



File System Analytics

ONTAP 9

NetApp
May 25, 2022

Table of Contents

- File System Analytics 1
 - File System Analytics overview 1
 - Enable File System Analytics 3
 - View file system activity 3
 - Enable Activity Tracking 5
 - Take corrective action based on analytics 6
 - Considerations for File System Analytics 8

File System Analytics

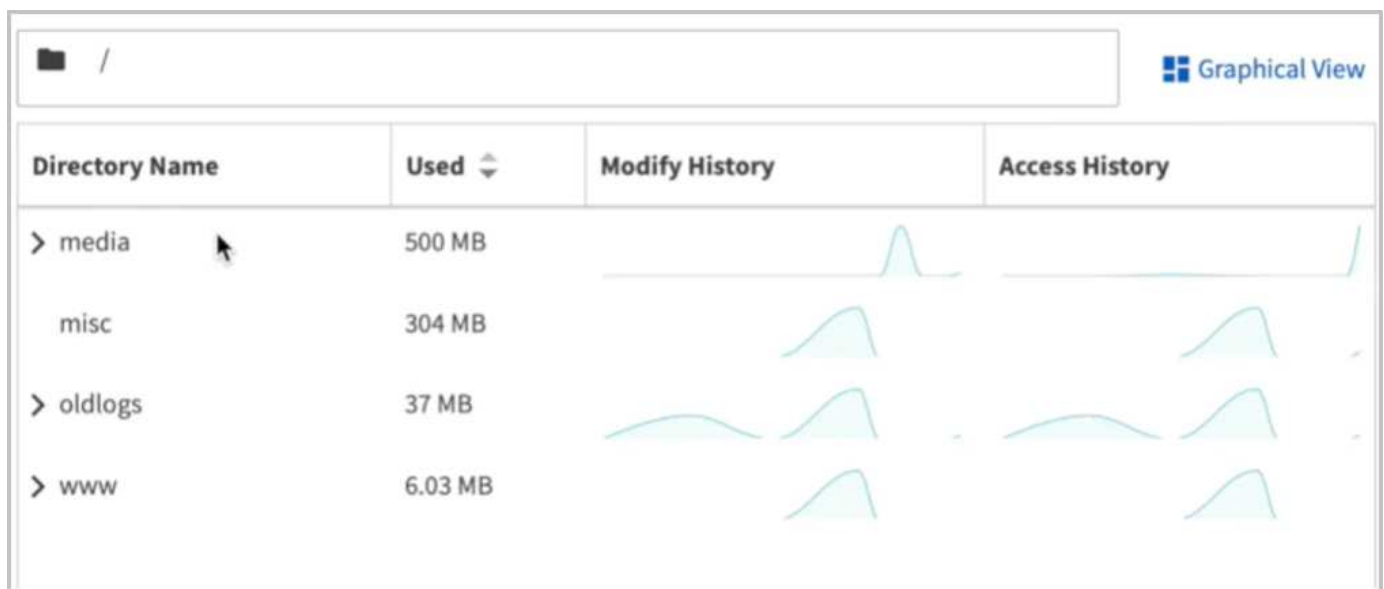
File System Analytics overview

Beginning with ONTAP 9.8, File System Analytics (FSA) is a framework for collecting and visualizing key data points about the contents of ONTAP FlexGroup or FlexVol volumes. FSA helps solve organizational data challenges around data locality, storage consumption, storage optimization, and data protection. By providing real-time insights into data usage, FSA helps you with decision-making in effective data management and operations, such as quality-of-service changes to throughput and moving volumes to a different controller or aggregate.

FSA presents detailed information at each level of the volume's file system hierarchy, allowing you to:

- [Assess capacity usage and trends](#)
- [Monitor file and directory counts](#)
- [Evaluate file activity and history](#)
- [Capture information on inactive data](#)
- [Take corrective action based on insights \(beginning with ONTAP 9.9.1\)](#)
- [Analyze and download information on high traffic and high throughput activity \(beginning with ONTAP 9.10.1\)](#)

File System Analytics offers a list and graphical view of your data. In the graphical view, variably sized blocks and color shades provide a graphical representation of the directories. In the list view shown below, you can see how much space each directory uses as well modification and access histories in numerical form and in interactive graphs for each directory.



File System Analytics data is displayed using ONTAP System Manager. You can also use the ONTAP REST API to access the data programmatically.

Feature availability by release

	ONTAP 9.8	ONTAP 9.9.1	ONTAP 9.10.1	ONTAP 9.11.1
Visualization in System Manager	X	X	X	X
Capacity analytics	X	X	X	X
Inactive Data information	X	X	X	X
Availability for volumes transitioned from 7-mode systems		X	X	X
Option to change Inactive period in System Manager		X	X	X
Activity Tracking (volume-level)			X	X
Download to CSV (Activity Tracking, volume-level)			X	X
Activity Tracking (SVM-level)				X
Download to CSV (Activity Tracking, SVM-level)				X
Timeline				X

Learn more about File System Analytics

ONTAP File System Analytics



Daniel Tennant
Director of Software Engineering
December 13, 2020

© 2020 NetApp, Inc. All rights reserved. — NETAPP CONFIDENTIAL —




Further Reading

[TR 4687: Best-practice guidelines for ONTAP File System Analytics](#)

Enable File System Analytics

To collect and display usage data such as capacity analytics, you must enable File System Analytics.

You can enable File System Analytics when you create a new volume or when you upgrade a system with volumes to ONTAP 9.8 or later. After upgrading, ensure that all upgrade processes have completed before enabling analytics.

Procedure

Depending on the size and contents of the volume, enabling analytics may take time while ONTAP processes existing data in the volume. System Manager displays progress and presents analytics data when complete. If you need more precise information about initialization progress, you can use the ONTAP CLI command `volume analytics show`.

You can enable File System Analytics with ONTAP System Manager, the ONTAP CLI, or the ONTAP REST API.

System Manager

In ONTAP 9.8 and 9.9.1	Beginning in ONTAP 9.10.1
<ol style="list-style-type: none">1. Select Storage > Volumes.2. Select the desired volume, then select Explorer.3. Select Enable Analytics or Disable Analytics.	<ol style="list-style-type: none">1. Select Storage > Volumes.2. Select the desired volume. From the individual volume menu, select File System > Explorer.3. Select Enable Analytics or Disable Analytics.

CLI

To enable File System Analytics with the CLI:

1. Run the following command:

```
volume analytics on -vserver svm_name -volume volume_name [-foreground {true|false}]
```

By default, the command runs in the foreground; ONTAP displays progress and presents analytics data when complete. If you need more precise information, you can run the command in the background by using the `-foreground false` option and then use the `volume analytics show` command to display initialization progress in the CLI.

2. After successfully enabling File System Analytics, use ONTAP System Manager or the ONTAP REST API to display the analytic data.

View file system activity

After File System Analytics (FSA) is enabled, by default you can view the root directory contents of a selected volume sorted by the space used in each subtree.

Clicking on any file system object allows you to browse the file system and to display detailed information about each object in a directory. Information about directories can also be displayed graphically. Over time, historical data is displayed for each subtree. Space used is not sorted if there are more than 3000 directories.

Explorer

The File System Analytics **Explorer** screen consists of three areas:

- Tree view of directories and subdirectories; expandable list showing name, size, modify history, and access history.
- Files; showing name, size, and accessed time for the object selected in the directory list.
- Active and inactive data comparison for the object selected in the directory list.

Beginning with ONTAP 9.9.1, you can customize the range to be reported. The default value is one year. Based on these customizations, you can take corrective actions, such as moving volumes and modifying the tiering policy.

Accessed time is shown by default. However, if the volume default has been altered from the CLI, by setting the `-atime-update` option to `false` with the `volume modify` command, only last modified time is shown. For example:

- The tree view will not display the **access history**.
- The files view will be altered.
- The active/inactive data view will be based on modified time (`mtime`).

Using these displays, you can examine the following:

- File system locations consuming the most space
- Detailed information about a directory tree, including file and subdirectory count within directories and subdirectories
- File system locations that contain old data (for example, scratch, temp, or log trees)

Keep the following points in mind when interpreting FSA output:

- FSA show where and when your data is in use, not how much data is being processed. For example, large space consumption by recently accessed or modified files does not necessarily indicate high system processing loads.
- The way that the **Volume Explorer** tab calculates space consumption for FSA might differ from other tools. In particular, there could be significant differences compared to the consumption reported in the **Volume Overview** if the volume has storage efficiency features enabled. This is because the **Volume Explorer** tab does not include efficiency savings.
- Due to space limitations in the directory display, it is not possible to view a directory depth greater than 8 levels in the *List View*. To view directories more than 8 levels deep, you must switch to *Graphical View*, locate the desired directory, then switch back to *List View*. This will allow additional screen space in the display.

Procedure

In ONTAP 9.8 and 9.9.1	Beginning in ONTAP 9.10.1
Click Storage > Volumes , select the desired volume, then click Explorer .	Select Storage > Volumes , select the desired volume. From the individual volume menu, select File System > Explorer .

Enable Activity Tracking

Beginning with ONTAP 9.10.1, File System Analytics includes an Activity Tracking feature that allows you to identify hot objects and download them as a CSV file. Beginning with ONTAP 9.11.1, Activity Tracking is expanded to the SVM scope. Also beginning in ONTAP 9.11.1, System Manager features a timeline for Activity Tracking, allowing you to look through up to five minutes of Activity Tracking data.

Activity Tracking enables monitoring in four categories:

- Directories
- Files
- Clients
- Users

For each category monitored, Activity Tracking will display read IOPs, write IOPs, read throughputs, and write throughputs. Queries on Activity Tracking refresh every 10 to 15 seconds pertaining to hot spots seen in the system over the previous five-second interval.

Activity tracking information is approximate, and the accuracy of the data depends on the distribution of the incoming I/O traffic.

When viewing Activity Tracking in System Manager at the volume level, only the menu of the expanded volume will actively refresh. If the view of any volumes are collapsed, they will not refresh until the volume display is expanded. You can stop the refreshes with the **Pause Refresh** button. Activity data can be downloaded in a CSV format that will display all the point-in-time data captured for the selected volume.

With the timeline feature available beginning in ONTAP 9.11.1, you can keep a record of hotspot activity on a volume or SVM, continuously updating approximately every five seconds and retaining the previous five minutes of data. Timeline data is only retained for fields that are visible area of the page. If you collapse a tracking category or scroll so the timeline is out of view, the timeline will stop collecting data. By default, timelines are disabled and will automatically be disabled when you navigate away from the Activity tab

Enable Activity Tracking for a volume

If you use RBAC with the ONTAP REST API or System Manager, you will need to create custom roles to manage access to Activity Tracking. See [Considerations for File System Analytics](#) for this process.

Procedures

You can enable Activity Tracking with ONTAP System Manager, the ONTAP CLI, or the ONTAP REST API.

System Manager

1. Select **Storage > Volumes**. Select the desired volume. From the individual volume menu, select File System and then select the Activity tab.
2. Ensure **Activity Tracking** is turned on to view individual reports on top directories, files, clients, and users.
3. To analyze data in greater depth without refreshes, select **Pause Refresh**. You can download the data to have a CSV record of the report as well.

CLI

1. Enable Activity Tracking:
`volume activity-tracking on -vsverver svm_name -volume volume_name`
2. You can check if the Activity Tracking state for a volume is on or off with the command:
`volume activity-tracking show -vsverver svm_name -volume volume_name -state`
3. Once enabled, use ONTAP System Manager or the ONTAP REST API to display Activity Tracking data.

Enable Activity Tracking for multiple volumes

You can enable Activity Tracking for multiple volumes at once through System Manager or the ONTAP REST API.

For specific volumes

1. Select **Storage > Volumes**. Select the desired volume. From the individual volume menu, select File System and then select the Activity tab.
2. Select the volumes that you want to enable Activity Tracking on. At the top of the volume list, select the **More Options** button. Select **Enable Activity Tracking**.
3. To view Activity Tracking at the SVM level, select the specific SVM you would like to view from **Storage > Volumes**. Navigate to the File System tab then Activity and you will see data for the volumes that have Activity Tracking enabled.

For all volumes in an SVM

1. Select **Storage > Volumes**. Select an SVM from the menu.
2. Navigate to the **File System** tab, choose the **More** tab to enable Activity Tracking on all volumes in the SVM.

Take corrective action based on analytics

Beginning with ONTAP 9.9.1, you can take corrective actions directly from File System Analytics displays based on current data and desired outcomes.

When analytics are enabled, you can take the following actions:

- Delete directories and files

In the Explorer display, you can select directories or individual files to delete. Directories are deleted with

low-latency fast directory delete functionality. (Fast directory delete is also available beginning in ONTAP 9.9.1 without analytics enabled.)

- Assign media cost in storage tiers to compare costs of inactive data storage locations


Media cost is a value that you assign based on your evaluation of storage costs, represented as your choice of currency per GB. When set, System Manager uses the assigned media cost to project estimated savings when you move volumes.

The media cost you set is not persistent; it can only be set for a single browser session.

- Move volumes to reduce storage costs
Based on analytics displays and media cost comparisons, you can move volumes to less expensive storage in local tiers.

Only one volume at a time can be compared and moved.

Steps

To perform this action...	Take these steps...
Delete directories or files	<ol style="list-style-type: none">1. Click Storage > Volumes, then click Explorer. <p>When you hover over a file or folder, the option to delete appears. You can only delete one object at a time.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px;"> When directories and files are deleted, the new storage capacity values are not displayed immediately.</div>
Enable media cost comparison	<ol style="list-style-type: none">1. Click Storage > Tiers, then click Set Media Cost in the desired local tier (aggregate) tiles. Be sure to select active and inactive tiers to enable comparison.2. Enter a currency type and amount. <p>When you enter or change the media cost, the change is made in all media types.</p>

To perform this action...	Take these steps...
Move volumes to a less expensive tier	<ol style="list-style-type: none"> 1. After enabling media cost display, click Storage > Tiers, then click Volumes. 2. To compare destination options for a volume, click : for the volume, then click Move. 3. In the Select Destination Local Tier display, select destination tiers to display the estimated cost difference. 4. After comparing options, select the desired tier and click Move.

Considerations for File System Analytics

You should be aware of certain usage limits and potential performance impacts associated with implementing File System Analytics.

Supported volume types

File System Analytics is designed to provide visibility on volumes with active NAS data, with the exception of FlexCache caches and SnapMirror destination volumes.

Performance considerations

Enabling File System Analytics may incur a performance impact during the initial metadata collection. Do not enable analytics on the systems that are at a maximum utilization.

ONTAP System Manager provides [performance monitoring tools](#) to manage cluster health and metrics.

SVM protected relationships

If you have enabled File System Analytics on volumes whose containing SVM is in a protection relationship, the analytics data is not replicated to the destination SVM. If the source SVM must be resynchronized in a recovery operation, you must manually reenables analytics on desired volumes after recovery.

RBAC

If you are user of System Manager or the ONTAP REST API [RBAC](#) in your ONTAP set up, you will need to create a dedicated role to moderate access to Activity Tracking in File System Analytics.

High Level

1. Create a role that has access to all accept the “volume file show-disk-usage” command directory
2. Create web access for the new role
3. Create a user with the new role
4. Login to System Manager or make REST calls using the new user

Control access to Activity Tracking

1. Create a default role to have access to all features. This needs to be done before creating the restrictive role to ensure the role is only restrictive on the Activity Tracking:

```
security login role create -cmddirname DEFAULT -access all -role storageAdmin
```

2. Create the restrictive role:

```
security login role create -cmddirname "volume file show-disk-usage" -access none -role storageAdmin
```

3. Authorize roles to access the SVM's web services:

- rest for REST API calls
- security for password protection
- sysmgr for System Manager access

```
vserver services web access create -vserver svm-name -name_ -name rest -role storageAdmin  
vserver services web access create -vserver svm-name -name security -role storageAdmin  
vserver services web access create -vserver svm-name -name sysmgr -role storageAdmin
```

4. Create a user. You must issue a distinct create command for each application you would like to apply to the user. Calling create multiple times on the same user simply applies all the applications to that one user and does not create a new user each time. The `http` parameter for application type applies for the ONTAP REST API and System Manager.

```
security login create -user-or-group-name storageUser -authentication-method password -application http -role storageAdmin
```

Copyright Information

Copyright © 2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.