

Foglight® Capacity Director 5.7.7
User Guide



© 2017 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. "Apache HTTP Server", Apache, "Apache Tomcat" and "Tomcat" are trademarks of the Apache Software Foundation. Google is a registered trademark of Google Inc. Android, Chrome, Google Play, and Nexus are trademarks of Google Inc. Red Hat, JBoss, the JBoss logo, and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the U.S. and other countries. CentOS is a trademark of Red Hat, Inc. in the U.S. and other countries. Fedora and the Infinity design logo are trademarks of Red Hat, Inc. Microsoft, .NET, Active Directory, Internet Explorer, Hyper-V, Office 365, SharePoint, Silverlight, SQL Server, Visual Basic, Windows, Windows Vista and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. AIX, IBM, PowerPC, PowerVM, and WebSphere are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Java, Oracle, Oracle Solaris, PeopleSoft, Siebel, Sun, WebLogic, and ZFS are trademarks or registered trademarks of Oracle and/or its affiliates in the United States and other countries. SPARC is a registered trademark of SPARC International, Inc. in the United States and other countries. Products bearing the SPARC trademarks are based on an architecture developed by Oracle Corporation. OpenLDAP is a registered trademark of the OpenLDAP Foundation. HP is a registered trademark that belongs to Hewlett-Packard Development Company, L.P. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. MySQL is a registered trademark of MySQL AB in the United States, the European Union and other countries. Novell and eDirectory are registered trademarks of Novell, Inc., in the United States and other countries. VMware, ESX, ESXi, vSphere, vCenter, vMotion, and vCloud Director are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Sybase is a registered trademark of Sybase, Inc. The X Window System and UNIX are registered trademarks of The Open Group. Mozilla and Firefox are registered trademarks of the Mozilla Foundation. "Eclipse", "Eclipse Foundation Member", "EclipseCon", "Eclipse Summit", "Built on Eclipse", "Eclipse Ready" "Eclipse Incubation", and "Eclipse Proposals" are trademarks of Eclipse Foundation, Inc. IOS is a registered trademark or trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. Apple, iPad, iPhone, Mac OS, Safari, Swift, and Xcode are trademarks of Apple Inc., registered in the U.S. and other countries. Ubuntu is a registered trademark of Canonical Ltd. Symantec and Veritas are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. OpenSUSE, SUSE, and YAST are registered trademarks of SUSE LCC in the United States and other countries. Citrix, AppFlow, NetScaler, XenApp, and XenDesktop are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. AlertSite and DéjàClick are either trademarks or registered trademarks of Boca Internet Technologies, Inc. Samsung, Galaxy S, and Galaxy Note are registered trademarks of Samsung Electronics America, Inc. and/or its related entities. MOTOROLA is a registered trademark of Motorola Trademark Holdings, LLC. The Trademark BlackBerry Bold is owned by Research In Motion Limited and is registered in the United States and may be pending or registered in other countries. Quest is not endorsed, sponsored, affiliated with or otherwise authorized by Research In Motion Limited. Ixia and the Ixia four-petal logo are registered trademarks or trademarks of Ixia. Opera, Opera Mini, and the O logo are trademarks of Opera Software ASA. Tevron, the Tevron logo, and CitraTest are registered trademarks of Tevron, LLC. PostgreSQL is a registered trademark of the PostgreSQL Global Development Group. MariaDB is a trademark or registered trademark of MariaDB Corporation Ab in the European Union and United States of America and/or other countries. Vormetric is a registered trademark of Vormetric, Inc. Intel, Itanium, Pentium, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Debian is a registered trademark of Software in the Public Interest, Inc. OpenStack is a trademark of the OpenStack Foundation. Amazon Web Services, the "Powered by Amazon Web Services" logo, and "Amazon RDS" are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. Infobright, Infobright Community Edition and Infobright Enterprise Edition are trademarks of Infobright Inc. POLYCOM®, RealPresence® Collaboration Server, and RMX® are registered trademarks of Polycom, Inc. All other trademarks and registered trademarks are property of their respective

owners.

Legend

- **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

- ⚠ **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

- ℹ **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Foglight Capacity Director User Guide
Updated - November 2018
Software Version - 5.7.7

Contents

Using Foglight Capacity Director	7
Installation requirements	7
Performance Agent Configuration for Foglight Capacity Director	8
Hyper-V Agent Configuration for Foglight Capacity Director	9
Dashboard location and UI elements	9
Group selector	10
Actions bar	10
Menu bar	13
Quick view	13
Getting Started Tab	14
Capacity Planning	14
Capacity Allocation	15
Capacity Optimization	16
Capacity Maintenance	16
Current Capacity Tab	18
Planned Changes pane	19
VirtualCenter/Hosts Overview pane	20
Calculating Slot Size Based on VM	20
Calculating Slot Size Using Custom Size	21
Clusters/Hosts pane	21
Resource Utilization Tab	22
Planned Changes pane	23
Clusters/Hosts and Workload pane	23
Calculating Host Size based on specific host	24
Resource Utilization related views	25
Scenario Modeler Tab	28
Optimal Placement	29
Balance Score	30
Modeling Placement	31
Deploying Placement	33
Manual Placement	34
Hog Score	36
Server Purchase Tab	38
Replace Existing Servers	39
Purchase Additional Servers	39
Managing Server List	40
Adding a new server	40
Editing existing servers	41
Deleting existing servers	41

Importing server list	41
Exporting server list	42
Infrastructure Organizations Tab	43
Organizations tab	43
Adding a new organization	44
Editing existing organizations	44
Removing existing organizations	45
Clusters tab	45
Performance Tuning Tab	46
Select a Cluster pane	47
Select a Host pane	47
Performance Tuning pane	47
Changes and DRS Rules Tab	49
Adding a Planned Change	50
Modeling a Cluster Addition	50
Modeling a Host Addition	51
Adding Virtual Machines	52
Modeling a Host Decommission	53
Decommissioning Virtual Machines	54
Modeling a Host Movement	54
Moving Virtual Machines	55
Adding a Datastore (Preview)	57
Modeling a Datastore Decommission (Preview)	58
Modeling a Datastore Expansion (Preview)	58
Modeling a Datastore Shrink (Preview)	58
Editing a Planned Change	59
Deleting a Planned Change	59
Cloud Model Tab	61
Cloud Instances pane	62
Adding a Cloud instance	63
Editing existing Cloud instance	63
Delete existing Cloud instance(s)	64
Importing Cloud instances	64
Importing Cloud Models from Portal	64
Exporting Cloud instances	65
Running Cloud Model	65
Model Outcome pane	65
Report Tab	67
Printing Reports from Current Capacity	67
Hosts Report	68
Clusters Report	68
Printing Reports from Resource Utilization	69
Clusters Report	70
Hosts and Workload Report	70

About Us	72
We are more than just a name	72
Our brand, our vision. Together.	72
Contacting Quest	72
Technical support resources	72

Using Foglight Capacity Director

Foglight® Capacity Director delivers complete virtual infrastructure capacity life cycle management. Get deep insight into your environment, including the availability and estimated time left for CPU, memory, storage, and Input/Output Operations Per Second (IOPs). Use historical trends to better forecast resource requirements for clusters, hosts, and workloads—and proactively predict and budget for capital expenditures. Optimize resource utilization and reduce operational expenses by rebalancing the distribution of clusters and workloads. Model your infrastructure changes (add, move, or remove clusters, hosts, and virtual machines) before deployment, to avoid bottlenecks, performance issues, or reaching the system's capacity limits.

The *Foglight Capacity Director User Guide* is intended for:

- *VMware*: users who have been assigned the *System Administrator*, *Advanced Operator*, *Capacity Management Administrator*, *VMware Administrator*, or *VMware QuickView User* role.
- *Hyper-V*: users who is a member of a local group administrators.

This section introduces you to the Foglight Capacity Director environment, and provides you with essential information.

For more information, see the following topics:

- [Installation requirements](#)
- [Performance Agent Configuration for Foglight Capacity Director](#)
- [Hyper-V Agent Configuration for Foglight Capacity Director](#)
- [Dashboard location and UI elements](#)

Installation requirements

Foglight Capacity Director comes installed on Foglight for Virtualization, Enterprise Edition and can be installed on a Foglight Management Server.

Foglight Capacity Director requires the following cartridges for data collection:

- 1 `DRP-5.7.7.car`
- 2 `Virtual-VMware-5.7.7.car`
 - **NOTE:** After installing the `Virtual-VMware-5.7.7.car` cartridge, existing vCenter agents (VC agents) must be updated using the Update Agent button in the *VMware Environment > Administration* tab.
- 3 `Virtual-VMware-Admin-5.7.7.car`
- 4 `Virtual-HyperV-5.7.7.car`
- 5 `SCVMM.car` (optional)
- 6 `OptimizerAutomation-5.7.7.car`
- 7 `CommonAnalytics-5.7.7.car`
- 8 `Capacity-Director-5.7.7.car`

While Foglight for Virtualization, Enterprise Edition comes with these cartridges pre-installed and enabled, a stand-alone Foglight release requires that these components be installed on the Foglight Management Server. The

sequence of cartridge installation is important because of their dependencies. For more information about installing Foglight Capacity Director, and for details about system requirements and version compatibility, see the *Foglight Capacity Director Release Notes*.

Performance Agent Configuration for Foglight Capacity Director

Each VMware Performance Agent monitors a single vCenter®. When you create a VMware Performance Agent instance, the Performance Agent starts importing historical data into Foglight Capacity Director. To fully enable Foglight Capacity Director functionalities in the VMware Environment and to enable the configuration of the **Automated** attribute when you create a planned change using the Add/Move/Decommission wizard in Foglight Capacity Director, a user with sufficient privileges must be configured.

On the vCenter® server, create a new role with the following privileges:

- Datastore > Allocate space
- Datastore > Browse Datastore
- Datastore > Low-level file operations
- Scheduled Task > Create task
- Scheduled Task > Delete task
- Scheduled Task > Run task
- Network > Assign network
- Resource > Assign Virtual machine to resource pool
- Resource > Migrate Powered off Virtual Machine
- Resource > Migrate Powered on Virtual Machine
- Virtual Machine > Configuration > Add existing disk
- Virtual Machine > Configuration > Add new disk
- Virtual Machine > Configuration > Add or remove device
- Virtual Machine > Configuration > Advanced
- Virtual Machine > Configuration > Change CPU count
- Virtual Machine > Configuration > Change resource
- Virtual Machine > Configuration > Extend virtual disk
- Virtual Machine > Configuration > Memory
- Virtual Machine > Configuration > Modify device settings
- Virtual Machine > Configuration > Remove disk
- Virtual Machine > Configuration > Settings
- Virtual Machine > Interaction > Configure CD Media
- Virtual Machine > Interaction > Device Connection
- Virtual Machine > Interaction > Power Off
- Virtual Machine > Interaction > Power On
- Virtual Machine > Inventory > Create from existing
- Virtual Machine > Inventory > Create new

- Virtual Machine > Provisioning > Deploy template

Refer to “*Configuring monitoring agents for data collection*” in the *Using Foglight for VMware* guide for instructions about configuring agent privileges.

Hyper-V Agent Configuration for Foglight Capacity Director

Each Hyper-V Agent monitors a single virtual machine. When you create a Hyper-V Agent instance, the agent starts importing historical data into Foglight Capacity Director. To fully enable Foglight Capacity Director functionalities in the Hyper-V environment and to enable the configuration of the **Automated** attribute when you create a planned change using the Add/Move/Decommission wizard in Foglight Capacity Director, the user must be a member of a local group administrators.

Dashboard location and UI elements

After installing Foglight Capacity Director, the **Capacity Director** dashboard appears in the Foglight Management Server.

To access the Capacity Director dashboard:


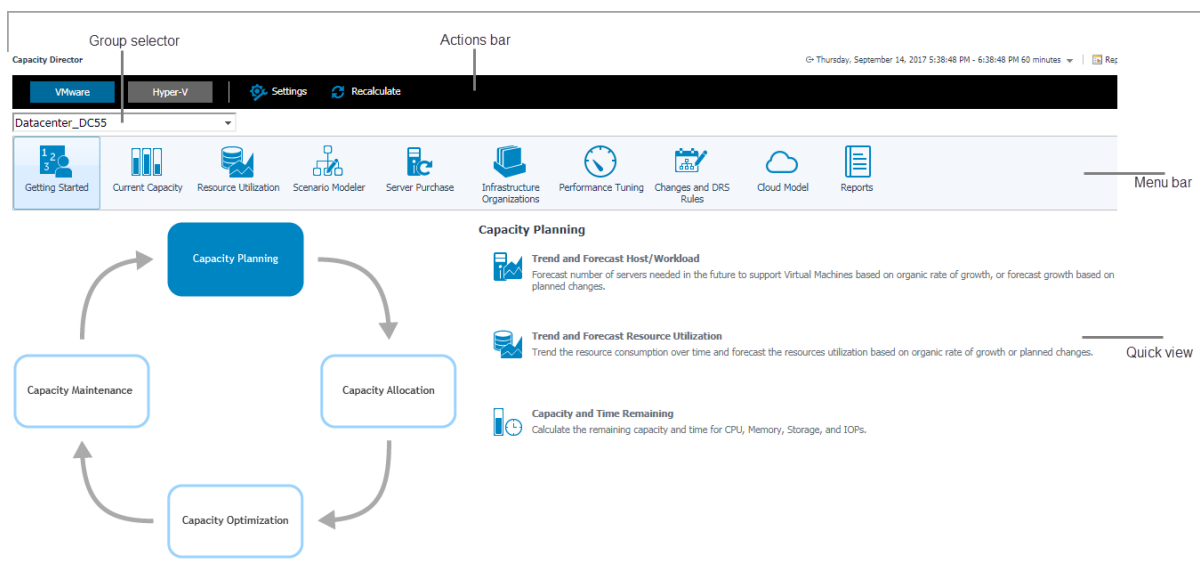
- 1 Log in to the Foglight browser interface.
- 2 Ensure that the navigation panel is open.
To open the navigation panel, click the right-facing arrow  on the left.
- 3 On the navigation panel, under *Dashboards*, click **Capacity Director**.
The **Capacity Director** dashboard opens.

Figure 1. Capacity Director dashboard



The **Capacity Director** dashboard consists of the following UI elements:

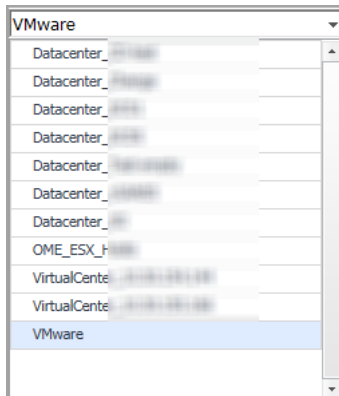
- **Group selector**

- [Actions bar](#)
- [Menu bar](#)
- [Settings Menu](#)

Group selector

The Group selector is located at the top of the dashboard and allows you to select the virtual environment that you want to monitor.

Figure 2. Group Selector



Actions bar

The actions bar at the top of the screens contains the following options, [Settings Menu](#) and [Recalculating Results](#).

Figure 3. Actions bar



Settings Menu

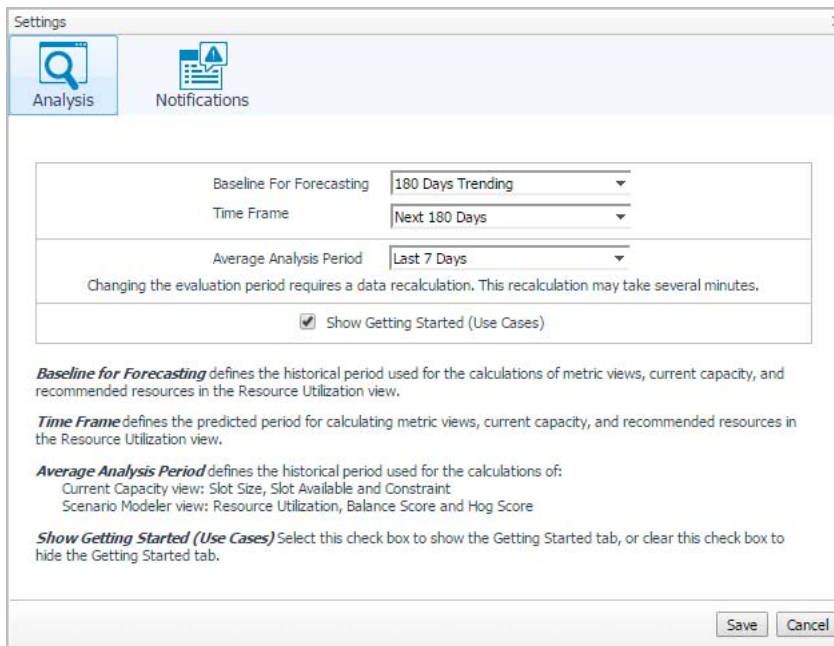
Click the Settings icon on the actions bar to open the Settings dialog box. For details about the settings that you can configure in this dialog, see the [“Configuring Analysis”](#) and [“Configuring Notifications”](#) sections.

Configuring Analysis

The Settings menu > Analysis tab allows you to configure the settings to be taken into calculation by Foglight Capacity Director for analyzing the current capacity and resource utilization, and for predicting future resource requirements.

To configure the analysis settings:

- 1 In the Settings dialog box, click the Analysis tab.



2 Specify the following values as needed.

- **Baseline For Forecasting:** Defines the historical period used for the calculations of metric views, current capacity, and recommended resources in the Resource Utilization view. The following options are available: 30 Days Trending (default option), 180 Days Trending, and 365 Days Trending.
- **Time Frame:** Defines the predicted period for calculating metric views, current capacity, and recommended resources in the Resource Utilization view. The following options are available: Next 30 Days (default option), Next 60 Days, Next 90 Days, and Next 180 Days.

For example, when Next 30 Days is selected, the Resource Utilization view displays metric views, recommended resources, and additional servers required in next 30 days.

- **Average Analysis Period:** Defines the historical period used for the calculations of:
 - Current Capacity view: Slot Size, Slot Available, and Constraint
 - Scenario Modeler view: Resource Utilization, Balance Score, and Hog Sore

For details about Balance Score and Hog Score, see the “[Threshold of Balance Score](#)” and “[Threshold of Hog Score](#)” sections.

The following options are available: Last 1 Day, Last 3 Days, Last 7 Days (default option), and Last 30 Days.

- **Show Getting Started (Use Cases):** Select this check box to show the *Getting Started* tab, or clear this check box to hide the *Getting Started* tab.

3 Click **Save**.

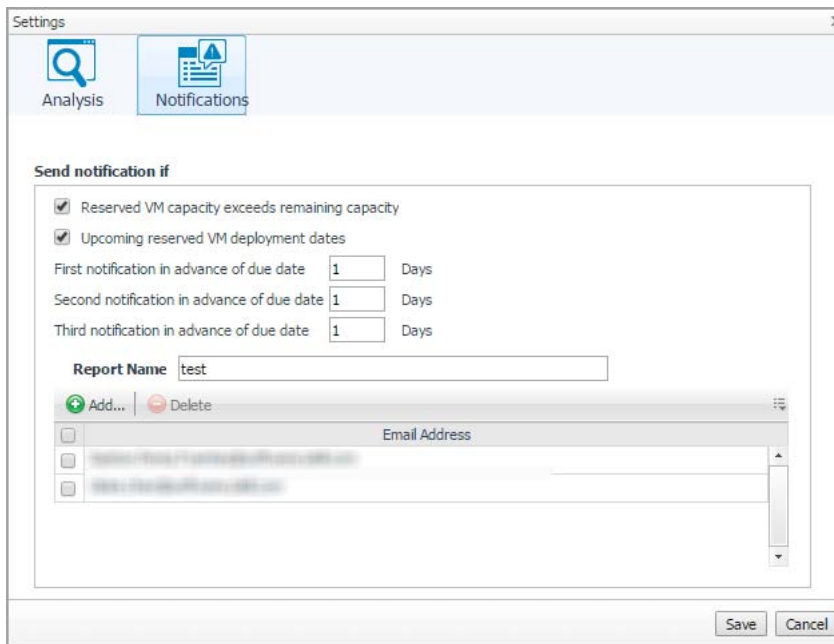
The settings are saved and applied next time when Foglight Capacity Director performs calculations.

Configuring Notifications

The Settings menu > Notifications tab allows you to configure the settings to be taken into calculation by Foglight Capacity Director when sending notifications.

To configure the notifications settings:

1 In the Settings dialog box, click the Notifications tab.



- 2 Select either or both of the following options, as needed:
 - Reserved VM capacity exceeds remaining capacity: Sends a notification email when the reserved VM capacity exceeds the remaining capacity.
 - Upcoming reserve VM deployment dates: Sends a notification email when the reserved VM deployment date gets closer.
- 3 Specify the first notification due date, if needed.
- 4 Specify the second notification due date, if needed.
- 5 Specify the third notification due date, if needed.
- 6 Name the report name in the **Report Name** field. For more information, see [Report Tab](#) on page 67.
- 7 Click **Add**, type the email address in the prompted dialog, and click **Add**.
Or
Select one or more email address(es) as needed, and click **Delete**.
- 8 Click **Save**.
The settings are saved and applied next time when Foglight Capacity Director performs calculations.

Recalculating Results

Foglight Capacity Director allows you to get the latest analytical results for the Current Capacity view, the Resource Utilization view, and the Scenario Modeler view, based on the data collected by the VMware Performance Agents. To get the latest analytic results, click the **Recalculate** menu on the actions bar, and click Recalculate in the Recalculate Confirmation dialog box.

i **NOTE:** Results populated in **Current Capacity > Clusters**, **Resource Utilization > Clusters**, **Resource Utilization > Hosts and Workload**, and **Scenario Modeler** are impacted if the High Availability (HA) policy is enabled in the vCenter®.

The Recalculation Results dialog box includes the following recalculating results:


- *Current Capacity:* Capacity of the vCenter® or datacenter, average/maximum slot size for clusters, available slots, and constraint.
- *Resource Utilization:* History data and forecasting data

- *Scenario Modeler*: Average utilization data of resources, including CPU, Memory, and Storage

Menu bar

The Menu bar contains the following tabs: *Getting Started Tab*, *Current Capacity Tab*, *Resource Utilization Tab*, *Scenario Modeler Tab*, *Server Purchase Tab*, *Infrastructure Organizations Tab*, and *Changes and DRS Rules Tab*.

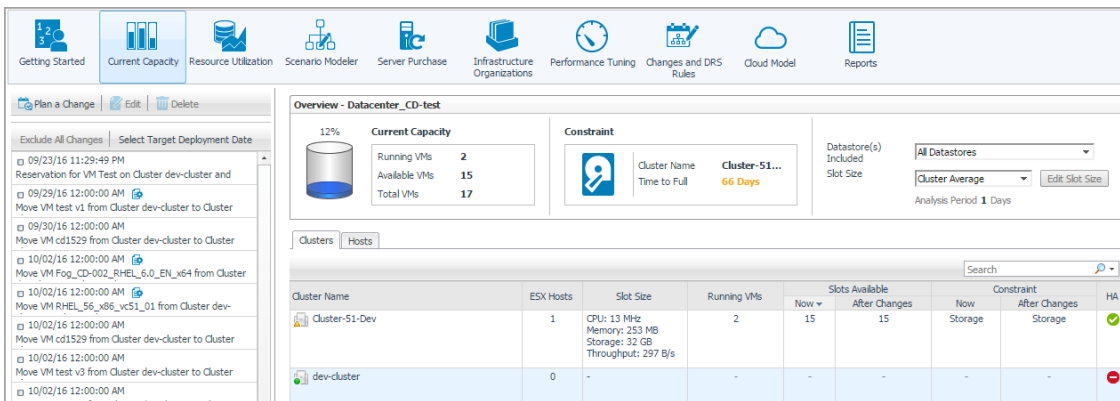
Figure 4. Menu bar



Quick view

The quick view is located on the lower part of the Capacity Director dashboard, which is updated based on the tab selected on the Menu bar.

Figure 5. Quick view in the Capacity Director Dashboard

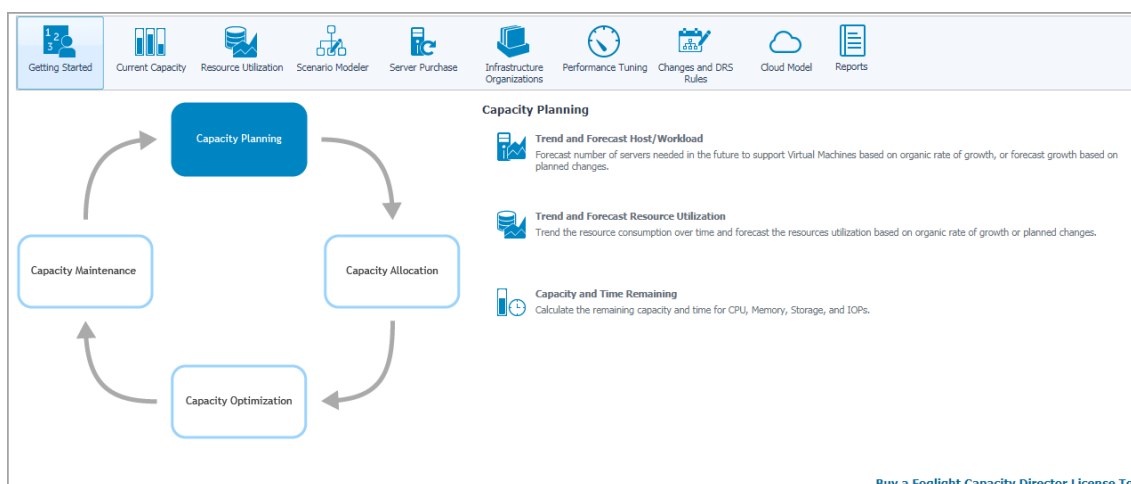


Getting Started Tab

This section provides an overview of the *Capacity Director dashboard > Getting Started* view, which walks you through the typical tasks performed in a datacenter for managing the system's capacity.

The *Getting Started* view consists of two panes. The left pane includes a diagram showing the typical workflow for managing the capacity in a datacenter. The right pane displays the *Capacity Planning* view, *Capacity Allocation* view, *Capacity Optimization* view, or *Capacity Maintenance* view, depending on the element selected on the left pane.

Figure 6. Getting Started tab



For more information, see the following topics:

- [Capacity Planning](#)
- [Capacity Allocation](#)
- [Capacity Optimization](#)
- [Capacity Maintenance](#)

Capacity Planning

Capacity planning is the first step in capacity life cycle management. It provides a long-term view of resource requirements, and allows you to build an effective capacity management strategy.

To access the *Capacity Planning* view:

- 1 On the navigation panel, under *Dashboards*, click **Capacity Director**.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the **Getting Started** tab.

The *Getting Started* view appears at the bottom of the *Capacity Director* dashboard.

- 4 On the *Getting Started* view > left pane, click **Capacity Planning**.

The *Capacity Planning* view appears on the right pane of the *Getting Started* view.

The *Capacity Planning* view provides a list of tasks that system administrators typically perform when planning the capacity of a virtual environment. Clicking a task directs you to the view where you can perform the selected task:

- *Trend and Forecast Host/Workload*: Allows you to forecast number of servers needed in the future to support Virtual Machines based on organic rate of growth, or to forecast growth based on planned changes. Click this link to navigate to the *Resource Utilization* view > *Hosts and Workload* tab. For more information, see [Clusters/Hosts and Workload pane](#) on page 23.
- *Trend and Forecast Resource Utilization*: Allows you to trend the resource consumption over time and forecast the resources utilization based on organic rate of growth or planned changes. Click this link to navigate to the *Resource Utilization* view > *Clusters* tab. For more information, see [Clusters/Hosts and Workload pane](#) on page 23.
- *Capacity and Time Remaining*: Allows you to calculate the remaining capacity and time for CPU, Memory, Storage, and IOPs. Click this link to navigate to the *Resource Utilization* view > *Clusters* tab. For more information, see [Clusters/Hosts and Workload pane](#) on page 23.

Capacity Allocation

When the resources are in place, the next step is to create the environment and to allocate the available capacity to support the workloads.

To access the *Capacity Allocation* view:

- 1 On the navigation panel, under *Dashboards*, click **Capacity Director**.

The *Capacity Director* dashboard opens.

- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the *Getting Started* tab.

The *Getting Started* view appears at the bottom of the *Capacity Director* dashboard.

- 4 On the *Getting Started* view > left pane, click **Capacity Allocation**.

The *Capacity Allocation* view appears on the right pane of the *Getting Started* view.

The *Capacity Allocation* view provides a list of tasks that system administrators typically perform when allocating resources in a virtual environment. Clicking a task directs you to the view where you can perform the selected task:

- *Plan Infrastructure Changes*: Walks you through the process of modeling future changes (add, move, or remove virtual machines, hosts, and clusters) or plan infrastructure changes (add, move, or decommission virtual machines). Click this link to navigate to the *Planned Changes* view. For more information, see [Changes and DRS Rules Tab](#) on page 49.
- *Manage Scheduled Infrastructure Changes*: Allows you to review the status of upcoming planned changes and modify the deployment date. Click this link to navigate to the *Planned Changes* view. For more information, see [Changes and DRS Rules Tab](#) on page 49.
- *Deploy Virtual Machines*: Walks you through the process of creating and deploying a virtual machine. Click this link to navigate to the *Add Virtual Machines Wizard*. For more information, see [Adding Virtual Machines](#) on page 52.
- *Decommission Virtual Machines*: Walks you through the process of decommissioning virtual machines that are no longer needed in your environment. The virtual machines are just turned off, but not deleted from the inventory. Click this link to navigate to the *Decommission Virtual Machines Wizard*. For more information, see [Decommissioning Virtual Machines](#) on page 54.

Capacity Optimization

Resource utilization can easily become unbalanced, as organizations and virtual environments grow over time. This leads to some clusters that are over-utilized, while others are under-utilized. The result is poor performance on the over-utilized clusters, and wasted resources and investment on clusters that are under-utilized. Optimal workload placement is critical to Return on Investment (ROI) in the virtual infrastructure. Foglight Capacity Director provides methods for both automatic and manual rebalancing of existing workloads, to ensure optimal resource utilization.

To access the Capacity Optimization view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Getting Started tab.
The Getting Started view appears at the bottom of the Capacity Director dashboard.
- 4 On the Getting Started view > left pane, click Capacity Optimization.
The Capacity Optimization view appears on the right pane of the Getting Started view.

The Capacity Optimization view provides a list of tasks that system administrators typically perform for optimizing virtual environments. Clicking a task directs you to the view where you can perform the selected task:

- *Optimal Placement*: Allows you to automatically rebalance the workload from clusters, by using an automatic algorithm that calculates and recommends the optimal distribution of virtual machines across clusters. Click this link to navigate to the Scenario Modeler > Optimal Placement view. For more information, see [Optimal Placement](#) on page 29.
- *Manual Placement*: Allows you to manually rebalance the workload from clusters, by moving virtual machines between clusters. Click this link to navigate to the Scenario Modeler > Manual Placement view. For more information, see [Manual Placement](#) on page 34.

Capacity Maintenance

Foglight Capacity Director allows you to monitor fully functional and optimized datacenters, and proactively identify any short-term resource requirements, to ensure the system's continuous availability.

To access the Capacity Maintenance view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Getting Started tab.
The Getting Started view appears at the bottom of the Capacity Director dashboard.
- 4 On the Getting Started view > left pane, click Capacity Maintenance.
The Capacity Maintenance view appears on the right pane of the Getting Started view.

The Capacity Maintenance view provides a list of tasks that system administrators typically perform for monitoring virtual environments. Clicking a task directs you to the view where you can perform the selected task:

- *Available Capacity*: Allows you to review the available capacity on the system and identify any potential bottlenecks. Click this link to navigate to the Current Capacity view. For more information, see [Current Capacity Tab](#) on page 18.

- *Resource Utilization*: Allows you to review the resources currently available in the system and identify potential bottlenecks. Click this link to navigate to the Resource Utilization view > Clusters tab. For more information, see [Clusters/Hosts and Workload pane](#) on page 23.
- *Slot Size*: Allows you to configure the size of the virtual machine (the slot size) that Foglight Capacity Director uses in calculations for future planned changes. Click this link to navigate to the Current Capacity view. For more information, see [Current Capacity Tab](#) on page 18.

Current Capacity Tab

The Current Capacity view displays information about the current capacity of a selected virtual center or datacenter. It includes the following three panes:

- **Planned Changes pane:** Lists all valid planned changes before the deployment date, and provides the capability to create/edit/remove one or more planned changes.
- **VirtualCenter/Hosts Overview pane:** Provides details about Current Capacity and Constraint, and allows you to select the Datastore information and Slot Size.
- **Clusters/Hosts pane:** Consists of two tabs:
 - **Clusters:** Provides detailed information about clusters that are available in the selected virtual center or datacenter.
 - **Hosts:** Provides detailed information about hosts that are available in the selected cluster.

To access the Current Capacity view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Current Capacity tab.

The Current Capacity view appears at the bottom of the Capacity Director dashboard.

NOTE: Enabling or disabling the HA policy, which is configured in vCenter®, is only applicable for the **Clusters** tab in the Current Capacity view. If the HA policy is enabled for a cluster, Foglight Capacity Director calculates the cluster resources and reserve the cluster resources that are required to satisfy the HA policy. A green check mark appears in the HA column when the HA policy is enabled.


The screenshot shows the Foglight Capacity Director interface. The top navigation bar includes 'VMware', 'Hyper-V', 'Settings', and 'Recalculate'. The main area is divided into three panes:

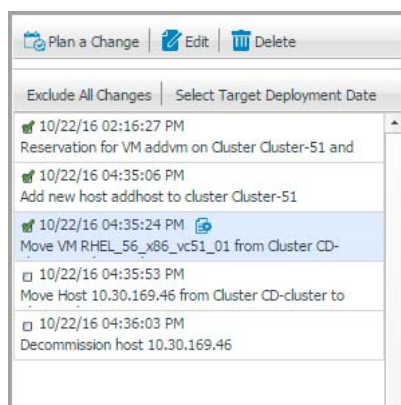
- Planned Changes pane:** Located on the left, it shows a list of planned changes with columns for 'Exclude All Changes' and 'Select Target Deployment Date'. The list includes actions like 'Add new cluster', 'Move VM testVM', and 'Move VM test4'.
- VirtualCenter/Hosts Overview pane:** Located in the middle, it displays 'Current Capacity' (83%) and 'Constraint' information. It includes a 'Datastore(s)' dropdown menu and a 'Slot Size' dropdown menu.
- Clusters/Hosts pane:** Located at the bottom, it displays a table of clusters with columns for Cluster Name, ESX Hosts, Slot Size, Running VMs, Slots Available, Constraint, and HA status.

Cluster Name	ESX Hosts	Slot Size	Running VMs	Slots Available		Constraint		HA
				Now	After Changes	Now	After Changes	
cluster51	2	CPU: 389 MHz Memory: 1.0 GB Storage: 4.4 GB Throughput: 8.7 KB/s	7	8	8	CPU	CPU	✓
Cluster-51-QA	2	CPU: 199 MHz Memory: 382 MB Storage: 22 GB Throughput: 5.5 KB/s	1	7	7	Storage	Storage	✗
Cluster-51-Dev	2	CPU: 27 MHz Memory: 269 MB Storage: 29 GB Throughput: 5.4 KB/s	21	-9	-9	Memory	Memory	✓
test	0	-	-	-	-	-	-	✗

Planned Changes pane

The Planned Changes pane on the left bottom of the Capacity Director dashboard offers the following menus:

- Plan a Change menu: Allows you to add a new planned change. For more information, see [Changes and DRS Rules Tab](#) on page 49.
- Edit menu: Allows you to make changes to existing planned changes. For more information, see [Changes and DRS Rules Tab](#) on page 49.
- Delete menu: Allows you to delete any existing planned changes if needed. For more information, see [Changes and DRS Rules Tab](#) on page 49.
- Exclude All Changes menu: Allows you to clear the selected planned changes in the Planned Changes table at a time.
- Select Target Deployment Date menu: Allows you to select the deployment date until which the information is displayed in the Planned Changes table.
- Planned Changes table: Includes a list of valid planned changes sorted chronologically up to the deployment date, which unifies the user interface and enables you to observe the result expected after executing planned changes. If you set the value of **Automated** to Yes when creating a planned change with the Add/Move/Decommission wizard, an icon  appears, indicating that the change is planned to be automatically deployed.



To identify the valid planned changes before a target deployment date:

- 1 Click **Select Target Deployment Date**.
- 2 In the **Select a Planned Deployment Date** dialog box, click the date icon and select the deployment date as needed.
- 3 Click **Select**.

The Select a Planned Deployment Date dialog box closes and the Planned Changes table is updated to indicate all valid changes planned before the selected deployment date (the check box for these changes is selected).

OR

In the Planned Changes table, select the check box in front of a planned change.

The Planned Changes table is updated to indicate all valid changes planned before the deployment date associated with the selected change (the check box for these changes is selected).

For details about how to edit or delete existing change(s), refer to [Changes and DRS Rules Tab](#) on page 49.

VirtualCenter/Hosts Overview pane

After selecting a virtual center or datacenter by using the Group selector located at the top of the dashboard, Foglight Capacity Director calculates the capacity of the selected center in the background and displays the following information on the VirtualCenter/Hosts Overview pane:

- **Current Capacity:** Includes information about the following metrics: Running VMs, Available VMs, and Total VMs.
- **Constraint:** Includes information about the following metrics: Cluster Name and Time to Full.
- **Datastore(s) Included:** Allows you to select from the drop-down list the datastores that Foglight Capacity Director should use in calculations for future planned changes.

The following options are available: All Datastores (default option) and particular datastores that varies depending on the virtual center/datacenter selected by the group selector.

- All Datastores: The Slot Size and available slots are calculated based on all available datastores.
- Particular datastore: The Slot Size and available slots are calculated based on the cluster/host connected to the selected datastore.

i | NOTE: The Clusters/Hosts pane is automatically updated to reflect the selection made.

- **Slot Size:** Allows you to select from the drop-down list the size of the virtual machine (the slot size) that Foglight Capacity Director should use in calculations for future planned changes.

The following options are available: Based on VM, Cluster Average (default option), Cluster Maximum, and Custom Size.


For details, see the “Calculating Slot Size Based on VM” and “Calculating Slot Size Using Custom Size” sections.

i | NOTE: The Clusters/Hosts pane is automatically updated to reflect the selection made.

Calculating Slot Size Based on VM

To recalculate the slot size Based on VM:

- 1 In the VirtualCenter/Hosts Overview pane, select Based on VM from the Slot Size drop-down list. The Select a VM button appears on the right side of the Slot Size drop-down list.




- 2 Click Select a VM. The VM Selection Dialog box appears.
- 3 Select a virtual machine from the table, or by typing its name in the Search box.
- 4 Click OK. The VM Selection Dialog box closes and the Clusters/Hosts pane is updated to reflect this change.

Calculating Slot Size Using Custom Size

To recalculate the slot size based on a custom size:

- 1 In the VirtualCenter/Hosts Overview pane, select Custom Size from the Slot Size drop-down list.
The Edit Slot Size button appears on the right side of the Slot Size drop-down list.



- 2 Click Edit Slot Size.
The Enter custom values dialog box appears.
- 3 Specify the values for CPU, Memory, Storage, and Throughput for the custom virtual machine.
- 4 Click Save.
The Enter custom values dialog box closes and the Clusters/Hosts pane is updated to reflect this change.

Clusters/Hosts pane

The **Clusters/Hosts** pane displays the details related to the selected virtual center or datacenter. This information is automatically updated when the Slot Size settings are modified and when including/applying planned changes. For more information, see [Calculating Slot Size Based on VM](#) on page 20 and [Calculating Slot Size Using Custom Size](#) on page 21.

The **Clusters/Hosts** pane includes the following capacity data:

- *Filtered by cluster (Hosts tab only):* Filter the hosts that contain the selected cluster.
- *Cluster Name (Clusters tab) or Name (Hosts tab):* Name of the selected cluster or host.
- *ESX Hosts (Clusters tab only):* Number of ESX[®] hosts in the selected cluster.
- *Slot Size:* CPU, Memory, Storage, and Throughput resources available on the virtual machine.
- *Running VMs:* Number of running virtual machines in the selected cluster or host.
- *Slots Available:*
 - *Now:* Number of available slots in the selected cluster or host
 - *After Changes:* Number of available slots after including/excluding planned changes. This value should be identical to the value of Now if there is no planned change selected on the Planned Changes pane. For more information, see [Changes and DRS Rules Tab](#) on page 49.
- *Constraint:*
 - *Now:* Available Constraint in the selected cluster or host
 - *After Changes:* Available slots after including/excluding planned changes.
- *HA (Clusters tab only):* If the HA policy is enabled for a cluster, Foglight Capacity Director calculates the cluster resources and reserves the cluster resources that are required to satisfy the HA policy. A green check mark appears in the HA column when the HA policy is enabled.

Resource Utilization Tab

The *Resource Utilization* view displays information about the resource utilization of a selected virtual center or datacenter. It includes the following panes:

- **Planned Changes pane:** Lists all valid planned changes before the deployment date, and provides the capability to create/edit/remove one or more planned changes.
- **Clusters/Hosts and Workload pane:** Consists of the following two tabs:
 - **Clusters:** Provides detailed information about clusters that are available in the selected virtual center or datacenter.
 - **Hosts and Workload:** Provides detailed information about hosts that are available in the selected cluster.

To access the Resource Utilization view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Resource Utilization tab.
The Resource Utilization view appears at the bottom of the Capacity Director dashboard.

NOTE: Enabling or disabling the HA policy, which is configured in vCenter®, is applicable for both the **Clusters** tab and the Hosts and Workload tab in the Resource Utilization view. If the HA policy is enabled for a cluster, Foglight Capacity Director calculates the cluster resources and reserve the cluster resources that are required to satisfy the HA policy. A green check mark appears in the HA column when the HA policy is enabled.

The screenshot displays the Foglight Capacity Director interface for the Resource Utilization view. On the left is the **Planned Changes pane**, listing various actions such as adding new clusters and moving VMs. The main area is the **Clusters/Hosts and Workload pane**, which shows the **Resource Utilization - cluster51** details. This includes a table for selecting a cluster, where 'cluster51' is highlighted with a green checkmark in the HA column. Below the table are four charts: **Projected CPU Usage** (showing utilization at 3.7 GHz and total capacity at 9.2 GHz), **Projected Memory Usage** (showing 12.0 GB consumed and 30.0 GB capacity), **Projected Storage Usage**, and **Projected Storage IOPs Usage**. The CPU chart also indicates a growth rate of 208.5 MHz per week and 99 days to full capacity. The Memory chart indicates a growth rate of 221.9 MB per week and 320 days to full capacity.

Planned Changes pane

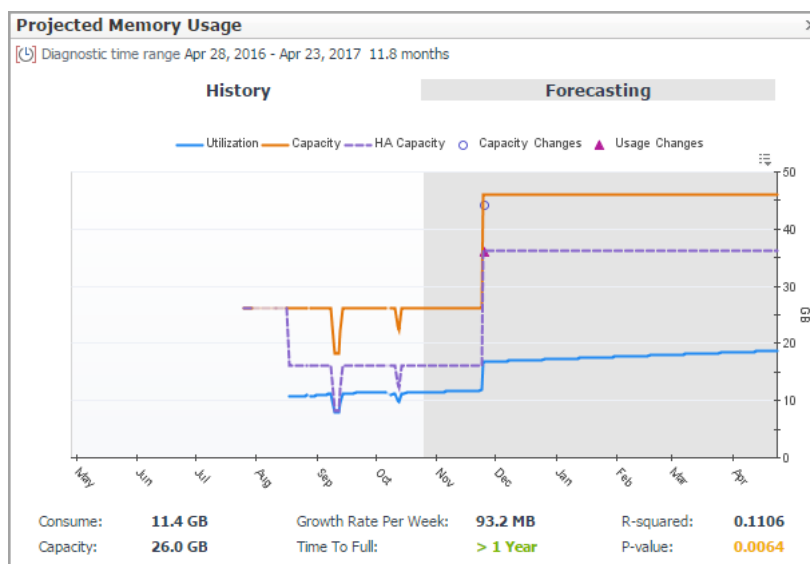
For more information, see [Planned Changes pane](#) on page 19.



Clusters/Hosts and Workload pane

The **Clusters** tab and the **Hosts and Workload** tab display the details related to the selected virtual center or datacenter. This information is automatically updated after selecting a cluster and when including/applying planned changes.

The **Clusters** tab includes the following data:

- *Select a Cluster*: Lists available clusters in the selected virtual center or datacenter.
- *Resource Utilization*
 - *Current Capacity*: Displays the current capacity of the selected cluster, which includes the following data: CPU, Memory, Storage, Hypervisors, and VMs and Templates. These values are not affected by including/excluding planned changes on the Planned Changes pane.
 - *Peak CPU Performance*: Displays the current CPU data of the selected cluster, which includes the following data: vCPU/Core Ratio, Physical CPU Cores, Powered On vCPUs, and Average % Ready. These values are not affected by including/excluding planned changes.
 - *Recommendations*: Displays recommendations on the resources (CPU, Memory, and Storage) to be added in the future to ensure that the resource utilization is balanced. Use the **Settings** dialog box to edit the Baseline for Forecasting. For more information, see on page 17. These values are automatically updated when including/applying planned changes.
- *Resource utilization related views*: [Projected CPU Usage view](#), [Projected Memory Usage view](#), [Projected Storage Usage view](#), and [Projected Storage IOPs Usage view](#). The metric line in the white area is the history data, while the metric line in the gray area is the predicted data. These views are refreshed automatically after including planned changes. The following is an example of the Projected Memory Usage view. For more information, see [Resource Utilization related views](#) on page 25.



Click  or  on the Projected Memory Usage view to display the **Planned Changes Summary** dialog box that summarizes the planned changes.

Planned Changes Summary	
Deploy Date	Description
9/30/16 12:00 AM	Move VM cd1529 from Cluster dev-cluster to Cluster Cluster-51-Dev

The **Hosts and Workload** tab includes the following data:

- *Select a Cluster:* Lists available clusters in the selected virtual center or datacenter.
- *Overview:* Summarizes the CPU capacity and memory capacity of the selected cluster. Both capacities are calculated based on the value selected from the Host Size drop-down list. The Host Size drop-down list includes the following options: average host size (default value) and based on specific host. The Host Requirement - Average Host Size view is refreshed automatically after changing the Host Size value. For more information, see [Calculating Host Size based on specific host](#) on page 24.

The value of Baseline for Forecasting on the display area can be edited using the **Settings** dialog box. For more information, see [Settings Menu](#) on page 10.

- *Recommendation on Additional Servers:* Displays how many new servers should be added in the future (namely the value of Time Frame). This value is calculated based on the selected planned changes and the projected resource usage. Use the **Settings** dialog box to edit Time Frame. For more information, see on page 17.

i | **NOTE:** The new servers need to be added in the **Server Purchase** tab. Click **>>Link to Server Purchase** to add a new server. For more information, see [Server Purchase Tab](#) on page 38.

- *Host Requirement - Average Host Size:* Displays the number of Required Hosts and Available Hosts, if the existing server is used as the calculation benchmark.
 - *Required Hosts:* Shows with the blue line, indicating how many hosts are required to support the current number of virtual machines. This number is populated based on the value selected from the Host Size drop-down list.
 - *Available Hosts:* Shows with the amber line, indicating the number of hosts, available in the selected cluster.
- *Host Requirement - <New Servers>:* Displays the number of Required Hosts and Available Hosts, if the new server is used as the calculation benchmark.
 - *Required Hosts:* Indicates how many hosts are required to support the usage of virtual machines. This number is populated based on the *Server Details*.
 - *Available Hosts:* Indicates the number of hosts, available in the selected cluster.

Calculating Host Size based on specific host

To recalculate the host size based on specific host:

- 1 In the Hosts and Workload Overview pane, select based on specific host from the Host Size drop-down list. The Browse button becomes available on the right side of the Host Size drop-down list.

- 2 Click Browse to select a host in the Specific Host Selector dialog box, and click OK.

The Specific Host Selector dialog box closes and the Host Requirement - Average Host Size view is refreshed automatically.

Resource Utilization related views

This section provides a detailed description of the following metrics presented in the embedded views in the Resource Utilization > Clusters tab.

- [Projected CPU Usage view](#)
- [Projected Memory Usage view](#)
- [Projected Storage Usage view](#)
- [Projected Storage IOPs Usage view](#)

Projected CPU Usage view

This embedded view is part of the Resource Utilization > Clusters view. To display this view, click the Clusters tab in the Resource Utilization view, and select a cluster from the Cluster Name table.

Table 1. Projected CPU Usage view

Description	<p>This view displays CPU utilization metrics on the selected cluster.</p> <ul style="list-style-type: none">• Utilization: This chart shows the actual CPU usage (in GHz).• Capacity: This chart shows the total CPU capacity (in GHz).• HA Capacity: If the HA policy is enabled on the selected cluster, this chart is plotted to show the actual CPU usage (in GHz).• Capacity Changes: Shows the changes on CPU capacity (in GHz) after including planned changes.• Usage Changes: Shows the changes on CPU usage (in GHz) after including planned changes.• Used Hz: Shows the current CPU usage (in GHz).• Growth Rate Per Week: Shows the growth rate of CPU capacity per week.• R-squared: Indicates the accuracy of the forecasting data. This value is ranged from 0 to 1, and a higher value indicates that the forecasting data of utilization is more accurate.
Data displayed	<ul style="list-style-type: none">• Total Hz: Shows the current CPU capacity (in GHz).• Time To Full:<ul style="list-style-type: none">HA enabled: Shows the period when the actual CPU usage (the Utilization metric) exceeds the HA capacity metric. This metric is calculated based on the Growth Rate Per Week.HA disabled: Shows the period when the actual CPU usage (the Utilization metric) exceeds the total CPU capacity (the Capacity metric). This metric is calculated based on the Growth Rate Per Week.• P-value: Indicates the dependency between the utilization and history time. This value is ranged from 0 to 1, and a lower value indicates that the utilization has a more significant dependency with the history time.

Projected Memory Usage view

This embedded view is part of the Resource Utilization > Clusters view. To display this view, click the Clusters tab in the Resource Utilization view, and select a cluster from the Cluster Name table.

Table 2. Projected Memory Usage view

Description	<p>This view displays memory utilization metrics on the selected cluster.</p> <p>NOTE: Unless otherwise indicated, the units of measurement for the following metrics are automatically converted for better plotting.</p> <ul style="list-style-type: none">• Utilization: This chart shows the actual memory usage.• Capacity: This chart shows the total memory capacity.• HA Capacity: If the HA policy is enabled on the selected cluster, this chart is plotted to show the actual memory usage.• Capacity Changes: Shows the changes on memory capacity after including planned changes.• Usage Changes: Shows the changes on memory usage after including planned changes.• Consume: Shows the current memory usage.• Growth Rate Per Week: Shows the growth rate of memory capacity per week.• R-squared: Indicates the accuracy of the forecasting data. This value is ranged from 0 to 1, and a higher value indicates that the forecasting data of utilization is more accurate.
Data displayed	<ul style="list-style-type: none">• Capacity: Shows the current memory capacity.• Time To Full:<ul style="list-style-type: none">HA enabled: Shows the period when the actual memory usage (the Utilization metric) exceeds the HA capacity metric. This metric is calculated based on the Growth Rate Per Week.HA disabled: Shows the period when the actual memory usage (the Utilization metric) exceeds the total memory capacity (the Capacity metric). This metric is calculated based on the Growth Rate Per Week.• P-value: Indicates the dependency between the utilization and history time. This value is ranged from 0 to 1, and a lower value indicates that the utilization has a more significant dependency with the history time.

Projected Storage Usage view

This embedded view is part of the Resource Utilization > Clusters view. To display this view, click the Clusters tab in the Resource Utilization view, and select a cluster from the Cluster Name table.

Table 3. Projected Storage Usage view

Description	<p>This view displays storage utilization metrics on the selected cluster.</p> <p>NOTE: Unless otherwise indicated, the units of measurement for the following metrics are automatically converted for better plotting.</p> <ul style="list-style-type: none">• Utilization: This chart shows the actual storage usage.• Capacity: This chart shows the total storage capacity.• Capacity Changes: Shows the changes on storage capacity after including planned changes.• Usage Changes: Shows the changes on storage usage after including planned changes.• Consume: Shows the current Storage usage.• Growth Rate Per Week: Shows the growth rate of storage capacity per week.
Data displayed	<ul style="list-style-type: none">• R-squared: Indicates the accuracy of the forecasting data. This value is ranged from 0 to 1, and a higher value indicates that the forecasting data of utilization is more accurate.• Capacity: Shows the current Storage capacity.• Time To Full: Shows the period when the actual storage usage (the Utilization metric) exceeds the total storage capacity (the Capacity metric). This metric is calculated based on the Growth Rate Per Week.• P-value: Indicates the dependency between the utilization and history time. This value is ranged from 0 to 1, and a lower value indicates that the utilization has a more significant dependency with the history time.

Projected Storage IOPs Usage view

This embedded view is part of the Resource Utilization > Clusters view. To display this view, click the Clusters tab in the Resource Utilization view, and select a cluster from the Cluster Name table.

Table 4. Projected Storage IOPS Usage view

Description	<p>This view displays storage Input/Output Operations Per Second (IOPS) utilization metrics on the selected cluster.</p> <ul style="list-style-type: none">• iops: This chart shows the actual storage IOPS in the unit of connections per second (c/s).• Usage Changes: Shows the changes on storage IOPS usage (in c/s) after including planned changes.• Consumed: Shows the current storage IOPS (in c/s).• Growth Rate Per Week: Shows the growth rate of storage IOPS usage per week.
Data displayed	<ul style="list-style-type: none">• R-squared: Indicates the accuracy of the forecasting data. This value is ranged from 0 to 1, and a higher value indicates that the forecasting data of utilization is more accurate.• Capacity: This metric is not collected for the Projected Storage IOPS Usage view.• Time To Full: This metric is not collected for the Projected Storage IOPS Usage view.• P-value: Indicates the dependency between the utilization and history time. This value is ranged from 0 to 1, and a lower value indicates that the utilization has a more significant dependency with the history time.

Scenario Modeler Tab

The *Scenario Modeler* view is intended to rebalance the workload among virtual centers or datacenters, to achieve optimal resource utilization. The balance metrics include CPU utilization, Memory utilization, and Storage Capacity utilization. A metric is defined as 'balanced' if the resource utilization is averagely distributed across all clusters and standalone nodes; otherwise this metric can be regarded as 'non-balanced'. Foglight Capacity Director provides two methods to help you optimize the use of clusters resources, by using the *Scenario Modeler* view:

- **Optimal Placement:** Automatically rebalances the workload from clusters, by using an automatic algorithm that calculates and recommends the optimal distribution of virtual machines across clusters.
- **Manual Placement:** Manually rebalances the workload from clusters, by moving virtual machines between clusters.

NOTE: Results populated in the **Scenario Modeler** view are impacted if the HA policy is enabled in a virtual center.

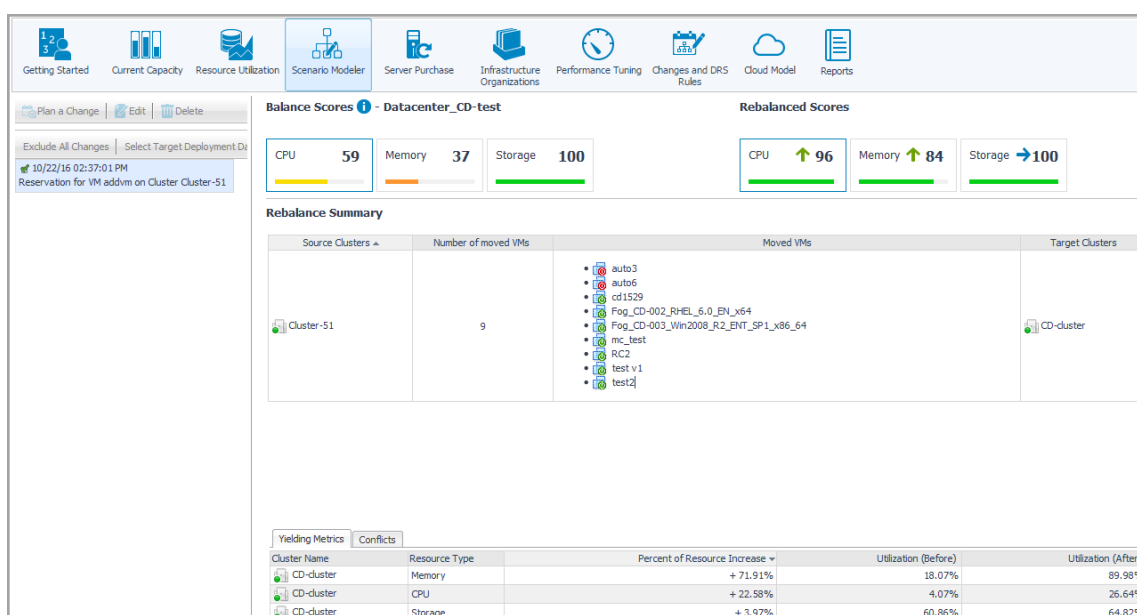
To access the Scenario Modeler view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Scenario Modeler tab.

The Scenario Modeler view appears at the bottom of the Capacity Director dashboard.

The left side of the Scenario Modeler view includes the **Planned Changes pane** and the right side is the Scenario Modeler pane that walks you through the placement procedure.

Figure 7. Scenario Modeler tab



Optimal Placement

To perform an optimal placement:

- 1 In the Scenario Modeler pane, select a resource that you want to balance under Step 1: Select a resource for balance.

Balance Scores - VirtualCenter

Step 1: Select a resource for balance

CPU 74 Memory 98 Storage 74

Select a scope for balance

Step 2: Select a placement approach

Optimal Placement Manual Placement

Step 3: Select an organization >> [Link to Infrastructure Organizations](#)

Organizations	Balance Scores			Clusters	ESX Hosts	VMs
	CPU	Memory	Storage			
<input checked="" type="radio"/> Unassigned	74	98	74	3	6	66

Step 4: Select at least two clusters to be balanced

Cluster Name	Utilization			Datacenters	ESX Hosts	VMs	HA
	CPU	Memory	Storage				
<input type="checkbox"/> vSAN55	30.07%	71.05%	6.76%	vSAN55	3	14	✓
<input type="checkbox"/> cluster55	22.60%	72.81%	55.58%	DC55	2	40	✗
<input type="checkbox"/> XD	2.92%	57.57%	17.05%	XD	1	12	✗

- 2 Select **Optimal Placement** under Step 2: Select a placement approach.

- 3 Select an organization from the available organizations.

Or

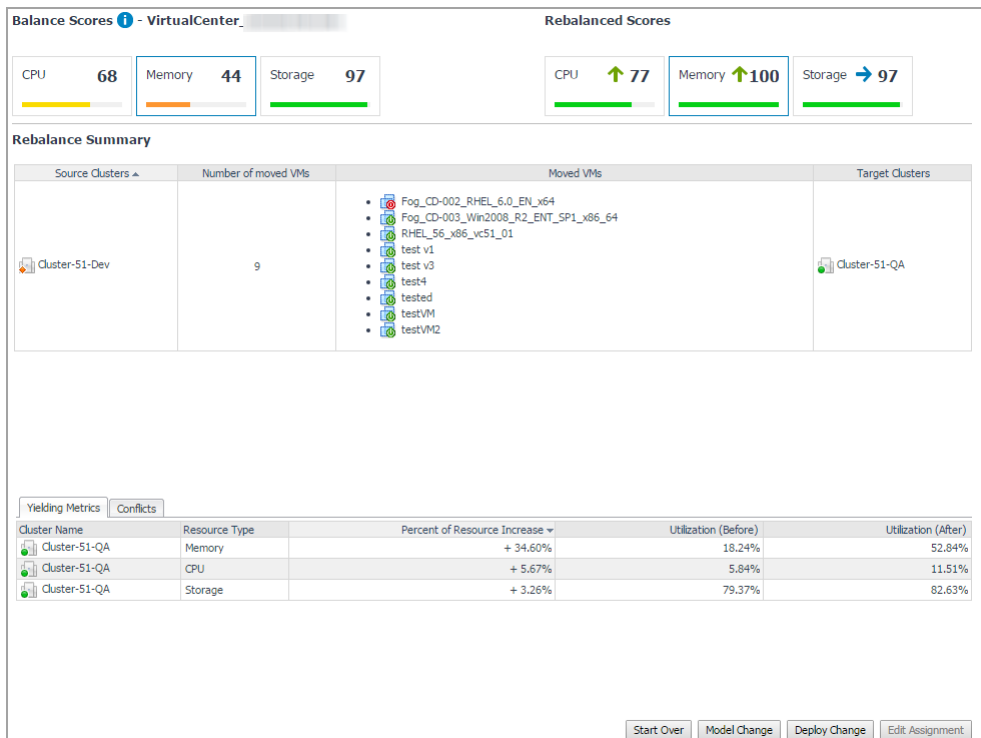
Click >> **Link to Infrastructure Organizations** to add a new organization as needed. For more information, see [Infrastructure Organizations Tab](#) on page 43.

The available clusters list is refreshed automatically.

- 4 Select the cluster(s) that you want to rebalance at the bottom of the page, and click **Next**.

NOTE: A minimum of two clusters must be selected so that the rebalancing and the VM replacement take effects. The Next button is disabled when only one cluster is available, which means the current organization cannot be rebalanced.

- 5 Foglight for VMware calculates in the background and displays the following optimal placement recommendations:



- **Balance Scores:** The original Balance Score before the optimal placement. For more information, see [Balance Score](#) on page 30.
- **Rebalance Scores:** The updated Balance Score if the optimal placement is deployed.
- **Source Clusters:** The clusters from which virtual machines are moved.
- **Number of moved VMs:** The number of the virtual machines to be moved to the Target Clusters.
- **Moved VMs:** The list of virtual machines to be moved to the Target Clusters.
- **Target Clusters:** The clusters to which virtual machines are moved.
- **Yielding Metrics:** The resource utilization before/after the optimal placement, and the increase percent of resources.
- **Conflicts:** The list of conflicts to occur if the optimal placement is deployed. The types of conflicts include affinity rules violation, organization conflicts, and unavailable network accesses.

6 Click **Start Over** if you want to redo the resource rebalance.

Click **Model Change** to model the optimal placement. For more information, see [Modeling Placement](#) on page 31.

Click **Deploy Change** to execute the optimal placement as Foglight Capacity Director recommended. For more information, see [Deploying Placement](#) on page 33.

NOTE: The **Edit Assignment** button is disabled for the *Optimal Placement* method.

Balance Score

Balance Score is scaled from 0 (non-balanced) to 100 (balanced). The same resource utilizations across all containers achieve a perfect Balance Score, for example, **Balance Score** is 100 if the CPU/Memory/Storage Capacity utilizations are 30% for all clusters. The **Balance Score** is approximately 0 if the resource utilization is 100% for cluster 1 and cluster 2, while resource utilization for cluster 3 is 0%.

Balance Score is divided into the following four ranges:

Table 5. Threshold of Balance Score

Balance Score state	Default range	Default threshold
Fatal (in red)	0 ~ 25	Default value is 25. To change this value, search for and edit <i>CapacityHealthScoreFatal</i> from Dashboards > Administration > Rules & Notifications > Manage Registry Variables .
Critical (in orange)	26 ~ 50	Default value is 50. To change this value, search for and edit <i>CapacityHealthScoreCritical</i> from Dashboards > Administration > Rules & Notifications > Manage Registry Variables .
Warning (in yellow)	51 ~ 75	Default value is 75. To change this value, search for and edit <i>CapacityHealthScoreWarning</i> from Dashboards > Administration > Rules & Notifications > Manage Registry Variables .
Normal (in green)	76 ~ 100	N/A

Modeling Placement

After Foglight Capacity Director calculates the result of optimal placement or manual placement, you can only model the virtual machines movement without deploying the movement in the realtime environment.

To model an optimal placement:

- 1 After [Step 5 in To perform an optimal placement:](#) or [Step 7 in To perform a manual placement:](#), the placement recommendations appear in the Scenario Modeler pane.

The screenshot displays the 'VirtualCenter' interface with the following components:


- Balance Scores:** CPU 68, Memory 44, Storage 97.
- Rebalanced Scores:** CPU 77, Memory 100, Storage 97.
- Rebalance Summary Table:**

Source Clusters	Number of moved VMs	Moved VMs	Target Clusters
Cluster-51-Dev	9	<ul style="list-style-type: none"> Fog_CD-002_RHEL_6.0_EN_x64 Fog_CD-003_Win2008_R2_ENT_SP1_x86_64 RHEL_56_x86_vc51_01 test v1 test v3 test4 tested testVM testVM2 	Cluster-51-QA
- Yielding Metrics Table:**

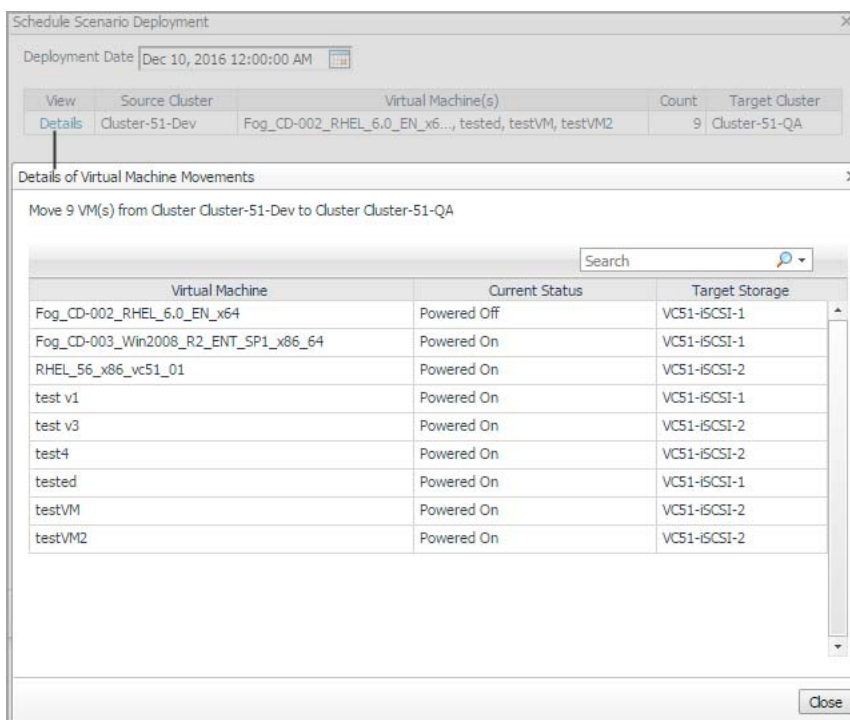
Cluster Name	Resource Type	Percent of Resource Increase	Utilization (Before)	Utilization (After)
Cluster-51-QA	Memory	+ 34.60%	18.24%	52.84%
Cluster-51-QA	CPU	+ 5.67%	5.84%	11.51%
Cluster-51-QA	Storage	+ 3.26%	79.37%	82.63%

- 2 Click **Model Change**.

The Schedule Scenario Deployment dialog box appears.



- (Optional) In the Schedule Scenario Deployment dialog box, click **Details** to view details of Virtual Machine Movements. Click Close and return to the Schedule Scenario Deployment dialog box.



- In the **Schedule Scenario Deployment** dialog box, select the **Deployment Date**, and click **Schedule Scenario Deployment**.

The Planned Changes pane is automatically updated to include the new modeled deployment.

Deploying Placement

After Foglight Capacity Director calculates the result of optimal placement or manual placement, you can deploy the virtual machines movement in the realtime environment, as needed.

To deploy an optimal placement:

- 1 After [Step 5 in To perform an optimal placement:](#) or [Step 7 in To perform a manual placement:](#), the placement recommendations appear in the Scenario Modeler pane.

The screenshot shows the 'VirtualCenter_' interface with 'Rebalanced Scores' for CPU (68), Memory (44), and Storage (97). The 'Rebalance Summary' table is as follows:

Source Clusters	Number of moved VMs	Moved VMs	Target Clusters
Cluster-51-Dev	9	<ul style="list-style-type: none"> Fog_CD-002_RHEL_6.0_EN_x64 Fog_CD-003_Win2008_R2_ENT_SP1_x86_64 RHEL_56_x86_vc51_01 test v1 test v3 test4 tested testVM testVM2 	Cluster-51-QA

Cluster Name	Resource Type	Percent of Resource Increase	Utilization (Before)	Utilization (After)
Cluster-51-QA	Memory	+ 34.60%	18.24%	52.84%
Cluster-51-QA	CPU	+ 5.67%	5.84%	11.51%
Cluster-51-QA	Storage	+ 3.26%	79.37%	82.63%

- 2 Click **Deploy Change**.

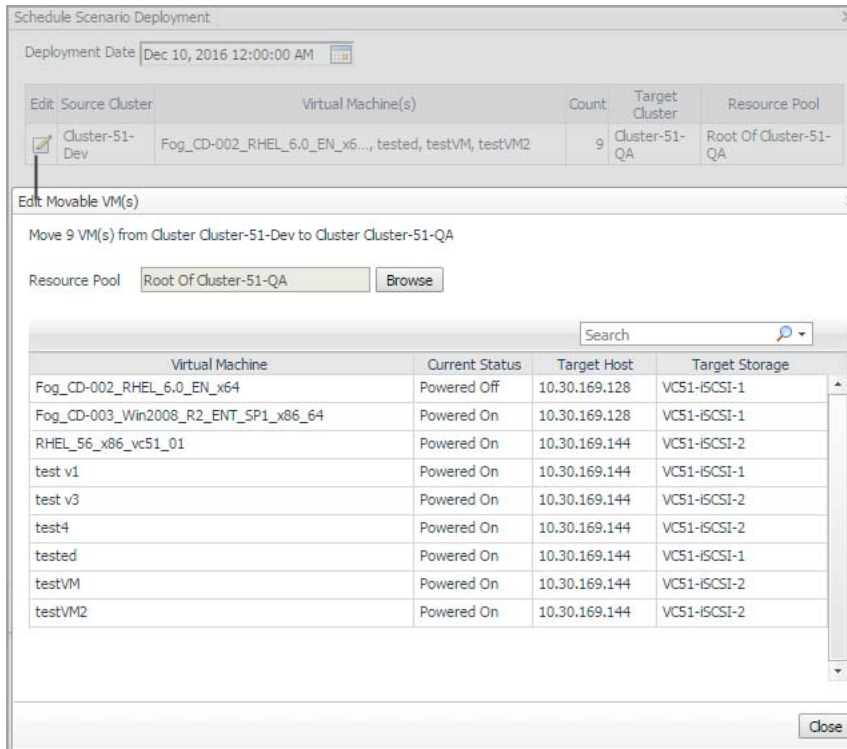
The Schedule Scenario Deployment dialog box appears.

The 'Schedule Scenario Deployment' dialog box shows a deployment date of 'Dec 10, 2016 12:00:00 AM'. The table below details the deployment plan:

Edit	Source Cluster	Virtual Machine(s)	Count	Target Cluster	Resource Pool
<input checked="" type="checkbox"/>	Cluster-51-Dev	Fog_CD-002_RHEL_6.0_EN_x6..., tested, testVM, testVM2	9	Cluster-51-QA	Root Of Cluster-51-QA

Buttons at the bottom: Schedule & Deploy Scenario Deployment, Cancel

- (Optional) In the Schedule Scenario Deployment dialog box, click the **Edit** icon to open the **Edit Movable VM(s)** dialog box, and then click **Browse** next to Resource Pool to make changes as needed. Click Close to save changes and return to the Schedule Scenario Deployment dialog box.



- In the Schedule Scenario Deployment dialog box, select the Deployment Date, and click Schedule & Deploy Scenario Deployment.

The placement is executed in the realtime environment, and Planned Changes pane is automatically updated to include the new modeled deployment.

Manual Placement

To perform a manual placement:

- In the Scenario Modeler pane, select a resource that you want to balance under Step 1: Select a resource for balance.

Balance Scores - VirtualCenter_

Step 1: Select a resource for balance

CPU **68** Memory **44** Storage **97**

Select a scope for balance

Step 2: Select a placement approach

Optimal Placement Manual Placement

Step 3: Select an organization >> [Link to Infrastructure Organizations](#)

Organizations	Balance Scores			Clusters	ESX Hosts	VMs
	CPU	Memory	Storage			
<input checked="" type="radio"/> Unassigned	68	44	97	3	6	61
<input type="radio"/> Test_1	100	100	100	1	0	0

Step 4: Select a Cluster for balancing

Cluster Name	Utilization			Datacenters	ESX Hosts	VMs	HA
	CPU	Memory	Storage				
<input checked="" type="radio"/> Cluster-51-Dev	27.88%	133.63%	83.70%	CD-test	2	24	✓
<input type="radio"/> cluster51	55.33%	53.25%	69.75%	DCS1	2	36	✓
<input type="radio"/> Cluster-51-QA	5.84%	18.24%	79.37%	CD-test	2	1	✗

- 2 Select **Manual Placement** under Step 2: Select a placement approach.
- 3 Select an organization from the available organizations.

Or

Click >> **Link to Infrastructure Organizations** to add a new organization as needed. For more information, see [Infrastructure Organizations Tab](#) on page 43.

The available clusters list is refreshed automatically.

- 4 Select a cluster that you want to rebalance as needed, and click **Next**.

The Scenario Modeler pane is updated to display the following information:

Balance Scores - Datacenter_CD-test

CPU **65** Memory **72** Storage **100**

Rebalanced Scores

CPU **↑ 99** Memory **↑ 82** Storage **→ 100**

Manual Virtual Machine Assignment

Select Movable VM(s) from **Cluster-51**

Name	Hog Score
test2	83
auto3	94
mc_test	95
RC2	95
auto2	96
auto5	96
cd1529	96
Fog_CD-002_RHEL_6.0_EN_x64	96
auto6	97
mc_test2	100
test1	100

Select Target Cluster

CD-cluster

Plan to Move 2 VM(s)

Movements	Conflicts
CD-cluster(2)	
test v1	
Fog_CD-003_Win2008_R2_ENT_SP1_x86_64	

Utilization (Before/After) of **Cluster-51**

CPU	Memory	Storage
Before: 39.05% / After: -6.45%	Before: 70.38% / After: -28.71%	Before: 57.34% / After: 0.00%
Before: 32.61% / After: 32.61%	Before: 41.67% / After: 41.67%	Before: 57.34% / After: 57.34%

Utilization (Before/After) of **CD-cluster**

CPU	Memory	Storage
Before: 4.07% / After: +19.82%	Before: 18.07% / After: +63.25%	Before: 60.86% / After: 0.00%
Before: 23.88% / After: 23.88%	Before: 81.32% / After: 81.32%	Before: 60.86% / After: 60.86%

- *Balance Scores*: The original Balance Score before the manual placement. For more information, see [Balance Score](#) on page 30.
 - *Rebalance Scores*: The updated Balance Score if the manual placement is deployed.
 - List of virtual machines available in the source cluster. For more information, see [Hog Score](#) on page 36.
 - Resource utilization before/after the movement in the source cluster.
 - List of virtual machines available in the target cluster.
 - Resource utilization before/after the movement in the target cluster.
- 5 From the list of virtual machines available in the source cluster, select the virtual machines that you want to move to the target cluster, and click **Move In**.

Or

From the list of virtual machines available in the target cluster, select the virtual machines that you want to move back to the source cluster, and click **Move Out**.

- 6 Click **Proceed to Summary**.
- 7 The Scenario Modeler pane is updated to display the following information:
- The Rebalance Summary view
 - *Balance Scores*: The original Balance Score before the manual placement.
 - *Rebalance Scores*: The updated Balance Score if the manual placement is deployed.
 - *Source Clusters*: The clusters from which virtual machines are moved.
 - *Number of moved VMs*: The number of the virtual machines to be moved to the Target Clusters.
 - *Moved VMs*: The list of virtual machines to be moved to the Target Clusters.
 - *Target Clusters*: The clusters to which virtual machines are moved.
 - *Yielding Metrics*: The resource utilization before/after the manual placement, and the increase percent of resources.
 - *Conflicts*: The list of conflicts that will be occurred if the manual placement is deployed. The types of conflicts include affinity rules violation, organization conflicts, and unavailable network accesses.

- 8 Click **Start Over** if you want to redo the resource rebalance.

Click **Model Change** to model the manual placement. For more information, see [Modeling Placement](#) on page 31.

Click **Deploy Change** to execute the manual placement as Foglight Capacity Director recommends. For more information, see [Deploying Placement](#) on page 33.

Click **Edit Assignment** to return the **Manual Virtual Machine Assignment view**, and redo the manual placement from [Step 4](#).

Hog Score

Hog Score can be regarded as the health score. It is scaled from 0 to 100 and helps achieve more predictable results when performing the rebalance. If Hog Score of a virtual machine is too low and is displayed in red, Foglight for VMware recommends moving this virtual machine to other clusters.

Resource utilization of a virtual machine and free capacity of a cluster are two factors that impact the **Hog Score**. For example, the **Hog Score** is low (unhealthy) if CPU utilization of a virtual machine is much higher than other resource utilizations in the same cluster. The **Hog Score** is high (healthy) if the CPU utilization of a virtual machine is much lower than the free capacity of a cluster on which this virtual machine is deployed.

Hog Score is divided into the following four ranges:

Table 6. Threshold of Hog Score

Hog Score state	Default range	Default threshold
Fatal (in red)	0 ~ 25	Default value is 25. To change this value, search for and edit <i>CapacityHealthScoreFatal</i> from Dashboards > Administration > Rules & Notifications > Manage Registry Variables .
Critical (in orange)	26 ~ 50	Default value is 50. To change this value, search for and edit <i>CapacityHealthScoreCritical</i> from Dashboards > Administration > Rules & Notifications > Manage Registry Variables .
Warning (in yellow)	51 ~ 75	Default value is 75. To change this value, search for and edit <i>CapacityHealthScoreWarning</i> from Dashboards > Administration > Rules & Notifications > Manage Registry Variables .
Normal (in green)	76 ~ 100	N/A

Server Purchase Tab

The *Server Purchase* view is intended to review the cost of replacing the existing infrastructure with new servers, and identify how many additional servers should be purchased in the specified Time Frame. The *Server Purchase* view contains two panes:

- **Planned Changes** (left pane): Lists all valid planned changes before the deployment date, and provides the capability to create/edit/remove one or more planned changes. For more information, see [Scenario Modeler Tab](#) on page 28.
 - **Server Purchase** (right pane): Includes the following elements:
 - **Select a Cluster**: Lists available clusters in the selected virtual center or datacenter.
 - **Weighting Options**: Includes the following weighting options to be used when calculating the cost of replacing the infrastructure: Cost Weight, Power Weight, and Space Weight. Click **Recalculate** button to perform the calculation.
 - **Replace Existing Servers** tab: Displays details about how many new servers are required to replace the existing servers, without including the planned changes.
 - **Purchase Additional Servers** tab: Displays details about how many new servers should be purchased, based on the calculation of resource utilization and planned changes in the specified Time Frame. For details about how to edit the Time Frame, see [Settings Menu](#) on page 10.
- NOTE:** Servers must be added or imported, to make them visible in the **Server List** table. For more information, see [Managing Server List](#) on page 40.

To access the Server Purchase view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Server Purchase tab.

The Server Purchase view appears at the bottom of the Capacity Director dashboard.

The screenshot displays the 'Server Purchase' tab in the Capacity Director dashboard. The top navigation bar includes icons for Getting Started, Current Capacity, Resource Utilization, Scenario Modeler, **Server Purchase**, Infrastructure Organizations, Performance Tuning, Changes and DRS Rules, Cloud Model, and Reports. Below the navigation bar, there are controls for 'Plan a Change', 'Edit', and 'Delete'. The main area is divided into two panes. The left pane shows a message: '11/30/16 03:46:06 PM Add new cluster Cluster_test to datacenter DCS1 with 1'. The right pane is titled 'Server Purchase' and contains a 'Select a Cluster' dropdown set to 'Cluster-51-Dev', 'Weighting Options' (Cost Weight: 80%, Power Weight: 10%, Space Weight: 10%), and a 'Recalculate' button. Below these are tabs for 'Replace Existing Servers' and 'Purchase Additional Servers'. The 'Top Recommendation' section features an image of a server and a table with the following data:

test test11111	
CPU	12.0 GHz
Memory	16.0 GB
Power	200.0 W
Space	2 U
Cost per Server	\$ 1.0
Number of Servers Required	2
Total Cost	\$ 2.0

The 'Server List' section at the bottom contains a table with the following data:

Manufacturer	Model Name	Ranking	Requirement	Total Cost	Unit Cost	CPU	Memory	Power
test	test11111	1	2	\$ 2.0	\$ 1.0	12.0 GHz	16.0 GB	200.0 W

Replace Existing Servers

After selecting a cluster and specifying the weighting options, Foglight Capacity Director populates the recommendation results for replacing the existing servers.

To get recommendations for replacing existing servers:

- 1 In the *Server Purchase* pane, click the Replace Existing Servers tab.
- 2 Select a cluster from the **Select a Cluster** drop-down list.
- 3 Type the percent of **Weighting Options** to be used when calculating the cost of replacing the infrastructure:
 - Cost Weight: The percent of server cost.
 - Power Weight: The percent of server power.
 - Space Weight: The percent of server space.
- 4 Click **Calculate**. This button is changed to **Recalculate** after executing the calculation the first time.

The **Top Recommendation** area is updated to show the summary of most recommended server, and the **Server List** table is updated to show all server types recommended as replacements.

i | **NOTE:** The sum of Cost Weight, Power Weight, and Space Weight must be equal to 100%. For more information, see [Managing Server List](#) on page 40.

Purchase Additional Servers

After specifying the *Time Frame* in the **Settings** dialog, Foglight for VMware populates the recommendation about how many new servers should be purchased in order to meet the requirement of planned changes in the specified Time Frame. For details about how to edit the Time Frame, see on page 17.

To get recommendations for purchasing additional servers:


- 1 In the *Server Purchase* pane, click the Purchase Additional Servers tab.
- 2 Select a cluster from the **Select a Cluster** drop-down list.
- 3 Type the percent of **Weighting Options** to be used when calculating the cost of replacing the infrastructure:
 - Cost Weight: The percent of server cost.
 - Power Weight: The percent of server power.
 - Space Weight: The percent of server space.
- 4 Click **Calculate**. This button is changed to **Recalculate** after executing the calculation first time.
- 5 The **Top Recommendation** area is updated to show the summary of most recommended server and how many servers should be purchased in the specified Time Frame.

The Server List table is updated to show all server types recommended to be purchased.

i | **NOTE:** You can review the recommendations provided and customize the results, as needed. For more information, see [Managing Server List](#) on page 40.

Top Recommendation

Based on the calculation of resource utilization and planned changes, you need to purchase additional 0 server(s) in next 180 days and will cost \$0.00.



CPU	12.0 GHz
Memory	16.0 GB
Power	200.0 W
Space	2 U
Cost per Server	\$ 1.0
Number of Servers Required	0
Total Cost	\$ 0.0

Server List

Add Edit Delete Import Export

Manufacturer	Model Name	Ranking	Requirement	Total Cost	Unit Cost	CPU	Memory	Power	Space
		1	0	\$ 0.0	\$1.0	12.0 GHz	16.0 GB	200.0 W	2 U

Managing Server List

Foglight Capacity Director allows you to manage the Server List table as follows:

- Adding a new server
- Editing existing servers
- Deleting existing servers
- Importing server list
- Exporting server list

Adding a new server

Before adding a new server, in the *Server Purchase* pane, click the *Replace Existing Servers* tab or the *Purchase Additional Servers* tab, and locate the *Server List* table at the bottom of the *Server Purchase* view.

To add a new server:

- 1 In the *Server List* table, click **Add**.
- 2 In the **Add Host Model** dialog box, specify the following values, as needed.
 - *Manufacturer*: The manufacturer of the new server.
 - *Model Name*: The model name of the new server.
 - *Unit Cost (\$)*: The unit cost (in dollar) of the new server.
 - *Number of CPU Cores*: The number of CPU cores of the new server.
 - *Physical CPU Speed*: The physical CPU speed of the new server. The unit is one of the following: MHz, GHz (default option), and THz.
 - *Amount of Physical Memory*: The physical memory of the new server. The unit is one of the following: MB, GB (default option), and TB.
 - *Power Requirement*: The power of the new server. The unit is one of the following: W (default option), kW, and MW.
 - *Space Requirement*: The space (in U) of the new server.
- 3 Click **Save**.

The *Server List* table is updated automatically to show the new server.

Editing existing servers

Before editing an existing server, in the *Server Purchase* pane, click the Replace Existing Servers tab or the Purchase Additional Servers tab, and locate the Server List table at the bottom of the Server Purchase view.

To edit an existing server:

- 1 Select an existing server from the **Server List** table, and click **Edit**.
- 2 In the **Edit Host Model** dialog box, change any of following values, as needed.
 - *Manufacturer*: The manufacturer of this server.
 - *Model Name*: The model name of this server.
 - *Unit Cost (\$)*: The unit cost (in dollar) of this server.
 - *Number of CPU Cores*: The number of CPU cores of this server.
 - *Physical CPU Speed*: The physical CPU speed of this server. The unit is one of the following: MHz, GHz (default option), and THz.
 - *Amount of Physical Memory*: The physical memory of this server. The unit is one of the following: MB, GB (default option), and TB.
 - *Power Requirement*: The power of this server. The unit is one of the following: W (default option), kW, and MW.
 - *Space Requirement*: The space (in U) of this server.
- 3 Click **Save**.

The Server List table is updated automatically to show the changes made to *this* server.

Deleting existing servers

Before deleting existing servers, in the *Server Purchase* pane, click the Replace Existing Servers tab or the Purchase Additional Servers tab, and locate the Server List table at the bottom of the Server Purchase view.

To delete one or more existing server(s):

- 1 In the Server List table, select the sever(s) that you want to remove, and click **Delete**.
- 2 In the **Confirm Delete Host Models** dialog box, click **Delete**.

The server is removed from the Server List table.

Importing server list

Before importing a server list, in the *Server Purchase* pane, click the Replace Existing Servers tab or the Purchase Additional Servers tab, and locate the Server List table at the bottom of the Server Purchase view.

To import an existing server list:

- 1 In the Server List table, click **Import**.
- 2 In the **Import Host Models** dialog box, select Replace to replace the existing servers with the imported server list.
OR
Select Add to add the imported servers to the existing server list.
- 3 Click **Browse** to browse your local drive and select a .csv file, and click Open.

- 4 Click **Import**.

The Server List table is updated automatically to include the imported servers.

Exporting server list

Before exporting a server list, in the *Server Purchase* pane, click the Replace Existing Servers tab or the Purchase Additional Servers tab, and locate the Server List table at the bottom of the Server Purchase view.

To export a server list:

- 1 In the Server List table, click **Export**.
- 2 In the File Download dialog box that appears, choose one of the following options:
 - Open the servers list for viewing.
 - a Click Open.
 - b The server list is opened automatically using Microsoft Excel for viewing.
 - Download the server list as a .csv file to a location of your choice.
 - a Click Save.
 - b In the dialog box that appears, select a location for saving the file, then click Save.

The server list is saved as a .csv file to the selected location.

Infrastructure Organizations Tab

The *Infrastructure Organizations* view shows detailed information about organizations and clusters, and provides an approach to manage them. The *Infrastructure Organizations* view contains two tabs:

- **Organizations tab:** Lists and manages the infrastructure organization.
- **Clusters tab:** Lists the cluster information and assigns the clusters, as needed.

To access the *Infrastructure Organizations* view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Infrastructure Organization tab.

The *Infrastructure Organizations* view appears at the bottom of the Capacity Director dashboard.

Name of Host or Cluster	Organization	Total Cost	Geographical Location
Cluster-51-Dev	Unassigned	n/a	Unassigned
Cluster-51-Test	Unassigned	n/a	Unassigned
cluster51	Unassigned	n/a	Unassigned
test	Unassigned	n/a	Unassigned
Cluster_test (Planned)	Unassigned	n/a	Unassigned

Organizations tab

The *Organizations* tab includes:

- A table showing details about the Organizations in your infrastructure. For more information, see [Editing existing organizations](#) on page 44.
 - Organization Name: The name of your organization
 - Location: The geographical location of your organization
 - Number of Clusters: The number of clusters existing in your organization
 - CPU Unit Cost: The CPU unit cost
 - Memory Unit Cost: The Memory unit cost
 - Storage Unit Cost: The storage unit cost
 - Host License Fee: The license fee of the host
- Two buttons allow you to manage the organizations in your infrastructure.
 - Add: Allows you to add a new organization. For more information, see [Adding a new organization](#) on page 44.

- **Remove:** Allows you to remove existing organizations. For more information, see [Removing existing organizations](#) on page 45.

Adding a new organization

To add a new organization:

- 1 In the *Infrastructure Organizations view*, click **Organizations**, and click **Add**.
- 2 In the **Add Organization** dialog box, specify the following fields, as needed.
 - **Organization Name:** The name of your organization
 - **Geographical Location:** The geographical location of your organization
 - **Unit Cost (\$):**
 - **CPU:** The unit cost (in dollar) of CPU
 - **Memory:** The unit cost (in dollar) of memory
 - **Storage:** The unit cost (in dollar) of storage
 - **Host License Fee (\$):** The license fee (in dollar) of the host
- 3 Click **Save**. The new organization is added to the Organizations table.

Or

Click **Save** and add another to save the current organization, then add another organization.

The Add Organization dialog box remains open. Continue with [Step 2](#).

Editing existing organizations

To edit an existing organization:

- 1 In the *Infrastructure Organizations view*, click **Organizations**, and select an organization to be updated.
- 2 Click **Edit** next to **Organization Name**.
- 3 In the **Edit Organization** dialog box, change the following fields, as needed.
 - **Organization Name:** The name of your organization
 - **Geographical Location:** The geographical location of your organization
 - **Unit Cost (\$):**
 - **CPU:** The unit cost (in dollar) of CPU
 - **Memory:** The unit cost (in dollar) of memory
 - **Storage:** The unit cost (in dollar) of storage
 - **Host License Fee (\$):** The license fee (in dollar) of the host
- 4 Click **Save**. The organization changes are updated in the Organizations table.

Or

Click **Save** and add another. The **Edit Organization** dialog box remains open but all fields become blank. Continue with [Step 2](#) in [Adding a new organization](#) on page 44 to add a new organization.

Removing existing organizations

To remove one or more existing organization(s):

- 1 In the *Infrastructure Organizations* view, click **Organizations**, and select the organizations to be removed.
- 2 Click **Delete**.
- 3 In the **Confirm Delete Organization** Dialog box, click **Delete**.
The organization is removed from the Organizations table.

Clusters tab

The *Clusters* tab allows you to assign the available cluster to an organization. Each cluster can be assigned only to one organization. The *Clusters* tab includes:

- A table showing details about the clusters in your infrastructure
 - Name of Host or Cluster: The name of your host or cluster
 - Organization: The name of your organization
 - Total Cost: The cost of your host or cluster
 - Geographical Location: The geographical location of your host or cluster
- An Assign button allowing you to assign a cluster to an organization

To assign a cluster to one organization:

- 1 In the *Infrastructure Organizations* view, click **Clusters**, and select a cluster that you want to assign.
- 2 Click **Assign**.
- 3 In the **Select Organization** dialog box, select an organization from the list, and click **Save**.
The Cluster/Host table is updated to show the updated assignment.

i **NOTE:** After assigning a cluster to an organization, the virtual machine movement between different organizations in the Scenario Modeler pane is not allowed.

NOTE: The cost of cluster is re-calculated when assigning this cluster to an organization.

Performance Tuning Tab

The *Performance Tuning* view helps to identify and resolve performance issues in your virtual environment by analyzing contention as well as the resources assigned to a virtual machine and then automating the recommended changes. It includes the following elements:

- **Select a Cluster pane:** Lists available clusters in the selected virtual center or datacenter, and provides the capability to search the cluster as needed.
- **Select a Host pane:** Lists available hosts in the selected cluster, and provides the capability to search the host as needed.
- **Information icon:** Click this icon to pop up the dialog box with details explaining what is Performance Tuning.
- **Performance Tuning pane:** Consists of two tabs:
 - **VMs Performance:** Design to show the virtual machines on the given host or all hosts in a cluster to quickly highlight the worst performing. This data is then used to drive the recommendations in the Resource Optimization tab.
 - **Resource Optimization:** Provides three automated recommendation engines (Aggressive, Moderate, and Top 10) built to move virtual machines between hosts to reduce the contention and improve the performance.

To access the Performance Tuning view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
- The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Performance Tuning tab.

The Performance Tuning view appears at the bottom of the Capacity Director dashboard.

Virtual Machine	Performance State	Performance Score	CPU Utilization	CPU Latency	CPU Max Limited	Memory Utilization	Memory Swapped	Memory Balloon	Virtual Disk Latency
test1	N/A	n/a	0.89 %	0.50 %	0.00 %	4.5 %	275 MB	1.2 GB	0ms
tested	N/A	n/a	0.31 %	0.27 %	0.00 %	0.00 %	0.00 B	0.00 B	0ms
Fog_CD-002_RHEL_6.0_ENL_x64	N/A	n/a	99 %	1.6 %	0.00 %	20 %	0.00 B	0.00 B	0ms
mic_test	N/A	n/a	0.30 %	0.30 %	0.00 %	0.00 %	0.00 B	0.00 B	0ms
CD-Test-3ay-1	N/A	n/a	11 %	5.3 %	0.00 %	5.6 %	32 MB	245 MB	11ms
kxlong	N/A	n/a	4.2 %	1.5 %	0.00 %	6.7 %	11 MB	0.00 B	12ms

Select a Cluster pane

The Select a Cluster pane on the left top of the Capacity Director dashboard offers the following menus:

- Search menu: Provides the search functionality that helps you easily find the cluster, as needed.
- Select a Cluster table: Lists all clusters that are available in the selected virtual center or datacenter.

Description	Lists the name and the performance score of all clusters that are available in the selected virtual center or datacenter.
Data displayed	<ul style="list-style-type: none">• Cluster Name. The name of the cluster.• Performance Score. The performance score of the cluster.

Select a Host pane

The Select a Host pane on the left bottom of the Capacity Director dashboard offers the following menus:

- Search menu: Provides the search functionality that helps you easily find the host, as needed.
- Select a Host table: Lists all hosts that are available in the selected cluster.

Description	Lists the name and the performance score of all hosts that are available in the selected cluster.
Data displayed	<ul style="list-style-type: none">• Host Name. The name of the host.• Performance Score. The performance score of the host.

Performance Tuning pane

The VMs Performance tab and Resource Optimization tab provide a view of the selected cluster and hosts.

The VMs Performance tab includes the following elements:

- Filter menu: Allows you to filter the virtual machines based on the following options: All VMs (default option), Good, Degrading, and Bad.
- Virtual Machine: Displays the name of the virtual machine.
- CPU: Shows the CPU Utilization, CPU Latency, and CPU Max Limited in percent for the virtual machine.
- Memory: Shows the memory Utilization, memory Swapped, and memory balloon for the virtual machine.
- Virtual Disk Latency: Shows the data transfer latency (in ms) of the virtual machine.
- Performance State: Indicates the performance state of the virtual machine, which can be one of the following: Good, Degrading, and Bad.
- Performance Score: Shows the performance score of the virtual machine.

The Resource Optimization tab includes the following elements:

- Select a mode for recommendation: Provides the optimization modes that help you obtain the recommended resource optimization, as needed.
- Resources Optimization table: Lists the resource optimization result of virtual machines, which includes the name of Virtual Machine, Performance Score, Resource Type, Recommendation, and Modify Recommendation.

This table is refreshed automatically after selecting a mode from the **Select a mode for recommendation** drop-down list.

- Migration Recommendation table: Lists the resource migration result of virtual machines, which includes the name of Virtual Machine, Performance Score, Source Host, Destination, and Status.

This table is refreshed automatically after selecting a mode from the **Select a mode for recommendation** drop-down list.

Changes and DRS Rules Tab

The *Changes and DRS Rules* view displays the list of planned changes scheduled before the deployment date, as well as the list of planned changes that have been executed. It includes the following elements:

- Planned Changes tab: Includes the Planned Changes table and Add/Delete menus.
 - Planned Changes table: Includes the following elements:
 - The Edit button: Allows you to edit planned changes. For more information, see [Editing a Planned Change](#) on page 59.
 - Deploy Date: The date when the planned change is deployed.
 - Description: The description of the planned change.
 - Validation:
 - Automation: Indicates whether the planned change will be deployed automatically or not.
 - **NOTE:** Foglight Capacity Director does not support the following automated deployments in Hyper-V environments: *Add Virtual Machines, Decommission Virtual Machines, and Move Virtual Machines.*
 - Performance: Indicates whether the planned change impacts the performance of the selected virtual center or not.
 - Conflict: Indicates whether the planned change has conflicts or not.
 - Add/Delete menus: Allow you to manage planned changes. For more information, see [Adding a Planned Change](#) on page 50 and [Deleting a Planned Change](#) on page 59.
- DRS Rules tab: Allow you to select a cluster to view the DRS status of the selected cluster.

To access the Changes and DRS Rules view:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Changes and DRS Rules tab.

The Changes and DRS Rules view appears at the bottom of the Capacity Director dashboard.

Edit	Deploy Date	Description	Automation	Validation Performance	Conflict
	5/5/17 11:52 PM	Add new host newESX to cluster Cluster-51-QA			
	4/19/17 11:49 PM	Reservation for VM BigVM on Cluster Cluster-51-QA and Datastore VCS1-HSCSI-1			

For more information, see the following topics:

- [Adding a Planned Change](#)
- [Editing a Planned Change](#)
- [Deleting a Planned Change](#)

Adding a Planned Change



To add a planned change, do one of the following:

- [Modeling a Cluster Addition](#)
- [Modeling a Host Addition](#)
- [Adding Virtual Machines](#)
- [Modeling a Host Decommission](#)
- [Decommissioning Virtual Machines](#)
- [Modeling a Host Movement](#)
- [Moving Virtual Machines](#)
- [Adding a Datastore \(Preview\)](#)
- [Modeling a Datastore Decommission \(Preview\)](#)
- [Modeling a Datastore Expansion \(Preview\)](#)
- [Modeling a Datastore Shrink \(Preview\)](#)

Modeling a Cluster Addition

Foglight Capacity Director allows you to model a cluster addition without deploying the planned change in the realtime environment.

To schedule a cluster addition:

- 1 In the Planned Changes view, click Add.
The wizard of scheduled planned changes appears.
- 2 Select **Model a Cluster Addition**, and click **Next**.
- 3 In the **Model a Cluster Addition Wizard** dialog box, specify the following fields, as needed.
 - Cluster Name: Type the name of the cluster.
 - Organization: Select an organization from the drop-down list.
 - Deploy Date: Click the icon  and select the deployment date.
 - Time: Click the icon  to select the accurate time on the selected deployment date.
 - Host Quantity: Type the number of host.
 - Host Configuration: Specify the following options:
 - CPU: Type the CPU frequency. The unit is one of the following: GHz (default option), MHz, and THz.
 - Memory: Type the physical memory. The unit is one of the following: GB (default option), MB, and TB.
 - Storage: Type the physical storage. The unit is one of the following: GB (default option), MB, and TB.
 - HA Policy: Specify whether the HA policy is enabled or disabled on clusters, as needed.
 - Host failures the cluster tolerates: Select this option to enable the “Host failures the cluster tolerates” HA policy on clusters, then type the number of host that enables this HA policy.

- Percentage of cluster resources reserved as failover spare capacity: Select this option to enable the “Percentage of cluster resources reserved as failover spare capacity” HA policy on clusters, then type the following percentages:
 - CPU (%): Type the percent of the reserved CPU resource.
 - Memory (%): Type the percent of the reserved Memory resource.
 - Disable HA Policy: Select this option to disable the HA policy on clusters.
- Datacenter: Click Browse to select a cluster in the Cluster - Datacenter Selector dialog box, and click Save.
- Like Cluster: Select a cluster from the drop-down list. The Average VM Utilization fields are automatically filled in based on your selection.
- Average VM Utilization: Type the average utilization of virtual machines. These fields are automatically filled in after you select a cluster from the Like Cluster drop-down list. You can edit this value if needed.
 - CPU: Type the CPU frequency. The unit is one of the following: GHz (default option), MHz, and THz.
 - Memory: Type the physical memory. The unit is one of the following: GB (default option), MB, and TB.
 - Storage: Type the physical storage. The unit is one of the following: GB (default option), MB, and TB.
 - Throughput: Type the value of throughput. The unit includes the following: MB/s (default option), B/s, KB/s, GB/s, and TB/s.

4 Click **Finish**.

The Planned Changes view is updated to include this new change.

Modeling a Host Addition

Foglight Capacity Director allows you to model a host addition without deploying the planned change in the realtime environment.



To schedule a host addition:

- 1 In the Planned Changes view, click **Add**.

The wizard of scheduled planned changes appears.

- 2 Select **Model a Host Addition**, and click **Next**.

- 3 In the **Model a Host Addition Wizard** dialog box, specify the following fields, as needed:

- Host Name: Type the name of the host.
- Deploy Date: Click the icon  and select the deployment date.
- Time: Click the icon  to select the accurate time on the selected deployment date.
- Change Type: Includes the following options:
 - Add Host To Cluster: The planned change type is “Add Host To Cluster”.
 - Add Standalone Host: The planned change type is “Add Standalone Host”.
- Cluster: Click Browse to select a cluster in the Host - Cluster Selector dialog box, then click Select.
- Like Host: Select a host from the drop-down list. The Configuration fields are automatically filled in based on your selection.

- Configuration: Specify the following options. These fields are automatically filled in after you select a host from the Like Host drop-down list. You can edit them as needed.
 - CPU: Type the CPU frequency. The unit is one of the following: GHz (default option), MHz, and THz.
 - Memory: Type the physical memory. The unit is one of the following: GB (default option), MB, and TB.
 - Storage: Type the physical storage. The unit is one of the following: GB (default option), MB, and TB.

4 Click **Finish**.

The Planned Changes view is updated to include this new change.

Adding Virtual Machines

Foglight Capacity Director allows you to create and deploy a new change for adding virtual machines in the realtime environment.

To schedule a change for adding virtual machines:

- 1 In the Planned Changes view, click **Add**.

The wizard of scheduled planned changes appears.

- 2 Select **Add Virtual Machines**, and click **Next**.

- 3 In the **Add Virtual Machines Wizard** dialog box, specify the following fields, as needed:

- VM Name: Type the name of the virtual machine.

i **NOTE:** If the value of Automated is set to Yes when deploying multiple virtual machines in the environment, the VM Name must end with one of the following: #, ##, and ###. In other words, if ten virtual machines are to be deployed to the environment, the VM Name must end with ##. For example, you plan to add 15 virtual machines and define the VM Name as VM_##, the new virtual machines' names will be VM_01, VM_02,... VM_15.

- Deploy Date: Click the icon  and select the deployment date.

- Time: Click the icon  to select the accurate time on the selected deployment date.

- Automated: Specify whether this planned change is deployed automatically or not.

- VM Quantity: Type the number of virtual machines.

i **NOTE:** VM Quantity cannot exceed the maximum value of the following customized patterns, if you plan to deploy multiple virtual machines and the value of Automated is set to Yes.

- #: The value of this customized pattern can be any number from 1 to 9.
- ##: The value of this customized pattern can be any number from 1 to 99.
- ###: The value of this customized pattern can be any number from 1 to 500.

- VM Size (available only if Automated is set to No): Includes the following options:

- Existing VM (default option): Specify the size of virtual machine based on an existing virtual machine. The Browse button is enabled after selecting this option. Click Browse to select a virtual machine in the VM Selector dialog box, and click Select. The Average Utilization values are filled in automatically after selecting a virtual machine.

- Default: Set the size of virtual machine to default values (CPU 2.0 GHz, Memory 4.0 GB, Storage 25.0 GB, and Throughput 1.0 MB/s). The Browse button is disabled and the Average Utilization values are filled in automatically after selecting this option.

- Cluster / Host: Click Browse to select a virtual machine in the VM - Cluster / Host Selector dialog box, and click Save.

- Resource Pool (available only if Automated is set to Yes): Click Browse to select a resource pool in the Resource Pool Selector dialog box, and click Select.
- Storage: Click Browse to select a virtual machine in the VM - Storage Selector dialog box, and click Save.
- Template (available only if Automated is set to Yes): Click Browse to select a virtual machine template in the VM - Template Selector dialog box, and click Save.
- Allocation (available only if Automated is set to Yes): These values are automatically filled in after you select a virtual machine template.
 - vCPU: Type the number of vCPU.
 - Memory: Type the physical memory. The unit is one of the following: GB (default option), MB, and TB.
- Average Utilization: These values are automatically filled in after you select a value from the VM Size drop-down list. You can edit these values, if needed.
 - CPU: Type the CPU frequency. The unit is one of the following: GHz (default option), MHz, and THz.
 - Memory: Type the physical memory. The unit is one of the following: GB (default option), MB, and TB.
 - Storage: Type the physical storage. The unit is one of the following: GB (default option), MB, and TB.
 - Throughput: Type the value of throughput. The unit includes the following: MB/s (default option), B/s, KB/s, GB/s, and TB/s.

4 Click **Finish**.

The Planned Changes view is updated to include this new change.

Modeling a Host Decommission

Foglight Capacity Director allows you to model a host decommission without deploying the planned change in the realtime environment.



To schedule a change for decommissioning a host:

1 In the Planned Changes view, click **Add**.

The wizard of scheduled planned changes appears.

2 Select **Model a Host Decommission**, and click **Next**.

3 In the **Model a Host Decommission Wizard dialog box**, specify the following fields, as needed:

- Host Name: Click Browse to select the host to be decommissioned in the Host Selector dialog box, and click Select.
- Decommission Date: Click the icon  and select the decommission date.
- Time: Click the icon  to select the accurate time on the selected decommission date.



4 Click **Finish**.

The Planned Changes view is updated to include this new change.

Decommissioning Virtual Machines

Foglight Capacity Director allows you to create and deploy a new change for decommissioning virtual machines from the realtime environment.



To schedule a change for decommissioning a virtual machine:

- 1 In the Planned Changes view, click **Add**.
The wizard of scheduled planned changes appears.
- 2 Select **Decommission Virtual Machines**, and click **Next**.
- 3 In the **Decommission Virtual Machines Wizard dialog box**, specify the following fields, as needed:
 - VM Name: Click Browse to select a virtual machine to be decommissioned in the VM Selector dialog box, and click Select.
 - Automated: Specify whether this planned change is deployed automatically or not.
 - Schedule Decommission: Specify the details about the virtual machine decommission.
 - Now / Later (available only if Automated is set to Yes): Specify this planned change is deployed now or later.
 - Date: Click the icon  and select the decommission date.
 - Time: Click the icon  to select the accurate time on the selected decommission date.
- 4 Click **Finish**.
The Planned Changes view is updated to include this new change.

Modeling a Host Movement

Foglight Capacity Director allows you to model a host movement without deploying the planned change in the realtime environment.



To schedule a change for moving a host:

- 1 In the Planned Changes view, click **Add**.
The wizard of scheduled planned changes appears.
- 2 Select **Model a Host Movement**, and click **Next**.
- 3 In the **Model a Host Movement Wizard dialog box**, specify the following fields, as needed:
 - Host Name: Click Browse to select a host to be moved in the Host Selector dialog box, and click Select.
 - Source: This value is the original destination where the host is deployed and is automatically filled in after you select a host.
 - Destination: Click Browse to select a destination to which the host is moved in the Destination Selector dialog box, and click Select.
 - Move Date: Click the icon  and select the movement date.
 - Time: Click the icon  to select the accurate time on the selected movement date.
- 4 Click **Finish**.
The Planned Changes view is updated to include this new change.

Moving Virtual Machines

Foglight Capacity Director allows you to create and deploy a new change for moving virtual machines between various destinations.

To schedule a change for moving a virtual machine:

- 1 In the Planned Changes view, click **Add**.
The wizard of scheduled planned changes appears.
- 2 Select **Move Virtual Machines**, and click **Next**.
- 3 In the **Move Virtual Machines** Wizard dialog box, specify the following fields, as needed:
 - **VM Name:** Click Browse to select a virtual machine to be moved in the VM Selector dialog box, and click Select.
 - **Source:** This value is the original destination where the virtual machine is deployed and is automatically filled in after you select a virtual machine.
 - **Automated:** Specify whether this planned change is deployed automatically or not.
 - **Destination:** Click Browse to select a destination to which the virtual machine is moved in the VM Movement - Destination Selector dialog box, and click Save.
 - **Specific Host** (available only if Automated is set to Yes): Click Browse to select a host in the VM Movement - Specific Host Selector dialog box, and click Save.
 - **Resource Pool** (available only if Automated is set to Yes): Click Browse to select a resource pool in the Resource Pool Selector dialog box, and click Select.
 - **Storage:** Click Browse to select a storage in the VM Movement - Storage Selector dialog box, and click Save.
 - **Move Date:** Click the icon  and select the movement date.
 - **Time:** Click the icon  to select the accurate time on the selected movement date.
- 4 Click **Finish**.
The Planned Changes view is updated to include this new change.

For migrating Hyper-V VMs, please do the following:

- 1 Set all `isInitiator=true` for `module.com.sun.security.jgss.initiate` (For AIX, set the `com.ibm.security.jgss.initiate`) from `FGLAM_HOM\state\default\config\auth.login.config`.

```
10 com.sun.security.jgss.login {  
11   com.sun.security.auth.module.Krb5LoginModule required useTicketCache=false isInitiator=true debug=false;  
12 };  
13  
14 com.sun.security.jgss.initiate {  
15   com.sun.security.auth.module.Krb5LoginModule required useTicketCache=false isInitiator=true debug=false;  
16 };  
17  
18 com.sun.security.jgss.accept {  
19   com.sun.security.auth.module.Krb5LoginModule required useTicketCache=false isInitiator=true debug=false;  
20 };
```

- 2 Set `forwardable=true` in `libdefaults` from `FGLAM_HOM\state\default\config\krb5.config`.

```

# Copyright 2018 Quest Software Inc.
# ALL RIGHTS RESERVED.

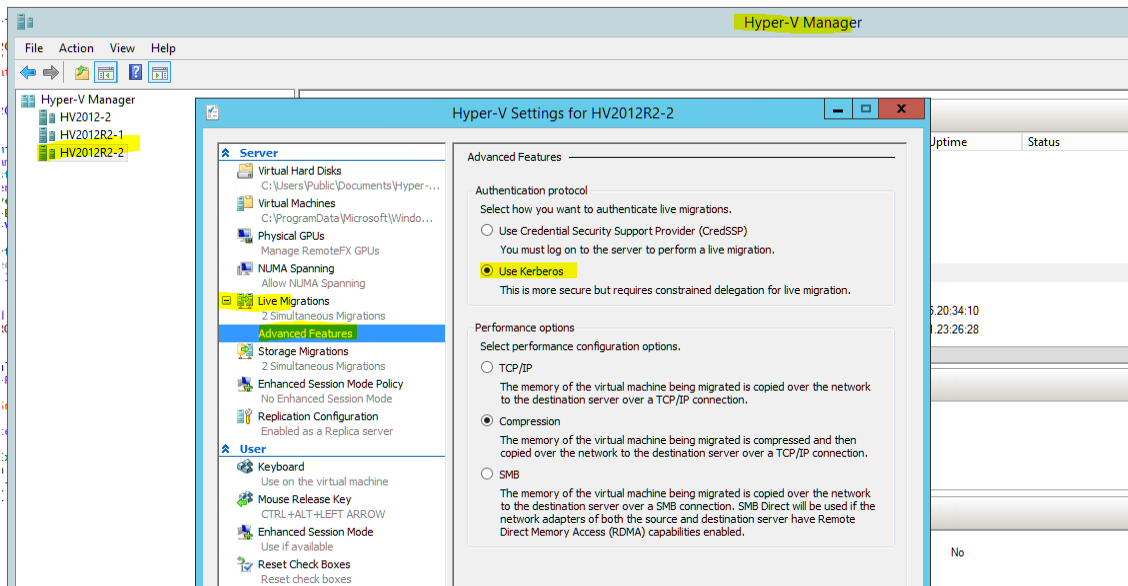
[libdefaults]
    default_realm = VFOG.LOCAL
    forwardable=true

[realms]
    VFOG.LOCAL = {
        kdc = ZHUVM-FOG-3326.VFOG.LOCAL
    }


[domain_realm]
    .vfog.local = VFOG.LOCAL

```

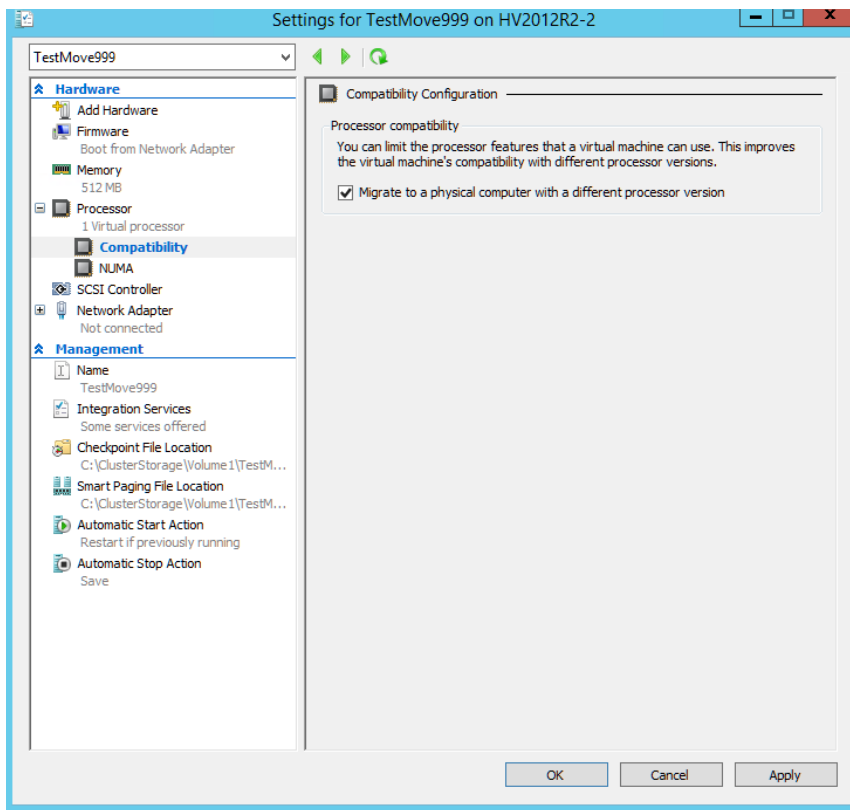
- 3 In Hyper-V Manager, select the **Hyper-V server and settings**. Then, go to **Live Migration and Advanced Features**. Then, select **Use Kerberos** under **Authentication Protocol**.



- 4 In **Active Directory Users and Computers**, find the Hyper-V server and go to **properties**. Then, find **Delegation** tab and add **CIFS** and **Microsoft Virtual System Migration Service**.





- 5 For Live Migration, In Hyper-V Manager, find the VM and go to **properties**. Then, expand **Processor** and select **Compatibility** and check **Migrate to physical computer with a different processor version**.



Adding a Datastore (Preview)

Foglight Capacity Director allows you to add a datastore without deploying the planned change in the realtime environment.



To schedule a datastore addition:

- 1 In the Planned Changes view, click **Add**.
The wizard of scheduled planned changes appears.
- 2 Select **Add a Datastore (Preview)**, and click **Next**.
- 3 In the **Model a Datastore Addition Wizard** dialog box, specify the following fields, as needed.
 - **Datastore Name:** Type the name of the datastore.
 - **Deploy Date:** Click the icon  and select the deployment date.
 - **Time:** Click the icon  to select the accurate time on the selected deployment date.
 - **Capacity:** Type the number of capacity. The unit is one of the following: GB (default option), B, KB, MB, TB, and b.
 - **Select Hosts:** Select a host from the Select Hosts table.
- 4 Click **Finish**.
The Planned Changes view is updated to include this new change.

Modeling a Datastore Decommission (Preview)

Foglight Capacity Director allows you to decommission a datastore without deploying the planned change in the realtime environment.

To schedule a datastore decommission:

- 1 In the Planned Changes view, click Add.
The wizard of scheduled planned changes appears.
- 2 Select **Model a Datastore Decommission (Preview)**, and click **Next**.
- 3 In the **Model a Datastore Decommission Wizard** dialog box, specify the following fields, as needed.
 - Datastore Name: Click Browse to select a datastore in the Datastore Selector dialog box, and click Select.
 - Deploy Date: Click the icon  and select the deployment date.
 - Time: Click the icon  to select the accurate time on the selected deployment date.
- 4 Click **Finish**.
The Planned Changes view is updated to include this new change.

Modeling a Datastore Expansion (Preview)

Foglight Capacity Director allows you to expand a datastore without deploying the planned change in the realtime environment.

To schedule a datastore expansion:

- 1 In the Planned Changes view, click Add.
The wizard of scheduled planned changes appears.
- 2 Select **Model a Datastore Expansion (Preview)**, and click **Next**.
- 3 In the **Model a Datastore Expansion Wizard** dialog box, specify the following fields, as needed.
 - Datastore Name: Click Browse to select a datastore in the Datastore Selector dialog box, and click Select.
 - Current Utilization: These values are automatically filled in after you select a datastore.
 - Capacity: The capacity of the selected datastore.
 - Allocated: The allocated capacity on the selected datastore and the relevant percent.
 - In Use: The capacity that is being used and the relevant percent.
 - Available Space: The available capacity on the selected datastore.
 - Extent Capacity: Type the number of capacity to be expanded. The unit is one of the following: GB (default option), B, KB, MB, TB, and b.
- 4 Click **Finish**.
The Planned Changes view is updated to include this new change.

Modeling a Datastore Shrink (Preview)

Foglight Capacity Director allows you to shrink a datastore without deploying the planned change in the realtime environment.

To schedule a datastore shrink:

- 1 In the Planned Changes view, click **Add**.
The wizard of scheduled planned changes appears.
- 2 Select **Model a Datastore Shrink (Preview)**, and click **Next**.
- 3 In the **Model a Datastore Shrink Wizard** dialog box, specify the following fields, as needed.
 - **Datastore Name:** Click **Browse** to select a datastore in the Datastore Selector dialog box, and click **Select**.
 - **Current Utilization:** These values are automatically filled in after you select a datastore.
 - **Capacity:** The capacity of the selected datastore.
 - **Allocated:** The allocated capacity on the selected datastore and the relevant percent.
 - **In Use:** The capacity that is being used and the relevant percent.
 - **Available Space:** The available capacity on the selected datastore.
 - **Shrink Capacity:** Type the number of capacity to be shrunk. The unit is one of the following: GB (default option), B, KB, MB, TB, and b.
- 4 Click **Finish**.
The Planned Changes view is updated to include this new change.

Editing a Planned Change

Foglight Capacity Director allows you to edit the existing planned changes, as needed.

To edit an existing planned change:

- 1 In the Planned Changes view, select a planned change that you want to edit, and click the **Edit** icon.
- 2 The **Settings** dialog box appears.
i **NOTE:** The **Settings** dialog box varies from the following change types: Model a Cluster Addition, Model a Host Addition, Add Virtual Machines, Model a Host Decommission, Decommission Virtual Machines, Model a Host Movement, and Move Virtual Machines. For more information, see [Adding a Planned Change](#) on page 50.
- 3 Edit the values as needed, and click **Save**. For detailed information about the available fields, refer to [Adding a Planned Change](#) on page 50.
The Planned Changes view is updated to show the updates of this planned change.

Deleting a Planned Change

Foglight Capacity Director allows you to remove the existing planned changes that you do not want to deploy.

To delete one or more planned change(s):

- 1 In the Planned Changes view, select one or more planned change(s) that you want to remove, and click **Delete**.
- 2 In the **Confirm delete planned changes** dialog box, click **Confirm**.
The selected planned change is removed from the Planned Changes view.

Cloud Model Tab

The *Cloud Model* view helps you understand costs and risks caused by migrating virtual machine workloads to any MSP, Cloud, or IaaS offering. Simply enter the proposed destinations that are under review, like Azure A2 Small, select the virtual machines that have been designated for a cloud migration and input the level of acceptable performance risk for CPU, Memory, Disk IO and Network. Select the cloud IaaS offering, select the virtual machines to be modeled, and then click **Run Model**. The *Cloud Model* view includes the following elements:

- *Cloud Instances pane*: Allow to add or edit the hardware specification and unit cost of your proposed IaaS subscriptions, such as Azure, AWS EC2. IaaS subscriptions are supported to import or export as a .CSV file.
- *Select Virtual Machines*: Lists available virtual machines in the selected virtual center or datacenter, and provides the capability to search VMs as needed.
- *Select Virtual Machines*: Lists available virtual machines in the selected virtual center or datacenter, and provides the capability to search VMs as needed.
- *Run Model button*: Click this button to start the recommendation analysis based on historical performance metrics for the specified values (including selected virtual machines, performance risks, selected Cloud Instances, and the modeling scenario).
- *Model Outcome pane*: Lists migration recommendations and resources fit analysis. In addition, this pane shows migration risks caused by restricted hardware resources, including *CPU Demand*, *Memory Consumed*, *Disk IOPs*, *Disk Size*, and *Unable to Model* information.

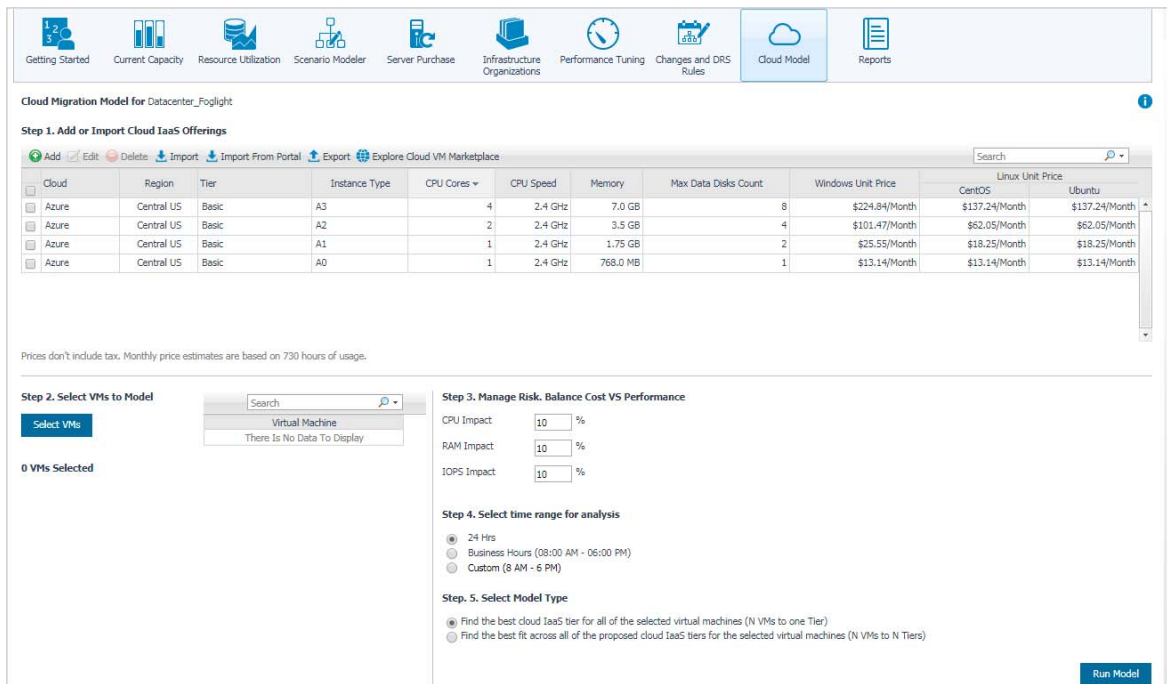
To access the *Cloud Model* view:

- 1 On the navigation panel, under **Dashboards**, click **Capacity Director**.

The *Capacity Director* dashboard opens.

- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the **Cloud Model** tab.

The *Cloud Model* view appears at the bottom of the *Capacity Director* dashboard.



Cloud Instances pane

The Cloud Instance table locates in the middle of the *Cloud Model* view offers the following menus:

- The Add button: Provides the functionality to add a new Cloud instance. For more information, see [Adding a Cloud instance](#) on page 63.
- The Edit button: Provides the functionality to edit existing Cloud instances. For more information, see [Editing existing Cloud instance](#) on page 63.
- The Delete button: Provides the functionality to delete existing Cloud instances. For more information, see [Delete existing Cloud instance\(s\)](#) on page 64.
- The Import button: Provides the functionality that helps you easily import Cloud instances in batches, as needed. For more information, see [Importing Cloud instances](#) on page 64.
- The Import From Portal button: Provides the functionality that helps you easily import Cloud instances from the portal, as needed. For more information, see [Importing Cloud Models from Portal](#) on page 64.
- The Export button: Provides the functionality that helps you easily import Cloud instances in batches, as needed. For more information, see [Exporting Cloud instances](#) on page 65.
- The Run Model button: Provides the functionality that starts the analysis. For more information, see [Running Cloud Model](#) on page 65.
- Cloud instance table: Lists all cloud instances that are available in the selected virtual center or datacenter.

Description

Lists the detailed information about all Cloud instances that are available in the selected vCenter or datacenter.

Data displayed

- **Name.** The provider name of the Cloud instance.
- **Tier.** The name of tier.
- **Instance Type.** The type of Cloud instance.
- **VMs to Migrate.** The number of VMs that can be migrated to the Cloud. This value is not available until you click **Run Model** and get analysis result.

- **Total Cost.** The total cost associated with VM migration. This value is not available until you click **Run Model** and get analysis result.
- **Unit Cost.** The cost per unit.
- **CPU Cores.** The number of CPU cores.
- **CPU Speed.** The CPU speed.
- **Memory.** The memory size.
- **Disk Size.** The size of disk.
- **IO Rate(IOPs).** The measurement of input/output performance that is used to characterize computer storage devices.

Adding a Cloud instance

To add a new Cloud instance:

- 1 In the Cloud Instance pane, click **Add**.
The **Add Cloud Model** dialog box opens.
- 2 In the Add Cloud Model dialog box, specify the following values, as needed.
 - **Cloud:** Type the provider name of the Cloud instance.
 - **Tier:** Type the name of the Cloud tier.
 - **Instance Type:** Type the type of the Cloud instance.
 - **Unit Cost:** Type the cost per unit.
 - **CPU Cores:** Type the number of CPU cores available in the Cloud instance.
 - **CPU Speed:** Type the CPU speed. The unit is one of the following: GHz (default option), MHz, and THz.
 - **Memory:** Type the physical memory. The unit is one of the following: GB (default option), MB, and TB.
 - **Disk Size:** Type the physical disk size. The unit is one of the following: GB (default option), MB, and TB.
 - **IO Rate(IOPs):** Type the IO Rate.
- 3 Click **Save**.
The *Cloud Instance* table refreshes automatically and the new Cloud instance is added into this table.

Editing existing Cloud instance

To edit an existing Cloud instance:

- 1 In the Cloud Instance pane, select a Cloud instance that you want to edit, and then click **Edit**.
The **Edit Cloud Model** dialog box opens.
- 2 In the Edit Cloud Model dialog box, modify the following values, as needed.
 - **Cloud:** Type the provider name of the Cloud instance.
 - **Tier:** Type the name of the Cloud tier.
 - **Instance Type:** Type the type of the Cloud instance.
 - **Unit Cost:** Type the cost per unit.

- CPU Cores: Type the number of CPU cores available in the Cloud instance.
 - CPU Speed: Type the CPU speed. The unit is one of the following: GHz (default option), MHz, and THz.
 - Memory: Type the physical memory. The unit is one of the following: GB (default option), MB, and TB.
 - Disk Size: Type the physical disk size. The unit is one of the following: GB (default option), MB, and TB.
 - IO Rate(IOPs): Type the IO Rate.
- 3 Click **Save**.

The *Cloud Instance* table refreshes automatically and the Cloud instance is updated.

Delete existing Cloud instance(s)

To delete existing Cloud instance(s):

- 1 In the Cloud Instance pane, select one or more Cloud instances that you want to delete, and then click **Delete**.

The **Confirm Delete Cloud Models Dialog** box opens.

- 2 In the Confirm Delete Cloud Models dialog box, click **Delete**.

The *Cloud Instance* table refreshes automatically and the selected Cloud instances are removed.

Importing Cloud instances

To import Cloud instances:

- 1 In the Cloud Instance pane, click **Import**.

The **Import Host Models** dialog box opens.

- 2 In the Import Host Models dialog box, click **Choose File**.

- 3 In the prompted dialog box, browse to select a .csv file, and then click **Open**.

- 4 Click **Import**.

The *Cloud Instance* table refreshes automatically and imported Cloud instances are added into this table.

Importing Cloud Models from Portal

To import Cloud models from the portal:

- 1 In the Cloud Instance pane, click **Import From Portal**.

The **Import Cloud Models From Portal** dialog box opens.

- 2 In the Import Cloud Models From Portal dialog box, select a cloud provider and a region, the instances table refreshes automatically.

- 3 In the instances table, select the cloud models as needed, and then click **Save**.

The *Cloud Instance* table refreshes automatically and imported Cloud instances are added into this table.

Exporting Cloud instances

To export Cloud instances:

- 1 In the Cloud Instance pane, click **Export**.
The existing Cloud instances are downloaded as a .csv file automatically.

Running Cloud Model

To run Cloud model:

- 1 In the Cloud Model view, click **Select Virtual Machines**.
The **Select Virtual Machines** dialog box opens.
- 2 In the **Select Virtual Machines** dialog box, select the VMs that you want to migrate to the Cloud instance, and then click **Save**.
- 3 In the **Performance Risk** area, specify the following:
 - NOTE:** If the virtual machines that have been selected for migration must always perform at their peak, enter 0% against each of CPU, Memory, Disk and Network. If Cost is the most important factor and performance is not a concern, enter 100 against CPU, Memory, Disk and Network; that way, no matter how often the virtual machine will exceed the destinations available resources it will always be included in the cloud model outcome and costs calculation.
 - CPU: The default value is 10%.
 - Memory: The default value is 10%.
 - Disk IO: The default value is 10%.
- 4 In the *Cloud Instances* table, select the Cloud instances to which VMs will be migrated, and then click **Run Model**.
The **Scenarios** dialog box opens.
- 5 In the **Scenarios** dialog box, select either of the following:
 - Find the best cloud IaaS tier for all of the selected virtual machines (N VMs to one Tier): Move all VMs to the same tier.
 - Find the best fit across all of the proposed cloud IaaS tiers for the selected virtual machines (N VMs to N Tiers): Moves all VMs to different tiers.
- 6 Click **Go**.
The Foglight Capacity Director runs the analysis and shows recommended values in the *VMs to Migrate* and *Total Cost* columns. For more information, see [Model Outcome pane](#) on page 65.

Model Outcome pane

The Model Outcome pane located on the bottom of the *Cloud Model* view offers the following menus:

- The *Model Outcome for* selector: Lists the Cloud instances selected in the *Cloud Instances* table.
- CPU Demand: Lists the VMs that are failed to migrate because of CPU restrictions.
- Memory Consumed: Lists the VMs that are failed to migrate because of memory restrictions.
- Disk IOPS: Lists the VMs that are failed to migrate because of disk IOPs restrictions.
- Disk Size: List the VMs that are failed to migrate because of disk size restrictions.

- **Unable to Model:** Shows the virtual machines that are unable to model due to missing or outdated VMware Tools. Check for VMware Tools status with virtual machine.
- **VM table:** Lists the detailed information about all VMs that can be migrated to the selected Cloud instance, including *VM Name*, *CPU Fit*, *Memory Fit*, *Disk Space Fit*, *IOPs Fit*, and *Recommendation*.

Report Tab

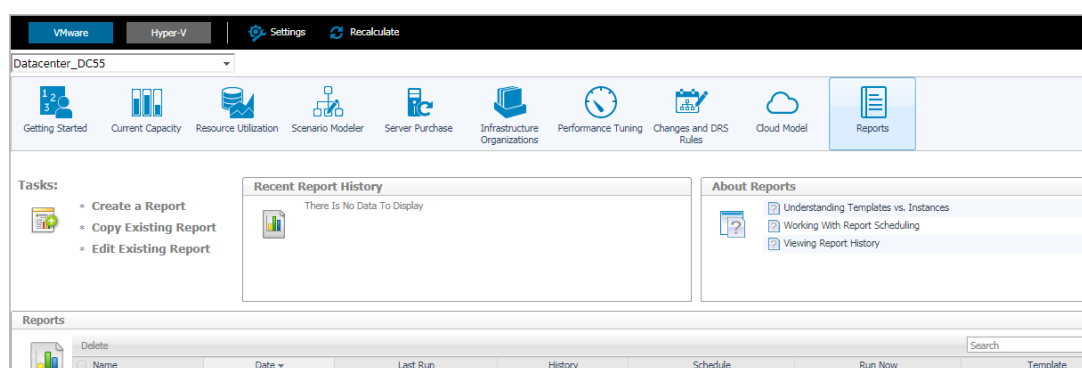
Foglight Capacity Director includes a report generation ability. This allows you to create reports using a set of predefined templates to report on the various aspects of your cloud environment. Foglight Capacity Director includes a collection of predefined report templates.

You can generate, copy, and edit reports using the Reports tab on the *Report* dashboard, or alternatively the *Reports* dashboard included with the Management Server.

To access the Report dashboard:

- 1 On the navigation panel, under Dashboards, click Capacity Director.
The Capacity Director dashboard opens.
- 2 Use the Group selector located at the top of the dashboard to select the virtual environment that you want to monitor.
- 3 On the Menu bar, click the Reports tab.

The Reports view appears at the bottom of the Capacity Director dashboard.



For complete information about this tab, see the *Managing Capacity in Virtual Environments User Guide*. For more information about the *Reports* dashboard, see the *Foglight User Help*.

Foglight Capacity Director allows you to print reports from the *Current Capacity* view and the *Resource Utilization* view. For more information, see the following topics:

- [Printing Reports from Current Capacity](#)
- [Printing Reports from Resource Utilization](#)

Printing Reports from Current Capacity


Foglight Capacity Director allows you to print the following reports from the Current Capacity view.

- **Hosts Report:** Walks through the **Current Capacity - Hosts Report** wizard to print a host report that includes the information related to Service, Slot Size, Datastore Include, Selected cluster, and selected Planned Changes.
- **Clusters Report:** Walks through the **Current Capacity - Clusters Report** wizard to print a cluster report that includes the information related to Service, Slot Size, Datastore Include, and selected Planned Changes.

Hosts Report

To print a hosts report from the Current Capacity tab:

- 1 Click **Reports** on the top right corner of the *Foglight* dashboard.



- 2 Click **Current Capacity - Hosts Report** from the context menu.

The Current Capacity - Hosts Report wizard appears.


- 3 In the Set Input Parameters page, supply the values for the report input parameters, and click **Next**.
- 4 In the Set Properties page, select the report properties:
 - Name: Type the name of this new report.
 - Output Format: Select an output format from the drop-down list. The output format includes the following options: PDF (default option), Excel, and XML.
 - Email Recipients: Type the email address, and separate the email addresses using a comma if there are multiple recipients.
 - Schedule This Report: Select this check box to enable this new report.
 - Retain: Type the number of days for which this report is retained. This value is only available after selecting the Schedule This Report option.
- 5 Click **Next**.
- 6 Select the scheduled date for executing this report, as needed, then click **Next**. This configuration is only available after selecting the Schedule This Report option.
- 7 Review the Summary page to ensure that the provided information are correct, and click **Finish**.
- 8 In the File Download dialog box that appears, choose one of the following options:
 - Open the report for viewing.
 - a Click **Open**.
 - b The report is opened automatically for viewing.
 - Download the report to a location of your choice.
 - a Click **Save**.
 - b In the dialog box that appears, select a location for saving the file, then click **Save**.

The report is saved as the output format specified in [Step 4](#) to the selected location.

Clusters Report

To print a clusters report from the Current Capacity view:

- 1 Click **Reports** on the top right corner of the *Foglight* dashboard.



- 2 Click **Current Capacity - Clusters Report** from the context menu.
The Current Capacity - Cluster Report wizard appears.
- 3 In the Set Input Parameters page, supply the values for the report input parameters, and click **Next**.
- 4 In the Set Properties page, select the report properties:
 - Name: Type the name of this new report.
 - Output Format: Select an output format from the drop-down list. The output format includes the following options: PDF (default option), Excel, and XML.
 - Email Recipients: Type the email address, and separate the email addresses using a comma if there are multiple recipients.
 - Schedule This Report: Select the check box to enable this new report.
 - Retain: Type the number of days for which this report is retained. This value is only available after selecting the Schedule This Report option.
- 5 Click Next.
- 6 Select the scheduled date for executing this report, as needed, then click Next. This configuration is only available after selecting the Schedule This Report option.
- 7 Review the Summary page to ensure that the provided information are correct, and click **Finish**.
- 8 In the File Download dialog box that appears, choose one of the following options:
 - Open the report for viewing.
 - a Click Open.
 - b The report is opened automatically for viewing.
 - Download the report to a location of your choice.
 - a Click Save.
 - b In the dialog box that appears, select a location for saving the file, then click Save.

The report is saved as the output format specified in [Step 4](#) to the selected location.

Printing Reports from Resource Utilization


Foglight for VMware enables you to print the following reports from the Resource Utilization view.

- [Clusters Report](#): Walks through the **Resource Utilization - Clusters Report** wizard to print a cluster report that includes the information related to Service, selected cluster, and selected Planned Changes.
- [Hosts and Workload Report](#): Walks through the **Resource Utilization - Hosts and Workload Report** wizard to print a host/workload report that includes the information related to Service, Slot Size, Datastore Include, and selected Planned Changes.

Clusters Report

To print a clusters report from the Resource Utilization view:

- 1 Click **Reports** on the top right of the *Foglight* dashboard.



- 2 Click **Resource Utilization - Clusters Report** from the context menu.

The Resource Utilization - Clusters Report wizard appears.

- 3 In the Set Input Parameters page, supply the values for the report input parameters, and click **Next**.

- 4 In the Set Properties page, select the report properties:

- Name: Type the name of this new report.
- Output Format: Select an output format from the drop-down list. The output format includes the following options: PDF (default option), Excel, and XML.
- Email Recipients: Type the email address, and separate the email addresses using a comma if there are multiple recipients.
- Schedule This Report: Select the check box to enable this new report.
- Retain: Type the number of days for which this report is retained. This value is only available after selecting the Schedule This Report option.

- 5 Click **Next**.

- 6 Select the scheduled date for executing this report, as needed, then click **Next**. This configuration is only available after selecting the Schedule This Report option.

- 7 Review the Summary page to ensure that the provided information are correct, and click **Finish**.

- 8 In the File Download dialog box that appears, choose one of the following options:


- Open the report for viewing.
 - a Click **Open**.
 - b The report is opened automatically for viewing.
- Download the report to a location of your choice.
 - a Click **Save**.
 - b In the dialog box that appears, select a location for saving the file, then click **Save**.

The report is saved as the output format specified in [Step 4](#) to the selected location.

Hosts and Workload Report

To print a clusters report from the Resource Utilization view:

- 1 Click **Reports** on the top right of the *Foglight* dashboard.



- 2 Click **Resource Utilization - Hosts and Workload Report** from the context menu.
The Resource Utilization - Hosts and Workload Report wizard appears.
- 3 In the Set Input Parameters page, supply the values for the report input parameters, and click **Next**.
- 4 In the Set Properties page, select the report properties:
 - Name: Type the name of this new report.
 - Output Format: Select an output format from the drop-down list. The output format includes the following options: PDF (default option), Excel, and XML.
 - Email Recipients: Type the email address.
 - Schedule This Report: Select the check box to enable this new report.
 - Retain: Type the number of days for which this report is retained.
- 5 Click Next.
- 6 Select the scheduled date for executing this report, as needed, then click Next. This configuration is only available after selecting the Schedule This Report option.
- 7 Review the Summary page to ensure that the provided information are correct, and click **Finish**.
- 8 In the File Download dialog box that appears, choose one of the following options:
 - Open the report for viewing.
 - a Click Open.
 - b The report is opened automatically for viewing.
 - Download the report to a location of your choice.
 - a Click Save.
 - b In the dialog box that appears, select a location for saving the file, then click Save.

The report is saved as the output format specified in [Step 4](#) to the selected location

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit www.quest.com/contact-us.aspx.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.