



## **View performance data**

SANtricity 11.8

NetApp  
April 05, 2024

This PDF was generated from <https://docs.netapp.com/us-en/e-series-santricity/sm-storage/overview-performance.html> on April 05, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- View performance data ..... 1
  - Performance overview ..... 1
  - Performance graphs and guidelines ..... 1
  - Performance terminology ..... 2
  - View graphical performance data ..... 4
  - View and save tabular performance data ..... 6
  - Interpret performance data ..... 7

# View performance data

## Performance overview

The Performance page provides easy ways for you to monitor the performance of your storage array.

### What can I learn from performance data?

The Performance graphs and tables show performance data in near real-time, which helps you determine whether a storage array is experiencing problems. You can also save performance data to construct a historical view of a storage array and identify when a problem started or what caused a problem.

Learn more:

- [Performance graphs and guidelines](#)
- [Performance terms](#)

### How can I view performance data?

Performance data is available from the Home page and from the Storage page.

Learn more:

- [View graphical performance data](#)
- [View and save tabular performance data](#)
- [Interpret performance data](#)

## Performance graphs and guidelines

The Performance page provides graphs and tables of data that enable you to assess the storage array's performance in several key areas.

Performance functions allow you to accomplish these tasks:

- View performance data in near real-time to help you determine whether a storage array is experiencing problems.
- Export performance data to construct a historical view of a storage array and identify when a problem started or what caused a problem.
- Select the objects, performance metrics, and time frame you want to view.
- Compare metrics.

You can view performance data in three formats:

- **Real-time graphical** — Plots performance data on a graph in near real-time.
- **Near real-time tabular** — Lists performance data in a table in near real-time.
- **Exported CSV file** — Allows you to save tabular performance data in a file of comma-separated values for further viewing and analysis.

## Characteristics of performance data formats

Type of performance monitoring	Sampling interval	Length of time displayed	Maximum number of objects displayed	Ability to save data
Real-time graphical, live  Real-time graphical, historical	10 sec (live)  5 min (historical)  Data points shown depend on selected time frame	Default time frame is 1 hour.  Choices: <ul style="list-style-type: none"><li>• 5 minutes</li><li>• 1 hour</li><li>• 8 hours</li><li>• 1 day</li><li>• 7 days</li><li>• 30 days</li></ul>	5	No
Near real-time tabular (table view)	10 sec -1 hr	Most current value	Unlimited	Yes
Comma-separated values (CSV) file	Depends on selected time frame	Depends on selected time frame	Unlimited	Yes

## Guidelines for viewing performance data

- Performance data collection is always on. There is no option to turn it off.
- Each time the sampling interval elapses, the storage array is queried and the data is updated.
- For graphical data, the 5-minute time frame supports 10-second updating averaged over 5 minutes. All other time frames are updated every 5 minutes, averaged over the selected time frame.
- Performance data in the graphical views is updated in real time. Performance data in the table view is updated in near real time.
- If a monitored object changes during the time data is collected, the object might not have a complete set of data points spanning the selected time frame. For example, volume sets can change as volumes are created, deleted, assigned, or unassigned; or drives can be added, removed, or failed.

## Performance terminology

Learn how the performance terms apply to your storage array.

Term	Description
Application	An application is a software program, such as SQL or Exchange.
CPU	CPU is short for "central processing unit." CPU indicates the percentage of the storage array's processing capacity being used.

Term	Description
Host	A host is a server that sends I/O to a volume on a storage array.
IOPS	IOPS stands for input/output operations per second.
Latency	Latency is the time interval between a request, such as for a read or write command, and the response from the host or the storage array.
LUN	<p>A logical unit number (LUN) is the number assigned to the address space that a host uses to access a volume. The volume is presented to the host as capacity in the form of a LUN.</p> <p>Each host has its own LUN address space. Therefore, the same LUN can be used by different hosts to access different volumes.</p>
MiB	MiB is an abbreviation for mebibyte (mega binary byte). One MiB is 2 <sup>20</sup> , or 1,048,576 bytes. Compare with MB, which signifies a base 10 value. One MB equals 1,024 bytes.
Object	<p>An object is any logical or physical storage component.</p> <p>Logical objects include volume groups, pools, and volumes. Physical objects include the storage array, array controllers, hosts, and drives.</p>
Pool	A pool is a set of drives that is logically grouped. You can use a pool to create one or more volumes accessible to a host. (You create volumes from either a pool or a volume group.)
Read	Read is short for "read operation," which occurs when the host requests data from the storage array.
Volume	<p>A volume is a container in which applications, databases, and file systems store data. It is the logical component created for the host to access storage on the storage array.</p> <p>A volume is created from the capacity available in a pool or a volume group. A volume has a defined capacity. Although a volume might consist of more than one drive, a volume appears as one logical component to the host.</p>
Volume name	A volume name is a string of characters assigned to the volume when it is created. You can either accept the default name or provide a more descriptive name indicating the type of data stored in the volume.
Volume group	A volume group is a container for volumes with shared characteristics. A volume group has a defined capacity and RAID level. You can use a volume group to create one or more volumes accessible to a host. (You create volumes from either a volume group or a pool.)

Term	Description
Workload	A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, the system configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.
Write	Write is short for "write operation," when data is sent from the host to the array for storage.

## View graphical performance data

You can view graphical performance data for logical objects, physical objects, applications, and workloads.

### About this task

The performance graphs show historical data as well as live data currently being captured. A vertical line on the graph, labeled Live updating, distinguishes historical data from live data.

### Home page view

The Home page contains a graph showing storage array level performance. You can select limited metrics from this view, or you can click **View Performance Details** to select all the available metrics.

### Detailed view

The graphs available from the detailed performance view are arranged under three tabs:

- **Logical View** — Displays performance data for logical objects grouped by volume groups and pools. Logical objects include volume groups, pools, and volumes.
- **Physical View** — Displays performance data for the controller, host channels, drive channels, and drives.
- **Applications & Workloads View** — Displays a list of logical objects (volumes) grouped by the application types and workloads you have defined.

### Steps

1. Select **Home**.
2. To select an array-level view, click the IOPS, MiB/s, or CPU button.
3. To see more details, click **View Performance Details**.
4. Select **Logical View** tab, **Physical View** tab, or **Applications & Workloads View** tab.

Depending on the object type, different graphs appear in each tab.

View tabs	Performance data displayed for each object type
Logical View	<ul style="list-style-type: none"> <li>• <b>Storage array:</b> IOPS, MiB/s</li> <li>• <b>Pools:</b> Latency, IOPS, MiB/s</li> <li>• <b>Volume groups:</b> Latency, IOPS, MiB/s</li> <li>• <b>Volumes:</b> Latency, IOPS, MiB/s</li> </ul>
Physical View	<ul style="list-style-type: none"> <li>• <b>Controllers:</b> IOPS, MiB/s, CPU, Headroom</li> <li>• <b>Host channels:</b> Latency, IOPS, MiB/s, Headroom</li> <li>• <b>Drive channels:</b> Latency, IOPS, MiB/s</li> <li>• <b>Drives:</b> Latency, IOPS, MiB/s</li> </ul>
Applications & Workloads View	<ul style="list-style-type: none"> <li>• <b>Storage array:</b> IOPS, MiB/s</li> <li>• <b>Applications:</b> Latency, IOPS, MiB/s</li> <li>• <b>Workloads:</b> Latency, IOPS, MiB/s</li> <li>• <b>Volumes:</b> Latency, IOPS, MiB/s</li> </ul>


5. Use the options to view the objects and information you need.

### Options

Options for viewing objects	Description
Expand a drawer to see the list of objects.	<p><i>Navigation drawers</i> contain storage objects, such as pools, volume groups, and drives.</p> <p>Click the drawer to view the list of objects in the drawer.</p>
Select objects to view.	Select the check box to the left of each object to choose the performance data you want to view.
Use Filter to find object names or partial names.	In the Filter box, enter the name or a partial name of objects to list just those objects in the drawer.
Click <b>Refresh Graphs</b> after selecting objects.	After selecting objects from the drawers, select <b>Refresh Graphs</b> to view graphical data for the items you have selected.
Hide or show graph	Select the graph title to hide or show the graph.

6. As needed, use the additional options for viewing performance data.

## Additional options

Option	Description
Time frame	<p>Select the length of time you want to view (5 minutes, 1 hour, 8 hours, 1 day, 7 days, or 30 days). The default is 1 hour.</p> <div><p>Loading performance data for a 30-day time frame can take several minutes. Do not navigate away from the web page, refresh the web page, or close the browser while data is loading.</p></div>
Data point details	Hover the cursor over the graph to see metrics for a particular data point.
Scroll bar	Use the scroll bar below the graph to view an earlier or later time span.
Zoom bar	<p>Below the graph, drag the zoom bar handles to zoom out on a time span. The wider the zoom bar, the less granular the details of the graph.</p> <p>To reset the graph, select one of the time frame options.</p>
Drag and drop	<p>On the graph, drag the cursor from one point in time to another to zoom in on a time span.</p> <p>To reset the graph, select one of the time frame options.</p>

## View and save tabular performance data

You can view and save performance graphs data in tabular format. This allows you to filter the data you want displayed.

### Steps

1. From any performance data graph, click **Launch table view**.

A table appears that lists all the performance data for the selected objects.

2. Use the object selection pull-down and the filter as needed.
3. Click the **Show/Hide Columns** button to select the columns you want to include in the table.

You can click each check box to select or deselect an item.

4. Select **Export** at the bottom of the screen to save the tabular view to a file of comma-separated values (CSV).

The Export Table dialog box appears, indicating the number of rows to be exported and the file format of the export (comma-separated values, or CSV format).

5. Click **Export** to proceed with the download, or click **Cancel**.



Depending on your browser settings, the file is either saved, or you are prompted to choose a name and location for the file.

The default file name format is `performanceStatistics-yyyy-mm-dd_hh-mm-ss.csv`, which includes the date and time when the file was exported.

## Interpret performance data

Performance data can guide you in tuning the performance of your storage array.

When interpreting Performance data, keep in mind that several factors affect the performance of your storage array. The following table describes the main areas to consider.

Performance data	Implications for performance tuning
Latency (milliseconds, or ms)	<p>Monitor the I/O activity of a specific object.</p> <p>Potentially identify objects that are bottlenecks:</p> <ul style="list-style-type: none"><li>• If a volume group is shared among several volumes, the individual volumes might need their own volume groups to improve the sequential performance of the drives and decrease latency.</li><li>• With pools, larger latencies are introduced and uneven workloads might exist between drives, making the latency values less meaningful and, in general, higher.</li><li>• Drive type and speed influence latency. With random I/O, faster spinning drives spend less time moving to and from different locations on the disk.</li><li>• Too few drives result in more queued commands and a greater period of time for the drive to process the command, increasing the general latency of the system.</li><li>• Larger I/Os have greater latency due to the additional time involved with transferring data.</li><li>• Higher latency might indicate that the I/O pattern is random in nature. Drives with random I/O will have greater latency than those with sequential streams.</li><li>• A disparity in latency among drives or volumes of a common volume group could indicate a slow drive.</li></ul>

Performance data	Implications for performance tuning
IOPS	<p>Factors that affect input/output operations per second (IOPS or IOs/sec) include these items:</p> <ul style="list-style-type: none"> <li>• Access pattern (random or sequential)</li> <li>• I/O size</li> <li>• RAID level</li> <li>• Cache block size</li> <li>• Whether read caching is enabled</li> <li>• Whether write caching is enabled</li> <li>• Dynamic cache read prefetch</li> <li>• Segment size</li> <li>• The number of drives in the volume groups or storage array</li> </ul> <p>The higher the cache hit rate, the higher I/O rates will be. Higher write I/O rates are experienced with write caching enabled compared to disabled. In deciding whether to enable write caching for an individual volume, look at the current IOPS and the maximum IOPS. You should see higher rates for sequential I/O patterns than for random I/O patterns. Regardless of your I/O pattern, enable write caching to maximize the I/O rate and to shorten the application response time.</p> <p>You can see performance improvements caused by changing the segment size in the IOPS statistics for a volume. Experiment to determine the optimal segment size, or use the file system size or database block size.</p>
MiB/s	<p>Transfer or throughput rates are determined by the application I/O size and the I/O rate. Generally, small application I/O requests result in a lower transfer rate but provide a faster I/O rate and shorter response time. With larger application I/O requests, higher throughput rates are possible.</p> <p>Understanding your typical application I/O patterns can help you determine the maximum I/O transfer rates for a specific storage array.</p>
CPU	<p>This value is a percentage of processing capacity that is being used.</p> <p>You might notice a disparity in the CPU usage of the same types of objects. For example, the CPU usage of one controller is heavy or is increasing over time while that of the other controller is lighter or more stable. In this case, you might want to change the controller ownership of one or more volumes to the controller with the lower CPU percentage.</p> <p>You might want to monitor CPU across the storage array. If CPU continues to increase over time while application performance decreases, you might need to add storage arrays. By adding storage arrays to your enterprise, you can continue to meet application needs at an acceptable performance level.</p>

Performance data	Implications for performance tuning
Headroom	<p>Headroom refers to the remaining performance capability of the controllers, the controller host channels, and the controller drive channels. This value is expressed as a percentage and represents the gap between the maximum possible performance these objects are able to deliver and the current performance levels.</p> <ul style="list-style-type: none"> <li>• For the controllers, headroom is a percentage of maximum possible IOPS.</li> <li>• For the channels, headroom is a percentage of maximum throughput, or MiB/s. Read throughput, write throughput, and bidirectional throughput are included in the calculation.</li> </ul>

## 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.