



Understand FPolicy

ONTAP 9

NetApp
April 24, 2024

Table of Contents

- Understand FPolicy 1
 - What the two parts of the FPolicy solution are 1
 - What synchronous and asynchronous notifications are 1
 - FPolicy persistent stores 2
 - FPolicy configuration types 3
 - Roles that cluster components play with FPolicy implementation 4
 - How FPolicy works with external FPolicy servers 5
 - What the node-to-external FPolicy server communication process is 7
 - How FPolicy services work across SVM namespaces 8
 - How FPolicy passthrough-read enhances usability for hierarchical storage management 9

Understand FPolicy

What the two parts of the FPolicy solution are

FPolicy is a file access notification framework that is used to monitor and manage file access events on storage virtual machines (SVMs) through partner solutions. Partner solutions help you address various use cases such as data governance and compliance, ransomware protection, and data mobility.

Partner solutions include both Netapp supported 3rd party Solutions and NetApp products Workload Security and Cloud Data Sense.

There are two parts to an FPolicy solution. The ONTAP FPolicy framework manages activities on the cluster and sends notifications to Partner Application (aka External FPolicy Servers). External FPolicy servers process notifications sent by ONTAP FPolicy to fulfill customer use cases.

The ONTAP framework creates and maintains the FPolicy configuration, monitors file events, and sends notifications to external FPolicy servers. ONTAP FPolicy provides the infrastructure that allows communication between external FPolicy servers and storage virtual machine (SVM) nodes.

The FPolicy framework connects to external FPolicy servers and sends notifications for certain file system events to the FPolicy servers when these events occur as a result of client access. The external FPolicy servers process the notifications and send responses back to the node. What happens as a result of the notification processing depends on the application and whether the communication between the node and the external servers is asynchronous or synchronous.

What synchronous and asynchronous notifications are

FPolicy sends notifications to external FPolicy servers via the FPolicy interface. The notifications are sent either in synchronous or asynchronous mode. The notification mode determines what ONTAP does after sending notifications to FPolicy servers.

- **Asynchronous notifications**

With asynchronous notifications, the node does not wait for a response from the FPolicy server, which enhances overall throughput of the system. This type of notification is suitable for applications where the FPolicy server does not require that any action be taken as a result of notification evaluation. For example, asynchronous notifications are used when the storage virtual machine (SVM) administrator wants to monitor and audit file access activity.

If an FPolicy server operating in asynchronous mode experiences a network outage, FPolicy notifications generated during the outage are stored on the storage node. When the FPolicy server comes back online, it is alerted of the stored notifications and can fetch them from the storage node. The length of time the notifications can be stored during an outage is configurable up to 10 minutes.

Beginning with ONTAP 9.14.1, FPolicy allows you to set up a persistent store to capture file access events for asynchronous non-mandatory policies in the SVM. Persistent stores can help decouple client I/O processing from FPolicy notification processing to reduce client latency. Synchronous (mandatory or non-mandatory) and asynchronous mandatory configurations are not supported.

- **Synchronous notifications**

When configured to run in synchronous mode, the FPolicy server must acknowledge every notification before the client operation is allowed to continue. This type of notification is used when an action is required based on the results of notification evaluation. For example, synchronous notifications are used when the SVM administrator wants to either allow or deny requests based on criteria specified on the external FPolicy server.

Synchronous and asynchronous applications

There are many possible uses for FPolicy applications, both asynchronous and synchronous.

Asynchronous applications are ones where the external FPolicy server does not alter access to files or directories or modify data on the storage virtual machine (SVM). For example:

- File access and audit logging
- Storage resource management

Synchronous applications are ones where data access is altered or data is modified by the external FPolicy server. For example:

- Quota management
- File access blocking
- File archiving and hierarchical storage management
- Encryption and decryption services
- Compression and decompression services

FPolicy persistent stores

Beginning with ONTAP 9.14.1, FPolicy allows you to set up a persistent store to capture file access events for asynchronous non-mandatory policies in the SVM. Persistent stores can help decouple client I/O processing from FPolicy notification processing to reduce client latency. Synchronous (mandatory or non-mandatory) and asynchronous mandatory configurations are not supported.

This feature is only available in FPolicy external mode. The partner application you use needs to support this feature. You should work with your partner to ensure this FPolicy configuration is supported.

Best practices

Cluster administrators need to configure a volume for the persistent store on each SVM where FPolicy is enabled. When configured, a persistent store captures all matching FPolicy events, which are further processed in the FPolicy pipeline and sent to the external server.

The persistent store remains as it was when the last event was received when there is an unexpected reboot or FPolicy is disabled and enabled again. After a takeover operation, new events will be stored and processed by the partner node. After a giveback operation, the persistent store resumes processing any unprocessed events that might remain from when the node takeover occurred. Live events would be given priority over unprocessed events.

If the persistent store volume moves from one node to another in the same SVM, the notifications that are yet to be processed will also move to the new node. You will need to re-run the `fpolicy persistent-store`

`create` command on either node after the volume is moved to ensure the pending notification are delivered to the external server.

The persistent store volume is setup on a per SVM basis. For each FPolicy enabled SVM you will need to create a persistent store volume.

Create the persistent store volume on the node with LIFs that expect maximum traffic to be monitored by Fpolicy.

If the notifications accumulated in the persistent store exceed the size of the volume provisioned, FPolicy will start dropping the incoming notification with appropriate EMS messages.

The persistent Store volume name and the junction-path specified at the time of volume creation should match.

Have the snapshot policy set to `none` for that volume instead of `default`. This is to ensure that there is no accidental restore of the snapshot leading to loss of current events and to prevent possible duplicate event processing.

Make the persistent store volume inaccessible for external user protocol access (CIFS/NFS) to avoid accidental corruption or deletion of the persisted event records. To achieve this, after enabling FPolicy, unmount the volume in ONTAP to remove the junction path, this makes it inaccessible for the user protocol access.

For more information, see [Create persistent stores](#).

FPolicy configuration types

There are two basic FPolicy configuration types. One configuration uses external FPolicy servers to process and act upon notifications. The other configuration does not use external FPolicy servers; instead, it uses the ONTAP internal, native FPolicy server for simple file blocking based on extensions.

- **External FPolicy server configuration**

The notification is sent to the FPolicy server, which screens the request and applies rules to determine whether the node should allow the requested file operation. For synchronous policies, the FPolicy server then sends a response to the node to either allow or block the requested file operation.

- **Native FPolicy server configuration**

The notification is screened internally. The request is allowed or denied based on file extension settings configured in the FPolicy scope.

Note: File extension requests that are denied are not logged.

When to create a native FPolicy configuration

Native FPolicy configurations use the ONTAP internal FPolicy engine to monitor and block file operations based on the file's extension. This solution does not require external FPolicy servers (FPolicy servers). Using a native file blocking configuration is appropriate when this simple solution is all that is needed.

Native file blocking enables you to monitor any file operations that match configured operation and filtering events and then deny access to files with particular extensions. This is the default configuration.

This configuration provides a means to block file access based only on the file's extension. For example, to block files that contain `mp3` extensions, you configure a policy to provide notifications for certain operations with target file extensions of `mp3`. The policy is configured to deny `mp3` file requests for operations that generate notifications.

The following applies to native FPolicy configurations:

- The same set of filters and protocols that are supported by FPolicy server-based file screening are also supported for native file blocking.
- Native file blocking and FPolicy server-based file screening applications can be configured at the same time.

To do so, you can configure two separate FPolicy policies for the storage virtual machine (SVM), with one configured for native file blocking and one configured for FPolicy server-based file screening.

- The native file blocking feature only screens files based on the extensions and not on the content of the file.
- In the case of symbolic links, native file blocking uses the file extension of the root file.

Learn more about [FPolicy: Native File Blocking](#).

When to create a configuration that uses external FPolicy servers

FPolicy configurations that use external FPolicy servers to process and manage notifications provide robust solutions for use cases where more than simple file blocking based on file extension is needed.

You should create a configuration that uses external FPolicy servers when you want to do such things as monitor and record file access events, provide quota services, perform file blocking based on criteria other than simple file extensions, provide data migration services using hierarchical storage management applications, or provide a fine-grained set of policies that monitor only a subset of data in the storage virtual machine (SVM).

Roles that cluster components play with FPolicy implementation

The cluster, the contained storage virtual machines (SVMs), and data LIFs all play a role in an FPolicy implementation.

- **cluster**

The cluster contains the FPolicy management framework and maintains and manages information about all FPolicy configurations in the cluster.

- **SVM**

An FPolicy configuration is defined at the SVM level. The scope of the configuration is the SVM, and it only operates on SVM resources. One SVM configuration cannot monitor and send notifications for file access requests that are made for data residing on another SVM.

FPolicy configurations can be defined on the admin SVM. After configurations are defined on the admin SVM, they can be seen and used in all SVMs.

- **data LIFs**

Connections to the FPolicy servers are made through data LIFs belonging to the SVM with the FPolicy configuration. The data LIFs used for these connections can fail over in the same manner as data LIFs used for normal client access.

How FPolicy works with external FPolicy servers

After FPolicy is configured and enabled on the storage virtual machine (SVM), FPolicy runs on every node on which the SVM participates. FPolicy is responsible for establishing and maintaining connections with external FPolicy servers (FPolicy servers), for notification processing, and for managing notification messages to and from FPolicy servers.

Additionally, as part of connection management, FPolicy has the following responsibilities:

- Ensures that file notification flows through the correct LIF to the FPolicy server.
- Ensures that when multiple FPolicy servers are associated with a policy, load balancing is done when sending notifications to the FPolicy servers.
- Attempts to reestablish the connection when a connection to an FPolicy server is broken.
- Sends the notifications to FPolicy servers over an authenticated session.
- Manages the passthrough-read data connection established by the FPolicy server for servicing client requests when passthrough-read is enabled.

How control channels are used for FPolicy communication

FPolicy initiates a control channel connection to an external FPolicy server from the data LIFs of each node participating on a storage virtual machine (SVM). FPolicy uses control channels for transmitting file notifications; therefore, an FPolicy server might see multiple control channel connections based on SVM topology.

How privileged data access channels are used for synchronous communication

With synchronous use cases, the FPolicy server accesses data residing on the storage virtual machine (SVM) through a privileged data access path. Access through the privileged path exposes the complete file system to the FPolicy server. It can access data files to collect information, to scan files, read files, or write into files.

Because the external FPolicy server can access the entire file system from the root of the SVM through the privileged data channel, the privileged data channel connection must be secure.

How FPolicy connection credentials are used with privileged data access channels

The FPolicy server makes privileged data access connections to cluster nodes by using a specific Windows user credential that is saved with the FPolicy configuration. SMB is the only supported protocol for making a privileged data access channel connection.

If the FPolicy server requires privileged data access, the following conditions must be met:

- A SMB license must be enabled on the cluster.
- The FPolicy server must run under the credentials configured in the FPolicy configuration.

When making a data channel connection, FPolicy uses the credential for the specified Windows user name. Data access is made over the admin share ONTAP_ADMIN\$.

What granting super user credentials for privileged data access means

ONTAP uses the combination of the IP address and the user credential configured in the FPolicy configuration to grant super user credentials to the FPolicy server.

Super user status grants the following privileges when the FPolicy server accesses data:

- Avoid permission checks

The user avoids checks on files and directory access.

- Special locking privileges

ONTAP allows read, write, or modify access to any file regardless of existing locks. If the FPolicy server takes byte range locks on the file, it results in immediate removal of existing locks on the file.

- Bypass any FPolicy checks

Access does not generate any FPolicy notifications.

How FPolicy manages policy processing

There might be multiple FPolicy policies assigned to your storage virtual machine (SVM); each with a different priority. To create an appropriate FPolicy configuration on the SVM, it is important to understand how FPolicy manages policy processing.

Each file access request is initially evaluated to determine which policies are monitoring this event. If it is a monitored event, information about the monitored event along with interested policies is passed to FPolicy where it is evaluated. Each policy is evaluated in order of the assigned priority.

You should consider the following recommendations when configuring policies:

- When you want a policy to always be evaluated before other policies, configure that policy with a higher priority.
- If the success of requested file access operation on a monitored event is a prerequisite for a file request that is evaluated against another policy, give the policy that controls the success or failure of the first file operation a higher priority.

For example, if one policy manages FPolicy file archiving and restore functionality and a second policy manages file access operations on the online file, the policy that manages file restoration must have a higher priority so that the file is restored before the operation managed by the second policy can be allowed.

- If you want all policies that might apply to a file access operation to be evaluated, give synchronous policies a lower priority.

You can reorder policy priorities for existing policies by modifying the policy sequence number. However, to have FPolicy evaluate policies based on the modified priority order, you must disable and reenable the policy with the modified sequence number.

What the node-to-external FPolicy server communication process is


To properly plan your FPolicy configuration, you should understand what the node-to-external FPolicy server communication process is.

Every node that participates on each storage virtual machine (SVM) initiates a connection to an external FPolicy server (FPolicy server) using TCP/IP. Connections to the FPolicy servers are set up using node data LIFs; therefore, a participating node can set up a connection only if the node has an operational data LIF for the SVM.

Each FPolicy process on participating nodes attempts to establish a connection with the FPolicy server when the policy is enabled. It uses the IP address and port of the FPolicy external engine specified in the policy configuration.

The connection establishes a control channel from each of the nodes participating on each SVM to the FPolicy server through the data LIF. In addition, if IPv4 and IPv6 data LIF addresses are present on the same participating node, FPolicy attempts to establish connections for both IPv4 and IPv6. Therefore, in a scenario where the SVM extends over multiple nodes or if both IPv4 and IPv6 addresses are present, the FPolicy server must be ready for multiple control channel setup requests from the cluster after the FPolicy policy is enabled on the SVM.

For example, if a cluster has three nodes—Node1, Node2, and Node3—and SVM data LIFs are spread across only Node2 and Node3, control channels are initiated only from Node2 and Node3, irrespective of the distribution of data volumes. Say that Node2 has two data LIFs—LIF1 and LIF2—that belong to the SVM and that the initial connection is from LIF1. If LIF1 fails, FPolicy attempts to establish a control channel from LIF2.



How FPolicy manages external communication during LIF migration or failover

Data LIFs can be migrated to data ports in the same node or to data ports on a remote node.

When a data LIF fails over or is migrated, a new control channel connection is made to the FPolicy server. FPolicy can then retry SMB and NFS client requests that timed out, with the result that new notifications are sent to the external FPolicy servers. The node rejects FPolicy server responses to original, timed-out SMB and NFS requests.

How FPolicy manages external communication during node failover

If the cluster node that hosts the data ports used for FPolicy communication fails, ONTAP breaks the connection between the FPolicy server and the node.

The impact of cluster failover to the FPolicy server can be mitigated by configuring the failover-policy to migrate the data port used in FPolicy communication to another active node. After the migration is complete, a new connection is established using the new data port.

If the failover-policy is not configured to migrate the data port, the FPolicy server must wait for the failed node to come up. After the node is up, a new connection is initiated from that node with a new Session ID.



The FPolicy server detects broken connections with the keep-alive protocol message. The timeout for purging the session ID is determined when configuring FPolicy. The default keep-alive timeout is two minutes.

How FPolicy services work across SVM namespaces

ONTAP provides a unified storage virtual machine (SVM) namespace. Volumes across the cluster are joined together by junctions to provide a single, logical file system. The FPolicy server is aware of the namespace topology and provides FPolicy services across the namespace.

The namespace is specific to and contained within the SVM; therefore, you can see the namespace only from the SVM context. Namespaces have the following characteristics:

- A single namespace exists in each SVM, with the root of the namespace being the root volume, represented in the namespace as slash (/).
- All other volumes have junction points below the root (/).
- Volume junctions are transparent to clients.
- A single NFS export can provide access to the complete namespace; otherwise, export policies can export specific volumes.
- SMB shares can be created on the volume or on qtrees within the volume, or on any directory within the namespace.
- The namespace architecture is flexible.

Examples of typical namespace architectures are as follows:

- A namespace with a single branch off of the root
- A namespace with multiple branches off of the root

- A namespace with multiple unbranched volumes off of the root

How FPolicy passthrough-read enhances usability for hierarchical storage management

Passthrough-read enables the FPolicy server (functioning as the hierarchical storage management (HSM) server) to provide read access to offline files without having to recall the file from the secondary storage system to the primary storage system.

When an FPolicy server is configured to provide HSM to files residing on a SMB server, policy-based file migration occurs where the files are stored offline on secondary storage and only a stub file remains on primary storage. Even though a stub file appears as a normal file to clients, it is actually a sparse file that is the same size of the original file. The sparse file has the SMB offline bit set and points to the actual file that has been migrated to secondary storage.

Typically when a read request for an offline file is received, the requested content must be recalled back to primary storage and then accessed through primary storage. The need to recall data back to primary storage has several undesirable effects. Among the undesirable effects is the increased latency to client requests caused by the need to recall the content before responding to the request and the increased space consumption needed for recalled files on the primary storage.

FPolicy passthrough-read allows the HSM server (the FPolicy server) to provide read access to migrated, offline files without having to recall the file from the secondary storage system to the primary storage system. Instead of recalling the files back to primary storage, read requests can be serviced directly from secondary storage.



Copy Offload (ODX) is not supported with FPolicy passthrough-read operation.

Passthrough-read enhances usability by providing the following benefits:

- Read requests can be serviced even if the primary storage does not have sufficient space to recall requested data back to primary storage.
- Better capacity and performance management when a surge of data recall might occur, such as if a script or a backup solution needs to access many offline files.
- Read requests for offline files in Snapshot copies can be serviced.

Because Snapshot copies are read-only, the FPolicy server cannot restore the original file if the stub file is located in a Snapshot copy. Using passthrough-read eliminates this problem.

- Policies can be set up that control when read requests are serviced through access to the file on secondary storage and when the offline file should be recalled to primary storage.

For example, a policy can be created on the HSM server that specifies the number of times the offline file can be accessed in a specified period of time before the file is migrated back to primary storage. This type of policy avoids recalling files that are rarely accessed.

How read requests are managed when FPolicy passthrough-read is enabled

You should understand how read requests are managed when FPolicy passthrough-read is enabled so that you can optimally configure connectivity between the storage virtual machine (SVM) and the FPolicy servers.

When FPolicy passthrough-read is enabled and the SVM receives a request for an offline file, FPolicy sends a notification to the FPolicy server (HSM server) through the standard connection channel.

After receiving the notification, the FPolicy server reads the data from the file path sent in the notification and sends the requested data to the SVM through the passthrough-read privileged data connection that is established between the SVM and the FPolicy server.

After the data is sent, the FPolicy server then responds to the read request as an ALLOW or DENY. Based on whether the read request is allowed or denied, ONTAP either sends the requested information or sends an error message to the client.

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.