



Managing clusters

OnCommand Unified Manager 9.5

NetApp
October 23, 2024

Table of Contents

- Managing clusters 1
 - How the cluster discovery process works 1
 - Viewing the list of monitored clusters 2
 - Adding clusters 2
 - Editing clusters 4
 - Removing clusters 4
 - Rediscovering clusters 5
- Page descriptions for data source management 5

Managing clusters


You can manage the ONTAP clusters by using Unified Manager to monitor, add, edit, and remove clusters.

How the cluster discovery process works

After you have added a cluster to Unified Manager, the server discovers the cluster objects and adds them to its database. Understanding how the discovery process works helps you to manage your organization's clusters and their objects.

The monitoring interval for collecting cluster configuration information is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster objects in the Unified Manager UI. This time frame is also true when making changes to a cluster. For example, if you add two new volumes to an SVM in a cluster, you see those new objects in the UI after the next polling interval, which could be up to 15 minutes.

The following image illustrates the discovery process:



After all the objects for a new cluster are discovered, Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.



Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time.

Viewing the list of monitored clusters

You can use the Configuration/Cluster Data Sources page to view your inventory of clusters. You can view details about the clusters, such as their name or IP address and communication status.

Before you begin

You must have the Operator, OnCommand Administrator, or Storage Administrator role.

About this task

The list of clusters is sorted by the collection state severity level column. You can click a column header to sort the clusters by different columns.

Steps

1. In the left navigation pane, click **Configuration > Cluster Data Sources**.

Adding clusters

You can add a cluster to OnCommand Unified Manager so that you can monitor the cluster. This includes the ability to obtain cluster information such as the health, capacity, performance, and configuration of the cluster so that you can find and resolve any issues that might occur.

Before you begin

- You must have the OnCommand Administrator role or the Storage Administrator role.
- You must have the host name or cluster management IP address (IPv4 or IPv6) for the cluster.

When using the host name, it must resolve to the cluster management IP address for the cluster management LIF. If you use a node management LIF, the operation fails.

- You must have the user name and password to access the cluster.

This account must have the *admin* role with Application access set to *ontapi*, *ssh*, and *http*.

- You must know the type of protocol (HTTP or HTTPS) that is to be configured on the cluster and the port number use to connect to the cluster.
- You must have adequate space on the Unified Manager server. You are prevented from adding a cluster to the server when greater than 90% of space is already consumed.



You can add clusters which are behind a NAT/firewall by using the Unified Manager NAT IP address. Any connected Workflow Automation or SnapProtect systems must also be behind the NAT/firewall, and SnapProtect API calls must use the NAT IP address to identify the cluster.

About this task

- Each cluster in a MetroCluster configuration must be added separately.
- A single instance of Unified Manager can support a specific number of nodes. If you need to monitor an environment that exceeds the supported node count, you must install an additional instance of Unified Manager to monitor some of the clusters.
- You can monitor a single cluster by two instances of Unified Manager provided that you have configured a second cluster-management LIF on the cluster so that each instance of Unified Manager connects through a different LIF.

Steps

1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
2. On the **Configuration/Cluster Data Sources** page, click **Add**.
3. In the **Add Cluster** dialog box, specify the values as required, and then click **Submit**.
4. If HTTPS is selected, perform the following steps:
 - a. In the **Authorize Host** dialog box, click **View Certificate** to view the certificate information about the cluster.
 - b. Click **Yes**.

Unified Manager checks the certificate only when the cluster is added initially. Unified Manager does not check the certificate for each API call to ONTAP.

If the certificate has expired, you cannot add a new cluster. You must first renew the SSL certificate and then add the cluster.

Results

After all the objects for a new cluster are discovered (about 15 minutes), Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.



Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time. Additionally, if you restart Unified Manager during the data continuity collection period, the collection will be halted and you will see gaps in the performance charts for the missing timeframe.



If you receive an error message that you cannot add the cluster, check to see if the following issues exist:

- If the clocks on the two systems are not synchronized and the Unified Manager HTTPS certificate start date is later than the date on the cluster. You must ensure that the clocks are synchronized using NTP or a similar service.
- If the cluster has reached the maximum number of EMS notification destinations the Unified Manager address cannot be added. By default only 20 EMS notification destinations can be defined on the cluster.

Editing clusters

You can modify the settings of an existing cluster, such as the host name or IP address, user name, password, protocol, and port, by using the Edit Cluster dialog box.

Before you begin

You must have the OnCommand Administrator role or the Storage Administrator role.

About this task



If you change the IP address of a cluster to an IP address of an existing monitored cluster, all data for the existing cluster is lost when the former cluster is discovered. An error message is not displayed to warn you.

Steps

1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
2. On the **Configuration/Cluster Data Sources** page, select the cluster you want to edit, and then click **Edit**.
3. In the **Edit Cluster** dialog box, modify the values as required.
4. Click **Submit**.

Removing clusters

You can remove a cluster from Unified Manager by using the Configuration/Cluster Data Sources page. For example, you can remove a cluster if cluster discovery fails, or when you want to decommission a storage system.

Before you begin

You must have the OnCommand Administrator role or the Storage Administrator role.

About this task

This task removes the selected cluster from Unified Manager. After a cluster is removed, it is no longer monitored. The instance of Unified Manager registered with the removed cluster is also unregistered from the cluster.

Removing a cluster also deletes all its storage objects, historical data, storage services, and all associated events from Unified Manager. These changes are reflected on the inventory pages and the details pages after the next data collection cycle.

Steps

1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
2. On the **Configuration/Cluster Data Sources** page, select the cluster that you want to remove and click **Remove**.
3. In the **Remove Data Source** message dialog, click **Remove** to confirm the remove request.

Rediscovering clusters

You can manually rediscover a cluster from the Configuration/Cluster Data Sources page in order to obtain the latest information about the health, monitoring status, and performance status of the cluster.

About this task

You can manually rediscover a cluster when you want to update the cluster—such as by increasing the size of an aggregate when there is insufficient space—and you want Unified Manager to discover the changes that you make.

When Unified Manager is paired with OnCommand Workflow Automation (WFA), the pairing triggers the reacquisition of the data cached by WFA.

Steps

1. In the left navigation pane, click **Configuration > Cluster Data Sources**.
2. On the **Configuration/Cluster Data Sources** page, click **Rediscover**.

Unified Manager rediscovers the selected cluster and displays the latest health and performance status.



You can obtain the monitoring status of the cluster from the right pane of the Dashboards/Cluster View page.

Page descriptions for data source management

You can view and manage your clusters, including adding, editing, rediscovering, and removing clusters, from a single page.

Configuration/Cluster Data Sources page

The Configuration/Cluster Data Sources page displays information about the clusters that Unified Manager is currently monitoring. This page enables you to add additional clusters, edit cluster settings, and remove clusters.

A message at the bottom of the page indicates how frequently Unified Manager collects performance data from clusters. The default collection interval is five minutes, but you can modify this interval through the maintenance console if you find that collections from large clusters are not completing on time.

Command buttons

- **Add**

Opens the Add Cluster dialog box, which enables you to add clusters.

- **Edit**

Opens the Edit Cluster dialog box, which enables you to edit the settings of the selected cluster.

- **Remove**

Removes the selected cluster and all the associated events and storage objects. After the cluster is removed, it is no longer monitored.



The cluster, its storage objects, and all associated events are removed, and the cluster is no longer monitored by Unified Manager. The instance of Unified Manager registered with the removed cluster is also unregistered from the cluster.




- **Rediscover**

Forces a rediscover operation of the cluster so you can update the collection of health and performance data.

Clusters list

The Clusters list displays the properties of all the discovered clusters. You can click a column header to sort the clusters by that column.

- **Status**

Displays the current discovery status of the data source. The status can be Failed () , Completed () , or In Progress ().

- **Name**

Displays the cluster name.

Note that the name might take fifteen minutes or more to appear after the cluster is first added.

- **Maintenance Mode**

Enables you to specify the timeframe, or “maintenance window”, when a cluster will be down for maintenance so that you do not receive a storm of alerts from the cluster while it is being maintained.

When maintenance mode is scheduled for the future this field displays “Scheduled”, and you can hover your cursor over the field to display the scheduled time. When the cluster is in the maintenance window this field shows “Active”.

- **Host Name or IP Address**

Displays the host name, fully qualified domain name (FQDN), short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.

- **Protocol**

Displays the type of protocol that can be configured on the cluster: HTTP or HTTPS (for a secure connection).

If a connection is established with the cluster by using both protocols, HTTPS is chosen over HTTP. The default is HTTPS.

- **Port**

Displays the port number of the cluster.

If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

- **User Name**

Displays the user name that can be used to log in to the cluster.

- **Operation**

Displays the current operation that is supported by the cluster data source.

The following operations are supported by the data source:

- Discovery

Specifies the operation when the data source is being discovered.

- Health Poll

Specifies the operation when the data source is successfully discovered and has started sampling data.

- Deletion

Specifies the operation when the data source (cluster) is deleted from the respective storage objects list.

- **Operation State**

Displays the state of the current operation. The state can be Failed, Completed, or In Progress.

- **Operation Start Time**

The date and time the operation started.

- **Operation End Time**

The date and time the operation ended.

- **Description**

Any message related to the operation.

Add Cluster dialog box

You can add an existing cluster so that you can monitor the cluster and obtain information about its health, capacity, configuration, and performance.

You can add a cluster by specifying the following values:

- **Host Name or IP Address**

Enables you to specify the host name (preferred) or the IP address (IPv4 or IPv6) of the cluster-management LIF that is used to connect to the cluster. By specifying the host name, you will be able to match the name of the cluster across the web UI, rather than trying to correlate an IP address on one page

to a host name on another page.

- **User Name**

Enables you to specify a user name that can be used to log in to the cluster.

- **Password**

Enables you to specify a password for the specified user name.

- **Protocol**

Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

- **Port**

Enables you to specify the port number used to connect to the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Edit Cluster dialog box

The Edit Cluster dialog box enables you to modify the connection settings of an existing cluster, including the IP address, port, and protocol.

You can edit the following fields:

- **Host Name or IP Address**

Enables you to specify the FQDN, short name, or the IP address (IPv4 or IPv6) of the cluster-management LIF that is used to connect to the cluster.

- **User Name**

Enables you to specify a user name that can be used to log in to the cluster.

- **Password**

Enables you to specify a password for the specified user name.

- **Protocol**

Enables you to specify the type of protocol that can be configured on the cluster. You can enable HTTP or HTTPS (for a secure connection). Connection is established with the cluster by using both protocols and HTTPS is chosen over HTTP. By default, HTTPS is enabled with the default port 443.

- **Port**

Enables you to specify the port number used to connect to the cluster. If the port is not specified, the default port for the selected protocol is used (80 for HTTP or 443 for HTTPS).

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.