



# **NetApp Clustered Data ONTAP data source**

## OnCommand Insight

NetApp  
October 24, 2024

This PDF was generated from <https://docs.netapp.com/us-en/oncommand-insight/config-admin/cdot-landing-page-terminology-storage.html> on October 24, 2024. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

|   |   |
|---|---|
| NetApp Clustered Data ONTAP data source ..... | 1 |
| Terminology .....                             | 1 |
| Requirements .....                            | 1 |
| Configuration .....                           | 1 |
| Advanced configuration .....                  | 2 |
| Clustered Data ONTAP Storage .....            | 2 |
| Clustered Data ONTAP Storage Pool .....       | 2 |
| Clustered Data ONTAP Storage Node .....       | 3 |

# NetApp Clustered Data ONTAP data source

This data source should be used for storage systems using Clustered Data ONTAP, and requires an administrator account used for read-only API calls.

## Terminology

OnCommand Insight acquires the following inventory information from the Clustered Data ONTAP data source. For each asset type acquired by Insight, the most common terminology used for this asset is shown. When viewing or troubleshooting this data source, keep the following terminology in mind:

| Vendor/Model Term | Insight Term    |
|-------------------|-----------------|
| Disk              | Disk            |
| Raid Group        | Disk Group      |
| Cluster           | Storage         |
| Node              | Storage Node    |
| Aggregate         | Storage Pool    |
| LUN               | Volume          |
| Volume            | Internal Volume |



These are common terminology mappings only and might not represent every case for this data source.

## Requirements

- Administrator account used for read-only API calls
- Target IP is the cluster management LIF
- Username (with read-only role name to ontapi application to the default Vserver) and password to log into NetApp cluster
- Port requirements: 80 or 443
- License requirements: FCP license and mapped/masked volumes required for discovery

## Configuration

| Field                | Description   |
|----------------------|---|
| NetApp Management IP | IP address or fully-qualified domain name of the NetApp cluster |

|           |                                  |
|-----------|----------------------------------|
| User Name | User name for the NetApp cluster |
| Password  | Password for the NetApp cluster  |

## Advanced configuration

| Field                           | Description  |
|---------------------------------|--|
| Inventory Poll Interval (min)   | Interval between inventory polls (default 20 minutes)    |
| Performance Poll Interval (sec) | Interval between performance polls (default 300 seconds) |

## Clustered Data ONTAP Storage

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage asset landing pages.

### Clustered Data ONTAP Storage Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage asset landing pages. Many of these terms apply to other data collectors as well.

- Model — A comma delimited list of the unique, discrete node model names within this cluster. If all the nodes in the clusters are the same model type, just one model name will appear.
- Vendor — same Vendor name you would see if you were configuring a new data source.
- Serial number — The array serial number. On cluster architecture storage systems like NetApp Clustered Data Ontap, this serial number may be less useful than the individual “Storage Nodes” serial numbers.
- IP — generally will be the IP(s) or hostname(s) as configured in the data source.
- Microcode version — firmware.
- Raw Capacity — base 2 summation of all the physical disks in the system, regardless of their role.
- Latency — a representation of what the host facing workloads are experiencing, across both reads and writes. Ideally, OCI is sourcing this value directly, but this is often not the case. In lieu of the array offering this up, OCI is generally performing an IOPs-weighted calculation derived from the individual internal volumes’ statistics.
- Throughput — aggregated from internal volumes.
- Management — this may contain a hyperlink for the management interface of the device. Created programmatically by the Insight data source as part of inventory reporting.

## Clustered Data ONTAP Storage Pool

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages.

## Clustered Data ONTAP Storage Pool Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- Storage — what storage array this pool lives on. Mandatory.
- Type — a descriptive value from a list of an enumerated list of possibilities. Most commonly will be “Aggregate” or “RAID Group”.
- Node — if this storage array’s architecture is such that pools belong to a specific storage node, its name will be seen here as a hyperlink to its own landing page.
- Uses Flash Pool — Yes/No value — does this SATA/SAS based pool have SSDs used for caching acceleration?
- Redundancy — RAID level or protection scheme. RAID\_DP is dual parity, RAID\_TP is triple parity.
- Capacity — the values here are the logical used, usable capacity and the logical total capacity, and the percentage used across these.
- Over-committed capacity — If by using efficiency technologies you have allocated a sum total of volume or internal volume capacities larger than the logical capacity of the storage pool, the percentage value here will be greater than 0%.
- Snapshot — snapshot capacities used and total, if your storage pool architecture dedicates part of its capacity to segments areas exclusively for snapshots. Ontap in MetroCluster configurations are likely to exhibit this, while other Ontap configurations are less so.
- Utilization — a percentage value showing the highest disk busy percentage of any disk contributing capacity to this storage pool. Disk utilization does not necessarily have a strong correlation with array performance — utilization may be high due to disk rebuilds, deduplication activities, etc in the absence of host driven workloads. Also, many arrays’ replication implementations may drive disk utilization while not showing as internal volume or volume workload.
- IOPS — the sum IOPs of all the disks contributing capacity to this storage pool.
- Throughput — the sum throughput of all the disks contributing capacity to this storage pool.

## Clustered Data ONTAP Storage Node

Terms applying to objects or references that you might find on NetApp Clustered Data ONTAP storage node asset landing pages.

## Clustered Data ONTAP Storage Node Terminology

The following terms apply to objects or references that you might find on NetApp Clustered Data ONTAP storage pool asset landing pages. Many of these terms apply to other data collectors as well.

- Storage — what storage array this node is part of. Mandatory.
- HA Partner — on platforms where a node will fail over to one and only one other node, it will generally be seen here.
- State — health of the node. Only available when the array is healthy enough to be inventoried by a data source.
- Model — model name of the node.
- Version — version name of the device.

- Serial number — The node serial number.
- Memory — base 2 memory if available.
- Utilization — On Ontap, this is a controller stress index from a proprietary algorithm. With every performance poll, a number between 0 and 100% will be reported that is the higher of either WAFL disk contention, or average CPU utilization. If you observe sustained values > 50%, that is indicative of undersizing — potentially a controller/node not large enough or not enough spinning disks to absorb the write workload.
- IOPS — Derived directly from Ontap ZAPI calls on the node object.
- Latency — Derived directly from Ontap ZAPI calls on the node object.
- Throughput — Derived directly from Ontap ZAPI calls on the node object.
- Processors — CPU count.

## 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—with 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.