



Working with Queries

Data Infrastructure Insights

NetApp

February 03, 2026

This PDF was generated from https://docs.netapp.com/us-en/data-infrastructure-insights/concept_querying_assets.html on February 03, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Working with Queries	1
Querying Assets and Metrics	1
Inventory Assets	1
Integration Metrics	1
Sharing Queries	1
Creating Queries	2
See it in Action	3
Choosing Aggregation, Units, Conditional Formatting	4
More on Filtering	5
What do I do now that I have query results?	7
Annotation Rules require query	7
Viewing queries	7
Exporting query results to a .CSV file	8
Asynchronous export	8
Commas in asset names	9
Time format, or not time format?	9
Modifying or Deleting a Query	9
Modifying a Query	9
Deleting a Query	10
Assigning multiple applications to or removing multiple applications from assets	10
Copying table values	11
Log Explorer	11
Filtering	12
Log Details	15
Troubleshooting	16
Identifying inactive devices	16

Working with Queries

Querying Assets and Metrics

Query your infrastructure's physical and virtual assets to monitor performance, troubleshoot issues, and perform granular searches based on custom criteria like annotations. Data Infrastructure Insights enables queries across diverse asset types—from storage arrays and hosts to applications and virtual machines—along with integration metrics from Kubernetes, Docker, and ONTAP Advanced Data for comprehensive visibility.

Note that annotation rules, which automatically assign annotations to assets, require a query that has been shared with everyone. See below for more on sharing queries.


You can query the physical or virtual inventory assets (and their associated metrics) on your tenant, or the metrics provided with integration such as Kubernetes or ONTAP Advanced Data.

Inventory Assets

All inventory (also called infrastructure) asset types (storage, switch, VM, Application, etc.) can be used in queries, dashboard widgets, and customized asset landing pages. The fields and counters available for filters, expressions, and display will vary among asset types.

Integration Metrics

In addition to querying for inventory assets and their associated performance metrics, you can query for **integration data** metrics as well, such as those generated by Kubernetes or Docker, or provided with ONTAP Advanced Metrics.



Sharing Queries

Control access to your queries by choosing who can view and edit them. By default, new queries are private and visible only to you, and you can choose to share them with specific users or your entire organization with flexible permission levels (Read-Only or Edit).

You can choose to share the query to Everyone or to select users, with either Edit or Read-Only permissions.



Users with Account Owner permission can see all queries, regardless of privacy setting.

Creating Queries

Queries enable you to search the assets on your tenant at a granular level, allowing to filter for the data you want and sort the results to your liking.


For example, you can create a query for *volumes*, add a filter to find particular *storages* associated with the selected volumes, add another filter to find a particular *annotation* such as "Tier 1" on the selected storages, and finally add another filter to find all storages with *IOPS - Read (IO/s)* greater than 25. When the results are displayed, you can then sort the columns of information associated with the query in ascending or descending order.

Note: When a new data collector is added which acquires assets, or any annotation or application assignments are made, you can query for those new assets, annotations, or applications only after the queries are indexed. Indexing occurs at a regularly scheduled interval or during certain events such as running annotation rules.

Creating a Query is very simple:

1. Navigate to **Queries > *+New Query**.
2. From the 'Select...' list, select the object type you want to query for. You can scroll through the list or you can start typing to more quickly find what you're searching for.

Scroll list:



awesome site / All Queries / **New Query**

Select...

- agent.node
- agent.node_diskio
- agent.node_fs
- agent.node_net
- Application
- DataStore
- Disk
- Fabric
- GenericDevice

Type-to-Search:

You can add filters to further narrow down your query by clicking the **+** button in the **Filter By** field. Group rows by object or attribute. When working with integration data (Kubernetes, ONTAP Advanced Metrics, etc.), you can group by multiple attributes, if desired.

5 items found

Table Row Grouping		Metrics & Attributes	
aggr_name		cp_read_blocks	cluster_name ↓
oci02sat0		0.59	oci-phonehome
oci02sat1		0.15	oci-phonehome
oci02sat2		212.64	oci-phonehome
oci01sat0		0.39	oci-phonehome
oci01sat1		48.89	oci-phonehome

The query results list shows a number of default columns, depending on the object type searched for. To add, remove, or change the columns, click the gear icon on the right of the table. The available columns vary based on the asset/metric type.

14 items found

Table Row Grouping		Metrics & Attributes	
aggr_name		cp_read_blocks	agent_version ↑
agg0_optimus_02		1.72	Apache-HttpClient/4.5.13 (Java/1.8.0_292)
agg1_optimus_02		408.84	Apache-HttpClient/4.5.13 (Java/1.8.0_292)
ocinaneqa1_04_aggr0		6.19	Apache-HttpClient/4.5.13 (Java/1.8.0_292)
ocinaneqa1_03_aggr0		6.48	Apache-HttpClient/4.5.13 (Java/1.8.0_292)
oci02sat0		1.04	Apache-HttpClient/4.5.13 (Java/1.8.0_292)

Search...

Show Selected Only

agent_version

aggr_name

cluster_location

cluster_name

cluster_serial_number

cluster_version

See it in Action

Explore and analyze with queries in Data Infrastructure Insights (Video)

Choosing Aggregation, Units, Conditional Formatting

Aggregation and Units

For "value" columns, you can further refine your query results by choosing how the displayed values are aggregated as well as selecting the units in which those values are displayed. These options are found by selecting the "three dots" menu at the top corner of a column.

Units

You can select the units in which to display the values. For example, if the selected column shows raw capacity and the values are shown in GiB, but you prefer to display them as TiB, simply select TiB from the Unit Display drop-down.

Aggregation

By the same token, if the values shown are aggregated from the underlying data as "Average", but you prefer to show the sum of all values, select "Sum" from either the *Group by* drop-down (if you want any grouped values to show the sums) or from the *Time Aggregate By* drop-down (if you want the row values to show sums of underlying data).

You can choose to aggregate grouped data points by *Avg, Max, Min, or Sum*.

You can aggregate individual row data by *Average, Last data point acquired, Maximum, Minimum, or Sum*.

Conditional Formatting

Conditional Formatting allows you to highlight Warning-level and Critical-level thresholds in the query results list, bringing instant visibility to outliers and exceptional data points.

143 items found	
Table Row Grouping	Metrics & Attributes
agent.node_diskio ↑	io_time (sec)
nvme0n1	20,604.96
nvme0n1	29,184.97
nvme0n1	4,642.68
nvme0n1	31,918.99
nvme0n1	29,258.26
nvme0n1	18,022.16
nvme0n1	28,483.30
nvme0n1	69,835.02
nvme0n1	15,952.78

- ⋮
- > Aggregation
- > Unit Display
- ✓ Conditional Formatting Reset
- If value is > (Greater than) ▾
- ⚠ Warning 10000 sec
- ❗ Critical 20000 sec
- > Rename Column

Conditional formatting is set separately for each column. For example, you can choose one set of thresholds for a capacity column, and another set for a throughput column.

Rename Column

Renaming a column changes the displayed name on the Query results list. The new column name is also

shown in the resulting file if you export the query list to .CSV.


Save

After you have configured your query to show you the results you want, you can click the **Save** button to save the query for future use. Give it a meaningful and unique name.

More on Filtering

Wildcards and Expressions

When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create **expressions** using NOT or OR, or you can select the "None" option to filter for null values in the field.



kubernetes.pod X ▾

Filter By pod_name ingest ▼ X + ?

Group pod_name X

Create wildcard containing "ingest"

ci-service-datalake-ingestion-85b5bdfd6d-2qbwr

service-foundation-ingest-767dfd5bfc-vxd5p

None

71 items found

Table Row Grouping

Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.

kubernetes.pod X ▼

Filter By pod_name *ingest* X ci-service-audit-5f775dd975-brfdc X X ▼ X ▼ + ?

Group pod_name X X ▼

3 items found

Table Row Grouping

pod_name
ci-service-audit-5f775dd975-brfdc
ci-service-datalake-ingestion-85b5bdfd6d-2qbwr
service-foundation-ingest-767dfd5bfc-vxd5p

Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans.

Refining Filters

You can use the following to refine your filter:

Filter	What it does	Example	Result
* (Asterisk)	enables you to search for everything	vol*rhel	returns all resources that start with "vol" and end with "rhel"
? (question mark)	enables you to search for a specific number of characters	BOS-PRD??-S12	returns BOS-PRD12-S12, BOS-PRD23-S12, and so on
OR	enables you to specify multiple entities	FAS2240 OR CX600 OR FAS3270	returns any of FAS2440, CX600, or FAS3270
NOT	allows you to exclude text from the search results	NOT EMC*	returns everything that does not start with "EMC"
<i>None</i>	searches for NULL values in all fields	<i>None</i>	returns results where the target field is empty
Not *	searches for NULL values in <i>text-only</i> fields	Not *	returns results where the target field is empty

If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example, filtering for `"*"` will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this

case. The operators OR and NOT will also be treated as literal strings when enclosed in double quotes.

Filtering for Boolean values

When filtering for a Boolean value, you may be presented with the following choices on which to filter:

- **Any:** This will return *all* results, including results set to "Yes", "No", or not set at all.
- **Yes:** Returns only "Yes" results. Note that DII shows "Yes" as a check mark in most tables. Values may be set to "True", "On", etc.; DII treats all of these as "Yes".
- **No:** Returns only "No" results. Note that DII shows "No" as an "X" in most tables. Values may be set to "False", "Off", etc.; DII treats all of these as "No".
- **None:** Returns only results where the value has not been set at all. Also referred to as "Null" values.

What do I do now that I have query results?

Querying provides a simple place to add annotations or assign applications to assets. Note that you can only assign applications or annotations to your inventory assets (Disk, Storage, etc.). Integration metrics cannot take on annotation or application assignments.

To assign an annotation or application to the assets resulting from your query, simply select the asset(s) using the check box column on the left of the results table, then click the **Bulk Actions** button on the right. Choose the desired action to apply to the selected assets.

The screenshot shows the DII interface with a query results table. At the top, there is a filter bar with 'Volume' selected, a 'Filter By' dropdown set to 'Name', and a 'Any' button. The table has columns: Name, Storage Pools, Capacity - Raw (GB), and Mapped Ports. There are five rows of data. The second and third rows have the 'Name' column checked, indicating they are selected. To the right of the table is a 'Bulk Actions' dropdown menu with options: Add Annotation, Remove Annotation, Add Application, Remove Application, and a dropdown menu item 'OS:windows_2008'. The table data is as follows:

Name	Storage Pools	Capacity - Raw (GB)	Mapped Ports
DmoESX_optimus:mc_Dm...	optimus-02:aggr1_optimu...	N/A	
<input checked="" type="checkbox"/> DmoSAN_optimus:hoffma...	optimus-02:aggr1_optimu...	N/A	
<input checked="" type="checkbox"/> DmoSAN_optimus:mc_D...	optimus-02:aggr1_optimu...	N/A	
oci-3070-01:/vol/vfiler_lun...	oci-3070-01:aggr5	N/A	OS:windows
spectrav1:sjimmylscsi/v...	ocinaneqa1-01:spectraaggr1	N/A	OS:linux

Annotation Rules require query

If you are configuring [Annotation Rules](#), each rule must have an underlying query to work with. But as you've seen above, queries can be made as broad or as narrow as you need.

Viewing queries

You can view your queries to monitor your assets and change how your queries display the data related to your assets.

Steps

1. Log in to your Data Infrastructure Insights tenant.

2. Click **Queries** and select **Show all queries**.

You can change how queries display by doing any of the following:

3. You can enter text in the filter box to search to display specific queries.

4. You can change the sort order of the columns in the table of queries to either ascending (up arrow) or descending (down arrow) by clicking the arrow in the column header.

5. To resize a column, hover the mouse over the column header until a blue bar appears. Place the mouse over the bar and drag it right or left.

6. To move a column, click on the column header and drag it right or left.

When scrolling through the query results, be aware that the results may change as Data Infrastructure Insights automatically polls your data collectors. This may result in some items being missing, or some items appearing out of order depending on how they are sorted.

Exporting query results to a .CSV file

You can export the results of any query to a .CSV file, which will allow you to analyze the data or import it into another application.

Steps

1. Log in to Data Infrastructure Insights.

2. Click **Queries** and select **Show all queries**.

The Queries page is displayed.

3. Click a query.

4. Click  to export the query results to a .CSV file.




Export to .CSV is also available in the "three dots" menu in dashboard table widgets as well as most landing page tables.

Asynchronous export

Exporting data to .CSV can take anywhere from a few seconds to a number of hours, depending on the amount of data to export. Data Infrastructure Insights exports that data asynchronously, so you can continue working while the .CSV is being compiled.

View and download your .CSV exports by selecting the "Bell" icon in the upper-right toolbar.



The screenshot shows the Data Infrastructure Insights interface. At the top, there is a blue header bar with a search icon, tenant name "Standard QA Main", and user "Tony Lavoie". Below the header, a notification box is displayed. The notification text is: "query_export_disk_1748360447802.csv is Ready to Download". It includes details: "Size: 13.29 KiB", "Initiated: 05/27/2025 11:40:47 AM", and "Ready: 05/27/2025 11:40:48 AM". There is a "Download" button and a timestamp "Just now".

The exported data will reflect the current filtering, columns, and column names displayed.

Commas in asset names

Note: When a comma appears in an asset name, the export encloses the name in quotes, preserving the asset name and the proper .csv format.

Time format, or not time format?

When opening an exported .CSV file with Excel, if you have an object name or other field that is in the format NN:NN (two digits followed by a colon followed by two more digits), Excel will sometimes interpret that name as a Time format, instead of Text format. This can result in Excel displaying incorrect values in those columns. For example, an object named "81:45" would show in Excel as "81:45:00".

To work around this, import the .CSV into Excel using the following steps:

1. Open a new sheet in Excel.
2. On the "Data" tab, choose "From Text".
3. Locate the desired .CSV file and click "Import".
4. In the Import wizard, choose "Delimited" and click Next.
5. Choose "Comma" for the delimiter and click Next.
6. Select the desired columns and choose "Text" for the column data format.
7. Click Finish.

Your objects should show in Excel in the proper format.

Modifying or Deleting a Query

You can change the criteria that are associated with a query when you want to change the search criteria for the assets that you are querying.

Modifying a Query

Steps

1. Click **Explore** and select **All Metric Queries**.

The Queries page is displayed.

2. Click the query name
3. To add a criteria to the query, click Columns icon and select a metric or attribute from the list.

When you have made all necessary changes, do one of the following:

- Click the **Save** button to save the query with the name that was used initially.
- Click the drop-down next to the **Save** button and select **Save As** to save the query with another name. This does not overwrite the original query.
- Click the drop-down next to the **Save** button and select **Rename** to change the query name that you had used initially. This overwrites the original query.
- Click the drop-down next to the **Save** button and select **Discard Changes** to revert the query back to the last saved changes.

Deleting a Query

To delete a query, click **Queries** and select **Show all queries**, and do one of the following:

1. Click on the "three dot" menu to the right of the query and click **Delete**.
2. Click on the query name and select **Delete** from the **Save** drop-down menu.

Assigning multiple applications to or removing multiple applications from assets

You can assign multiple [applications](#) to or remove multiple applications from assets by using a query instead of having to manually assign or remove them.



You can use these steps to add or remove [annotations](#) in the same way.

Before you begin

You must have already created a query that finds all the assets that you want to edit.

Steps

1. Click **Explore** and select **Metric Queries**.

The Queries page displays.

2. Click the name of the query that finds the assets.

The list of assets associated with the query displays.

3. Select the desired assets in the list or click the top checkbox to select All.

The Bulk Actions drop-down displays.

4. To add an application to the selected assets, click Bulk Actions and select **Add Application**.

5. Select one or more applications.

You can select multiple applications for hosts, internal volumes, qtrees, and virtual machines; however, you can select only one application for a volume or a share.

6. Click **Save**.

7. To remove an application assigned to the assets, click Bulk Actions and select **Remove Application**.

8. Select the application or applications you want to remove.

9. Click **Delete**.

Any new applications you assign override any applications on the asset that were derived from another asset. For example, volumes inherit applications from hosts, and when new applications are assigned to a volume, the new application takes precedence over the derived application.

After you click **Save** on a bulk add or **Remove** on a bulk delete action, Data Infrastructure Insights informs you that the action will take some time. You can dismiss this message; the action will continue in the background.



For environments with large amounts of related assets, inheritance of application assignments to those assets could take several minutes. Please allow more time for inheritance to occur if you have many related assets.

Copying table values

You can copy values in tables to the clipboard for use in search boxes or other applications.

About this task

There are two methods you can use to copy values from tables or query results to the clipboard.

Steps

1. Method 1: Highlight the desired text with the mouse, copy it, and paste it into search fields or other applications.
2. Method 2: For single-value fields, hover over the field and click the clipboard icon that appears. The value is copied to the clipboard for use in search fields or other applications.


Note that only values that are links to assets can be copied using this method. Only fields that include single values (i.e. non-lists) have the copy icon.

Log Explorer

The Data Infrastructure Insights Log Explorer is a powerful tool for querying system logs. In addition to helping with investigations, you can also save a log query in a Monitor to provide alerts when those particular log triggers are activated.

To begin exploring logs, click **Log Queries > +New Log Query**.

Select an available log from the list.




The types of logs available for querying may vary based on your environment. Additional log types may be added over time.

You can set filters to further refine the results of the query. For example, to find all log messages showing a failure, set a filter for *Messages* containing the word "failed".



You can begin typing the desired text in the filter field; Data Infrastructure Insights will prompt you to create a wildcard search containing the string as you type.


The results are displayed in a graph showing the number of log instances in each time period shown. Below the graph are the log entries themselves. The graph and the entries refresh automatically based on the selected time range.



Filtering

Include / Exclude

When filtering the logs, you can choose to **include** (i.e. "Filter to") or **exclude** the strings you type. Excluded strings are displayed in the completed filter as "NOT <string>".



Filters based on wildcards or expressions (e.g. NOT, OR, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.



At any point, you can click on *Create a Log Monitor* to create a new Monitor based on the current filter.


Advanced Filtering

When you are filtering for text or list values in queries or dashboard widgets, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create expressions using NOT, AND, or OR, or you can select the "None" option to filter for null values.



Be sure to Save your query early and often as you build your filtering. Advanced Querying is "free-form" string entry, and parsing mistakes may occur as you build.

Take a look at this screen image showing filtered results for an advanced query of the `logs.kubernetes.event` log. There is a lot going on in this page, which is explained below the image:



Log Entries					Last updated 08/30/2023 9:54:13 AM
timestamp	source	message	metadata.namespace	reason	
08/26/2023 8:40:28 AM	kubernetes_cluster:eg-stream;namespace:33994-monitoring;pod_name:event-exporter-5db67db995-bxmkf;	Error: context deadline exceeded	k3s-cm-monitoring	Failed	
08/26/2023 8:40:28 AM	kubernetes_cluster:eg-stream;namespace:ph-monitoring;pod_name:event-exporter-c4446976c-jxrdc;	Error: context deadline exceeded	k3s-cm-monitoring	Failed	
08/26/2023 8:40:29 AM	kubernetes_cluster:eg-stream;namespace:33994-monitoring;pod_name:event-exporter-5db67db995-bxmkf;	Error: failed to reserve	k3s-cm-monitoring	Failed	



1. This advanced query string filters for the following:

- Filter for log entries with a *reason* that includes the word "failed", but not anything with the specific reason of "FailedMount".
- Include any of those entries that also include a *metadata.namespace* including the word "monitoring", but exclude the specific namespaces of "cm-monitoring" or "eg-monitoring".

Note that in the case above, since both "cm-monitoring" and "eg-monitoring" contain a dash (" - "), the

strings must be included in double-quotes or a parsing error will be displayed. Strings that do not include dashes, spaces, etc. do not need to be enclosed in quotes. If in doubt, try putting the string in quotes.

2. The results of the current filter, including any "Filter By" values AND the Advanced Query filter, are displayed in the results list. The list can be sorted by any displayed columns. To display additional columns, select the "gear" icon.
3. The graph has been zoomed in to show only log results that occurred within a specific time frame. The time range shown here reflects the current zoom level. Select the *Reset Zoom* button to set the zoom level back to the current Data Infrastructure Insights time range.
4. The chart results have been Grouped By the *source* field. The chart shows results in each column grouped into colors. Hovering over a column in the chart will display some details about the specific entries.



Refining Filters

You can use the following to refine your filter:

Filter	What it does
* (Asterisk)	enables you to search for everything
? (question mark)	enables you to search for a specific number of characters
OR	enables you to specify multiple entities
NOT	allows you to exclude text from the search results
None	searches for NULL values in all fields
Not *	searches for NULL values in <i>text-only</i> fields


If you enclose a filter string in double quotes, Insight treats everything between the first and last quote as an exact match. Any special characters or operators inside the quotes will be treated as literals. For example,

filtering for "*" will return results that are a literal asterisk; the asterisk will not be treated as a wildcard in this case. The operators OR and NOT will also be treated as literal strings when enclosed in double quotes.

You can combine a simple filter with an advanced query filter; the resulting filter is an "AND" of the two.

The Chart Legend

The *Legend* below the chart has a few surprises as well. For each result (based on the current filter) shown in the Legend, you have an option to display only results for that line (Add Filter), or to display any results NOT for that line (Add Exclude Filter). The chart and the Log Entries list update to show results based on your selection. To remove this filtering, open the Legend again and select the [X] to clear the Legend-based filter.



Log Details

Clicking anywhere in a log entry in the list will open a detail pane for that entry. Here you can explore more information about the event.

Click on "Add Filter" to add the selected field to the current filter. The log entry list will update based on the new filter.

Note that some fields cannot be added as filters; in those cases, the *Add Filter* icon is not available.

Log Details



timestamp

09/20/2021 9:03:36 PM

message

2021-09-20T15:33:36Z E! [processors.execd] stderr: "Total time to process mountstats file: /hostfs/proc/1/mountstats, was: 0s"

id: 227814532095936770

node_name: ci-auto-dsacq-insights-1.cloudinsights-dev.netapp.com

Add Filter



source: telegraf-ds-dfcc5

type: logs.kubernetes

kubernetes

kubernetes.annotations.openshift.io_scc:

kubernetes.container_hash: ci-registry.name.openenglab.netapp.com:8022/telegraf@sha256:00h45a7cc0761c

Troubleshooting

Here you will find suggestions for troubleshooting problems with Log Queries.

Problem:	Try this:
I don't see "debug" messages in my log query	Debug log messaging is not collected. To capture messages you want, change the relevant message severity to <i>informational</i> , <i>error</i> , <i>alert</i> , <i>emergency</i> , or <i>notice</i> level.

Identifying inactive devices

Identifying the assets you have and who's using them is critical to "right-sizing" and freeing up unused infrastructure. You can easily reallocate or decommission underused resources and avoid unnecessary purchases.

Use the following steps to identify inactive assets.

Steps

- Navigate to **Observability > Explore → +New Metric Query**.

- Select *Storage* from the drop down.
- Click the gear and add *isActive* as a column.

Rows showing a Check are active. "X" indicates inactive devices.

To remove inactive devices, simply select the devices to remove and in the **Bulk Actions** drop-down, select *Delete Inactive Devices*.

Copyright information

Copyright © 2026 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.