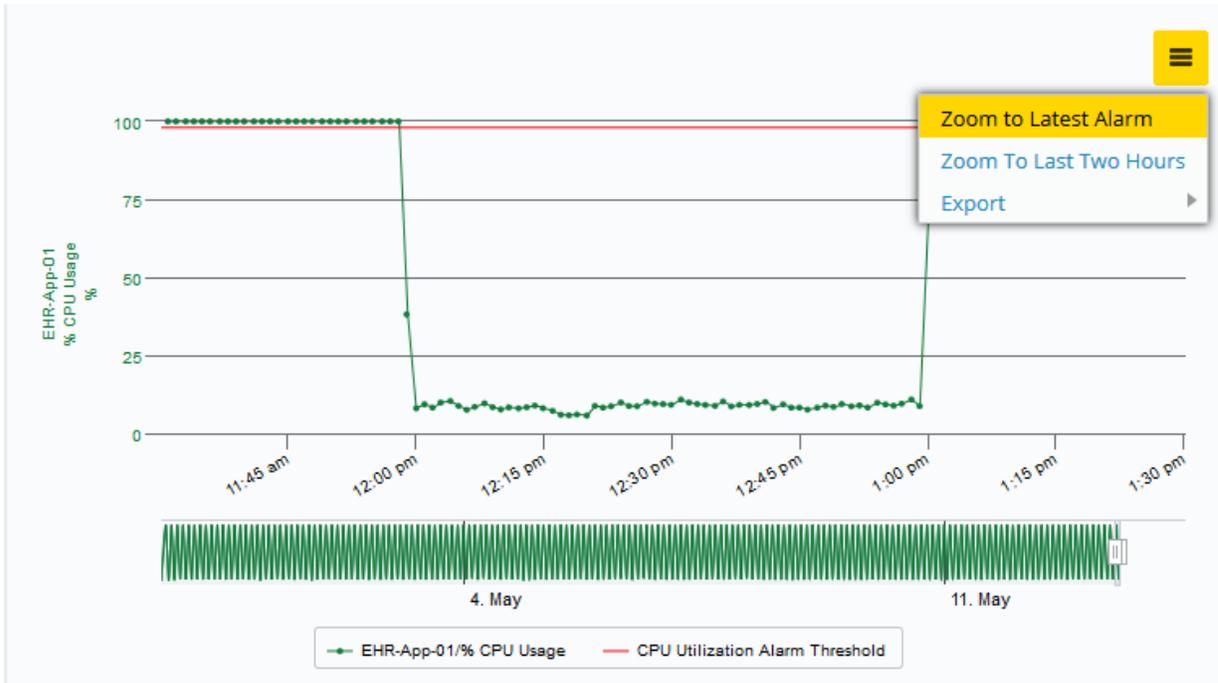
Overview Tab

The **Overview tab** displays summary information for the open case. It displays the impacted entity, applications and tiers. It shows you statistics such as total number of alarms and first and most recent alarm date and time. It also shows you the primary rule that triggered the alarm along with statistics for that alarm rule.

The screenshot shows the 'Overview' tab for a case titled '[21387] CPU Utilization'. The interface includes several sections:

- Entity & Impacted Applications:** Lists 'ESX VM EHR-App-01 in Prod/Platinum' as the impacted entity and 'SQL Cluster 2 (Platinum)' and 'Snow Proxy Test 1 (Platinum)' as impacted applications.
- Alarm Statistics:** Shows 'Total Alarms: 21750 over a month', 'First: 03/30/2020 11:15:00 AM PDT', and 'Most Recent: 05/14/2020 11:59:00 AM PDT' (marked as 'a day ago').
- Primary Rule:** 'Tier 0 Compute ESX VM CPU Utilization'. Description: 'Monitors ESX Virtual Machine (VM) CPU Utilization. Triggers if the CPU exceeds 98% utilization for a 15-minute interval.' Alarm condition: 'Alarm when more than 98% for 15 mins'. Statistics: 'Total Alarms: 21750 over a month', 'First: 03/30/2020 11:15:00 AM PDT', 'Most Recent: 05/14/2020 11:59:00 AM PDT' (marked as 'a day ago').
- Master/Detail Trend Charts:** A line chart showing 'EHR-App-01 CPU Usage' (green) and 'CPU Utilization Alarm Threshold' (red) over time. The chart shows a spike in CPU usage that crosses the 98% threshold, triggering the alarm. The x-axis shows dates from 4 May to 11 May.
- External Case:** A section for 'External case information, e.g., ServiceNow' with a text input field for 'External Case ID' and 'Notes', and a 'Save' button.
- Case Severity:** Indicated as 'CRITICAL'.
- Case Watchers:** A button with a user icon and a notification bell icon.
- Investigations (2):** A section on the right with options for 'Runaway process?' and 'Insufficient vCPU for the workload on this VM?'.

To the right, the **Master/Detail Trend Chart** displays a chart of the last alarm event and a chart below that displaying a 2-week summary of the data.



The Master/Detail Trend Chart shows you the alarm data for the last alarm event. Two hours of data are displayed by default. Under this chart is another chart that displays the last two weeks of alarm data to help you see when the issue started. You can also set the chart to show the latest alarm or the last two hours of data by clicking the hamburger icon and selecting **Zoom to Latest Alarm** or **Zoom To Last Two Hours**.

Use **Export** to export the chart data as a PNG, JPG, PDF, SVG, CSV or copy it to the clipboard.

At the top, the case severity is displayed and a link to see which users are watching the case. Click on Case Watchers to add your name or another user as a case watcher. Information on watched cases will appear in your notifications pane (click on alarm bell).

The **Topology** button takes you to a topology view of the impacted entity while the Close Case button lets you close the case.



On the far right is a pane that displays investigation workflows designed to help you troubleshoot the alarm.

Investigations (2)

[Runaway process?](#)

Not Started

Created 11/05/2019 10:21:23 AM PST

[Insufficient vCPU for the workload on this VM?](#)

Not Started

Created 11/05/2019 10:21:24 AM PST

The external case section at the bottom provides information on incident tickets that may have been opened in an external system. VirtualWisdom opens tickets in ServiceNow and displays the external case information here if ServiceNow has been integrated.

External Case

External Case ID

Notes

Archived Cases

An archived case is a case that was closed by the user or archived by the VirtualWisdom platform.



NOTE

The default archival interval for open cases (when the case is archived) is 21 days, and an email notification is automatically sent to all VirtualWisdom Administrators.

Table 8. Archived Case List View Fields

Column ID	Definition
Severity	Alarm Severity
Case ID	VirtualWisdom Case ID
Case Type	VirtualWisdom Case Type
Entity	Entity affected by the case
Entity Type	Entity type affected by the case
External Case ID	Text entry box for tracking a ticket number from an external ticketing system
Status	Archived Case Status (Archived, Closed)
Open On	Date and time case was opened
Archived On	Date and time case was archived or closed

You can sort the list view by these fields: Severity, Case ID, Case Type, Status, Open On, and Archived On. Drill down on a row to view the archived case.

Archived Cases (833 items)

Sort by these fields

Severity	Case ID	Case Type	Entity	Entity Type	External Case ID	Status	Open On	Archived On
Warning	51617	ISCSI Write Payload Rate	ISCSI_LIF_2-C10.20.10.88	IP Address	INC0962779	Archived	07/27/2020 08:40:00 PM PDT	09/01/2020 12:00:00 AM PDT
Critical	21539	Exchange Performance	Billing	Application	INC0490165	Archived	11/06/2019 10:45:00 AM PST	08/26/2020 12:00:00 AM PDT
Critical	21444	CPU Utilization	syslab-esx04.lab.vi.local	ESX Host	INC0489823	Archived	11/05/2019 05:56:00 PM PST	08/09/2020 12:00:00 AM PDT
Critical	40965	CPU Utilization	syslab-esx06.lab.vi.local	ESX Host	INC0821585	Archived	02/07/2020 11:32:00 AM PST	07/08/2020 12:00:00 AM PDT
Critical	47732	% CPU Ready	Ordering System	Application	INC0920152	Archived	06/06/2020 12:41:00 AM PDT	07/08/2020 12:00:00 AM PDT
Warning	15894	Cluster % /ifs Unavailable	pm-islion-sim	Islion Cluster		Closed	09/15/2019 03:00:00 PM PDT	06/22/2020 09:46:25 AM PDT
Critical	46713	Link Errors	vmx1955_105	Storage Port	INC0903134	Archived	05/17/2020 12:40:00 AM PDT	06/12/2020 12:00:00 AM PDT
Warning	46714	Link Errors	VMAX1955-PG1	Storage I/O M...	INC0903135	Archived	05/17/2020 12:40:00 AM PDT	06/12/2020 12:00:00 AM PDT
Critical	46710	Link Transmission Errors	syslab-esx05	Host	INC0903123	Archived	05/17/2020 12:40:00 AM PDT	06/12/2020 12:00:00 AM PDT
Critical	46712	Link Errors	vmx1955_104	Storage Port	INC0903133	Archived	05/17/2020 12:40:00 AM PDT	06/12/2020 12:00:00 AM PDT
Critical	46708	Link Transmission Errors	syslab-esx06	Host	INC0903121	Archived	05/17/2020 12:40:00 AM PDT	06/12/2020 12:00:00 AM PDT
Critical	46709	Link Transmission Errors	syslab-esx04	Host	INC0903122	Archived	05/17/2020 12:40:00 AM PDT	06/12/2020 12:00:00 AM PDT
Critical	44523	Capacity Forecast Alarm	syslab-esx01-local	ESX Datastore	INC0884399	Archived	04/10/2020 12:00:00 PM PDT	06/12/2020 12:00:00 AM PDT
Critical	42668	Memory Utilization	syslab-esx04.lab.vi.local	ESX Host	INC0862541	Archived	03/08/2020 05:03:00 PM PDT	06/03/2020 12:00:00 AM PDT
Critical	45903	% CPU Ready	Ordering System	Application	INC0897036	Archived	05/03/2020 02:09:00 AM PDT	06/03/2020 12:00:00 AM PDT
Critical	43974	Avg NFSv3 Write Completion Time	Order Manager	Application	INC0880064	Archived	03/31/2020 04:28:00 PM PDT	05/23/2020 12:00:00 AM PDT
Critical	44425	Average Performance	10.20.14.107	IP Address	INC0883772	Archived	04/08/2020 04:10:00 PM PDT	05/23/2020 12:00:00 AM PDT
Critical	44424	Average Performance	10.20.14.104	IP Address	INC0883771	Archived	04/08/2020 04:10:00 PM PDT	05/23/2020 12:00:00 AM PDT
Critical	44423	Average Performance	10.20.14.106	IP Address	INC0883770	Archived	04/08/2020 04:10:00 PM PDT	05/23/2020 12:00:00 AM PDT
Critical	44419	Average Performance	Win10Ent_Mgmt_Vttoo:4000:10.20.12.69	IP Address	INC0883744	Archived	04/08/2020 03:33:00 PM PDT	05/23/2020 12:00:00 AM PDT

Archived cases cannot be deleted using the UI. Archived cases age out automatically after which they are no longer visible on the Archived Cases page.

Archived Case Page

Overview Tab

The **Overview tab** displays summary information for the archived case. It displays the impacted entity, applications and tiers. It shows you statistics such as total number of alarms and first and most recent alarm date and time. It also shows you the primary rule that triggered the alarm along with statistics for that alarm rule.

The screenshot shows the 'Overview' tab for a case titled '[22607] CPU Utilization (Closed)'. The case severity is 'CRITICAL' and there are 0 case watchers. The impacted entity is 'ESX Host syslab-esx05.lab.vi.local' in 'ProdPlatinum' tier. Impacted applications include 'eCommerce (Platinum)' and 'OrderProcessing_Conflict (Platinum)'. Alarm statistics show 18 total alarms over a year, with the most recent on 11/22/2019. The primary rule is 'Tier 0 Compute ESX Host CPU Utilization', which triggers if CPU exceeds 90% utilization for a 15-minute interval. The external case ID is INC0534858, created by VirtualWisdom. On the right, there are investigation workflows, including one for 'vSphere cluster imbalanced in CPU utilization?' which is closed and unrelated.

At the top, the case severity is displayed and a link to see which users are watching the case.

The **Topology** button takes you to a topology view of the impacted entity while the Close Case button lets you close the case.

The external case section at the bottom provides information on incident tickets that may have been opened in an external system. VirtualWisdom opens tickets in ServiceNow and displays the external case information here if ServiceNow has been integrated.

On the right is a pane that displays the investigation workflows that may have been used to troubleshoot the alarm.

Investigations (2)

[vSphere cluster imbalanced in CPU utilization?](#)

Closed - Unrelated

Closed 11/13/2019 06:16:12 PM PST
a year ago

[Are there VMs on this ESX host that have runaway processes?](#)

Not Started

Created 11/13/2019 06:12:45 PM PST

The Latest Alarms tab shows you a list of the most recent alarm events. It includes the alarm rule that triggered the alarm, the metric being monitored, the threshold, and the metric value that triggered the alarm.

[22607] CPU Utilization (Closed)

Overview		Latest Alarms		Case Severity: CRITICAL	Case Watchers: 0
Timestamp ↓	Rule triggering the alarm	Metric measured	Alarm threshold value	Measured value	
01/14/2021 08:16:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	70.396	
01/14/2021 08:15:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	72.5	
01/14/2021 08:14:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	71.852	
01/14/2021 08:13:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	75.604	
01/14/2021 08:12:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	70.984	
01/13/2021 06:18:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	72.006	
01/13/2021 06:17:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	74.172	
01/13/2021 06:16:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	76.884	
01/13/2021 06:15:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	78.632	
01/13/2021 06:14:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	77.89	
01/13/2021 06:13:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	77.974	
01/13/2021 06:12:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	75.104	
01/13/2021 06:11:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	73.742	
01/13/2021 06:10:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	73.648	
01/13/2021 06:09:00 PM PST	Tier 0 Compute ESX Host CPU Utilization	ESX Host CPU Utilization	70	73.548	

Investigations

Investigations help guide you through a troubleshooting exercise. Each investigation is designed to look for one root cause, and the title should be answered by the content/troubleshooting within.

Investigations are associated with VirtualWisdom standard alarm rules and are unique to the entity type and metric defined for the alarm rule.

Investigations are found on the Open Case page in a panel on the right side of the page.

[21387] CPU Utilization

Case Severity: **CRITICAL** Case Watchers: 0

Investigations (2)

Entity & Impacted Applications:
ESX VM EHR-App-01
In ProdPlatinum

Impacted Applications:
SQL Cluster 2 (Platinum) | Snow Proxy Test 1 (Platinum) | 10 More

Alarm Statistics for All Triggered Rules:
Total Alarms: 21750 over a month
First: 03/30/2020 11:15:00 AM PDT
Most Recent: 05/14/2020 11:59:00 AM PDT **a day ago**

Primary Rule
Tier 0 Compute ESX VM CPU Utilization
Monitors ESX Virtual Machine (VM) CPU Utilization. Triggers if the CPU exceeds 98% utilization for a 15-minute interval.
CPU Utilization Alarm when more than 98% for 15 mins
Total Alarms: 21750 over a month
First: 03/30/2020 11:15:00 AM PDT
Most Recent: 05/14/2020 11:59:00 AM PDT **a day ago**

External Case
External Case ID:
Notes:
Save

Show Primary Threshold for: CPU Utilization Alarm Threshold

Graph: EHR-App-01 % CPU Usage vs. CPU Utilization Alarm Threshold

Investigations (2):
Runaway process?
- Not Started
Created 11/05/2019 10:21:23 AM PST
Insufficient vCPU for the workload on this VM?
- Not Started
Created 11/05/2019 10:21:24 AM PST

The Investigations panel lists all available investigations for an open case and provides the investigation name, when it was created, and its status. Click on the investigation name to open the investigation window.

Investigations (2)

[vSphere cluster imbalanced in CPU utilization?](#)

Closed - Pertinent

Closed 03/21/2020 07:04:13 PM PDT
2 months ago

[Are there VMs on this ESX host that have runaway processes?](#)

Not Started

Created 11/05/2019 06:02:29 PM PST

There are two types of investigations: **Automatic** and **Manual**.

- An automatic investigation performs root-case analysis using Analytics like **VM Coordinator** or **Trend Matcher** and provides solutions to the issue being investigated. The results of the analysis are displayed in the investigation window. In some cases, automatic investigations recommend that you run additional analytics or perform specific actions to resolve the issue.

Closed - Pertinent
Closed 03/21/2020 07:04:13 PM PDT; 2 months ago

vSphere cluster imbalanced in CPU utilization?

Case: [21444, CPU Utilization](#)
 ESX Host: [syslab-esx04.lab.vi.local](#)

Automatic investigation

VirtualWisdom 03/21/2020 07:02:29 PM PDT

Common causes of High ESX Host CPU Utilization

- An imbalanced vSphere Cluster, where multiple CPU-intensive VMs are hosted on a single ESX Host
- One or more VMs have runaway processes consuming CPU resources

VirtualWisdom 03/21/2020 07:02:29 PM PDT

Updated status: Active Investigation

Analyzing allocation of VMs across the cluster ...

VirtualWisdom 03/21/2020 07:04:13 PM PDT

Updated status: Closed - Pertinent

Ran [VM Coordinator](#) to see if some VMs can be moved to another host in the cluster.

Ran VM Coordinator

VM Coordinator Recommendation

Move 5 VMs

10/29/2019 7:02 pm to 11/05/2019 6:02 pm PDT [view analysis](#)

VMs to Move	From	To
SYSLAB-VC2	syslab-esx06.lab.vi...	syslab-esx04.lab.vi...
VW-PM-Master	syslab-esx04.lab.vi...	syslab-esx06.lab.vi...
ERP-ETL-DB-01	syslab-esx05.lab.vi...	syslab-esx04.lab.vi...
FinanceJMS	syslab-esx06.lab.vi...	syslab-esx04.lab.vi...
supply-warehouse-etl	syslab-esx04.lab.vi...	syslab-esx06.lab.vi...

- Manual investigations provide guidance on the steps required to discover an issue's root cause. Background on the issue, common causes, relevant charts, and suggested steps are displayed in the manual investigation window.

○ Not Started Created 11/05/2019 06:02:29 PM PST ✕

Are there VMs on this ESX host that have runaway processes?

Case: [21444, CPU Utilization](#)

ESX Host: [syslab-esx04.lab.vi.local](#)

Manual investigation



VirtualWisdom 11/05/2019 06:02:29 PM PST

Common causes of High ESX Host CPU Utilization

- An imbalanced vSphere Cluster, where multiple CPU-intensive VMs are hosted on a single ESX Host
- One or more VMs have runaway processes consuming CPU resources

How to determine if there are there VMs on this ESX host that have runaway processes

Observe the following chart which plots all the virtual machine's CPU utilization on this ESX Host over the last 24 hours

Provides troubleshooting steps



If the CPU utilization trend for any single VM appears to stay fixed at one CPU level over a long period of time (a day or more), this is usually an indication of a runaway process present on this VM.

How to resolve VMs with runaway processes

- Contact the server administrator to ensure the process is truly runaway.
- Terminate the runaway process.

Link to Analytics

Manual investigations may display a link to a VirtualWisdom Analytic for further troubleshooting and resolution steps. Click on the link to open the analytic.

○ Not Started

Created 12/05/2019 05:28:06 PM PST ✕

Is this queue depth setting too high?

Case: [29599, Queue Depth](#)Host: [SQL-DB-001](#)

VirtualWisdom 12/05/2019 05:28:06 PM PST

What is HBA queue depth?

HBA queue depth is a configurable parameter on the host side that determines the number of concurrent requests that can be outstanding to a single Target for a single LUN at one point in time. This concurrency has a significant impact on the overall performance of the storage system. Too few outstanding items and the storage device will not be utilized fully and work will not be completed as quickly as it could be. Too many outstanding items and the response time will degrade rapidly causing slow response time.

Little's Law shows us that any queued system where the input grows at a faster rate than the output, the time waiting in queue grows toward infinity. This can be translated to storage port queues, in that as the number of I/O requests in the queue increase, their respected wait times increase exponentially. In any exponential growth graph, there is always a dogleg curve. Points before the curve experience (in the storage port queue example) minimal increases to wait times as the queue grows. Points on or past this curve exhibit rapidly deteriorating response times as the queue continues to grow linearly.

Since the queue depth setting is per Target - LUN, if the queue depth is set to 32, and the host has 10 LUNs, then the HBA can issue 320 outstanding I/O requests to that port. If each HBA is zoned and masked to two storage ports, and the host has two HBAs then the single host could have 1280 outstanding requests to each storage array it attaches to. Often times, a storage port will service multiple hosts, each with multiple LUNs. The risk of overrunning the storage port and impacting response time is a very real concern in a SAN environment.

How to determine the optimal HBA queue depth setting

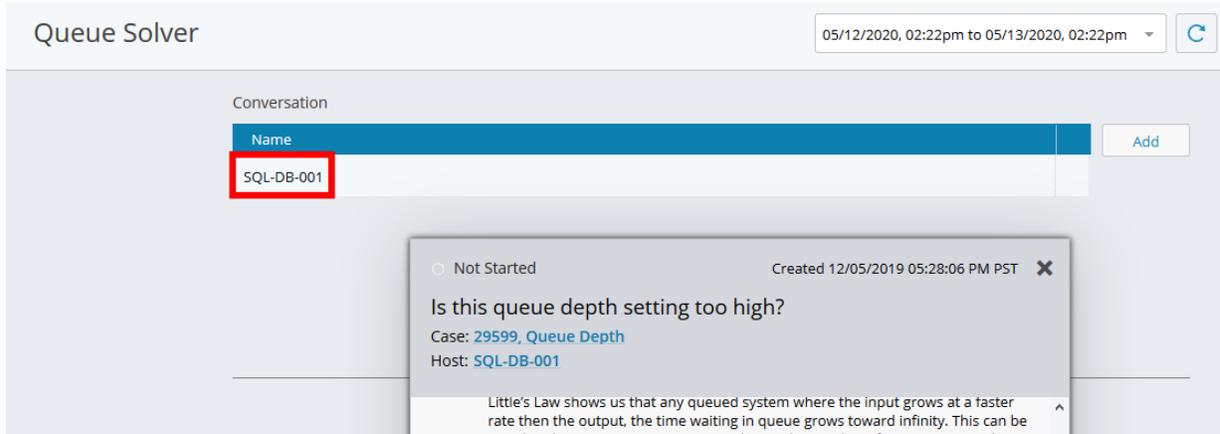
Run the Queue Solver analytic (button below) for the host (or hosts) zoned to the impacted storage array.

How to resolve a queue depth issue

The Queue Solver analytic will make suggestions on the proper HBA setting to optimize throughput and performance.

Queue Solver

The analytic window is displayed, and the analytic engine is populated with the entity associated with the case and investigation. The investigation window remains open as you navigate in the analytic and can be moved via drag/drop if necessary. You can also close the window using the x on the upper right corner.



Tracking Investigation Status

The status of an investigation can be tracked and changed through the troubleshooting process. The status is displayed at the top and bottom of the investigation window. Click on the down arrow next to the status to change the status.

Not Started Created 11/05/2019 06:02:29 PM PST ✕

Are there VMs on this ESX host that have runaway processes?

Case: [21444, CPU Utilization](#)
ESX Host: [syslab-esx04.lab.vi.local](#)

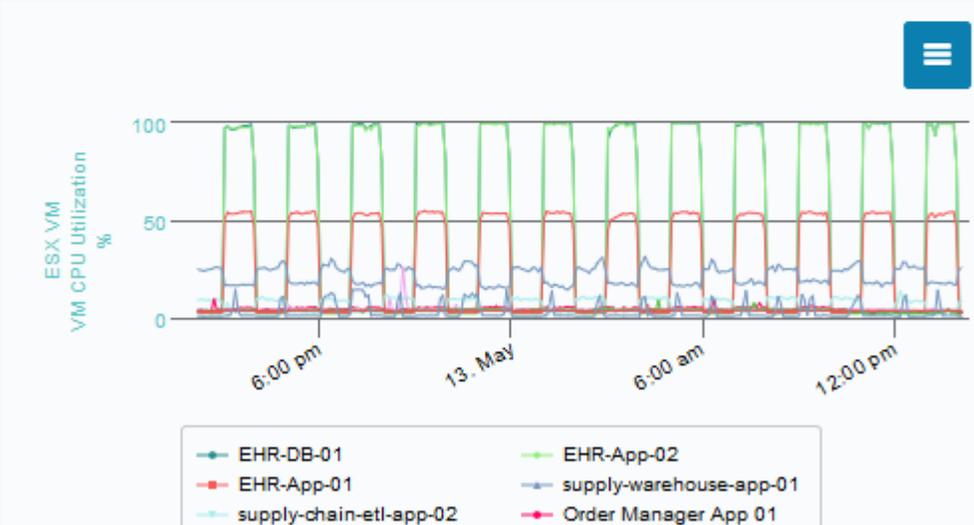
 **VirtualWisdom** 11/05/2019 06:02:29 PM PST

Common causes of High ESX Host CPU Utilization

- An imbalanced vSphere Cluster, where multiple CPU-intensive VMs are hosted on a single ESX Host
- One or more VMs have runaway processes consuming CPU resources

How to determine if there are there VMs on this ESX host that have runaway processes

Observe the following chart which plots all the virtual machine's CPU utilization on this ESX Host over the last 24 hours



The chart displays the CPU utilization percentage for several VMs over a 24-hour period. The y-axis represents 'ESX VM CPU Utilization %' from 0 to 100. The x-axis shows time from 6:00 pm to 12:00 pm on 13 May. The legend includes: EHR-DB-01 (green), EHR-App-02 (light green), EHR-App-01 (red), supply-warehouse-app-01 (blue), supply-chain-etl-app-02 (cyan), and Order Manager App 01 (magenta). EHR-App-02 shows the most significant utilization, with multiple peaks reaching 100%.

Set Status: Not Started Save

Choose from the following statuses:

- Not Started
- Active Investigation
- Waiting**
- Postponed
- Closed - Root Cause
- Closed - Pertinent
- Closed - Unrelated

Enter a note in the comments field then save your changes.

Set Status: Waiting ▼

Opened ticket in ServiceNow to download and runVM move script

Save

The activity is displayed in the Investigation pane.

 **vi.training** 05/13/2020 02:15:02 PM PDT
Updated status: Waiting
Opened ticket in ServiceNow to download and runVM move script

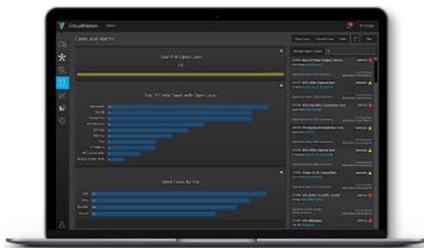
Infrastructure Assessment Workflow

You can use the following starting points to manage your infrastructure, identify and troubleshoot issues, manage capacity, balance workloads, and ensure availability:

1. [Topology \[361\]](#)



2. [Alarms \[364\]](#)



3. Analytics [365]

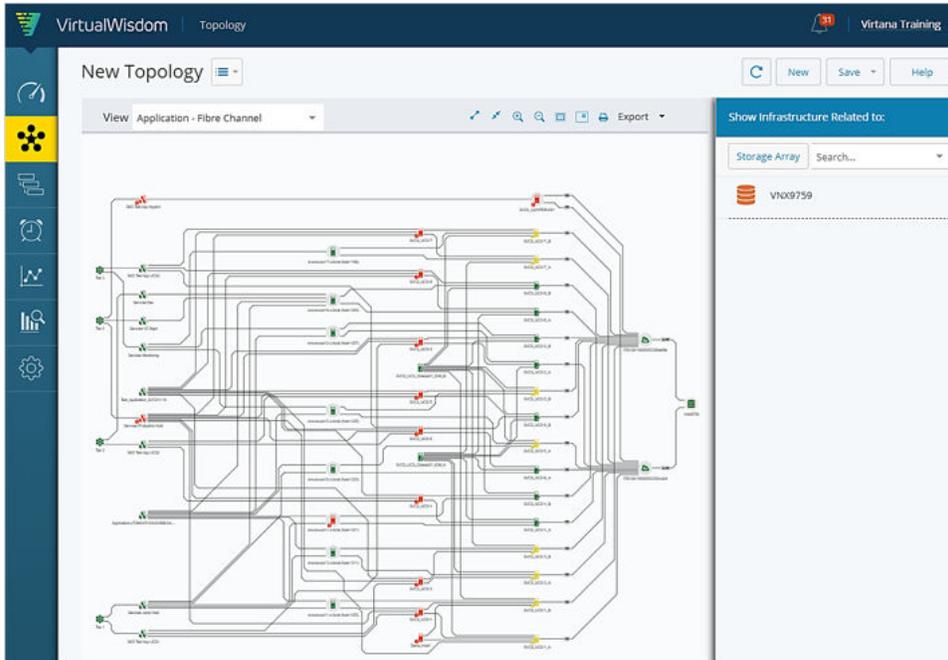


4. Dashboards [369]



Starting from Topology

Topology shows you the relationships between your entities, where there are problems, and lets you view basic metric trends.



Why start in Topology?

1. You are responsible for a defined infrastructure component, e.g., “these storage arrays, this ESX cluster”.
2. You prefer visual representations that show you the components rather than list views or report charts.
3. You care about what is going on in the infrastructure right now.

Getting started

If you're planning to start your workflow in Topology, it's recommended that you define Topology views that are aligned with the infrastructure you need to review.

Filter the view for a storage array

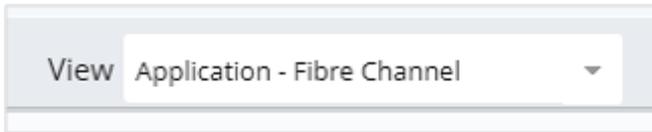
For example, as a Storage Administrator, you may need to determine if there are problems with a specific storage array.



From the Topology page, search for the storage array you want to filter for. The Topology view shows you the infrastructure related to that storage array.

Select an infrastructure type-specific view

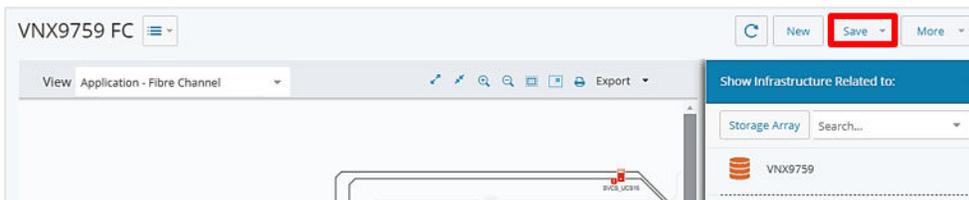
You may want to view only a specific infrastructure, e.g., "I only care about the Fibre Channel infrastructure supporting our applications":



In this case, you can select a view that shows you only that infrastructure type.

Save your topology

Once you have set your filters and view, save your Topology for future use:



Using the view to review your infrastructure

Now you can use the view to see where issues exist in your infrastructure:

Multiple storage ports have issues

Open Cases - VNX9759_A4

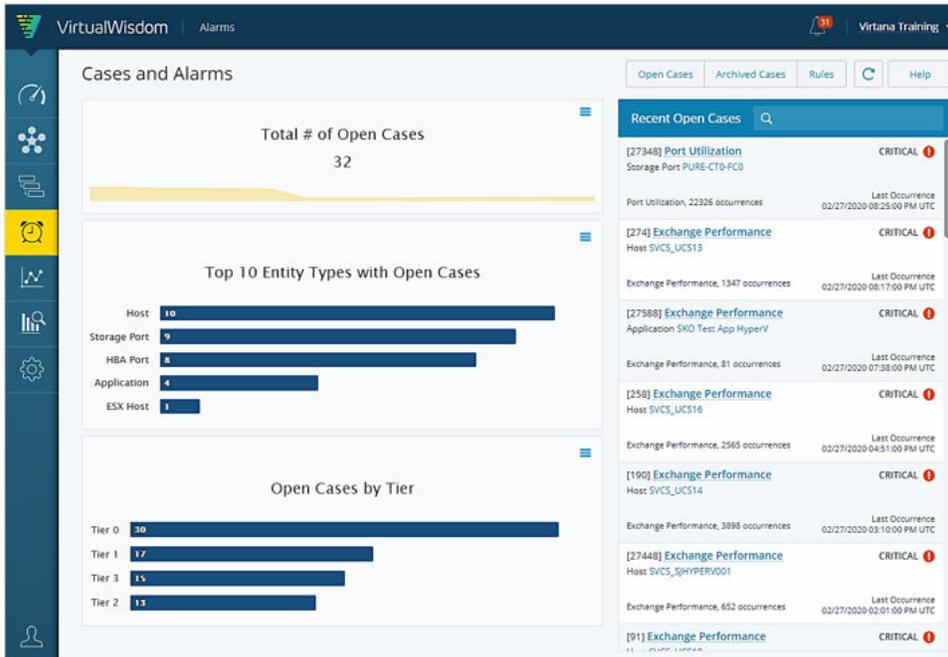
Rule Name	Entry Na...	Type	Occurren...	Last Occurrence
Port Utilization	VNX9759_A4	Port Utilization	139	02/27/2020 11:50:00 AM UTC
Link Buffer-to-buffer Credits	VNX9759_A4	Link Buffer-to-bu...	1	02/22/2020 11:27:00 AM UTC

Open Cases - VNX9759_A5

Rule Name	Entry Na...	Type	Occurren...	Last Occurrence
Port Utilization	VNX9759_A5	Port Utilization	93	02/27/2020 11:50:00 AM UTC
Link Buffer-to-buffer Credits	VNX9759_A5	Link Buffer-to-bu...	1	02/22/2020 11:27:00 AM UTC

Starting From Alarms

Alarms show you all the cases open against your infrastructure, organized by entity type, tier, and when it occurred.



Why start in Alarms?

- You have responsibility for a broad range of infrastructure, e.g., "I want to see all the alarms on all applications".
- You prefer a list view to a graphical view.
- You want to see historical data on alarms, or want to see the most recent alarms.

Getting started

The Alarms home page is useful for viewing alarms against entity types or viewing the most recent alarms.

- Drill down on an entity type to see all the alarms on that type.
- Use the Open Cases view to sort the alarms by severity, case type, entity and entity type, and when the case was opened.
- The search field can be used for filtering. Enter an entity type, entity name, case type, or severity.

Example: "Show me all cases based on the **Memory Utilization** rule type:

Open Cases

(2 items)

	Severity ↓	Case ID	Case Type	Entity
▶ <input type="checkbox"/>	! Critical	4262	Memory Utilization	vi-svcs-ucs12.vi.local
▶ <input type="checkbox"/>	! Critical	4196	Memory Utilization	vi-svcs-ucs16.vi.local

Review the open case details

Drill down to view the open case data, review topology, and conduct investigations.

[27592] Memory Utilization
Topology Close Case

Overview Latest Alarms ← **View alarm history**

Entity & Impacted Applications:

ESX Host vi-svcs-ucs14.vi.local (host-1221)
in SVCS_Prod2

Impacted Applications: ← **Impacted applications**

Services Jump Host (Tier 1)

Alarm Statistics for All Triggered Rules:

Total Alarms: 39 over an hour
First: 02/24/2020 04:20:00 PM UTC
Most Recent: 02/24/2020 05:18:00 PM UTC 3 days ago

Primary Rule

Tier 1 Compute ESX Host Memory Utilization
Monitors ESX Host Memory utilization. Triggers if the host exceeds 96% of its available memory for a 15-minute interval.

Memory Utilization Alarm when more than 96% for 15 mins

Total Alarms: 39 over an hour
First: 02/24/2020 04:20:00 PM UTC
Most Recent: 02/24/2020 05:18:00 PM UTC 3 days ago

External Case

External Case ID

Show Primary Threshold for: Memory Utilization Alarm Threshold

Investigations (2)

Is the vSphere Cluster balanced from a memory perspective?
 Not Started
Created 02/24/2020 04:26:16 PM UTC

Is there a VM Storm consuming excessive amounts of memory?
 Not Started
Created 02/24/2020 04:26:16 PM UTC

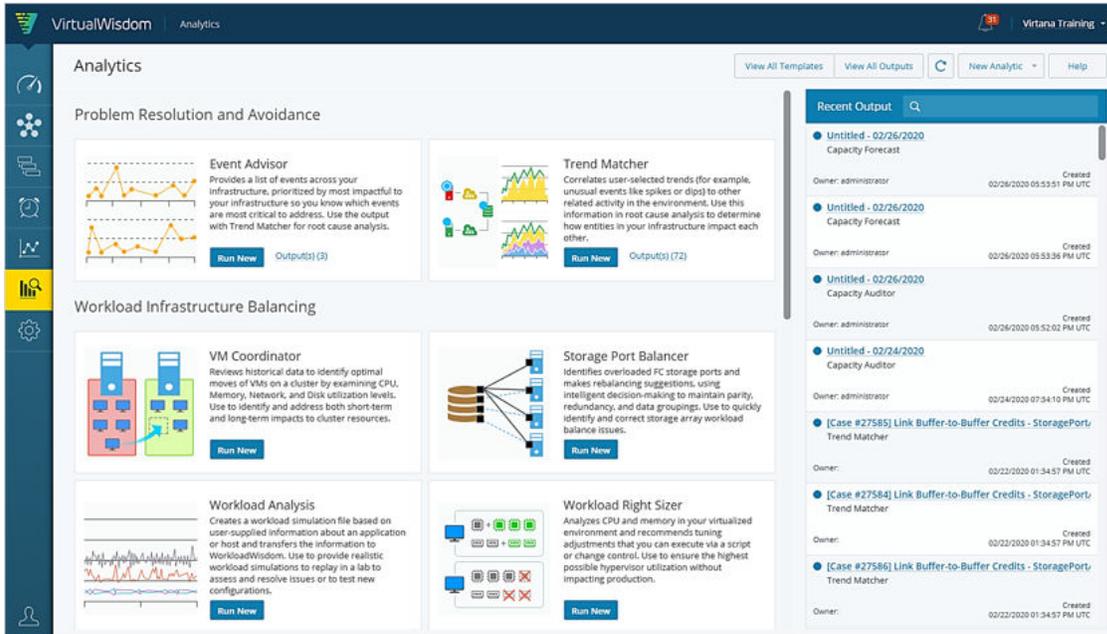
→ View topology

→ Use investigation to troubleshoot

→ View trends for the last event or the last 2 weeks

Starting From Analytics

Use Analytics to identify issues and assess infrastructure balance and capacity.



Why start with Analytics?

- You have a specific problem or situation in mind, e.g., multipathing failure, buffer credit starvation, workload right-sizing, optimal queue depth settings, capacity forecasting, etc.
- Alarms have not yet been configured in your VirtualWisdom environment.
- You want to run the analytic on a schedule and embed results in a report or dashboard.

Using Analytics to identify critical events and correlate with trends

Analytic	Objective
Event Advisor	<p>Use Event Advisor [259] to find events that may cause problems in your infrastructure:</p> <ul style="list-style-type: none"> • Slow draining devices • High utilization devices • Over-utilized virtualized resources • Aborts • Performance degradations
Trend Matcher	<p>Use Trend Matcher [261] to troubleshoot problems and perform root cause analysis:</p> <ul style="list-style-type: none"> • Accepts a source trend from Event Advisor or run standalone • Provides a topology view showing entities with correlating events and matching trends

Event Advisor Suggested Situations

Run Event Advisor for the past seven days. You can save as a template and set up a recurring schedule to run Event Advisor once a week.

Entities		
Type	Metrics	Filter
Storage Port	% Time at Zero Transmit Credits	

Buffer credit starvation on Fibre Channel storage ports

Entities		
Type	Metrics	Filter
HBA Port	Read Utilization	

High utilization on HBA and storage ports

Entities		
Type	Metrics	Filter
ESX VM	VM CPU Utilization	

High VM CPU utilization

Entities		
Type	Metrics	Filter
HBA Port	Abort Sequence Frames	

Entities		
Type	Metrics	Filter
Storage Port	Abort Sequence Frames	

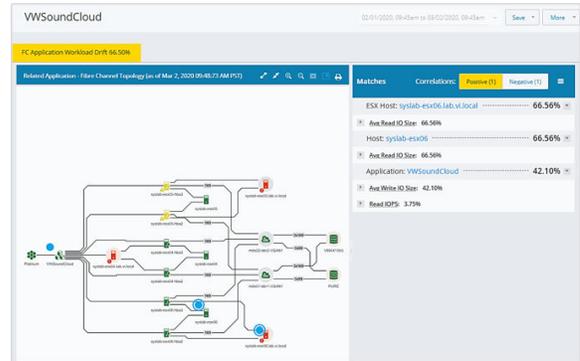
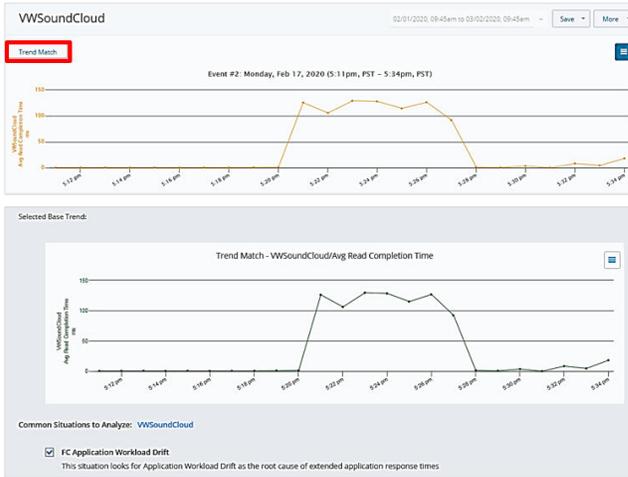
Aborts on Fibre Channel HBA ports and storage ports

Entities		
Type	Metrics	Filter
Application	Avg Read Completion Time	

Latency on the infrastructure supporting an application

Using Trend Matcher to discover correlating entities and metrics

After you've run Event Advisor, you can use Trend Matcher to find correlating entities and metrics.



Using other Analytics to assess your infrastructure

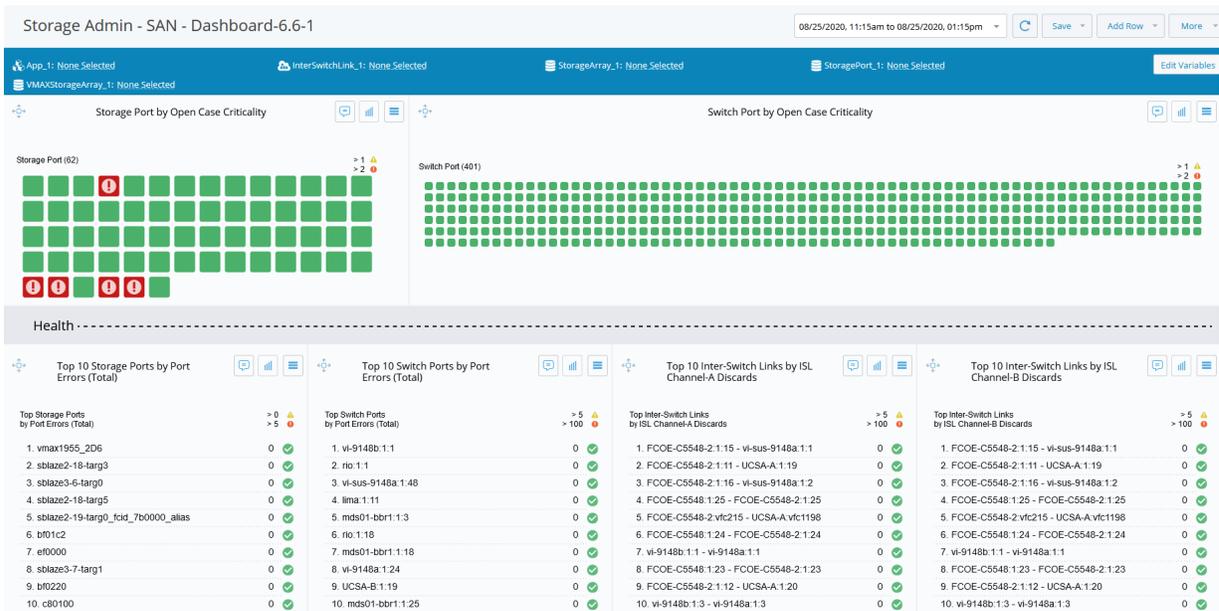
You can also use the following Analytics as part of a regular infrastructure assessment process:

Analytic	Objective
VM Coordinator	Use VM Coordinator [268] to identify optimal moves for VMs to avoid <ul style="list-style-type: none"> • Cluster degradation • Over-provisioning • Unnecessary rebalancing
Workload Right Sizer	Use Workload Right Sizer [278] to <ul style="list-style-type: none"> • Ensure the highest possible hypervisor utilization, without impacting production
Capacity Forecast	Use Capacity Forecast [297] to <ul style="list-style-type: none"> • Predict usage trends • Identify resource strain • Plan for growth
Capacity Auditor	Use Capacity Auditor [305] to <ul style="list-style-type: none"> • Identify where and when capacity adjustments should be made to VMAX and Isilon storage arrays

Analytic	Objective
Balance Finder	Use Balance Finder [313] to examine HBA throughput to <ul style="list-style-type: none"> • Verify host multi-pathing • Mitigate the risk of fabric failure • Increase availability and resiliency

Starting From Dashboards

Dashboards are designed to address the visibility requirements of a particular role, group, or infrastructure type.



Why start in Dashboards?

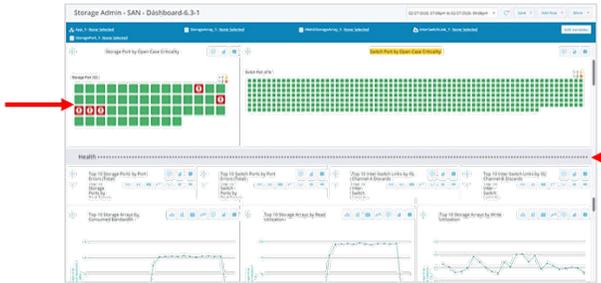
- You are responsible for the management and administration of a particular infrastructure types, e.g., SAN Administrator, NOC User/Operator
- You prefer visual representations that show you the current status of the components, as well as the historical data
- You care about what is going on in the infrastructure **now** but you also want the flexibility of being able to review past trends on the same infrastructure

Standard Admin Dashboards

Virtana Services has created a set of dashboards designed for use by infrastructure administrators. These admin dashboards can be used as a starting point for assessing your infrastructure, identifying issues, and performing troubleshooting exercises.

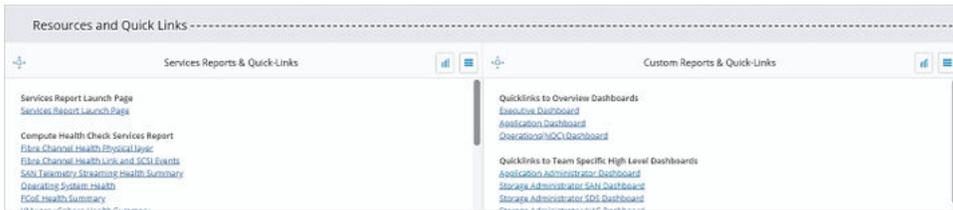
The dashboards are divided into sections that focus on infrastructure **health, utilization, and performance.**

Uses charts that show open case data.
Alarms must be configured for these charts to display data.



Divided into sections that focus on health, utilization, and performance

They also include a reference section with links to other Services reports and other dashboards.



Drill down on dashboard open case

Use the links to the infrastructure's open cases to drill down on identified issues.



The open case provides details and investigations to troubleshoot the issue.

The screenshot shows the 'Port Utilization' dashboard for case [27348]. The interface includes a top navigation bar with 'Overview' and 'Latest Alarms' tabs. The 'Latest Alarms' tab is highlighted in yellow and has an arrow pointing to it with the text 'View alarm history'. Below the navigation bar, the left sidebar contains 'Entity & Impacted Applications' (Storage Port PURE-CTD-FC0), 'Impacted Applications' (None), and 'Alarm Statistics for All Triggered Rules'. An arrow points to the 'Alarm Statistics' section with the text 'View impacted applications'. The main content area features a 'Show Primary Threshold for: Device Utilization Alarm Threshold' section with a line graph showing 'PURE-CTD-FC0 Real Utilization' and 'PURE-CTD-FC0 Time Utilization' over a period from 17 Feb to 24 Feb. A red arrow points to the graph with the text 'View trends for the last event or the last 2 weeks'. On the right side, there is an 'Investigations (4)' panel with several diagnostic questions and 'Not Started' buttons. A red arrow points to this panel with the text 'Use investigations to troubleshoot'. At the top right, a 'Topology' link is visible, with a red arrow pointing to it and the text 'View topology'. At the bottom, there is an 'External Case' section with input fields for 'External Case ID' and 'Notes', and a 'Save' button.

Contact Information

Sales Inquiries

To speak with a sales representative:

Complete the form at virtana.com/contact-us/.

Call us at +1-888-522-2557.

Support for VirtualWisdom Core and Integrations

VirtualWisdom support is available 24/7

Online Support

www.virtana.com/support

Technical Support

virtualwisdom.support@virtana.com

Weekend Severity 1 HOTLINE (For VirtualWisdom ONLY):

Toll Free: 1-888-988-9925

International: +1-408-579-4100

Feedback

We appreciate your input to help us improve the quality of our products and documentation. Send your suggestions, comments, and questions about Virtana products and documentation to:

Product Feedback

feedback@virtana.com

Documentation Feedback

techpubs@virtana.com

Legal

Copyright

Copyright © 2021 by Virtual Instruments Corporation (d/b/a Virtana). All rights reserved.

Virtual Instruments reserves the right to revise these specifications without notice or penalty.

Trademarks

Dell Technologies, Dell, EMC, Dell EMC, Isilon, Unisphere, VMAX, and other trademarks are trademarks of Dell Inc. or its subsidiaries.

IBM® and PowerVM® are registered trademarks of IBM Corporation in the United States, other countries, or both.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft®, Windows®, Windows Server®, and Hyper-V™ are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

NetApp®, OnCommand®, and ONTAP® are registered trademarks of NetApp, Inc., registered in the U.S. and/or other countries.

Oracle®, Java, and Solaris are registered trademarks or trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Pure Storage, the Pure Storage logo and the marks listed at <http://www.purestorage.com/legal/productenduserinfo.html> are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. or other countries.

VMware®, vCenter®, and vSphere® are registered trademarks of VMware, Inc. in the United States and other jurisdictions.

VirtualWisdom is a registered trademark of Virtual Instruments Corporation (d/b/a Virtana).