



# **What the Unified Manager server does**

## **OnCommand Unified Manager 9.5**

NetApp  
August 29, 2024

This PDF was generated from <https://docs.netapp.com/us-en/oncommand-unified-manager-95/online-help/concept-how-discovery-process-works.html> on August 29, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- What the Unified Manager server does ..... 1
- How the discovery process works ..... 1
- Cluster configuration and performance data collection activity ..... 1
- What a data continuity collection cycle is ..... 3
- What the timestamp means in collected data and events ..... 4

# What the Unified Manager server does

The Unified Manager server infrastructure consists of a data collection unit, a database, and an application server. It provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging.

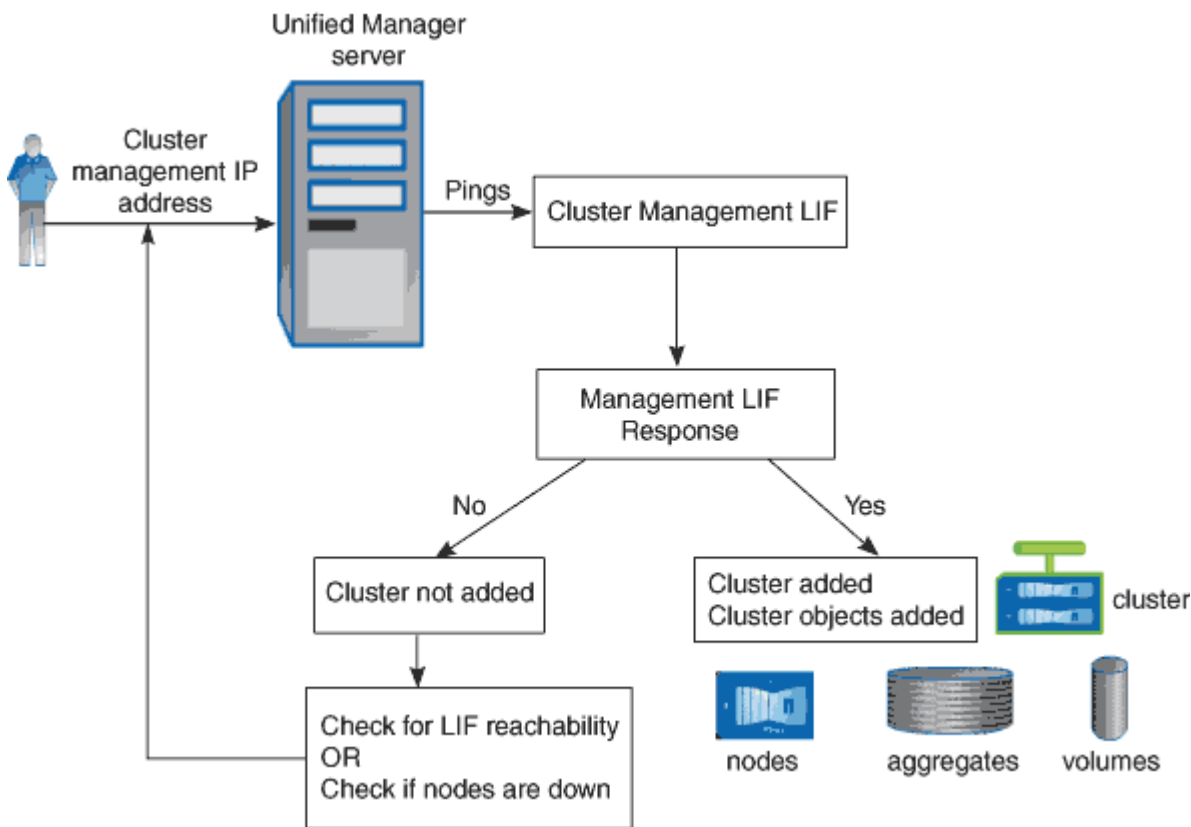
Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

## How the discovery process works

After you have added the 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 default monitoring interval is 15 minutes: if you have added a cluster to Unified Manager server, it takes 15 minutes to display the cluster details in the Unified Manager UI.

The following image illustrates the discovery process in OnCommand Unified Manager:



## Cluster configuration and performance data collection activity

The collection interval for *cluster configuration data* is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster details in the Unified

Manager UI. This interval applies when making changes to a cluster too.

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.

Unified Manager collects current *performance statistics* from all monitored clusters every five minutes. It analyzes this data to identify performance events and potential issues. It retains 30 days of five-minute historical performance data and 390 days of one-hour historical performance data. This enables you to view very granular performance details for the current month, and general performance trends for up to a year.

The collection polls are offset by a few minutes so that data from every cluster is not sent at the same time, which could affect performance.

The following table describes the collection activities that Unified Manager performs:

Activity	Time interval	Description
Performance statistics poll	Every 5 minutes	Collects real-time performance data from each cluster.
Statistical analysis	Every 5 minutes	After every statistics poll, Unified Manager compares the collected data against user-defined, system-defined, and dynamic thresholds.  If any performance thresholds have been breached, Unified Manager generates events and sends email to specified users, if configured to do so.
Configuration poll	Every 15 minutes	Collects detailed inventory information from each cluster to identify all the storage objects (nodes, SVMs, volumes, and so on).
Summarization	Every hour	Summarizes the latest 12 five-minute performance data collections into hourly averages.  The hourly average values are used in some of the UI pages, and they are retained for 390 days.

Activity	Time interval	Description
Forecast analysis and data pruning	Every day after midnight	Analyzes cluster data to establish dynamic thresholds for volume latency and IOPS for the next 24 hours.  Deletes from the database any five-minute performance data older than 30 days.
Data pruning	Every day after 2 a.m.	Deletes from the database any events and dynamic thresholds older than 390 days.
Data pruning	Every day after 3:30 a.m.	Deletes from the database any one-hour performance data older than 390 days.

## What a data continuity collection cycle is

A data continuity collection cycle retrieves performance data outside of the real-time cluster performance collection cycle that runs, by default, every five minutes. Data continuity collections enable Unified Manager to fill in gaps of statistical data that occur when it was unable to collect real-time data.

Data continuity collection is supported only on clusters installed with ONTAP version 8.3.1 or later software.

Unified Manager performs data continuity collection polls of historical performance data when the following events occur:

- A cluster is initially added to Unified Manager.

Unified Manager gathers historical performance data for the previous 15 days. This enables you to view two weeks of historical performance information for a cluster a few hours after it is added.

Additionally, system-defined threshold events are reported for the previous period, if any exist.



15 days of historical volume statistics are not currently collected.

- The current performance data collection cycle does not finish on time.

If the real-time performance poll goes beyond the five-minute collection period, a data continuity collection cycle is initiated to gather that missing information. Without the data continuity collection, the next collection period is skipped.

- Unified Manager has been inaccessible for a period of time and then it comes back online, as in the following situations:
  - It was restarted.
  - It was shut down during a software upgrade or when creating a backup file.

- A network outage is repaired.
- A cluster has been inaccessible for a period of time and then it comes back online, as in the following situations:
  - A network outage is repaired.
  - A slow wide area network connection delayed the normal collection of performance data.

A data continuity collection cycle can collect a maximum of 24 hours of historical data. If Unified Manager is down for longer than 24 hours, a gap in performance data appears in the UI pages.

A data continuity collection cycle and a real-time data collection cycle cannot run at the same time. The data continuity collection cycle must finish before the real-time performance data collection is initiated. When the data continuity collection is required to collect more than one hour of historical data, then you see a banner message for that cluster at the top of the Performance dashboard.

## What the timestamp means in collected data and events

The timestamp that appears in collected health and performance data, or that appears as the detection time for an event, is based on the ONTAP cluster time, adjusted to the time zone set on the web browser.

It is highly recommended that you use a Network Time Protocol (NTP) server to synchronize the time on your Unified Manager servers, ONTAP clusters, and web browsers.



If you see timestamps that look incorrect for a particular cluster, you might want to check that the cluster time has been set correctly.

## 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.