



**Hewlett Packard**  
Enterprise

**InfoSight User Guide for HPE Alletra 6000, Alletra 5000,  
HPE Nimble Storage**

# Legal Notices

Copyright © 2021 - 2021 by Hewlett Packard Enterprise Development LP

## Notices

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

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

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

Hewlett Packard Enterprise believes in being unconditionally inclusive. If terms in this document are recognized as offensive or noninclusive, they are used only for consistency within the product. When the product is updated to remove the terms, this document will be updated.

## Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

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

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

## Publication Date

Wednesday April 10, 2024 10:05:00

## Document ID

mtn1582564794393

## Support

All documentation and knowledge base articles are available on HPE InfoSight at <https://infosight.hpe.com>. To register for HPE InfoSight, click the *Sign up* link on the HPE Sign In page.

Email: [support@nimblestorage.com](mailto:support@nimblestorage.com)

For all other general support contact information, go to <https://www.hpe.com/us/en/services/nimble-storage.html>.



# Contents

<b>HPE InfoSight Portal Overview.....</b>	<b>6</b>
Register Your HPE Alletra 6000, Alletra 5000, HPE Nimble Storage Account.....	6
Set up Email Notifications.....	7
Select an Organization.....	7
Configure Your Home Page.....	8
Links Section.....	8
Landing Page Section.....	8
Bulletins Section.....	8
The Dashboards.....	9
Cross-Stack Performance Recommendations.....	10
Commonly Used HPE InfoSight Pages.....	13
The List Views and Details Pages.....	14
Data Availability FAQ.....	15
Data Availability Details.....	17
<b>Managing Your Arrays.....</b>	<b>18</b>
HPE InfoSight Roles.....	18
Configure Your Arrays.....	18
Configure Arrays to Send Data to HPE InfoSight.....	19
Use the GUI to Configure Arrays.....	19
Use the CLI to Configure Arrays.....	19
Plan Your Storage Resources.....	20
Add New Workloads.....	20
Scaling Workloads.....	21
Migrating Workloads.....	21
View Software Recommendations.....	22
Control Release Statuses for Software Recommendations .....	22
Wellness Issues and how to Manage Them.....	23
Configure Wellness Rules.....	23
Set-up Email Notifications for Alerts.....	24
Set a Snooze Period.....	24
<b>Displaying VMware Analytics.....</b>	<b>25</b>
Using VMware Virtualization Data to Evaluate Performance.....	25
Enable InfoSight to Collect VMware Information.....	25
Register the vCenter Plugin.....	26
Use the Array OS GUI to Register the vCenter Plugin .....	27
Use the Array OS CLI to Register the vCenter Plugin .....	27

Configure Arrays to Monitor your Virtual Environment.....	27
View Overprovisioned VMs.....	28
View VMs by IOPS and Latency.....	28
View VM Latency Due to the ESX Host.....	28
View Contentious VMs Based on the Datastore.....	28
View Contentious VMs Based on the ESX Host.....	29

## **Displaying Hyper-V Analytics.....30**

Analyzing Hyper-V Performance Issues.....	30
Enabling InfoSight to Collect Hyper-V Data.....	30
Key Data Provided by the Hyper-V Virtualization Pages.....	30
Hyper-V Cluster Analytics.....	31
Hyper-V Host Analytics.....	32
Hyper-V Storage Analytics.....	34
Hyper-V VM Analytics.....	36
Navigating Between Analytic Stacks .....	42
Displaying Data from Multiple Stacks.....	44
Current Behavior for Clustered/Standalone Hosts.....	44

## **Monitoring Your Array Environment.....46**

Monitor Equipment.....	46
Potential Impact Score.....	46
Monitor Volumes.....	46
Monitor Synchronously Replicated Volumes.....	46
The Executive Dashboard.....	47
Streaming Performance Telemetry Data.....	47
Enabling Stats Streaming.....	48
Viewing Arrays on which Stats Streaming is Enabled.....	49

## **Viewing dHCI Data.....50**

HPE InfoSight Analytics and dHCI .....	50
Enable InfoSight to Provide dHCI Information.....	50

## **Support of Your Arrays.....51**

Validate Contact and Install Location Information .....	51
Maintenance Windows.....	51
Schedule Maintenance Windows.....	52
Support Area Overview.....	52
Create a Support Case .....	53
Security of Customer-Originated Data .....	54
Information Security Practices .....	55
Core Data Dumps .....	55

Failed Disk Drives .....	56
Get Help.....	56

# HPE InfoSight Portal Overview

HPE InfoSight is a portal that enables you to manage, monitor, and troubleshoot HPE assets in various HPE [organizations](#). And while this document focuses on HPE Alletra 6000 and HPE Nimble Storage assets, you can access the documentation for other HPE products from the portal **Resources** menu, which you can access from the InfoSight Main Menu (☰)

Access to HPE InfoSight (available at <https://infosight.hpe.com>) is included in your support subscription package.



**Tip:** For an updated high-level introduction and review of HPE InfoSight, watch the [InfoSight Demo](#) on YouTube.

You can use the HPE InfoSight portal to perform the following tasks:

- Monitor your environment:
  - Manage storage health alerts and support cases
  - Analyze top causes of and possible corrections to latency issues
  - Monitor snapshots and replication
  - View data protection status and recovery point objective times
  - Monitor your virtual environment
  - Get a comprehensive overview on the dashboards
- Plan for efficient use of your infrastructure:
  - View capacity trends and forecasts
  - Assess CPU headroom on your storage assets and virtual infrastructure
  - View upgrade recommendations based on current utilization
  - Customize wellness summary emails
  - Customize array-originating events that require interaction with Support
  - Schedule maintenance windows
  - Plan array software updates
- Troubleshoot performance issues:
  - At the storage level, use detailed analysis of latency factors
  - At the per-VM level, troubleshoot with end-to-end visibility from the host to storage
  - Get an analysis of your virtual and physical infrastructure usage on a per-consumer basis through interactive heatmaps
- Access content:
  - View or download documentation
  - Manage issues and support cases
  - Download software and documentation

To watch a demo of HPE InfoSight, click the [InfoSight Demo](#) link on YouTube. To view a self-guided iterative demo of HPE InfoSight, click [HPE InfoSight Web Portal Demo](#).

## Register Your HPE Alletra 6000, Alletra 5000, HPE Nimble Storage Account

### Before you begin

You must have created an HPE account.



If you are not associated with any HPE Alletra 6000, Alletra 5000, HPE Nimble Storage account, you must be invited by another member of your organization. If you are a new customer and your email address is not recognized by our system, email [support@nimblestorage.com](mailto:support@nimblestorage.com) for assistance.

During the sales process, the following tasks are automatically completed for you:

- The HPE Alletra 6000, Alletra 5000, HPE Nimble Storage organization is created.
- The HPE Alletra 6000, Alletra 5000, HPE Nimble Storage system is added to the organization.
- An Admin user is created in the organization.

**Note:** The registration process is different for HPE Alletra 6000, Alletra 5000, HPE Nimble Storage than for other organizations. This following procedure is for HPE Alletra 6000, Alletra 5000, HPE Nimble Storage only.

### Procedure

1. Log into InfoSight at <https://infosight.hpe.com>.

If this is your first time logging in to HPE InfoSight, a verification email is sent to the email address associated with your HPE account. When you receive the verification email:

- 1 Follow the instructions in the verification email until the HPE InfoSight Terms of Use are displayed.
- 2 Accept the Terms of Use. The InfoSight Welcome page is displayed.

2. Either click **HPE Alletra 6000, HPE Nimble Storage Registration** on the InfoSight Welcome page or select **Settings > HPE Alletra 6000, HPE Nimble Storage > Customer Registration**.

The HPE Alletra 6000, Nimble Storage **Operational Dashboard** is displayed.

3. Add users to your organization.

For instructions on adding users, see [Associate New Users with an HPE InfoSight Account](#)

## Set up Email Notifications

### Procedure

Setup email notification for alerts.

- a) Click **Dashboards > Wellness**
- b) Click the Daily Summary Emails button on right-hand side of screen.
- c) Enter an email address for delivery of the email alerts.
  - 1 Click Add to add the email address as a recipient for the alerts.
  - 2 (Optional) Click **Test** to send a test email to the email address that you indicated.

## Select an Organization

The HPE InfoSight portal supports multiple organizations. An organization is a collection of users and systems that are grouped together in InfoSight. When you navigate within InfoSight, you must select the group of assets and the applications that you have been given access to by selecting an organization from the **Organization** selector.

The following are organizational models that are currently supported by HPE InfoSight. Each model is dedicated to managing a product domain:

- HPE Alletra 6000, HPE Nimble Storage
- HPE Alletra 9000, HPE 3PAR StoreServ, StoreOnce, Primera, and Recovery Manager Central
- HPE Compute/Servers
- HPE SimpliVity



While all HPE InfoSight users complete the same initial steps to create and verify their accounts, the procedures for registering devices, managing users, and navigating the user interface differ based on the product domain type.

## Configure Your Home Page

When you log into the HPE InfoSight portal, the HPE InfoSight Home page appears. The Home Page is designed to help you navigate the GUI and access key features and pages. You can display the Home Page from any other page in the GUI by clicking the HPE InfoSight logo at the top-left corner of every page.

---

**Note:** The Home Page appears only after you have registered with an organization or after you switch to a different organization.

---

The Home Page is composed of three sections:

- The Links section at the top left hand side
- The Landing Page section at the top right hand side
- The Bulletins section as the bottom of the screen.

### Links Section

The links section of the Home Page is composed of two different type of hyperlinks:

- Recent Pages
- Product Pages

The Recent Pages section shows recent pages you have visited in InfoSight. The list is updated each time you visit a new page. The list maintains a maximum of 10 pages, which will always be the last 10 pages you've visited. As you visit more pages, the older pages are pushed down and eventually disappear as new pages are displayed.

The Product Pages section enables you to access specific InfoSight areas dedicated to a platform that has been configured for InfoSight.

---

**Note:** Product pages are filtered on a per user and per organization selection basis. Consequently, switching to a different organization while logged into the same user account could show different product pages as different products can be included in different organizations.

---

### Landing Page Section

The Landing Page section is used to set your preferred landing page. Before you select a landing page, make sure you have first selected an organization. Once you select an organization, the newly selected organization appears in the Landing Page Organization field.

In the Landing Page drop-down list, the default page options are:

- The Home Page
- The products that are listed in the Product Pages section.

---

**Note:** The products in the drop-down list are not all the products that exist, but rather the pages for the type of product that you have configured to be managed by InfoSight.

---

To select a landing page, click on the page in the drop down list, and then click **Set as Default**. The page that you selected is displayed immediately, and the page becomes the default page that appears when you next log into InfoSight.

### Bulletins Section

By default, the Bulletins section shows all the InfoSight bulletins that have been released. Bulletins consist of notices, alerts, release notes, user tips and FAQs. You can dismiss a bulletin by clicking the **Dismiss** button that appears next to the bulletin or you can remove all bulletins by clicking **Dismiss all**.





---

**Note:** Once bulletins are dismissed, they can no longer be displayed in the Bulletins section of the Home Page. However, you can still view dismissed bullets at any time by clicking **My Account** > **Past Bulletins**.

---

## The Dashboards

### Operational

For owners of HPE Alletra 6000, HPE Nimble Storage assets, HPE InfoSight opens on the Operational Dashboard. This dashboard includes quick views of the performance of your arrays and the status of your system, protection, and capacity. It also shows pool and VMware object [performance recommendations](#).

### Recommendations

The Recommendations Dashboard displays cross-stack performance recommendations that can help you improve the performance of your virtual environment. These recommendations are generated by InfoSight's AI recommendation engine, which leverages machine learning to identify and diagnose a wide variety of performance issues. See [Cross-Stack Performance Recommendations](#) on page 10 for more information.

### Executive

The Executive Dashboard displays an overall view of the state of your arrays. The various graphs provide a quick picture of several aspects of your system, including space savings, operational efficiencies, data protection, and upgrade needs.

You can get more information on the Executive Dashboard in [The Executive Dashboard](#) on page 47.

### Wellness

The Wellness dashboard is where you can view issues against your arrays and details about those issues.

You can get more information on the Wellness Dashboard in [Wellness Issues and how to Manage Them](#) on page 23.

### Capacity Report

The Capacity Report displays the capacity of all your pools or a particular pool based on a time frame that you select. The information is presented as a summary in the headline bar or in graph format with the details listed at the bottom.

### Labs for Analytics

HPE InfoSight Labs is a framework that allows you to quickly view advanced analytics. Each application in Labs supports near-real time analytics and can be queried on demand. Labs also enables you to obtain instant feedback that you can use to refine your models to ensure maximum accuracy. All the applications can be aggregated by volumes, applications, performance policies, or volume collections.

The following is the current list of labs that you can access from the InfoSight portal.

- Capacity Consumers Forecast
- HPE Cloud Volumes Sizing Tool
- AI Performance Recommendations
- Pool Performance Telemetry
- Reflection Planning
- Replication Timeline
- Capacity Consumers Timeline
- Volume Performance
- Recurring Performance Patterns
- Inter-Volume Performance & Contention

For each application, the input parameters that you select are crucial to the results.

For information about a particular lab, open the lab and then click the information icon that appears next to the lab name. A **Lab User Guide** is displayed that explains the purpose of the lab as well as input and output information.



## Cross-Stack Performance Recommendations

HPE InfoSight automatically analyzes system telemetry to diagnose problems and recommend actions you can take to improve the performance of pools in your storage environment and in VMware hosts, datastores and VMs in your virtual environment.

The recommended actions appear in cards that you can access from:

- The Recommendations Dashboard
- The Operational Dashboard
- The Performance window for individual objects.

**Note:** To watch a video demonstration of this new feature, click [VM Recommendations](#).

### Display the Recommendations Dashboard

To display the Recommendations Dashboard, click **Dashboards** > **Recommendations**.

The Recommendations Dashboard displays information in cards that show the analysis period, the problem diagnosis, and recommended actions that you can take to improve the performance of the object.

The screenshot shows the 'Recommended Actions' dashboard. At the top right, it indicates the 'Analysis Period' as '24-Aug-2020 - 01-Sep-2020'. Below this is a filter bar with four dropdown menus: 'Object Type' (set to 'All'), 'Diagnosis Type' (set to 'All'), 'Snoozed' (set to 'No'), and a search box. The main content area contains two recommendation cards. The first card is for a 'Pool: Asup-array01-lvs (default)' with an analysis period of '8/24/2020 - 9/1/2020'. It identifies the issue as 'Elevated Random Read Latency' and recommends 'Random Read Cache Miss' actions: 'Consider QoS limits on random read workloads' and 'Consider staggering random read workloads'. The second card is for a 'Host: 10.175.1.15' with the same analysis period. It notes 'Performance issues have been identified affecting this host' and recommends 'Physical CPU Saturation' actions: 'Upgrade the physical CPU capabilities of this host' and 'Migrate some VMs to hosts with lower physical CPU utilization'. Both cards include a 'Show More' link and a snooze icon.

You can filter the display of objects in this dashboard by object type, diagnosis type, snoozed designation, and object name.

### Filter by Object Type

To filter objects in the dashboard by object type, select an object type from the **Object Type** drop-down list. You can select any of the following object type choices:

- ALL
- DATASTORE
- HOST
- POOL
- VM

**Note:** If VMware streaming is disabled or you don't have permission to view VMware data, then you can only display recommendations for pools.

### Filter by Diagnosis Type

To filter objects in the dashboard by diagnosis type, select one or more diagnosis types from the Diagnosis Type drop-down list. You can select any of the following diagnosis types:

**Table 1: Diagnosis Types**

All
Contention with Copy Offload
Elevated Latency
HDD Bandwidth Saturation
IOPS Limit Throttling
Low Compressibility
MBPS Limit Throttling
Network Bottleneck
Optimize Deduplication
Pending Deletes
Physical CPU Saturation
Random Read Cache Miss
SSD Bandwidth Saturation
Sequential IO Contention
Sequential Read Cache Miss
Storage CPU Saturation
Synchronous Replication
Unaligned IO
Unidentified Root Cause
Virtual CPU Overprovisioning
Virtual CPU Underprovisioning

Note that the diagnosis types that appear in the Diagnosis Types drop-down list are governed by the selected object type. For example, if you select **Datastore** as the object type, then only **Elevated Latency** will be available in the Diagnosis Type drop-down list.

### Change Object Snooze Settings

You can snooze the recommendation notification for an object using the Snooze option. When you snooze an object, the recommendation card is removed from the display for a specified time interval..

To change the snooze setting for an object, click the Snooze button (the alarm clock icon) next to the object name and select a specific snooze period. You can elect to snooze an object for one month, three months, or indefinitely.

You can filter snoozed objects in the dashboard by selecting either **Yes**, **No**, or **All** in the Snoozed drop-down list.



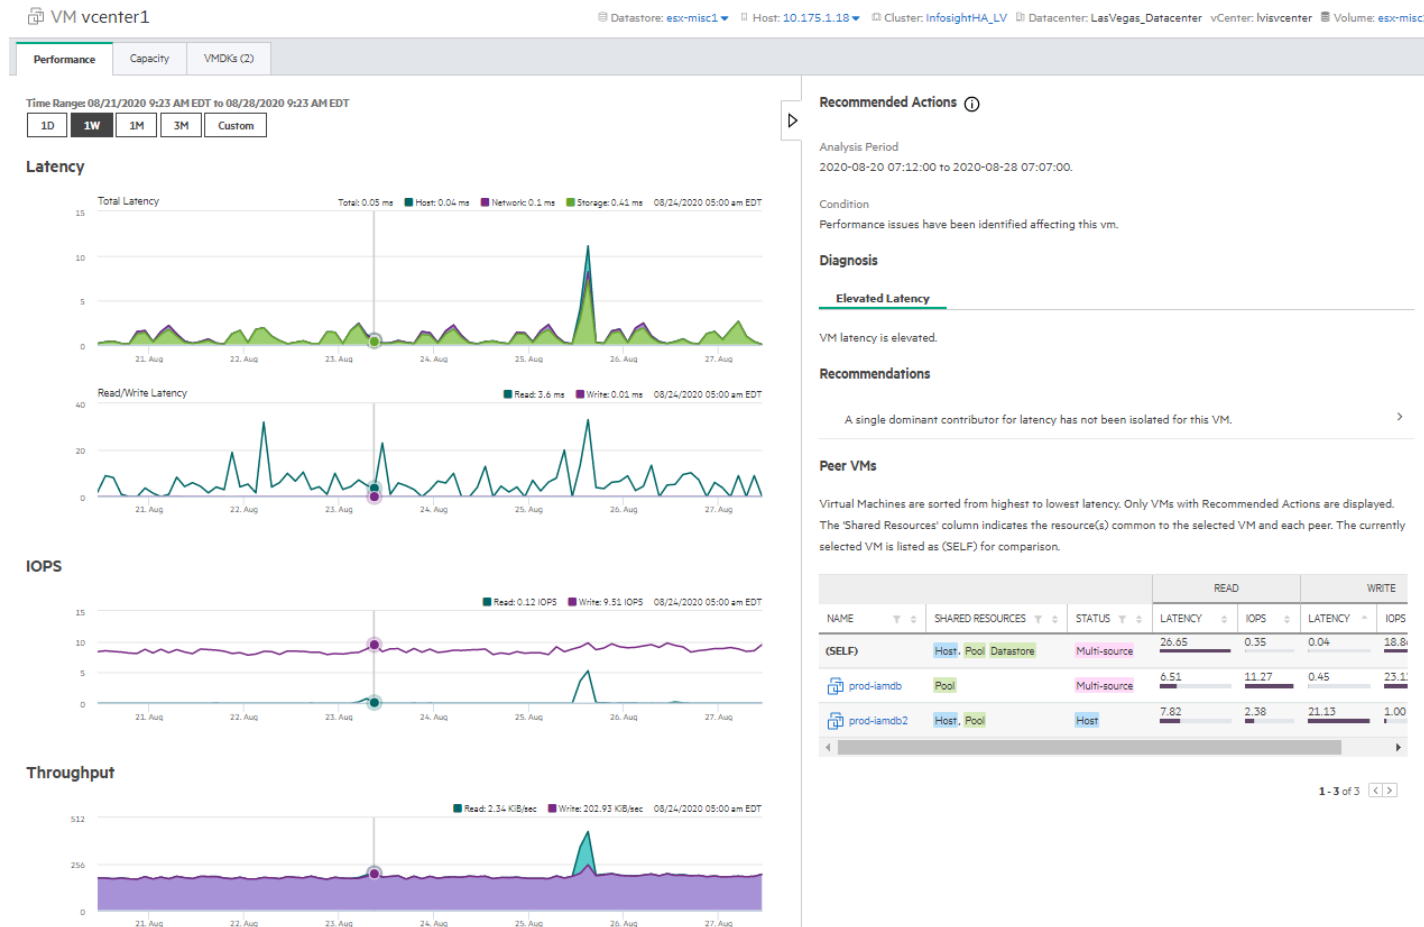
## Filter by Object Name

To filter objects in the dashboard by object name, type the name of the object in the Search field.

Note that the search results are governed by the inputs that you have selected in the Object Type, the Diagnosis Type, and the Snoozed lists. For example, if you search for VMs that have a diagnosis type of Elevated Latency, then the search will only return VMs with that diagnosis type.

## View Performance and Object Details

To view the details of an object in the Recommendations Dashboard, click the object. The Performance tab for the object appears with detailed Recommended Actions appearing in a collapsible panel on the right side of the screen. The following example shows the Performance tab for a VM object:



In the Performance tab, you can move the mouse pointer in the performance graph to correlate VM Latency with underlying IOPS and Throughput utilization. You can also click the graph and drag the mouse to zoom in on a subset of the performance data.

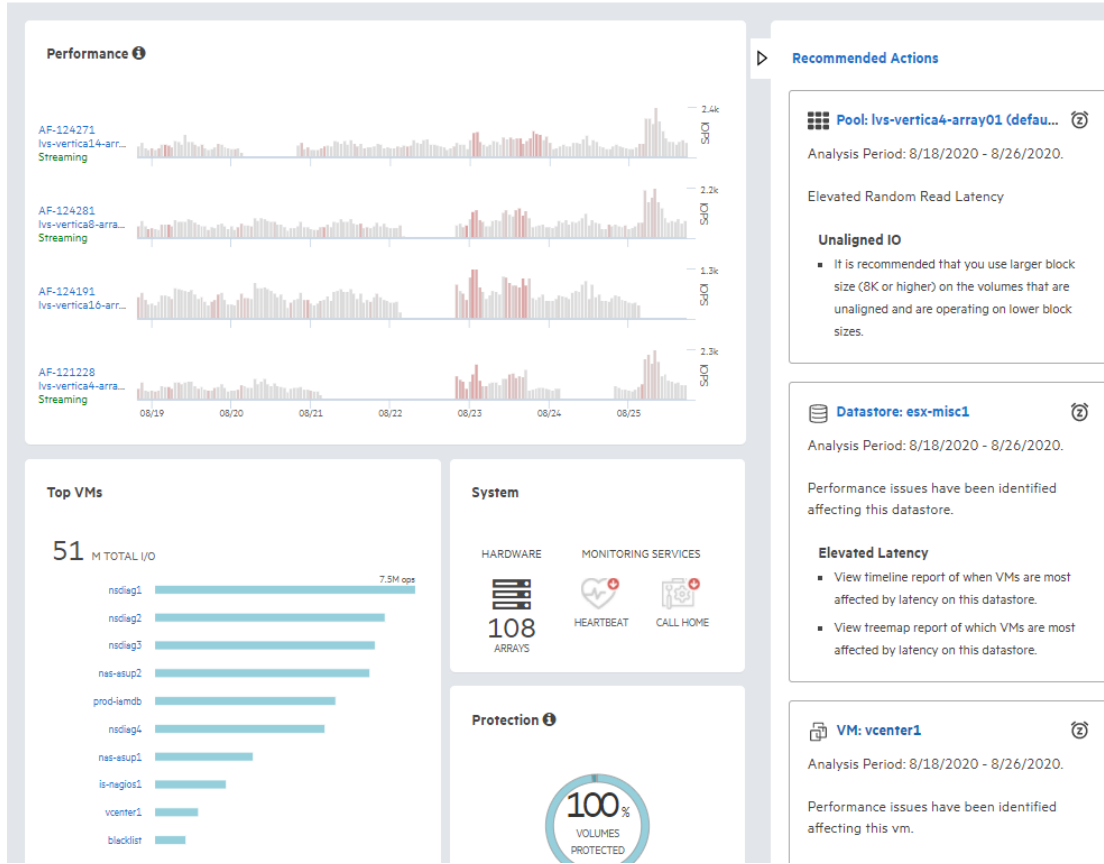
In the Recommended Actions panel, you can see a diagnosis of any discovered issues and view recommendations on how to fix the issues. In addition, the Recommended Actions panel displays performance analytics for Peer VMs, which enables you to analyze how peer objects might affect overall object performance.

For more detailed information about the Recommended Actions panel, click [VM Recommendations](#) to watch a video demonstration.

## Display the Operational Dashboard

To display the Operational Dashboard and view performance recommendations, click **Dashboards** > **Operational**.

## Operational Dashboard



The performance recommendations appear in collapsible cards on the right side of the display in the same manner as the cards in the [Recommendations Dashboard](#). For each analyzed pool, host, datastore, or VM, you can view the following information:

- The object name
- The dates the system telemetry was analyzed
- A brief description of any discovered performance issues
- Recommended actions you can take to address discovered performance issues

When working with performance recommendations in the Operational Dashboard, you can also:

- Snooze the objects in the Recommended Actions cards in the same manner as you can in the [Recommendations Dashboard](#).
- Navigate to the Recommendations Dashboard by clicking the **Recommended Actions** link.
- Display the performance tab and detailed recommended actions for an object in a card by clicking the object.
- Click a VM in the Top VMs list to display detailed recommended actions for that object.

## Commonly Used HPE InfoSight Pages

The following are some of the most commonly viewed pages in HPE InfoSight.


### Software Downloads page:

You can download software for the array OS, Integrations kits, and Solution kits from the Software Downloads page. To access the page, click the InfoSight Main Menu (☰), click **Resources** and then under the product name click **Software Downloads**.

### Documentation page:

You can view and download documentation for array software and hardware from the Documentation page. To access the page, click the InfoSight Main Menu (☰), click **Resources**, and then under the product name click **Documentation**.

**Support page:**

If you have admin privileges, you can go to the Support page to view and manage information about your arrays, cases, and contacts. To access the page, click the InfoSight Main Menu () , click **Resources**, and then under HPE Alletra 6000, HPE Nimble Storage click **Support**.

**The List Views and Details Pages**

The List Summary pages consist of a list of all the pools, arrays, volumes, VMs, datastores, or hosts that are visible in your account. The columns can all be sorted and the contents of several columns are links to the details pages for those items. The columns can also be filtered using the funnel icon on the column header.

On the Details pages, the headline bar indicates whether the information on display is for an array, VM, volume, datastore or host. Below the headline bar are graphs that display information based on a particular date range that you specify.



## Data Availability FAQ

Where can I view information on the latest data set that is being presented?

Page or View	Location	Details
Array Details Page	<ul style="list-style-type: none"> <li>Last Heartbeat and Last Data Sent columns</li> <li>The icons in the headline bar</li> </ul>	Indicates the time the last data set and heartbeat was received



Page or View	Location	Details
VM List Page	The column, Last Activity, on the Summary tab	Indicates the status and currency of the data
Wellness Dashboard	In the Issue Details table	Indicates the time and type of alert that was created

#### **How long after the installation of my array will I begin to see data and where will I see it?**

- The majority of new arrays appear in HPE InfoSight within 48 hours after receiving the first daily data set.
- For trending data such as capacity graphs, although data points might appear, expect two weeks before you have enough data for a meaningful forecast.

Although this data is updated on a weekly basis, it is not likely to be useful until there has been significant activity on the array. For example, the cache utilization metric is not useful until you have written enough data to fill the cache.

#### **After I enable VM Monitoring, how long will I have to wait before I can view the statistics?**

The statistics should be visible in HPE InfoSight within two to four hours. After the initial setup, which can take up to 48 hours, VMvision statistics are updated within four hours.

#### **After I enable the vCenter plugin, how much time does it take to start seeing associated data in HPE InfoSight?**

The VM related data in HPE InfoSight is based on a combination of the streaming statistics and VMware-related data that is part of the regular data set. Both of these sets of data must be received and processed (which requires at least 48 hours) before data is visible on HPE InfoSight.

#### **What is the average processing time for the streaming statistics?**

Usually about 4 hours.

#### **At what time are the nightly data sets processed (for example, are they processed based on the arrays time-zone or are they processed at the same time for all arrays?)**

Daily data sets are scheduled to be delivered a few minutes after midnight array local time. Data sets are processed in batches based on the order of arrival. Usually, the data sets are processed within 24 hours of their arrival.

#### **I upgraded my array. How long will it take for HPE InfoSight to show the version and other changes?**

Configuration changes, including upgrades and hardware changes, are updated in HPE InfoSight within 48 hours.

#### **When moving VMware virtual machines from one volume to another it generally takes time to reflect the change in HPE InfoSight correctly. How long should it take?**

You should see the changes within 48 hours, depending on the timing.

#### **If there is an HPE InfoSight recommendation, when does it appear and how long does it persist if the root cause is resolved?**

Preliminary recommendations appear within four days. Cache recommendations are only meaningful after the SSD cache has been fully written to at least once.

The recommendation is removed immediately after a data set event indicating a configuration change is received and processed, which could take up to 48 hours.

#### **If I have replication to an array that is in a different time zone, what is the impact of the time difference and location?**

We currently display time in local time of the browser unless where explicitly stated.

#### **How long does HPE InfoSight retain the Diagnostics for DNA? Can I request that the DNA be held longer?**

The DNA is held for one year. It is not possible to request that the data be held for a longer period. HPE InfoSight stops displaying the data of a particular array after the array stops sending heartbeats and data sets for seven consecutive days.





## Data Availability Details

Type of data	Notes	Quantity of data necessary for initial data visibility	Data currency (based on time change occurs)	Frequency of updates (visibility)
Wellness Dashboard, Issue Details (auto-opened case information)	Based on heartbeat, alerts and analysis of data sets	1 day worth of data		
Array details	Scroll over any graph and the analytics appear below the graph. The time the analytics were collected is in the far right.	1 week worth of data		
Capacity (raw data)		1 day worth of data		
Capacity forecasts		14 days worth of data		
Data protection information	Volumes List page	1 day worth of data	Updated within 48 hours	Every 24 hours
Replication Planning tool	Only available for volumes with snapshots	14 days worth of data from snapped volumes		
Performance graph	Data on IOPS, throughput, latency and its breakdown	7 days worth of data		
Performance information		1 day worth of data		
CPU and cache recommendations	Computed based on past 4 weeks	4 days worth of data		
Virtual environment	Assets inventory changes (upgrades and hardware changes)	1 day worth of data		
Virtual environment	Datastore, host, and VM statistics	1 day worth of data, post vCenter registration	Updated within 2 hours	Every 2 hours

## Managing Your Arrays

Before you can begin to monitor your arrays, you must perform some preparation and configuration tasks. These tasks ensure that you have access to the arrays, that you have registered your plugins, and that you can receive notifications and information on the health of your arrays.

---

**Note:** For information on how to verify software updates paths, and how to download HPE Alletra 6000, HPE Nimble Storage software, refer to the *GUI Administration Guide*.

---

### HPE InfoSight Roles

There are three roles that provide access to arrays on HPE InfoSight:

#### Admin

Users who have been assigned the **Admin** role can:

- Manage users
- Register systems
- Access all information within HPE InfoSight.

#### User

Users who have been assigned the **User** role can only view information in HPE InfoSight, with the exception of the Support page.

#### Third Party

This role has the same privileges as the User role, but is used for external company users, such as partners.

To see all the users in your organization and their assigned roles, go to the settings menu (the gear icon) and select **Users**.

### Configure Your Arrays

Configure your arrays to send data sets to Support.

It can take up to 48 hours for the array to show up in InfoSight after the first data set is sent.

#### Before you begin

These steps must be performed on the array GUI, not in InfoSight.

#### Procedure

1. From the **Administration** menu in the array GUI, select **Alerts and Monitoring > DNA > (DNA)**.
2. Select **Allow Support to collect analytics data from the array automatically**.
3. Click **Test DNA Settings** to confirm that DNA is set up correctly.

If errors appear after the test completes, check the proxy settings to verify whether a required HTTP proxy has been configured. If a proxy is not in use, the following external firewall ports will need to be opened from the array management and controller IPs to benefit from full InfoSight functionality:

`nsdiag.nimblestorage.com` (Port 443, HTTPS)

`nsstats.nimblestorage.com` (Port 443, HTTPS)

Further firewall ports are required for additional array features. See the appropriate Administration Guide available on the InfoSight documentation portal for more details.



#### 4. Click **Save**.

If errors appear after the autoSupport test completes, check the proxy settings to validate whether a required HTTP proxy has been configured. If a proxy is not in use, the following external firewall ports must be opened from the array management and controller IPs to benefit from complete InfoSight functionality:

`nsdiag.nimblestorage.com` (Port 443, HTTPS)

`nsstats.nimblestorage.com` (Port 443, HTTPS)

Additional firewall ports are required for additional array features. See the appropriate Administration Guide available on the InfoSight documentation portal for more details.

## Configure Arrays to Send Data to HPE InfoSight

HPE InfoSight recommendations and automatic fault detection are based on data that is sent from your arrays and processed by HPE InfoSight.

To take full advantage of the HPE InfoSight monitoring and analysis capabilities, configure your arrays to send data sets, statistics, and events to Support through DNA.

If you did not configure your arrays to send this data to Support during the initial setup, you can change the configuration at any time by following this procedure or by running the `group --edit` command in the CLI.

### Use the GUI to Configure Arrays

#### Before you begin

This procedure must be performed in the array OS GUI, version 4.x and later.

#### Procedure

1. In array OS, go to **Administration** > **Alerts and Monitoring** > **DNA**.
2. On the DNA page, select **Allow Support to collect analytics data from the array automatically**.
3. Click **Test Diagnostics for Analytics Settings** to confirm that DNA is set up correctly.
4. Click **Save**.
5. To enable event data, go to **Administration** > **Alerts and Monitoring** > **Email**.
6. On the **Email Alerts** page, select **Send event data to Support**.
7. Enter an SMTP host and port.  
Event data is relayed through HTTPS by default.
8. Click **Save**.

### Use the CLI to Configure Arrays

#### Before you begin

This procedure must be performed in the array OS CLI.

#### Procedure

1. Login to the CLI.
2. Enable sending data analytics to support:  
**group --edit --autosupport yes**
3. Enable sending event data to support:  
**group --edit --send\_event\_data yes**



## Plan Your Storage Resources

Resource Planner is a what-if scenario analysis tool that allow you to determine optimum workload placement. It helps IT administrators minimize the risk of downtime caused by the overuse of resources, maximize hardware utilization, therefore reduce costs while planning their hardware needs for their workloads. Resource Planner is for Alletra 6000 & Nimble Storage customers. Application-specific workload characteristics can be used in the modeling, helping to determine the workload's impact on array resources such as cache, capacity, I/Os, etc.

Resource Planner supports the following workflows:

- Adding new workloads  
Simulate a new workload that you configure to be placed on an existing system.
- Scaling workloads  
Model existing workloads that are running on your arrays to analyze what the results would be to scale or enlarge them on an existing system.
- Migrating workloads  
Create new, hypothetical workloads that model what would happen if you were to migrate existing workloads to other systems in your data center.

To access the Resource Planner, click on **Alletra 6000, Nimble Storage** from the InfoSight Main Menu, then select **Planning > Resource Planner**.

### Add New Workloads

Adding a workload can help you determine the effect of adding one or more application-specific workloads to your storage environment. Adding a workload allows you to define new workloads, simulate workloads on a new or existing array, and receive recommendations for scalability.




#### Procedure

1. From the Resource Planner, select **Start** in **Add Workloads**.
2. In the Define Workloads window, enter the **Name, Application Category, Max IOPs**, and **Dataset Size (TiB)**.


Advanced Options allows you to overwrite the default modeling for Performance, Capacity, and Data Protection.

3. Click **Run Analysis**.

The Analysis Results window opens and shows the Top Existing Hardware and Top New Hardware. The recommended

hardware based on headroom is indicated by , , or .

Based on the results, you can:

- Add another workload for analysis by selecting **Add**.
- Select a pool for more detailed information on a specific pool in you Top Existing Hardware by selecting **Select Pools**.
- Filter arrays based on Array Generation, Array Type, and Headroom Remaining by selecting the  icon in Top New Hardware.
- Display additional Top New Hardware by selecting **Show 5 More Array Models**.
- Export the Workload definitions as CSV or the Analysis Results as PDF by selecting **Export**.

**Note:** A maximum of five workloads can be analyzed at one time.



The results from the last analysis are saved. To view them, click **Return to your last analysis results** in **Add Workloads** from the Resource Planner. Adding a new Workload deletes the results from the last analysis.

## Scaling Workloads

Scaling workloads allows you to plan for infrastructure needs by simulating variability in an existing workload, providing a more granular analysis of a volume on an array.

### Procedure

1. From the Resource Planner, select **Start** in **Scale Workloads**.
2. In the Define Workload window, enter the **Name**, **Pool**, **Date Range**, **I/O Multiplier**, **Capacity Multiplier**, and **Add Volumes** field.




The available volumes are displayed in Selected Volume. You can remove a volume from the list by clicking the **X** under **ACTION**.

To view historical performance, select **Workload Performance**.


To view the past performance of another workload, select **Add**.

3. Select **Run Analysis**.

The Analysis Results window opens and shows the Current Hardware, Top Existing Hardware, and Top New Hardware.

The recommended hardware based on headroom is indicated by , , or .

Based on the results, you can:

- Add another workload to scale by selecting **Add**.
- Display the results of your Current Hardware.
- Select a pool for more detailed information on a specific pool in your Top Existing Hardware by selecting **Select Pools**.
- Filter arrays based on Array Generation, Array Type, and Headroom Remaining by selecting the  icon in Top New Hardware.
- Display additional Top New Hardware by selecting **Show 5 More Array Models**.
- Export the Workload definitions as CSV or the Analysis Results as PDF by selecting **Export**.

**Note:** A maximum of five workloads can be analyzed at one time.

The results from the last analysis are saved. To view them, click **Return to your last analysis results** in **Scale Workloads** from the Resource Planner. Scaling a new Workload deletes the results from the last analysis.

## Migrating Workloads

Migrating a workload can help you determine the optimal location in your existing environment to move a given workload.

### Procedure

1. From the Resource Planner, select **Start** in **Migrate Workloads**.
2. In the Define Workloads window, enter the **Name**, **Pool**, and **Group By** field.

The available volumes are displayed in Selected Volume. You can remove a volume from the list by clicking the **X** under **ACTION**.




To view historical performance, select **Workload Performance**.

To view the past performance of another workload, select **Add**.




### 3. Click **Run Analysis**.

The Analysis Results window opens and shows the Current Hardware, Top Existing Hardware, and Top New Hardware

that are best suited for the migration based on headroom, indicated by , , or .

Based on the results, you can:

- Add another workload for analysis by selecting **Add**.
- Select a pool for more detailed information on a specific pool in you Top Existing Hardware by selecting **Select Pools**.
- Filter arrays based on Array Generation, Array Type, and Headroom Remaining by selecting the  icon in Top New Hardware.
- Display additional Top New Hardware by selecting **Show 5 More Array Models**.
- Export the Workload definitions as CSV or the Analysis Results as PDF by selecting **Export**.

**Note:** A maximum of five workloads can be analyzed at one time.

The results from the last analysis are saved. To view them, click **Return to your last analysis results** in **Add Workloads** from the Resource Planner. Migrating a new Workload deletes the results from the last analysis.




## View Software Recommendations

HPE InfoSight provides software recommendations for updating the software on your arrays. You might want to update the current version of your software to incorporate a release that has new features or bug fixes.

To view recommended software updates for an array:

- 1 Log in to HPE InfoSight
- 2 Go to **Infrastructure** > **HPE Alletra 6000, HPE Nimble Storage** > **Arrays**
- 3 Click an array. The Software Recommendations window appears on the bottom right side of the page.

If the software release appears in red and is marked with a red dot, it means that the array has been added to the Deny List for that release. The Deny List entry is applied to an array when Support has determined that the update is not suitable for the array, and the software is unavailable for download for the array. Also, if there is information about an issue that might have caused an array to be added to the Deny List, a description of the issue appears under the release name:

Software Recommendations		
Release Statuses >		
Current Version: 4.5.6.0-726655-opt		
RELEASE	RELEASE STATUS	UPGRADE PATH
5.x		
<p> 5.1.4.200</p> <p>AS-94961 - Performance affecting firmware defect in a subset of 6TB drives. </p>	General Availability	Current →  5.1.4.200
5.0.10.0	General Availability	Current → 5.0.10.0

## Control Release Statuses for Software Recommendations

InfoSight now enables you to control the software release statuses that you want to make visible to an array for software download. The release statuses include:



- Long Term Support Release (LTSR)
- General Availability (GA)
- General Availability Candidate (GAC)
- Initial Production Release (IPR)

For a description of these software release statuses, refer to [Release Statuses](#) on the documentation portal.

To select the release statuses that you want to make visible to an array:

- 1 Click **Release Statuses** in the Software Recommendations window. The list of release statuses appears.
- 2 Select one or more release statuses.
- 3 Click **Save**.

The Software Recommendation list changes based on the release statuses you have selected.

---

**Note:** Filtering the release statuses in the Software Recommendation list does not affect arrays added to the Deny List, which are controlled by the support team. The release for an array on the Deny List will always be unavailable for downloading until the issue that caused the Deny List entry is addressed and the support team removes the array from the Deny List.

---

## Wellness Issues and how to Manage Them

---

**Note:** The Wellness Issues view is not meant for real time monitoring of the status of your arrays. Refer to the array GUI for real time monitoring.


---

Wellness issues arise when there are situations that you need to be aware of on any of your arrays. A wellness issue may involve latency, remaining space, issues with your power supply or any condition that is important to the health of your arrays.

You can track the wellness of your arrays on the Wellness Dashboard. The Wellness Dashboard is also where you can review wellness rules and tailor them to your configuration. You can configure rules to control the escalation of detected health conditions on the storage arrays.

You use the Create Issue checkbox to control the conversion of wellness issues for a given condition or object pair. If issue creation is disabled, case creation is automatically disabled as well.

To disable the alerts, you uncheck **Create Issue** on the Configure Wellness Rules page.

The support case options control the creation of support cases from wellness issues. When a case is created, subsequent occurrences of the condition do not cause the creation of additional cases until the snooze period expires. Most conditions (annotated with the  icon) allow only all or nothing control. In other words, you cannot disable case creation independently of issue creation.

## Configure Wellness Rules

### Procedure

On the Wellness Dashboard, click **Configure Wellness Rules**.

In the Configure Wellness Rules window, you can perform the following actions:

- To enable or disable case creation and notification for all arrays, check or uncheck **Create Issue** next to the wellness rule that you would like to modify.
- To enable or disable case creation and notification for a specific array, perform the following steps.
  - 1 Expand the wellness rule that you would like to enable or disable.
  - 2 Check or uncheck the **Create Issue** checkbox next to a specific serial number.

---

**Note:** After you disable automatic case generation for a certain condition, support cases are not created automatically until you manually re-enable case generation.

---



## Set-up Email Notifications for Alerts

You can subscribe to daily wellness emails that inform you of the status of your arrays.

### Procedure

1. On the Wellness Dashboard, click **Daily Summary Emails**.
2. Enter an email address for delivery of the email alerts.
3. (Optional) You can click **Test** to send a test email to the email address that you indicated.
4. Click **Add** to add the email address.
5. Click **Close** to conclude the email alerts setup.

## Set a Snooze Period

The snooze period is the number of days that a recurring condition updates an existing case before creating a new case. You are notified via email when a case is created. For most conditions, you are not notified when a case is updated.

### Before you begin

You must have both **Create Issue** and **Support Case Create** checked on the Configure Wellness Rules page before you can set a snooze period.

### Procedure

1. Make sure that **Support Case Created** is checked.  
You can only check the **Support Case Created** box if the **Create Issue** box is already checked.
2. Select the number of days that you want to snooze the support case in the **Snooze** column.  
You can also select "Do Not Snooze."
3. To disable the alerts, uncheck **Create Issue**.





# Displaying VMware Analytics

If you have set up VMware to work with your storage environment, you can access the VMware Virtualization pages on HPE InfoSight, which include pages for the following storage features:

- Datacenters
- Clusters
- ESXi Hosts
- Datastores
- VMs

From these pages, you can drill down to look at the most granular aspects of your data and then drill up to get a more global view.

---

**Note:** For information on how to set up VMware to work with your storage environment, refer to the *VMware Integration Guide* on the InfoSight documentation portal.

---



**Tip:** For an updated overview of Cross-Stack Analytics for VMware features and collection details, see [Cross Stack Analytics for VMware](#).

---

## Using VMware Virtualization Data to Evaluate Performance

You can view detailed information about virtual machines (VMs) hosted on arrays on the VMware Virtualization pages located under **Infrastructure > Virtualization:VMware**. You can use the analytics found on the pages in this menu to monitor latency and performance across the storage, network, and host stack. This information allows you to address potential issues and improve performance. For example, the data can show you where there are CPU bottlenecks and which hosts have the most CPU headroom. That way you can move a VM to a host that has more resources.

You can use the VMware Virtualization analytics to identify potential issues, such as:

- Contentious VMs
- Overloaded hosts (CPU and memory)
- Latent VMs by datastore
- Location of the volumes on VMware datastores

To collect this data, InfoSight uses the storage array and the VMware vSphere APIs. You do not need to install anything on the vCenter server.

You can quickly set up InfoSight to gather information about the storage array and VMware. From the array OS, you must register an HPE vCenter plugin or enable a VMware synchronized volume collection. From InfoSight, you must enable both the Streaming button for the array you want monitor and the VMware button. For detailed information about enabling this information, see [Enable InfoSight to Collect VMware Information](#) on page 25.

After InfoSight starts gathering this information, it can take up to 48 hours for the initial processing to complete and show data. After that, depending on the data, it can be updated as frequently as every 5 minutes. .

---

**Note:** InfoSight only provides analytics information about VMs and datastores that are hosted on an HPE storage array.

---

## Enable InfoSight to Collect VMware Information

You can enable HPE InfoSight to provide VMware data for your system by registering either the vCenter Server used with the HPE vCenter Plugin or a VMware synchronized snapshot volume collection and then enabling InfoSight to collect data



from the array and VMware. You perform steps in the array OS GUI and in InfoSight. You do not need to install anything on the vCenter Server.

### Before you begin

- Your array environment must be set up to work with VMware.
- At least one VM must be on a datastore.
- You must have an InfoSight account.
- Ports 80 and 443 must be open between the vCenter Server and the array.

### Procedure

1. Configure the vCenter Plugin on one or more arrays.
  - a) Log into the array.
  - b) From the array OS GUI, select **Administration** > **VMware Integration**.
  - c) Register the vCenter Server by entering the vCenter Server name and login credentials.

---

**Note:** If there is already a registered vCenter Server, select **Add Another vCenter** to add your vCenter Server.

---

**Optional:** If you do not want to use the HPE vCenter Plugin, you can enable a VMware synchronized snapshot volume collection and register it instead of a vCenter Server.

After you register the vCenter Server, it is a good practice to verify that the array can communicate with vCenter Server.

2. Configure InfoSight to access the VMware data.
  - a) Log into InfoSight.
  - b) From the settings menu (the gear icon) select **Telemetry Settings**.
  - c) Locate the array you want to monitor and click the button for **Streaming** to **On**.  
This button enables data streaming from the array.
  - d) In the same row, click the button for **VMware** to **On**.  
This button allows data to be collected from VMware.

---

**Note:** If you do not set the **VMware** button to **On**, you will only receive streaming data from the array. You will not receive any VMware data.

---

### What to do next

It takes up to 48 hours for InfoSight to process the vCenter registration and start streaming VMware and array data.

When the data streaming begins, you can view the analytics for the different objects by choosing **Infrastructure** > **Virtualization** > **VMware** > **[Datacenters | Clusters | ESXi Hosts | Datastores | VMs]**.

## Register the vCenter Plugin

The HPE vCenter plugin is part of the array OS. To take advantage of it, you must first register the plugin with a vCenter Server.




---

**Important:** The vCenter plugin is supported on ESX 5.5 update 1 and later.

---

**Note:** The plugin is not supported for:

- Multiple datastores located on one LUN.
- One datastore spanning multiple LUNs.



- LUNs located on a storage device not made by the array group.

You can register multiple vCenter plugins on the array. Each array that registers the vCenter plugin adds a tab to the vSphere client.

## Use the Array OS GUI to Register the vCenter Plugin

### Before you begin

Use a vCenter account that has sufficient privileges to install a plugin (for example, the root account, [administrator@vsphere.local](mailto:administrator@vsphere.local) or any administrator account). You need to know the vCenter hostname or IP address.

### Procedure

1. From array OS, select **Administration** > **VMware Integration**.
2. If the fields are not already filled, enter the vCenter name, subnet, vCenter server host name or IP address, port, description, username, and password; and check the appropriate boxes to register the web client, thin client, or VASA Provider.
3. Click **Save**.  
If a Security Warning message appears, click **Ignore**.
4. Click **Test Status** to ensure that the plugin has been registered.
5. Relaunch the vSphere client.

### Results

If the registration was successful, you can right click a datacenter, host cluster, host, or datastore and the Actions menu appears.

## Use the Array OS CLI to Register the vCenter Plugin

### Before you begin

Use a vCenter account that has sufficient privileges to install a plugin (for example, the root account, [administrator@vsphere.local](mailto:administrator@vsphere.local) or any administrator account). You need to know the vCenter hostname or IP address.

### Procedure

1. Log into your array using the array OS CLI.
2. At the command prompt, type:  
**vcenter --register <vcenter\_name> [--extension {web|vasa}]**  
Add the following command option if you want to run the web client:
  - **--client web** (for web client)
3. Relaunch the vSphere client.

### Results

If the registration was successful, you can right click a datacenter, host cluster, host, or datastore and the Actions menu appears.

## Configure Arrays to Monitor your Virtual Environment

### Procedure

1. Go to the Main Menu, and click **Alletra 6000, Nimble Storage** → **Settings** → **Telemetry Settings**



2. In the Array list, find the array for which you want to monitor the virtual environment.
3. In the Streaming column, make sure the On/Off box is set to **On**.  
You must have streaming enabled to send diagnostic content.
4. In the Diagnostics Content column, make sure the On/Off box is set to **On**.

## View Overprovisioned VMs

### Procedure

1. Go to **Infrastructure** > **Virtualization** > **VMs**.
2. To view the utilization compared to all other VMs, sort the list on the column headers for CPU usage or memory usage.
3. To view the utilization of a particular VM, click the VM in the list.  
In the headline bar, the allocation and usage for the VM is listed.

## View VMs by IOPS and Latency

### Procedure

1. Go to **Infrastructure** > **Virtualization** > **VMs**.
2. From the **View** list, select **VM I/O Contention Trend**.
3. Select a datastore from the Host list.

### Results

The top 10 VMs are shown in tile graphs where the date is depicted along the bottom of the graph and the number of IOPS or the degree of latency is represented by the depth of the color. The deeper the color of the tile, the higher the number of IOPS or degree of latency.

You can run your cursor along the graph to view details about a particular date.

## View VM Latency Due to the ESX Host

### Procedure

1. Go to **Infrastructure** > **Virtualization** > **VMs**.
2. From the summary list, select the VM for which you want to view latency.
3. On the Performance tab, under Latency, the average latency is indicated for the host, the network, and the storage.

## View Contentious VMs Based on the Datastore

### Procedure

1. Go to **Infrastructure** > **Virtualization** > **VMs**.
2. From the **View** dropdown list, select **VM I/O Contention Treemap**.
3. Select a vCenter or DataCenter from the Host list.



**Results**

A treemap graph can display two factors equally well in one graph. In this case the size of the boxes is determined by the amount of I/O. The higher the I/O, the bigger the box. The degree of latency determines the depth of the color. The darker the color, the higher the degree of latency.

## View Contentious VMs Based on the ESX Host

**Procedure**

1. Go to **Infrastructure** > **Virtualization** > **VMs**.
2. From the View list, select **VM CPU Contention Trend**.
3. Select a host from the Host list.

**Results**

The top VMs by CPU usage and CPU ready are shown in tile graphs where the date is depicted along the bottom of the graph and the CPU usage and CPU ready is represented by the depth of the color. The deeper the color of the tile, the higher the average usage or the higher the percent of CPU ready.

You can run your cursor along the graph to view details about a particular date.



## Displaying Hyper-V Analytics

If you have installed and are running the HPE InfoSight Hyper-V Collector Toolkit, you can view and analyze the virtualized Hyper-V infrastructure data that is sent from the array that you registered with the toolkit to HPE InfoSight. This data is displayed in virtualization pages for the following storage features:

- Clusters
- Hyper-V Hosts
- Hyper-V Storage
- VMs

From these pages, you can drill down to look at the most granular aspects of your data and then drill up to get a more global view.

The menu bar above the main data window enables you to tailor the display to specific information. Use the View menu to select the display you want for an object. InfoSight provides analytics for the entire stack as well as for individual objects.

The Virtualization display uses tables, graphs, and pop-up boxes to provide information. In some cases, you can specify custom date and time ranges for the information being displayed. In addition, some of the graphics display additional information when you hover the mouse cursor over a spot on a graph or the name of an object.

## Analyzing Hyper-V Performance Issues

To access the Hyper-V Virtualization pages, log on to InfoSight and go to **Infrastructure** > **Virtualization: Hyper-V**.

You can use the analytics found on these pages to monitor latency and performance across the storage and host stack. This information enables you to address potential issues and to improve performance. For example, the data can show you where there are CPU bottlenecks and which hosts have the most CPU headroom. This in turn can help you decide whether to move a VM to a host that has more resources.

The Virtualization pages include detailed analytics about Clusters, Hyper-V hosts, Hyper-V storage, and VMs. The analytics provide data about memory, latency, IOPS, and throughput.

You can use the analytics to identify potential issues, such as:

- Correlation of Hyper-V VMs to storage
- Noisy neighbors from a storage/host contention
- Overloaded hosts (CPU and memory)
- Underutilized VMs and hosts
- Average latency for VMs and storage

## Enabling InfoSight to Collect Hyper-V Data

To enable HPE InfoSight to stream Hyper-V telemetry data:

- 1 Log in to HPE InfoSight at <https://infosight.hpe.com>.
- 2 Go to the Main Menu, and click **Alletra 6000, Nimble Storage** > **Settings** > **Telemetry Settings**
- 3 In the list of arrays, locate the array on which you want to enable streaming.
- 4 In the Streaming column, click the streaming icon to change the streaming state from **Off** to **On**.

## Key Data Provided by the Hyper-V Virtualization Pages

The virtualization pages that appear under the **Virtualization:Hyper-V** section of HPE InfoSight provide both cross-stack analytics and drill-down analytics. You can look at performance and latency information across the storage and host stack for a specific object. You can use this data to determine how your system is working and whether there are issues within the stack



that could be improved. For example, if there is resource contention on one host, you might see improvement if you move a virtual machine (VM) to a different host.

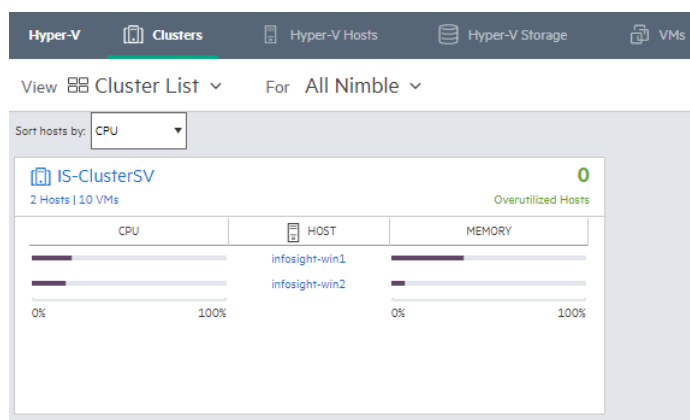
Some of the key data in the Virtualization pages includes details about VMs, which can help you locate and address potential performance problems. When you select the **VMs** tab under **Hyper-V**, you can then go to the **View** menu and select data categories that display the following key metrics.

VM View Menu	Key Metrics
VM Capacity Trend	Top VMs by Data Usage
VM I/O Contention Treemap	VM I/O and Latency Grouped by Hyper-V Storage
VM I/O Contention Trend	Top 10 VMs by IOPS and Top 10 VMs by Latency
VM Memory Contention Trend	Top 10 VMs by Memory Usage.
VM CPU Contention Trend	Top VMs by CPU Usage and Top VMs by CPU Ready

Refer to [Hyper-V VM Analytics](#) on page 36 for more information about capacity and contention trend analytics.

## Hyper-V Cluster Analytics

To display Hyper-V cluster analytics, click **Clusters** from the **Infrastructure** > **Virtualization** > **Hyper-V** menu. The Cluster List page appears in which you can view CPU, Host and Memory analytics for a specific cluster.



Note that Cluster is the highest level object for Hyper-V Analytics. That is, there is no concept of vCenter or datacenter as there is in VMware.

In the above example, only one cluster is shown, but if there were multiple clusters, they would also appear on this page with one cluster per card. You can sort hosts by CPU or memory, which is helpful if you have a large number of hosts or are interested in one particular resource over the other.

You can also see Overutilized Hosts, so if the top host is overutilized (for CPU for example) you can click on the Host name to drill down to further investigate any performance issues.

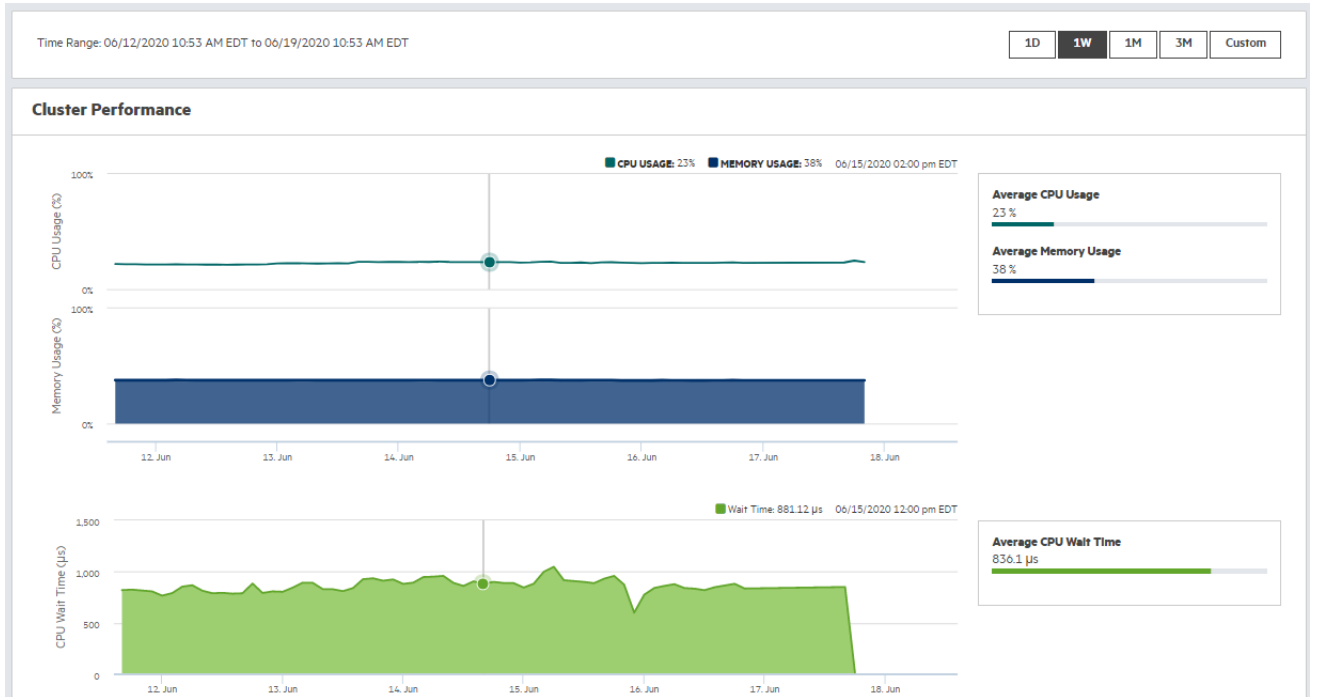
### Displaying Cluster Performance Analytics

To display Cluster Performance analytics, click a Cluster name in the Cluster List. The Cluster Performance page appears:

Hyper-V Clusters Hyper-V Hosts Hyper-V Storage VMs

Cluster IS-ClusterSV VHDs: 36 VMs: 22 Hyper-V Hosts: 2

2 CPU COUNT 511.4GiB TOTAL MEMORY 266.5TiB TOTAL CAPACITY



The Cluster Performance graph enables you to correlate Average CPU Usage and Average CPU Wait Time with Average Memory Usage. You can click the graph and drag the mouse to zoom in on a subset of the performance data, which is helpful if you want to examine a sudden CPU spike or any other performance issue. In addition, you can elect to display data collected over the last day, week, month, 3 months or over a customized date range that you can set by clicking the **Custom** button.

Note that when viewing the Cluster Performance page, you might find it especially useful to display the VM I/O Contention Treemap and the Host Performance Trend analytics pages, which you can select from the **VMs** and the **Hyper-V Host** pull-down menus at the top right corner of the page.

## Hyper-V Host Analytics

To display Hyper-V host analytics, click **Hyper-V Host** from the **Infrastructure > Virtualization > Hyper-V** menu. A detailed summary of Hyper-V host analytics is displayed along with details related to clusters, such as VHDs, VMs and Hyper-V Hosts.

Hyper-V Clusters **Hyper-V Hosts** Hyper-V Storage VMs

View Host List For All Hyper-V Hosts

**Summary** [Download CSV](#)

HOST NAME	CLUSTERS	CPU (PAST 24 HRS AVG)			MEMORY (PAST 24 HRS AVG)	
		USAGE	WAIT TIME	GUEST RUNTIME	USAGE	
infosight-win1	IS-ClusterSV	18%	60.50 µs	6.20%	5.70%	
infosight-win2	IS-ClusterSV	28%	792.06 µs	7.81%	69.35%	
INFOSIGHT-WIN3		N/A	N/A	N/A	N/A	
infosight-win4		30%	356.74 µs	7.90%	6.55%	

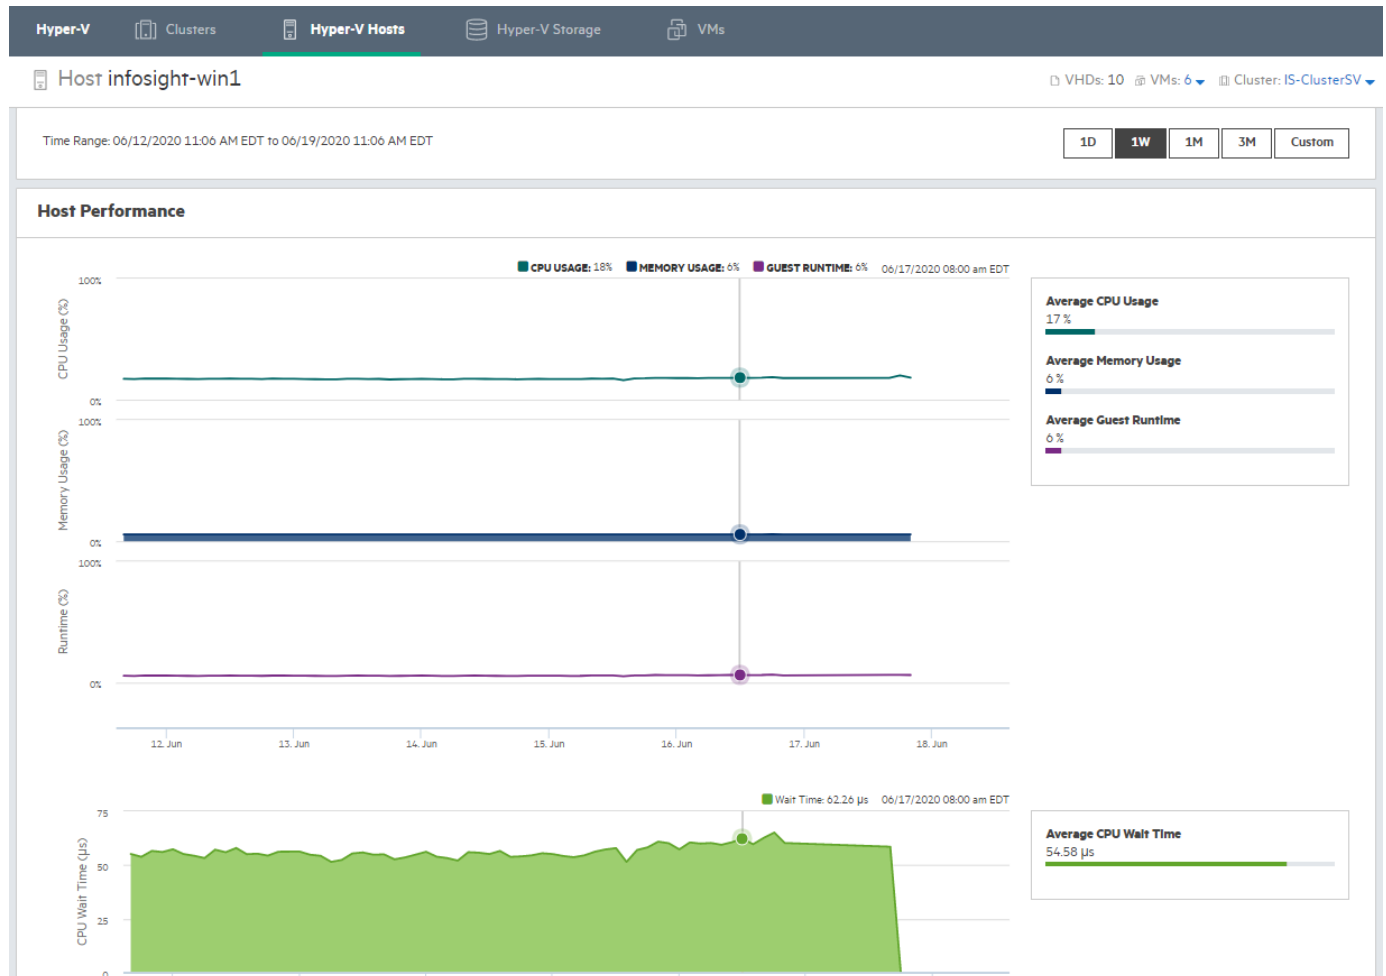




Note that the Hyper-V Host list Summary page employs some terminology that differs from that used in the VMware virtualization portal. For example, for Hyper-V hosts, the term Guest Runtime is used, which is the percentage of CPU time spent running cycles within a guest or within the host. For host infosight-win1 in the above example, you can see that 18% of total CPU resources are being used, but only 6.2% is being used for Hyper-V processing.

### Displaying Host Performance Analytics

To display Host Performance analytics, click a Host name in the Host List Summary page. The Host Performance window appears.



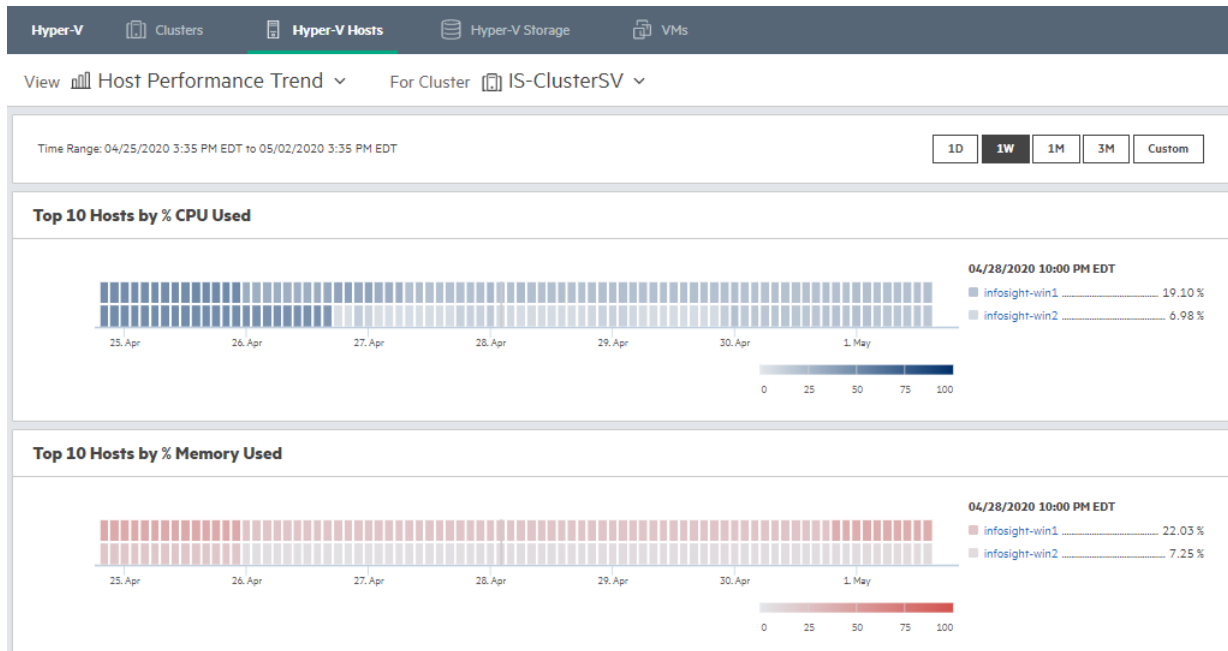
The Host Performance graph that is displayed enables you to correlate CPU and Memory usage with Runtime utilization for the host that you selected. You can click the graph and drag the mouse to zoom in on a subset of the performance data, which is useful when you want to examine a sudden CPU spike or any other performance issue. You can also display data collected over the last day, week, month, 3 months or over a customized date range that you can set by clicking the **Custom** button.

Note that from the **VMs** menu that appears in the top right corner of the page, you can elect to display VM Memory and CPU Contention Trend data for the host. Also, from the **Cluster** menu, you can elect to drill down and display Host Performance Trend and Cluster Performance analytics.

### Displaying Host Performance Trend Analytics

The main benefit of trend analysis is that it enables you to visualize average resource usage over time and can help you identify periods in which there were performance issues.

To display Host Performance Trend Analytics, click **Host Performance Trend** from the Host List View menu.



The Host Performance Trend window displays analytics showing the Top 10 Hosts by % CPU Used and the Top 10 Hosts by % Memory Used.

**Note:** Host Performance Trend only visualizes resources by cluster.

## Hyper-V Storage Analytics

To display Hyper-V storage analytics, click **Hyper-V Storage** from the **Infrastructure** > **Virtualization** > **Hyper-V** menu. The Storage List Summary page appears, which provides a detailed summary of Hyper-V storage statistics.

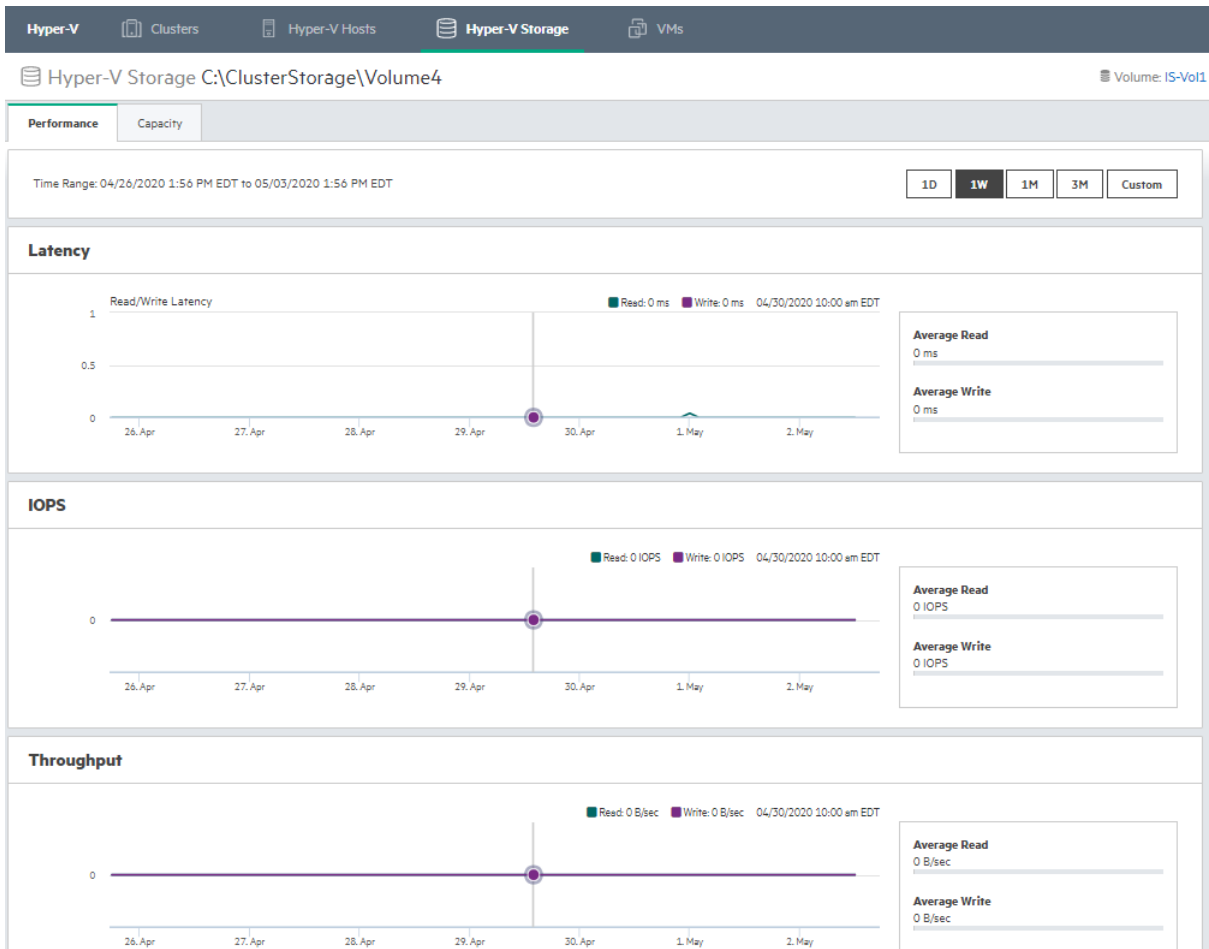
**Note:** Only initialized storage drives are listed. Pass-through drives are not displayed.

NAME	MAPPED TO	STORAGE TY...	VOLUME	ARRAY	FILE SYS...	BLOCK SIZE	USAGE	CAPACITY
\\?.\Volume{8a939c...	IS-ClusterSV	Cluster Volume	Qorum	sjc-array2040	NTFS	4.0 KIB	18.4% (560.6MiB)	3.0 GiB
W:\	infosight-win4	Standalone Volume	Madhan-SV-04011107-...	sjc-array2003	NTFS	4.0 KIB	1.7% (86.6MiB)	4.9 GiB
H:\	IS-ClusterSV	Cluster Volume	Madhan-CSV-04061448-...	c23-array7	NTFS	4.0 KIB	7.6% (611.2MiB)	7.9 GiB
Z:\	INFOSIGHT-WIN3	Standalone Volume	Madhan-SV-04011042-...	sjc-array2003	NTFS	4.0 KIB	0.3% (34.4MiB)	10.0 GiB

Storage Types can be either Cluster Shared Volumes, Standalone (non-clustered) Volumes, or Clustered Volumes. Cluster Shared Volumes are volumes that are mapped to all hosts in a cluster, while Standalone Volumes are mapped to only one host. Clustered Volumes are volumes that are clustered but not added to the CSV.

### Displaying Storage Performance and Capacity Analytics

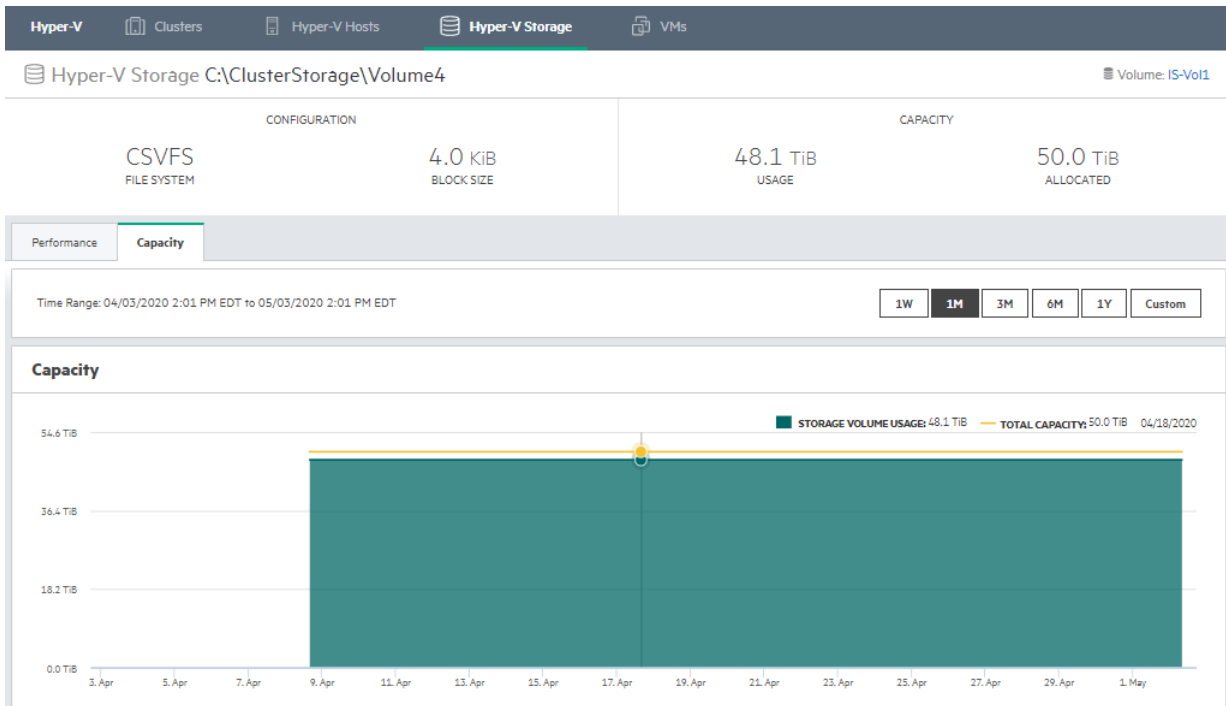
To display Storage Performance & Capacity, statistics, click a storage name in the Storage List Summary page. A window with Performance and Capacity tabs is displayed.



The performance graph that is displayed enables you to correlate Latency with IOPS and Throughput. As with other performance graphs, you can click the graph and drag the mouse to zoom in on a subset of the performance data, which is helpful when you want to examine a sudden performance spike or any other apparent performance issue. In addition, you can elect to display data collected over the last day, week, month, 3 months or over a customized date range that you can set by clicking the **Custom** button.

To display capacity data, click the Capacity tab.





The Capacity tab shows historical space utilization for the selected volume, which can help you discern growth patterns over time. You can also hover the mouse pointer over the capacity graph to display storage volume usage statistics over a specified time period.

## Hyper-V VM Analytics

To display Hyper-V VM analytics, click **VMs** from the **Infrastructure** > **Virtualization** > **Hyper-V** menu. The VM List Summary page appears, which displays an overview of your entire VM footprint:

The screenshot shows the Hyper-V VM List Summary page. The top navigation bar includes 'Hyper-V', 'Clusters', 'Hyper-V Hosts', 'Hyper-V Storage', and 'VMs'. Below the navigation bar, there are 'View' and 'For' dropdown menus. The main area has tabs for 'Summary' and 'Performance', with 'Summary' selected. A 'Download CSV' button is in the top right corner. The table below lists VMs with columns for Name, Host, Hyper-V Storage, Array, Total Usage, CPU (Past 24 Hrs), MEM (Past 24 Hrs), and Last Activity.

NAME	HOST	HYPER-V STORAGE	ARRAY	TOTAL USAGE	CPU (PAST 24 HRS)		MEM (PAST 24 HRS)		LAST ACTIVITY
					AVERAGE	PEAK	AVERAGE	PEAK	
Infosight-VM2	infosight-win1	C:\ClusterStorage\Volume3	c23-array7	120.2 GiB	0.00%	0.00%	0.0 B	0.0 B	an hour ago
IS-Move-VHD	infosight-win4	O:\	sjc-array2004	80.3 GiB	0.70%	1.42%	409.7 MiB	429.2 MiB	35 minutes ago
IS-VM7-New	infosight-win4	L:\	sjc-array2003	80.3 GiB	0.63%	1.12%	1.1 GiB	1.1 GiB	35 minutes ago
MemoryCPU2	infosight-win2	C:\ClusterStorage\Volume3	c23-array7	80.3 GiB	0.76%	1.18%	2.0 GiB	2.0 GiB	an hour ago
MemoryCPU3	infosight-win2	C:\ClusterStorage\Volume3	c23-array7	80.3 GiB	0.84%	1.28%	1.6 GiB	2.0 GiB	an hour ago
MemoryCPU...	infosight-win2	<b>2 Hyper-V Storages</b>	<b>2 Arrays</b>	204.7 GiB	0.24%	1.05%	2.7 GiB	3.4 GiB	an hour ago

As with other Hyper-V Summary pages, you can click on hyper-linked column item to drill down to other analytic pages. For example, if you click a Host name in the Summary page, you can display host performance information for that host. If you click a Storage name, you can display storage performance information for that host. For information about these pages, refer to [Hyper-V Host Analytics](#) on page 32 and to [Hyper-V Storage Analytics](#) on page 34.

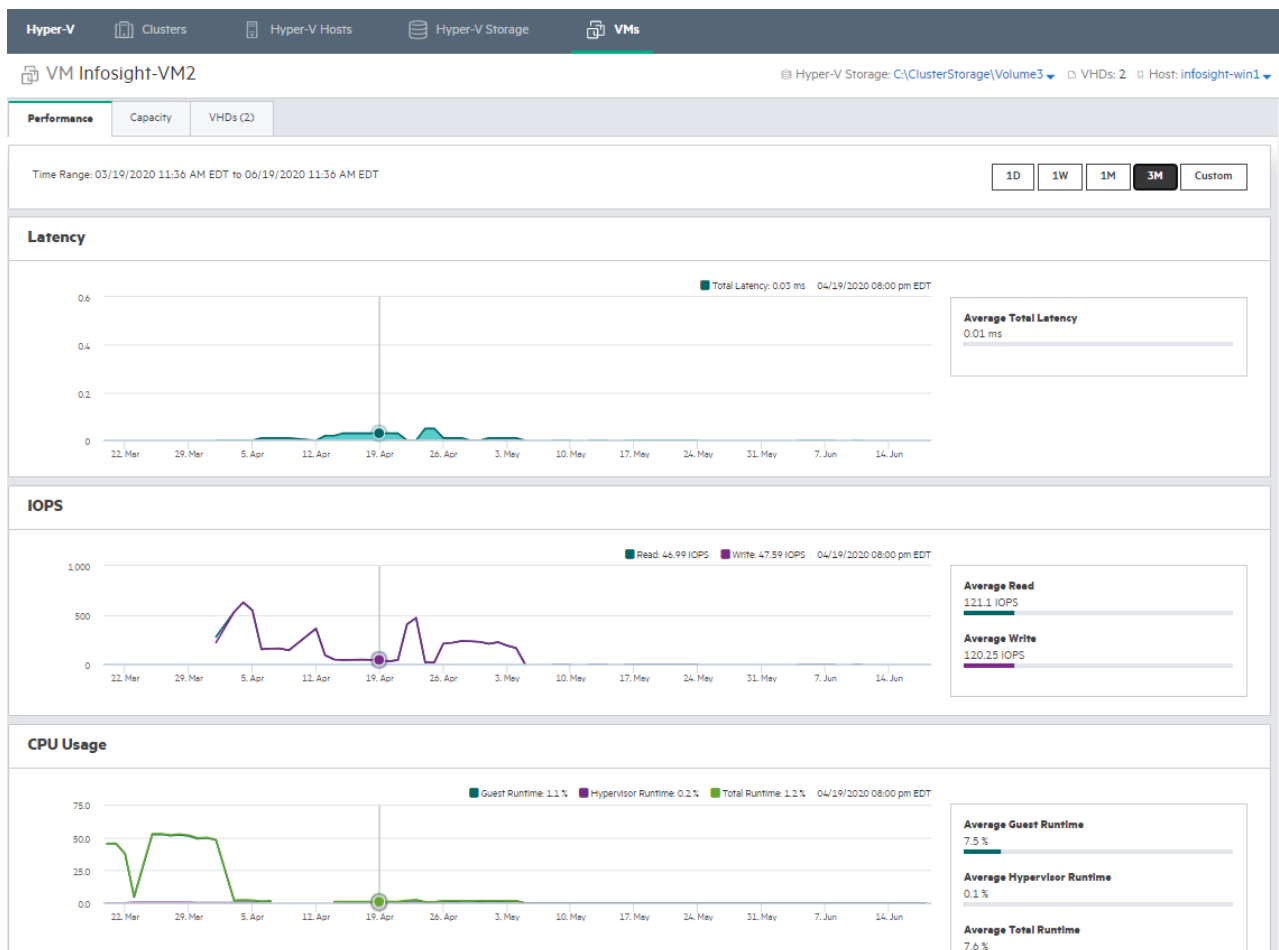
**Note:** For additional information about the VM List, see [Current Behavior for Clustered/Standalone Hosts](#) on page 44.

From the Performance tab, you can display performance data, such as Average IOPS, Average Latency, and Average Throughput for each VM:

Hyper-V Clusters Hyper-V Hosts Hyper-V Storage VMs								
View VM List For All VMs								
Summary Performance								Download CSV
NAME	AVG IOPS (PAST 24 HRS)			AVG LATENCY (PAST 24...		AVG THROUGHPUT (PAST 24 HRS)		
	READS	WRITES	OVERALL	OVERALL	READS	WRITES	OVERALL	
Infosight-VM2	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
IS-Move-VHD	0.0 IO/sec	0.0 IO/sec	0.0 IO/sec	0.1 ms	0.2 B/s	104.5 B/s	104.7 B/s	
IS-VM7-New	0.0 IO/sec	0.4 IO/sec	0.4 IO/sec	0.2 ms	0.1 B/s	1.5K B/s	1.5K B/s	
MemoryCPU2	0.0 IO/sec	0.1 IO/sec	0.1 IO/sec	7.7 ms	0.1 B/s	191.3 B/s	191.4 B/s	
MemoryCPU3	0.0 IO/sec	0.1 IO/sec	0.1 IO/sec	7.6 ms	0.1 B/s	192.2 B/s	192.3 B/s	
MemoryCPUTest1	0.0 IO/sec	0.1 IO/sec	0.1 IO/sec	7.8 ms	0.1 B/s	168.6 B/s	168.7 B/s	

### Displaying VM Performance, Capacity and VHD Analytics

To display VM Performance, Capacity, and VHD analytics, click a VM name in the VM List Summary page. A window appears with a tab for each of these metrics:

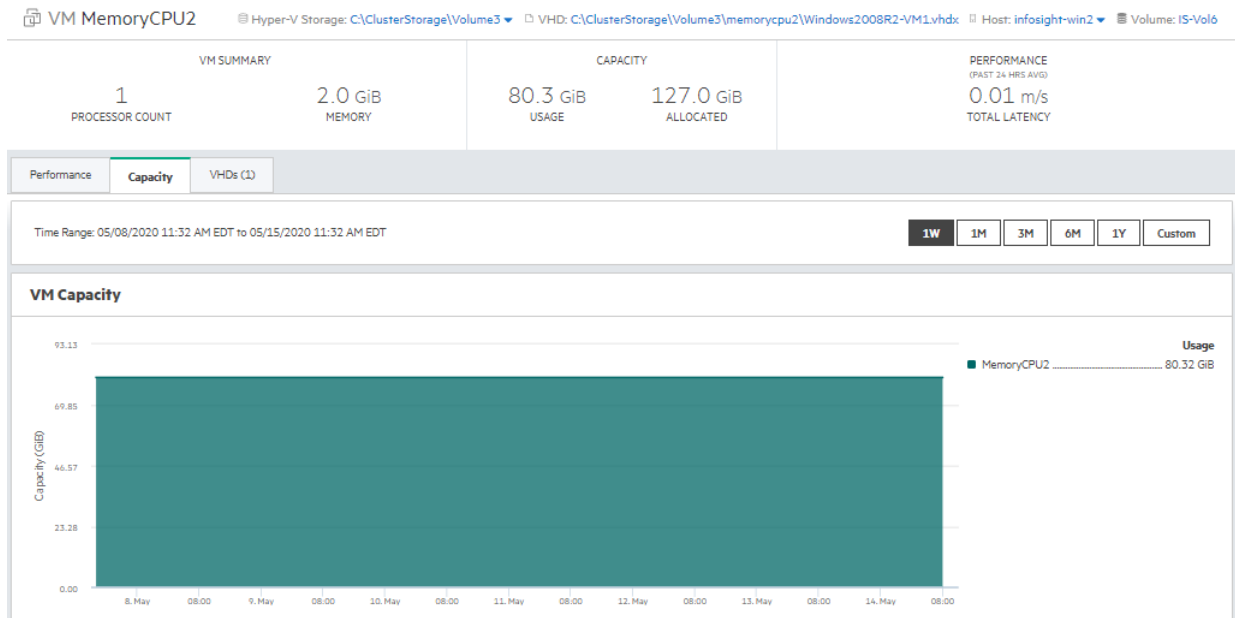


In the Performance tab, you can move the mouse pointer in the graph to correlate VM Latency with underlying IOPS and CPU utilization. You can also click the graph and drag the mouse to zoom in on a subset of the performance data, which is helpful when you want to examine a sudden CPU spike or any other performance issue.

Note that for CPU usage, there are some new Hyper-V specific terms that are not used in VMware Virtualization analytics:

- **Average Guest Runtime** is the percentage of CPU resources used to perform operations in the VM. This is the actual processing the servers are doing for their application workloads.
- **Average Hypervisor Runtime** is the percentage of CPU resources that are used to run the server in a virtualized environment. In other words, this is essentially the virtualization overhead or the virtualization tax on the server
- **Average Total Runtime** is the sum of the Guest and the Hypervisor Runtimes. This Guest and Hypervisor runtime breakdown enables you to visualize the penalty of running a workload on a VM versus running the workload on a physical server.

Click the Capacity tab, to track capacity consumption of the VM over a specified time frame.



Click the VHD tab to display the average hourly statistics for IOPS, Latency, and Throughput for all detected VHDs (Virtual Hard Drives) on your system:

The screenshot displays the 'VHDs (1)' tab for VM MemoryCPU2. The top summary section is identical to the Capacity tab. The main table shows the following data:

NAME	HYPER-V STORA...	HYPER-V VO...	TOTAL SIZE	AVG IOPS (PAST 24 HRS)			LATENCY (PAST 24...	AVG THROUGHPUT (PAST 24 HRS)		
				READS	WRITES	OVERALL	OVERALL	READS	WRITES	OVERALL
IDE_0_0	C:\ClusterStorage\Volum...	IS-Vol6	127.0 GiB	0.0 IO/sec	0.1 IO/sec	0.1 IO/sec	0.0 ms	0.0 MiB/s	0.0 MiB/s	0.0 MiB/s

**Tip:** When working in the VM Performance, Capacity, and VHD tabs, it can be useful to display related capacity and contention trend analytics. To access this data, select a choice from the **Hyper-V Storage** and the **Host** pull-down menus that appear at the top right corner of each page.

### Displaying Capacity and Contention Trend Analytics


From the VM List Summary page, you can display various Capacity and Contention Trend analytics pages by clicking the page from the VM List View menu.

The Contention Trend analytic pages enable you to identify Noisy Neighbor VMs that are using up the most IO, CPU and Memory resources, thereby reducing the availability of these resources to other VMs in the same resource domain.

The following are examples and brief descriptions of the capacity and contention analytic pages that are currently available:

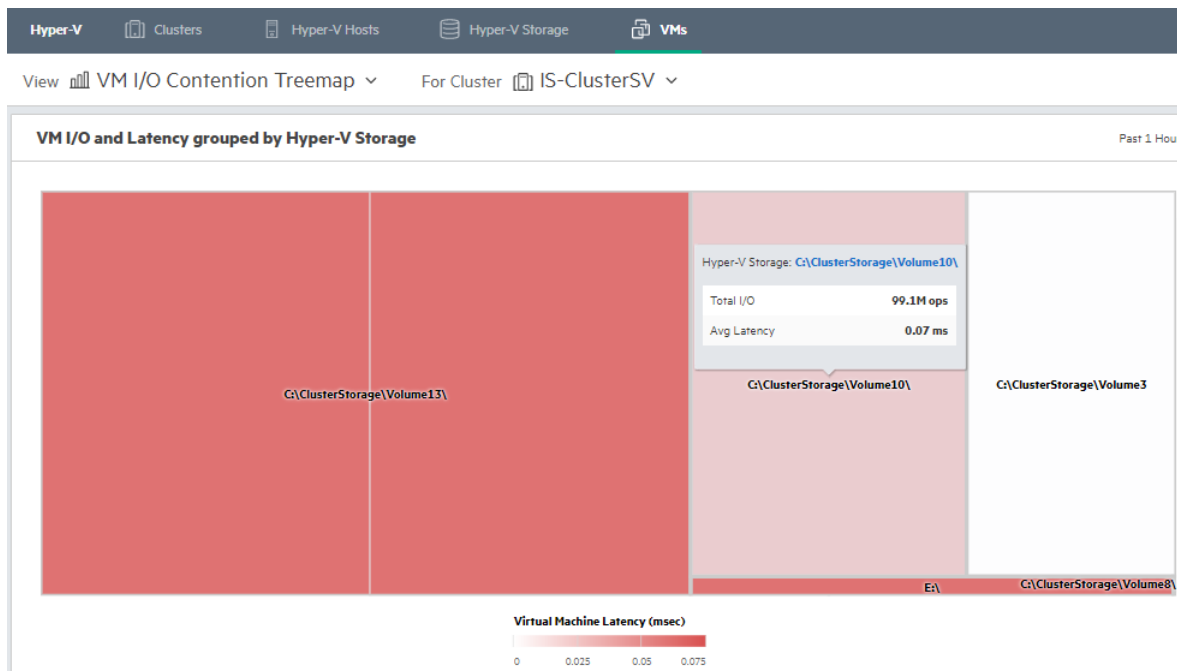


## VM Capacity Trend



This VM Capacity Trend page can help you identify the VMs that are using the most space over time. Like with most trend analysis, this information can help visualize average resource usage over time and can help identify performance issues or operational issues.

## VM I/O Contention Treemap



The VM I/O Contention Treemap shows VM I/O and Latency Grouped by Hyper-V Storage. You can use this page to visualize I/O and latency for a given host or cluster. The view is called a heat map. The larger the box, the more IOPS and the darker

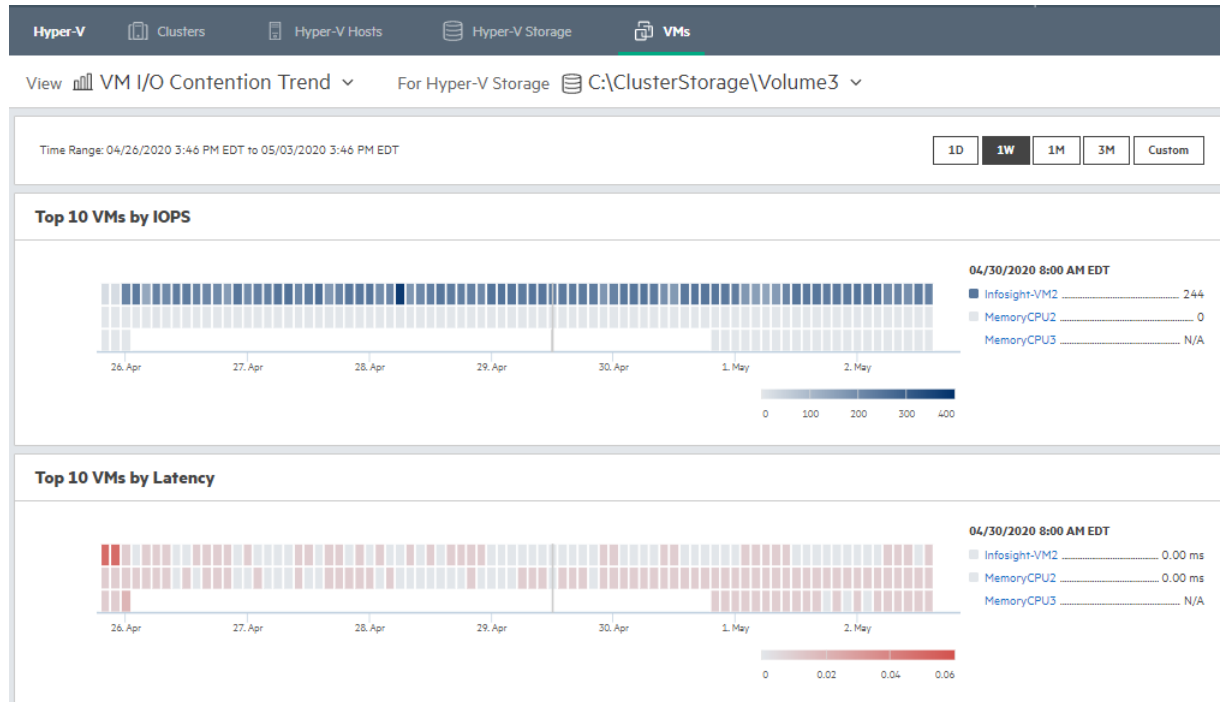


the color the more latency. The page is designed to show you hotspots with Hyper-V storage objects with respect to the volumes that are busy and the volumes that are experiencing the most latency

Note that there are two layers of contention accessible from the treemap for a given cluster or host, both shown in descending sequence of total I/O (size of box) along with values of total I/O and average latency. The outer layer shows storages/volumes, and within each storage (the inner layer) you can see contention of VMs. If you click a volume box, you will drill down to the inner layer where you can see VM contention within the volume. You can then navigate back to the outer layer by clicking a **Back** button.

**Note:** For additional information about the VM I/O Contention Treemap, see [Current Behavior for Clustered/Standalone Hosts](#) on page 44.

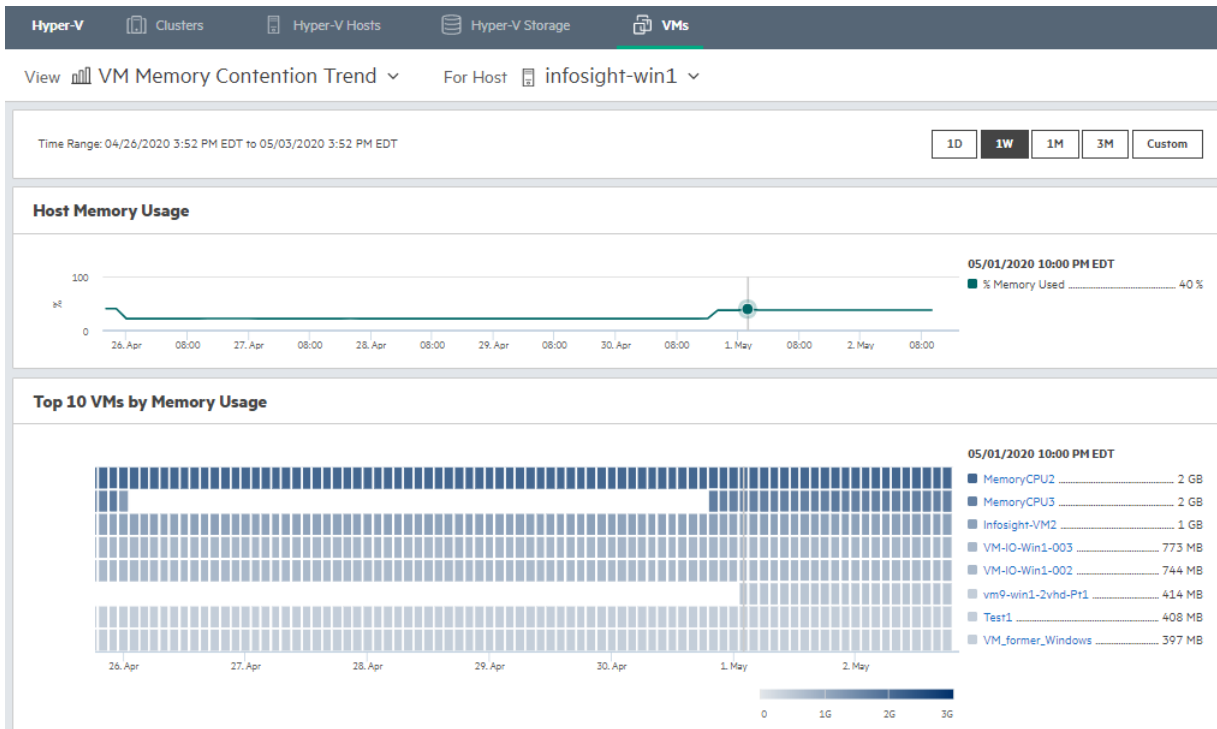
## VM I/O Contention Trend



The VM I/O Contention Trend page displays the Top 10 VMs by IOPS and Top 10 VMs by Latency. The Top 10 pane shows the busiest VMs in blue, the bottom pane shows VM latency in red. This page enables you to identify the VMs that are the busiest as well as those experiencing the highest latency (which can indicate potential performance issues). This page also enables you to identify noisy neighbor VMs that are using up the most IOPS relative to other VMs in the resource domain.

## VM Memory Contention Trend





The VM Memory Contention Trend page displays the VMs that are using the most memory. This page also enables you to identify noisy neighbor VMs that are using up the most memory relative to other VMs in the resource domain.




### VM CPU Contention Trend



View  VM CPU Contention Trend ▼ For Host  infosight-win2 ▼

Time Range: 05/10/2020 8:02 PM EDT to 05/17/2020 8:02 PM EDT

1D 1W 1M 3M Custom

**Host CPU Usage****Top VMs by CPU Usage****Top VMs by CPU Ready**

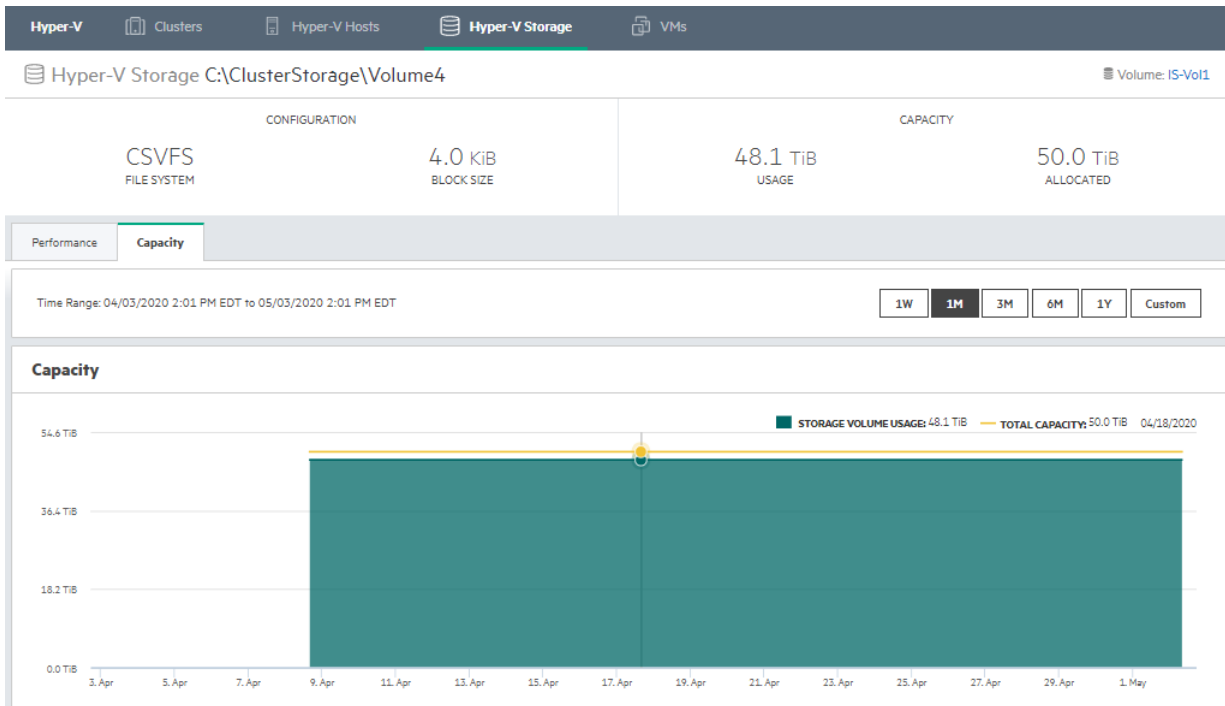
The VM CPU Contention Trend page identifies the VMs that are using the most CPU. For each Hyper-V Host, the top pane shows VM CPU usage in blue, and then any VM suffering from CPU-Ready in red at the bottom. CPU Ready means that there is processing to do but the VM is waiting for CPU resources from the Hyper-V host. Oftentimes CPU Ready is due to a noisy neighbor that is using a large share of the available resources.

## Navigating Between Analytic Stacks

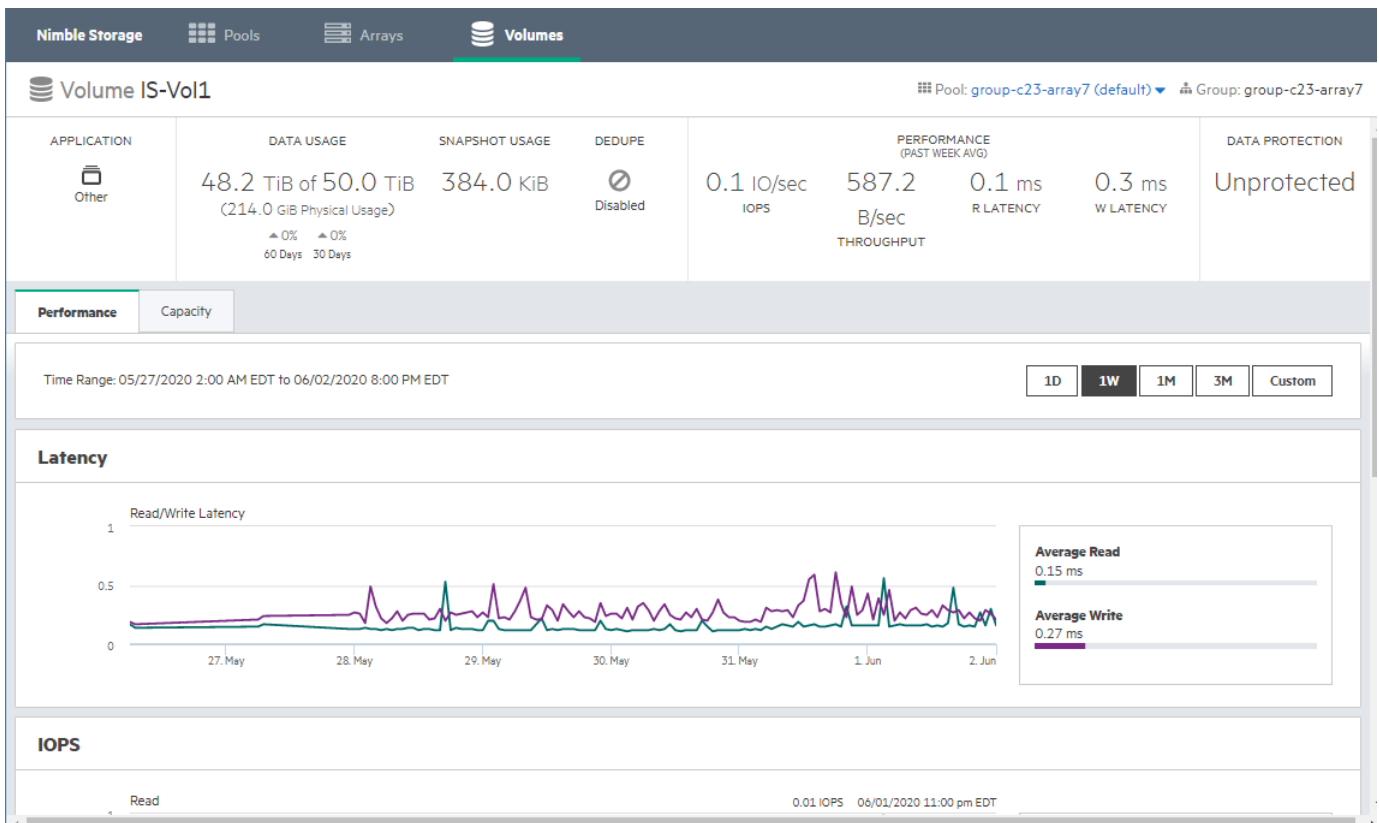
When working in the Hyper-V Virtualization pages you can easily jump from the Hyper-V stack to other analytics stacks.

For example, suppose you are in the Hyper-V Storage Capacity page and you notice that storage is filling up faster than you expect.





In this instance, you could click on the Volume link in the top right of the page, to jump to a storage centric view of the data in InfoSight:

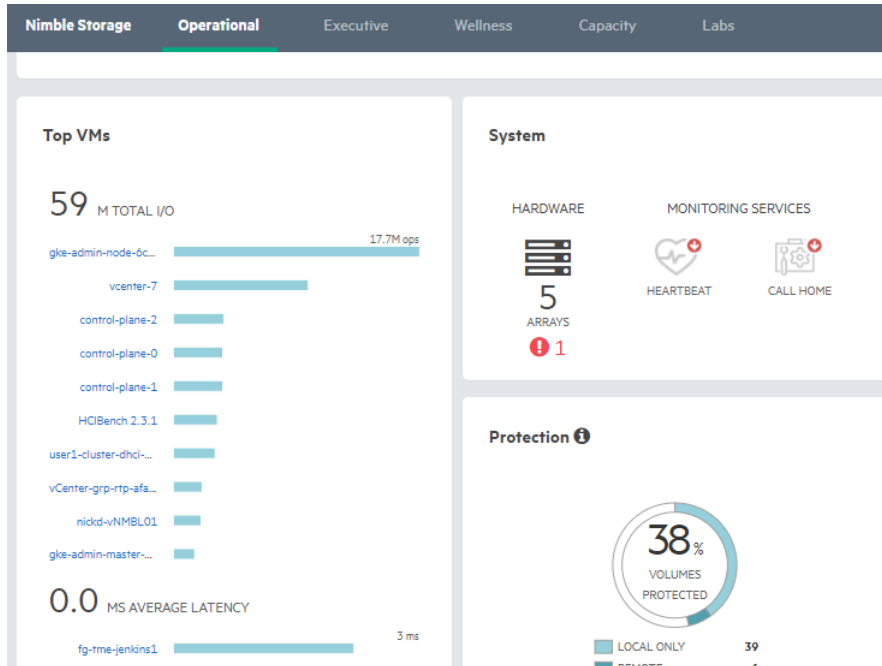


In this InfoSight storage view, you could then check to see whether DEDUPE is enabled or whether snapshots were set up for data protection. Then you could return to your Hyper-V Virtualization stack by simply clicking the browser back button.



## Displaying Data from Multiple Stacks

HPE InfoSight supports displaying data across multiple virtualization stacks. For example, if VMware or Hyper-V Analytics (or both) have been installed in an InfoSight organization, then the **Top VMs** list in the Operational Dashboard will include the VMs from each of these respective stacks.



If you click a VM in the **Top VMs** list, the respective VMware or Hyper-V Virtualization pages appear for that VM.

## Current Behavior for Clustered/Standalone Hosts

This topic provides information about InfoSight views of your Hyper-V environment with HPE Alletra 6000, HPE Nimble Storage in the event that one or more of the following conditions are true:

- You have a “standalone” host (defined below) in your environment
- You have a clustered host mapped to standalone physical storage

### Background

In a Hyper-V environment, hosts are typically grouped together in clusters for high availability. This allows the creation of highly available or clustered VMs. Clustered VMs need clustered storage. To achieve this, storage volumes are mapped to all hosts in the cluster. In other words, when we say storage is clustered, it is mapped to a “cluster”. If VMs are not highly available, their storage is typically only mapped to the host on which the VM runs. Such hosts are called “standalone” hosts.

In some cases, you can create a standalone VM on a host that is part of a cluster. Or, you can also map some storage volumes to only one host in the cluster. This causes some hosts to be designated as clustered and standalone, depending upon the context.

### Current Behavior:

The scenarios mentioned above can create several behaviors in the InfoSight GUI that can be considered atypical for following views:

#### **VM List View**

This view is meant to list the VMs and their basic attributes. You can also choose to limit your view to a particular ownership hierarchy (Cluster-> Host) using a scope selector menu.

Note the following:

- If you have a standalone host, it will be separately listed at the root hierarchy in the scope selector of VM list view.



- If you have a standalone disk mapped to a clustered host, it will be listed twice in the scope selector:
  - Under concerned cluster
  - As a standalone host

However, in the second scenario, regardless of whether you select host scope from cluster hierarchy or as standalone host, all VMs owned by that host will always be listed.

### **VM I/O contention Treemap View**

This view is intended to show the VMs and their contention for I/O (storage resources)

Also note the following

- If you have a standalone host, it will be listed separately at the root hierarchy in the scope selector of the VM I/O contention treemap. You will be able to see the VMs contending for the I/O on the standalone disks mapped to this host.
- Hosts that are exclusively clustered (e.g. no standalone disks mapped to them) will not be listed in scope selector. (Clustered storage resources are shared)
- If you have one or more standalone disk(s) mapped to a clustered host, it will be listed separately as a standalone host in the scope selector.
- A standalone host (or a clustered host with some standalone disks mapped to it) will always be listed in the scope selector, regardless of whether you have any application VMs deployed on this host.



# Monitoring Your Array Environment

The main purpose of HPE InfoSight is to enable you to better monitor your environment and determine when intervention is needed. To enable you to get the most detailed view of your arrays, you can use the various apps available on the Labs page.

## Monitor Equipment

You can monitor the overall health of all your arrays in HPE InfoSight. You can also view details about many aspects of the arrays.

### Procedure

1. On the Wellness Dashboard, go to the **Issue Details** section which lists all the alerts for all of your arrays.
2. Click to expand a particular alert to view the issue details.

## Potential Impact Score

Although latency is one of the key factors that you use to determine the health of your arrays, not all latency is the same and in some scenarios, higher latency is to be expected, such as when you perform a backup or transfer large files, but at other times, higher latency can indicate a problem.

To help you to determine when higher latency might be a problem, you can use the potential impact score. This score helps you to differentiate between when latency is normal and to be expected and when it is a sign of a potential problem.

Go to the performance graph for a particular pool. The graph indicates latency on an hourly basis for an entire week. The red bars indicate areas of concern. When you hover over the chart, a popup appears that indicates the read I/O per second, the latency, and the potential impact score at a particular point in time.

You can use the potential impact score to determine if further action is required on a particular array.

## Monitor Volumes

HPE InfoSight can be used to gain insight into your volume workloads, and retrieve detailed statistics. Data is collected from all installed arrays and can be filtered to show the required metrics as needed.

### Procedure

Go to **Infrastructure** > **Volumes**.

The Volume List page opens. To see the details about a particular volume, click the volume name.

## Monitor Synchronously Replicated Volumes

Synchronous Replication is a new feature in release V5.1.1 and higher that enables volumes replicas to be created synchronously. This feature helps protect against data loss in the event of an array or site failure. For information about synchronous replication, refer to the *GUI Administration Guide*.

When you create a volume collection in the array OS, you can specify whether volumes are protected locally or remotely. When volumes are protected locally, it means the volumes are protected by local snapshots that do not have replication schedules associated with them. When volumes are protected remotely, it means the volumes are protected by a replication schedule



that replicates snapshots to a replication partner based on a time frame that you specify in the schedule. Within the remote protection category, volumes can now have the following protection level subcategories:

- Remote Periodic
- Remote Synchronous

Remote Periodic indicates that volumes are backed up periodically, such as every 10 minutes or for some other user specified time period. Remote Synchronous indicates that the volumes are backed up synchronously, which means the volumes are backed up in near real time. HPE InfoSight enables you to visually identify and monitor volumes with these replication type protection levels by indicating the protection level subcategories in the Protection Level column of the Volume List window.

To display the Volume List window, log in to InfoSight and click **Dashboards** > **Infrastructure** > **Volumes**. The Volume List appears with the replication type protection method indicated in the Protection Level column.

Note that you can also see an indicator that a protection level is Remote when you click on an individual volume in the Volume List window.

## The Executive Dashboard

HPE InfoSight provides you with several different reports to keep you informed of the health of your arrays.

The Executive Dashboard is one place where you can go to get an overall view of the state of your arrays. The Executive Dashboard includes quick views of the following factors:

### Space Savings

This view includes the capacity used, compression savings, savings due to cloning, and capacity used (post savings)

### Operational Efficiencies

This view includes cases that were closed automatically as opposed to cases that were closed manually.

### Data Protection

This view includes graphs for RPO (recovery point objective), local retention configuration, and disaster recovery

### Upgrade Needs

This view provides upgrade recommendations, if any are needed, based on the last four weeks of activity.


## Streaming Performance Telemetry Data

InfoSight typically receives statistical information about monitored arrays at set time intervals. In some instances, these intervals can range from several hours to a full day. However, if you are monitoring arrays that run on release V4.x or higher and you enable stats streaming, the array performance telemetry data could be visible in InfoSight in as little as an hour.

You can see the this reduced data latency on an array by comparing its performance graph when stats streaming is disabled to the graph when stats streaming is enabled. An array with stats streaming enabled will have performance stats available within an hour in contrast to arrays that have stats streaming turned off .


For example, in the following arrays, which do not have stats streaming enabled, data is delayed by almost 24 hours.





SN048

However, in the 2nd, 3rd and 4th arrays below, in which streaming is enabled, you can see much more recent data:



SN049



**Tip:** The word **Streaming** appears in green next to the arrays that have stats streaming enabled.

### Enabling Stats Streaming

#### Before you begin

To enable stats streaming on an array, the array must be running release V4.x or higher.





**Procedure**

1. Log in to HPE InfoSight at <https://infosight.hpe.com>.
2. Go to the Main Menu and click **Alletra 6000, Nimble Storage**→**Settings**→**Telemetry Settings**
3. In the list of arrays, locate the array on which you want to enable streaming.
4. In the **Streaming** column, click the streaming icon to change the streaming state from **Off** to **ON**.

**Viewing Arrays on which Stats Streaming is Enabled**

You can see the reduction in data latency that results from enabling stats streaming in the following InfoSight pages:

- The Operational Dashboard
- The Pool Performance page, which displays when you click an array in the Operational Dashboard.
- The Volume Performance page, which displays when you click a volume in the Volume List window.
- The Volume Performance Contention page, which displays when you select **Volume Performance Contention** from the Volume List menu.

Note that in the Operational Dashboard, the word **Streaming** appears next to arrays that have stats streaming enabled.



## Viewing dHCI Data

If you have an HPE Alletra 6000, HPE Nimble Storage dHCI-enabled array, you can view data about your dHCI environment in the Infrastructure Solutions section of InfoSight.

For information on how to set up dHCI, see *Getting Started with dHCI* in the documentation portal.

### HPE InfoSight Analytics and dHCI

HPE InfoSight provides detailed information about the dHCI environment. You can use this information to identify potential problems and address them before they become problems. In addition, you can use the InfoSight data to make changes to your setup to make it more efficient.

InfoSight organizes information based on the dHCI clusters. It provides data about features such as the virtual machines (VMs), the host CPU, the host memory usage, the array capacity, and the group leader.

The main page is an overview of the dHCI environment. The **dHCI Details** page displays additional information about the servers, compute, storage, and VMs. The **Wellness** page provides an overview that shows whether the array has any server or storage issues.

To get to this information, log on to HPE InfoSight and select **Infrastructure** > **Solution** > **dHCI**.

---

**Note:** From InfoSight, you must enable streaming from the array where dHCI is hosted. You must also enable the VMware setting to collect data about VMware. See [Enable InfoSight to Provide dHCI Information](#) on page 50.

---

### Enable InfoSight to Provide dHCI Information

To take advantage of HPE InfoSight's proactive support functions, you must configure the HPE Alletra 6000, HPE Nimble Storage array where dHCI is installed to provide data to InfoSight.

#### Before you begin

You must have the following:

- dHCI must be installed.
- At least one VM must be on a datastore.
- You must have an InfoSight account.
- Ports 80 and 443 must be open between the vCenter Server and the array.

#### Procedure

1. Log in to InfoSight. You can enter [Infosight.hpe.com](https://infosight.hpe.com) in your browser or, from the dHCI dashboard, you navigate to the Configuration Checks page and click the HPE InfoSight link.
2. From the settings menu (the gear icon) select **Telemetry Settings**.
3. Locate the dHCI-enabled array and click the button for Streaming to On. This button enables data streaming from the array.
4. In the same row, click the button for VMware to change it to **On**. This button allows data to be collected from VMware.

---

**Note:** If you do not set the VMware button to **On**, you will receive diagnostic streaming data from the array only. You will not receive any VMware data. As a result, you will not be able to proactively identify any issues with the data from VMware.

---



# Support of Your Arrays

One of the strengths of HPE InfoSight is the ability to monitor your arrays and proactively resolve issues before they become serious problems. In those cases when a problem does arise, you can get the most out of your support experience by being prepared and using the tools that are provided in HPE InfoSight.

## Validate Contact and Install Location Information

If you are working with an HPE Alletra 6000, HPE Nimble Storage account, you should validate your contact and install location information to provide the support team with information that they can use to help you quickly resolve problems.

### Before you begin

This procedure must be performed by an admin.

### Procedure

Validate contact and install location information for each of your arrays.

- a) Click the Gear Icon in upper-right hand corner.
- b) Select **Asset Registry**.
- c) Click the **Address Book** Tab.
  - 1 To add a new address book entry, click the **New** button.
  - 2 To update an existing address, click the **Edit** button
- d) Click the **Asset Registry** tab to Review the Delivery address for the asset you want to validate to confirm that it is correct.
- e) Change the install Address drop down to the entry created in the Address book.
- f) Change the Delivery preference drop down to the preferred delivery preference.
- g) Click the **Confirm** button to Validate the address.

## Maintenance Windows

On the Wellness dashboard, you can use Manage Maintenance Windows to notify Support of any upcoming maintenance being carried out in your environment which may impact the array. Setting a maintenance window prevents maintenance-related automated support cases being generated for the duration of the window.

To access the Maintenance Windows page, go to **Dashboards > Wellness > Manage Maintenance Windows**.

You can disable automated support case generation during planned maintenance from the Manage Maintenance Windows page. After a maintenance window is set, no further communication with HPE Support is required.

Examples of maintenance activities when Support recommends setting the Maintenance Window are as follows:

### Array related activities:

- HPE Alletra 6000, HPE Nimble Storage Cache upgrades
- HPE Alletra 6000, HPE Nimble Storage Controller upgrades

### Datacenter related activities:

- Network related maintenance (VLAN changes, switch port configuration)
- Power related maintenance (UPS changes)
- Hardware relocation (Moving equipment)



- Air Conditioning (HVAC maintenance)

## Schedule Maintenance Windows

The time specified for the maintenance window should always be array local time based on the installation address. You can set up the system for repeatable occurrences so you only have to input your defined maintenance window one time.

---

**Note:** Automated support case generation is disabled only for maintenance related events that are encountered during the defined period on the HPE storage group. Support cases related to any unexpected array behavior continue to be generated during this period, the same as cases triggered through longer term analytical recommendations unrelated to the maintenance window.

---

### Procedure

1. Go to **Dashboards** > **Wellness** > **Manage Maintenance Windows**.
2. Click **New** and enter the appropriate information in the fields.
3. Select the arrays for which you want to create a maintenance window.
4. Click **Create**.
5. The Details view opens displaying the details about the newly-created maintenance window.
6. (Optional) You can edit the maintenance window by clicking **Edit**.

## Support Area Overview

The Support area of HPE InfoSight provides functionality to users with admin privileges to manage interactions with Support.

To access the page, click the InfoSight Main Menu () , click **Resources**, and then under Alletra 6000, HPE Nimble Storage click **Support**.

You can perform the following tasks in the Support area:

### Support Page

- File new support requests
- Search all case fields for all cases, including the description and resolution fields

The search does not include the following items:

- Case comments
- Emails
- Attachments
- Linked knowledge documents or articles
- Filter all cases

Add text in text boxes in column headings to create filters.

- Sort all columns
- Review status of support cases

To access the Case Details page, click on a particular case number.

### Case Details Pages

- Add comments to support cases

After a case is created, you can click the case number to show the case details page but you cannot edit any fields. You can only add comments or resolve the case.

Adding comments creates a conversation with a technical support engineer. The customer case contact is notified as responses are posted.



- Upload attachments to cases
- Resolve support cases

If you add notes on the case resolution, the notes are available to a technical support engineer.

**Note:** After a case is resolved, you cannot reopen it. You must call HPE Support to reopen a case.

## Create a Support Case

**Note:** If your case is critical (P1), you cannot create a case but must call Support.

When you open a case, it is important to include as much data as you can in the **Description** field to ensure that the assigned support engineer is able to attend to the issue as efficiently as possible. The following list details some of the data that you might want to include:


- Symptoms of the error condition, including any error messages or notifications
- Time when the issue began occurring and how often the issue arises
- Steps to reproduce the issue
- Whether the task worked previously and if so, when, or if the task in question was never successful
- Whether any changes occurred in the environment that may have triggered the issue
- Whether any troubleshooting has been carried out, either using existing self-service resources (such as support KB articles) or with other vendors (such as host OS vendor, switch vendor)
- Environmental data (for example, the host OS version, the number of hosts impacted, and the affected volume or volumes)

If you have supporting documentation for the case, you can upload attachments after you create the case.

### Before you begin

You must have admin privileges to create a support case.

### Procedure

1. From the InfoSight Main Menu (  ), go to **Resources** > **Support**.
2. Click **New Case**.
3. Select the severity from the **Severity** list.
4. Type in a brief summary of the problem including as much information as possible.  
While you type, suggestions are presented below the text field.  
After you click out of the **Summary** field, article titles appear on the right side of the screen.
5. Optional: Click one of the article suggestions to read the article.
6. Type in a description of the issue, including as much of the suggested information as possible.  
The list of articles continues to update as you include more information in the description field.
7. Select the affected array from the **Affected Array** list.  
You must select an array before you can submit the case.
8. Click **Submit Case**.

### Results

After you have submitted the case, a support case number is displayed. Use this number if you need to contact Support. The case is processed in accordance with your SLAs.

After a case is created, you can click on the case number to show the case details page, but you cannot edit any fields. You can only add comments, upload attachments, or resolve the case.



**Note:** If the case is assigned to an engineer, it is shown on HPE InfoSight as Assigned. If the case is in the support queue, it appears as Pending. If the case is an auto-closed case, it appears as N/A.

See <https://www.nimblestorage.com/customer-support> for more details.

## Security of Customer-Originated Data

HPE does not analyze or use customers' performance and system information in a customer-identifiable way other than for customer support purposes. However, support does review, store, analyze and use aggregated performance and system information, nor does it come into contact with any information that can be associated with the customer in any array support method.

The storage array products use up to four different support methods, outlined in the Array Support Methods table, to communicate with customer support staff and customer support automated systems to help improve the reliability and availability of storage arrays. One of those four array support methods is always enabled. The other three methods are opt-in. See [Configure Arrays to Send Data to HPE InfoSight](#) on page 19 to learn how to enable the remaining support methods.

Each of those opt-in array support methods is disabled by default until the customer uses the array OS to enable them. No substantive customer-originated data, such as files stored on the arrays or personally identifiable information, is communicated back to support.

The following table describes each array support method and the performance information that it communicates back support.

**Table 2: Array Support Methods**

	Heartbeat	DNA	Email Alerts	Tunnel
What is It?	A periodic transmission (via HTTPS) of a small set of packets from a given array to support. A heartbeat indicates that a given array is powered on and functioning appropriately.	A daily transmission (via HTTPS) of more detailed diagnostic information from a given array to HPE support. This diagnostic information might be reviewed periodically by Support.	Email notifications that are sent in the event of an error condition by a given array to internal customer email addresses designated by the customer, HPE Support, or both. The customer chooses who receives these notifications in the array OS. If email notifications go to HPE, these messages might be reviewed periodically by HPE Support.	An encrypted virtual private network tunnel ("VPN tunnel"), which might be established under customer control, outbound from a given array to HPE Support. Note that this VPN connection is a Secure Socket Layer ("SSL") tunnel with a key certification at HPE.
Information Communicated to HPE	Only the serial number and basic system health information for a given array.	System event logs and system configuration and operational statistics for a given array.	The serial number, customer-assigned name, and software version of a given array in addition to a brief description of the error condition on the array that requires attention.	An alert is logged in the array OS and HPE Support Staff uses Secure Shell (SSH) to securely log in to the operating software of the array.



	<b>Heartbeat</b>	<b>DNA</b>	<b>Email Alerts</b>	<b>Tunnel</b>
Benefit to Customer	Provides near-real time monitoring of basic system health.	Provides proactive support, for example shipping to the customer a replacement for an array component that might be close to its operating life span, before such component fails.	Alerts HPE Support staff of error conditions on a timely basis to help more rapidly create support cases and speed the provision of corrective action by HPE support to the customer.	Permits real-time, hands-on support by qualified HPE Support staff in the event of a configuration issue or error that is too serious to be otherwise resolved.
Opt-in	No, it is always enabled.	Yes, disabled by default until it is enabled by the customer using the array OS.	Yes, disabled by default until it is enabled by the customer using the array OS. Email addresses for customer notification are controlled through the array OS. Through the opt-in configuration, you can enable alerts to be emailed only to you or to you and HPE Support.	Yes; the customer must log in to the array OS for a given array and manually establish the VPN tunnel to HPE Support.
Trigger or Frequency	Sent automatically from a given array to HPE Support every five to eight minutes.	Sent automatically once per day by an array, or as triggered by the customer at any time from the array OS.	Sent by a given array when it experiences an error condition.	Per the above "Opt-in", manually enabled by the customer and then initiated in turn by HPE Support.

## Information Security Practices

Each time that Support receives customer-originated data through a core data dump, the following protections are in place:

- 1** You must have given your prior express written authorization, signed by a customer officer with appropriate authority to convey such consent.
- 2** While in transmission, the data must be encrypted using Secure Socket Layer (SSL).
- 3** After the data is received through a core data dump, the data will be in a secure environment reflected by the following general measures:
  - Strict Access Control Lists (ACLs) with logging to prevent outside access to internal services
  - Network level virus protection with regularly updated signatures
  - Physical security based on a locked server room and co-location facility
  - Tracking and historic monitoring of all network conversations
  - Role Based Access Control (RBAC) to limit the number of personnel allowed access to the data
  - A logging process to track which HPE support and engineering individuals and groups accessed the data and when the data was accessed

Whenever possible there will be no, or very minimal, exposure of customer data to HPE personnel.

## Core Data Dumps

Whenever possible, HPE Support staff de-bugs or otherwise troubleshoots a system without the need to access any customer-originated data. However, on occasion, Support is unable to successfully support your system without receiving a set of customer-originated data, called a core data dump. Often such a data set can be quite large and contain many various types of content.

Before Support receives a core data dump from you, you must identify a company officer or owner who is authorized to consent in writing to the secure, encrypted transmission to HPE of such information.



After HPE has received the signed written authorization form, the support staff arrange to receive the core data dump. Then, HPE maintains the confidentiality and security of the data by following the practices described in [Information Security Practices](#) on page 55.

## Failed Disk Drives

At times, one or more of the returnable disk drives on an array could fail. In the event of such a failure, you must return the affected media to HPE. Then, if the drive or drives cannot be repaired, HPE will send the drive or drives back to the manufacturer. The following list indicates your responsibilities when you send a drive back to HPE:

- If you send one or more drives back to HPE as part of an assessment, it is your responsibility to encrypt each drive before sending it.
- If you send one or more drives back to HPE for purposes other than for assessment, for example if it has been determined that the drive was or will be replaced, then it is your responsibility to securely wipe each drive before sending it.

## Get Help

Troubleshooting and other help is available to search from the HPE InfoSight Documentation portal at <https://infosight.hpe.com/InfoSight/#documentation>.

All support contact information is available and maintained at <https://www.nimblestorage.com/customer-support>. You can contact HPE Support in the following ways:

- The Support link, which you can access under **Resources** in the InfoSight Main Menu.
- Email: [support@nimblestorage.com](mailto:support@nimblestorage.com).
- Phone: The numbers are listed for each country

Location	Number
United States	Toll Free: 1-877-3NIMBLE (877-364-6253), x2; Local: 408-432-9600, x2
Belgium	Toll Free/Mobile Enabled: 0800-262-76
France	Toll Free: 0800-911093
Germany	Toll Free/Mobile Enabled: 0800-183-0882
Luxembourg	Toll Free: 800-8-8078, National: +352-20202368
Netherlands	Toll Free/Mobile Enabled: 0-800-020-0730
Norway	\Toll Free: 800-17-365
Sweden	Toll Free: 020-791-743, Local: (10) +46-101992444
Switzerland	Toll Free/Mobile Enabled: 0800-802-544, Zurich: +41-435082007
United Kingdom	Toll Free: 0-808-134-9962, Local: 0330 808 0085

