



Technical Report

FPolicy Solution Guide for ONTAP: IntraFind

Full-Text and Enterprise Searches with iFinder5 Elastic Edition for NetApp

Brahmanna Chowdary Kodavali, NetApp
Joerg Viechtbauer, Joerg Issel, Ralf Klinkhammer, IntraFind Software AG
April 2018 | TR-4670

TABLE OF CONTENTS

1	Introduction	5
1.1	Audience	5
1.2	Purpose and Scope	5
2	FPolicy Overview	5
2.1	Role of ONTAP Components in FPolicy Configuration	6
2.2	How FPolicy Works with External FPolicy Servers	6
3	FPolicy Solution Architecture	7
3.1	Components of FPolicy Framework on ONTAP	7
3.2	FPolicy Application Software: iFinder5 Elastic Edition for NetApp	8
4	Installing and Configuring iFinder5 Elastic Edition for NetApp	9
4.1	Software Requirements and Installation	9
4.2	Prerequisites	9
4.3	Installing IntraFind Software	10
4.4	Configure IntraFind Software	13
4.5	Complete IntraFind Software Installation	17
4.6	Additional Information	17
5	Configuring IntraFind FPolicy Server	17
5.1	Service Description	17
5.2	Prerequisites	18
5.3	Update Behavior	18
5.4	Dependencies with Other IntraFind Components	18
5.5	Install IntraFind FPolicy Infrastructure	18
6	Configuring IntraFind Policy Servers	20
6.1	Configure FPolicy Server Mode	21
7	Configuring Event Processing	21
7.1	Configure Multiple Event Logs	22
7.2	Configure in the config.cfg File	22
8	Reconfiguring Share Indexer	25
9	Miscellaneous	26
9.1	Scaling and Failover	26
9.2	Memory Settings	26
9.3	Licensing	26

9.4	Logging	26
9.5	Compression of Event Logs	26
10	Configuring FPolicy on ONTAP	27
10.1	FPolicy Configuration Workflow	27
10.2	Create an FPolicy Event	28
10.3	Create an FPolicy External Engine	29
10.4	Create an FPolicy Policy	29
10.5	Create an FPolicy Scope	30
10.6	Enable an FPolicy Policy	30
11	Configuring Security Login for FPolicy Server	30
12	ONTAP Best Practices	30
12.1	Policy Configuration	31
12.2	Network Configuration	31
12.3	Hardware Configuration	31
12.4	Multiple Policy Configuration	31
12.5	Managing FPolicy Workflow and Dependency on Other Technologies	31
12.6	Sizing Considerations	32
13	iFinder5 Elastic Edition for NetApp Best Practices	32
14	Troubleshooting	32
14.1	Problem 1: The FPolicy Server Is Disconnected	32
14.2	Problem 2: The FPolicy Server Does Not Connect	32
14.3	Problem 3: The External Engine Is Not Native for the Policy	33
14.4	Problem 4: Notifications Are Not Being Received for the File Operations on Volume, Share, and Export	33
15	Performance Monitoring	34
15.1	Collect and Display FPolicy Counters	34
15.2	Counters to Monitor	34
15.3	Performance Monitoring for iFinder5 Elastic Edition for NetApp	35
Appendix: FPolicy Server Cheat Sheets		35
	FPolicy Server Installation Cheat Sheet	35
	FPolicy Server Troubleshooting Cheat Sheet	36
Where to Find Additional Information		36
Version History		37

LIST OF TABLES

Table 1) Options for configuring in the config.cfg file.	22
Table 2) Options for configuring in the fpolicy.table file.	23
Table 3) FPolicy counters.	34
Table 4) FPolicy server troubleshooting.	36

LIST OF FIGURES

Figure 1) FPolicy solution architecture.	7
Figure 2) iFinder5 elastic search application.	8
Figure 3) FPolicy events.	17
Figure 4) FPolicy mode configuration.	19
Figure 5) FPolicy port configuration.	20
Figure 6) IntraFind policy servers for ONTAP configuration.	20
Figure 7) Event processing configuration.	21
Figure 8) IntraFind share indexer.	25
Figure 9) Compression of event logs.	27
Figure 10) FPolicy configuration workflow.	28

1 Introduction

NetApp® FPolicy™ is a file access notification framework that allows users to monitor file access over NFS and CIFS protocols. This feature was introduced in NetApp clustered Data ONTAP® 8.2, a scale-out architecture that enables a rich set of use cases working with partners. The FPolicy framework requires that all the nodes in the cluster are running Data ONTAP 8.2 or later. FPolicy supports all SMB versions such as SMB 1.0 (also known as CIFS), SMB 2.0, SMB 2.1, and SMB 3.0. It also supports major NFS versions such as NFS v3 and NFS v4.0.

The FPolicy framework natively supports a simple file-blocking use case, which enables administrators to restrict end users from storing unwanted files. For example, an administrator can block audio and video files from being stored in data centers, which saves precious storage resources. This feature blocks files based on only extension. For more advanced features, partner solutions must be considered.

The FPolicy framework enables partners to develop applications catering to a diverse set of use cases. The use cases include, but are not limited to, the following:

- File screening
- File access reporting
- User and directory quotas
- HSM and archiving solutions
- File replication
- Data governance

1.1 Audience

The target audience for this document is customers implementing IntraFind's iFinder5 elastic, a full-text search product, with the FPolicy server to index files stored on ONTAP® software. By using the FPolicy service, administrators can define the rules, which subsequently leads to very fast indexing of the respective documents. These documents are then searchable in iFinder5 elastic, enabling customers to tap into the wealth of information stored in the ONTAP shares. iFinder5 elastic is an enterprise-ready insight engine/enterprise search engine that delivers a single point of information access for all file service data and other data repositories such as intranets, wikis, collaboration platforms, or emails. Full-text content, metadata, and access rights are indexed. iFinder5 elastic includes ready-to-use graphical user interfaces (GUIs), a feature-rich knowledge worker GUI, lightweight integration components such as search bar and hitlist, accessible GUI, and even an iOS and Android app.

1.2 Purpose and Scope

The purpose of this document is to provide an understanding of FPolicy framework and define steps to deploy the iFinder5 elastic. The scope of the document includes the required deployment steps and best practices for the solution.

2 FPolicy Overview

The ONTAP FPolicy framework creates and maintains the FPolicy configuration, monitors file events resulting from client access, and sends notifications to external FPolicy servers. The communication between the storage node and the external FPolicy servers is either asynchronous or synchronous.

Asynchronous or synchronous communication depends on whether the FPolicy framework expects a notification response from the FPolicy server:

- Asynchronous notification is suitable for use cases for which ONTAP does not act based on the notification response from the FPolicy server. Therefore, it won't wait for the response from the FPolicy server.

Monitoring and auditing file access activities are examples of asynchronous notification use cases.

- In contrast to asynchronous notification, synchronous notification is suitable for the use cases for which ONTAP must allow or deny the client access based on the notification response from the FPolicy server.

Quotas, file screening, file archiving recall, replication, and so on are examples of synchronous notification use cases.

2.1 Role of ONTAP Components in FPolicy Configuration

The administrator storage virtual machine (SVM) (cluster), data SVMs, and data LIFs associated with SVM play a role in an FPolicy configuration.

Administrator SVM (Cluster)

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

Data SVMs

FPolicy configuration can be defined at the cluster or the SVM. The scope option in the FPolicy configuration defines the resources that are monitored in the SVM context; the scope option operates on only the SVM resources. One SVM configuration cannot monitor and send notifications for the data (shares) that belong to another SVM.

FPolicy configurations defined on the administrator SVM can be leveraged in all data SVMs.

Data LIFs

Connections to the FPolicy servers are made through data LIFs belonging to the data 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.

2.2 How FPolicy Works with External FPolicy Servers

The FPolicy framework runs on every node in the cluster. This framework is responsible for establishing and maintaining connections with the external FPolicy servers. As part of the connection management, the FPolicy framework manages the following tasks:

- Make sure that file notification flows through the correct LIF to the FPolicy server.
- Load balances the notifications to the FPolicy server when multiple FPolicy servers are associated with a policy.
- Attempts to reestablish the connection when a connection to an FPolicy server is broken.
- Sends the notifications to FPolicy servers over an authenticated session.
- Establishes the connection with the data LIFs on all the nodes participating in the SVM.

For synchronous use cases, the FPolicy server accesses data on the SVM through a privileged data access path. To make this privileged data access path secure, ONTAP uses a combination of user credentials and the IP address of the FPolicy server configured as part of the FPolicy configuration on ONTAP. After the FPolicy server is enabled, the user credentials used in the FPolicy configuration are granted the following special privileges on the file system:

- The ability to bypass permissions checks while accessing the data. The user avoids checks on files and directory access.
- Special locking privileges. ONTAP allows the FPolicy server to read, write, or modify access to any file regardless of existing locks.

Note: If the FPolicy server takes byte range locks on the file, it results in immediate removal of existing locks on the file.

- The ability to bypass any FPolicy checks. File access over a privileged data path does not generate FPolicy notification.

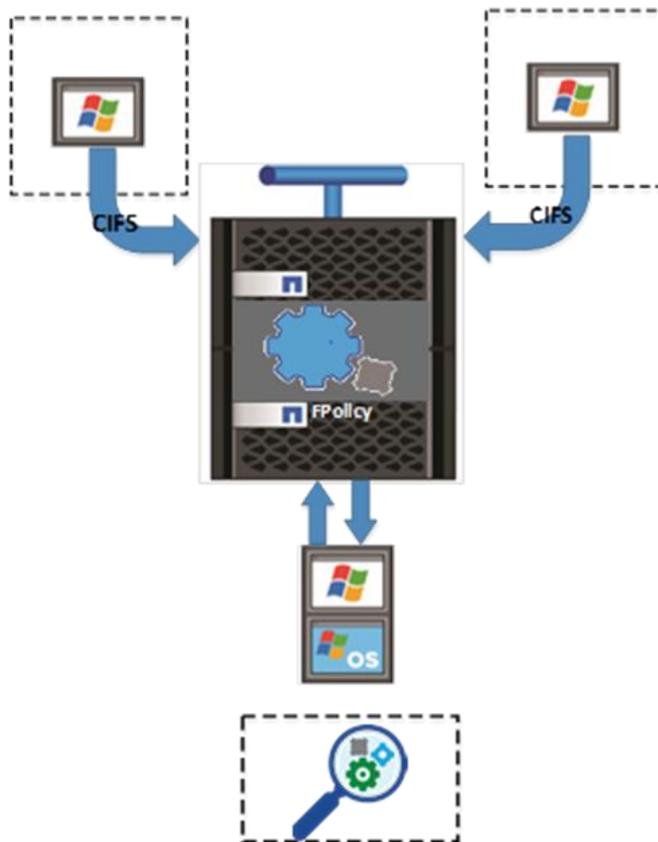
For more information about FPolicy functionality, see the “File Access Management Guide for CIFS” on the [NetApp Support site](#).

3 FPolicy Solution Architecture

The FPolicy solution consists of the following components, as shown in Figure 1:

- ONTAP FPolicy framework
- FPolicy application: iFinder5 elastic

Figure 1) FPolicy solution architecture.



FPolicy application software runs on an external server. The FPolicy framework is part of ONTAP software. The FPolicy framework connects to external FPolicy servers and sends notifications for certain file system events to the FPolicy servers when these events occur because of client access. The external FPolicy servers process the notifications and send responses back to the node.

3.1 Components of FPolicy Framework on ONTAP

The FPolicy framework on ONTAP includes the following components:

- **External engine.** This container manages external communications with the FPolicy server application, the iFinder5 elastic.
- **Events.** This container captures information about protocols and file operations monitored for the policy.
- **Policy.** This main container associates different constituents of the policy and provides the platform for policy management, such as policy enabling and disabling.
- **Scope.** This container defines the storage objects on which the policy acts: for example, volumes, sharers, exports, and file extensions.

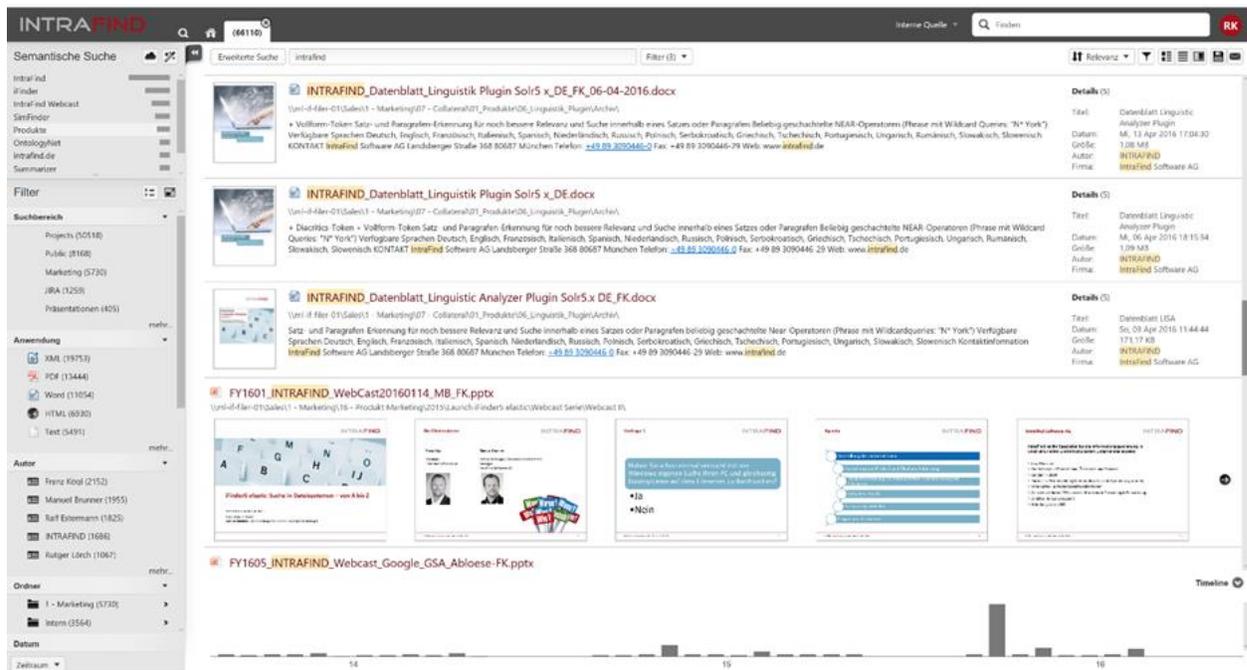
3.2 FPolicy Application Software: iFinder5 Elastic Edition for NetApp

The iFinder5 elastic is an enterprise search insight engine application. Through its connector framework, data from various sources is ingested into the search index of iFinder5 elastic, providing superfast and easy access to data and documents. These types of applications are also known as unified search, application-based search, insight engine, 360° view, and other names.

These applications all serve the same purpose: to allow users to find documents based on metadata and content, while still maintaining security on the documents: in other words, only documents for which the individual user is authorized to be part of the hitlist result from an individual search.

Figure 2 shows the iFinder5 elastic search application.

Figure 2) iFinder5 elastic search application.



Companies currently encounter many challenges with regard to information retrieval. Data sources and data diversity are increasing, but at the same time, employees want a flexible and independent process to gain access to relevant corporate information at any time. For many years, IntraFind has been dealing with these challenges. It now offers an intelligent solution with the product iFinder5 elastic.

Innovation is the driving force for successful companies to open new markets, satisfy customers, and generate revenue. Innovation takes place when historical findings are matched with current and future requirements. IntraFind supports its customers by always providing knowledge in the form of on-hand information, regardless of the application in which it was created or stored.

With the search solution, iFinder5 elastic enables companies to gain full insight into their enterprise information. Through extensive and powerful connectors, IntraFind taps important data sources (NetApp ONTAP, standard file systems, portals, websites, and so on), which allows companywide documents and information to be quickly searchable and checked for access rights. Content is processed, evaluated, and connected based on the latest text and content analytics methods.

iFinder5 elastic goes far beyond classic full-text search. At any given time, employees can gain insight into important and relevant corporate data, making it possible to combine existing knowledge with new findings and requirements and to create innovation.

The iFinder5 elastic edition for NetApp includes the FPolicy server that connects to ONTAP cluster-mode FPolicy service. Without the benefit of FPolicy, content from file shares is typically crawled. Crawling features periodically analyze the file share and identify new, updated, or deleted documents. However, the more documents that are stored on the file share, the more cumbersome this task becomes. Instead, with the FPolicy server now part of the iFinder5 elastic, events are processed almost in real time. These events include creating, deleting, and updating files as well as updates to security information. With the FPolicy server, these events are reflected almost instantly in the search index. This enhancement provides users and administrators comfort while searching for data and documents, knowing that the search index is always up to date.

iFinder5 elastic offers several key features. One of the most popular features is linguistic processing. The IntraFind linguistic processes content with a very powerful function set. As a result, users do not need to know how the original content was written. For example, if the word "policies" is included in a document, the IntraFind linguistic allows the user to also find the word "policy." While going beyond using stemmer, the IntraFind lemmatizer especially shows its strength while processing European languages, where stemming is not sufficient. For example, for the German word "laufen" (which means "to run"), all grammatical variations of this word would also be found, including the past tense ("lief"). This simple example demonstrates how stemming is not sufficient. With its state-of-the-art technology, the IntraFind linguistics solves these issues for almost 30 languages.

4 Installing and Configuring iFinder5 Elastic Edition for NetApp

4.1 Software Requirements and Installation

iFinder5 elastic is a Java-based application and runs on multiple platforms. If file security is required, Windows Server 2008 R2 and later is currently supported. If file security (Active Directory support) is not required, Ubuntu 1604, CentOS, and SLES 11 and later are supported.

For more information, see section **Error! Reference source not found.**, "**Error! Reference source not found.**"

4.2 Prerequisites

This section describes how to install and configure the IntraFind software.

Linux

The following tools must be installed on your Linux server:

- unzip
- gcc
- make
- `sudo apt-get install unzip make gcc`

User

The default user for the software installation is `intrafind`.

The default directory path for the software installation is `/home/intrafind/`.

You can provide a different user and different installation directory by passing the information by parameters or by entering the values in the installation dialog box.

IntraFind recommends creating and using the `intrafind` user name, which should be a sudoer for the installation. After successfully installing the software, the user no longer needs to use the sudoer and does not need a login bash. You can install the software and service start scripts with one attempt. Otherwise, you must use a sudo user or root to create the service scripts in `/etc/init.d/` by running a script for every service, which can be five or more.

Environment

Before you begin the installation process, edit the `/etc/environment` file, add three lines at the end of the file, and then change the path.

1. Add the following red text to the path:

Note: Your default path can look different.

```
PATH="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/usr/games:/usr/local/games:/home/intrafind/jdk/bin:/home/intrafind/tomcat8/bin"

JAVA_HOME=/home/intrafind/jdk
INTRAFIND_LICENSE=/home/intrafind/intrafind.lic
CATALINA_HOME=/home/intrafind/tomcat8
```

Note: You cannot use variables in `/etc/environment`.

2. In the `~/.profile` file, the path variable must also be defined. Append the following lines to the end of the file:

```
export JAVA_HOME=/home/intrafind/jdk
export INTRAFIND_LICENSE=/home/intrafind/intrafind.lic
export CATALINA_HOME=/home/intrafind/tomcat8
PATH="$JAVA_HOME/bin:$CATALINA_HOME/bin:$PATH"
```

3. Log out and log in again to use the new settings.

Note: `su - intrafind` is not sufficient to activate the settings; you must log in again with the current user.

4.3 Installing IntraFind Software

This section provides detailed information about installing IntraFind software.

Install with User Intrafind and Without sudo Permissions

If the user is not in the sudoers file, press Enter when the installation process asks for a password. After three attempts, you receive an error message; however, the software is installed.

```
Sorry, try again.
[sudo] password for intrafind:
sudo: 3 incorrect password attempts
```

Install with User Intrafind and sudo Permissions

It is assumed that your installation files are in the `/home/intrafind/install` directory.

Make sure that all `*.sh` and `*.bin` files are executable (`chmod +x *.bin`).

IntraFind recommends installing the services in subdirectories: for example, `/home/intrafind/services/` directory for the services and `/home/intrafind/apps` directory for the apps, which is the default.

The following installer packages are part of the default installation (single-node installation):

1. `if-meta-jdk-installer-1.8.0.91.bin`
2. `if-elasticsearch-installer-2.4.4.2.bin`
3. `if-sv-search-installer-2.4.4.4-40981.bin`
4. `if-sv-configstore-installer-1.8.0.10-38193.bin`
5. `if-app-admin-ui-installer-5.0.2.5-41081.bin`
6. `if-sv-access-installer-1.9.0.0-38193.bin`
7. `if-sv-converter-installer-3.6.1.3-41367.bin`
8. `if-sv-store-lucene4-installer-4.11.0.4-38193.bin`
9. `if-sv-thesaurus-installer-1.7.0.2-38193.bin`
10. `if-app-eindexer-soa-installer-1.4.1.2-43129.bin`
11. `if-app-sitemapcrawler-installer-1.0.2.2-44352.bin`
12. `apache-tomcat-8.0.14.zip`
13. `if-app-search-ui-5.0.2.5-SP3-44489.war`

The following subsections provide instructions for installing each package.

Install License

Copy the IntraFind product license file to the `/home/intrafind/intrafind.lic` directory.

Installer Package 1: Install jdk

```
intrafind@localhost:~/install$ ./if-meta-jdk-installer-x.x.x.x.bin
```

To install the `jdk` package, complete the following steps:

1. Create a symbolic link to the `jre` directory in your home directory.

```
intrafind@localhost:~/install$ cd ..
intrafind@localhost:~$ ln -s /home/intrafind/jdk/jre jre
```

2. After `jdk` is successfully installed, if you haven't already, edit the `~/.profile` file of the `intrafind` user and add the following three lines at the end of the file:

```
export JAVA_HOME=/home/intrafind/jdk
export INTRAFIND_LICENSE=/home/intrafind/intrafind.lic
PATH="$JAVA_HOME/bin:$CATALINA_HOME/bin:$PATH"
```

3. Log out and log back in to use the new settings.

Installer Package 2: Install Elasticsearch and Linguistics

To install the `elasticsearch` and `linguistics` package, complete the following steps:

1. Info: If you install multiple instances of `elasticsearch` in the same environment, change the cluster name. The cluster name can be passed by the option `-c`, but must be manually changed in `if-sv-search/config.cfg`.

```
intrafind@localhost:~/install$ ./if-elasticsearch-installer-x.x.x.x.bin
```

The following output is displayed:

```
IntraFind Installer 1.0
```

```

Installing "elasticsearch"

Summary of selected options:
INSTALLATION DIR : /home/intrafind/services/if-elasticsearch
JRE               : /home/intrafind/jdk
PORT              : 9200
USERNAME          : intrafind

Continue? [Y/n]

```

2. Press Enter to install elasticsearch.

Note: The linguistic features are installed with the elasticsearch base installation.

Installer Packages 3 Through 11: Install Services One by One

1. To install each service or app (packages 3 through 11), run the following installation script (here: search service):

```
intrafind@localhost:~$ ./if-sv-search-installer-1.7.5.5-40981.bin
```

After execution, the installation process starts:

```

IntraFind Installer 1.0

Installing "search"

Summary of selected options:
INSTALLATION DIR : /home/intrafind/services/if-sv-search
JRE               : /home/intrafind/jdk
LICENSE           : /home/intrafind/intrafind.lic
PORT              : 9605
USERNAME          : intrafind

Continue? [Y/n]

```

Note: To make changes, abort the installation process by entering `n` or by pressing `Ctrl + C`. Otherwise, process the installation by pressing `Enter` (`Y` is the default).

2. After the service or app is installed, you are prompted to enter the user's password to check sudo permissions for the services.

```

Extracting skeleton into /home/intrafind/services/if-sv-search
Extracting if-libs
Extracting 3rdparty-libs
Root privileges required to register daemon.
If your user is not in the sudoers file, please register the daemon manually
Do this by executing the installer with --no-daemon and exec afterwards
INSTDIR/bin/installDaemonNoPriv.sh with a user of the sudoers file or root
[sudo] password for intrafind:

```

If the user is not in the sudoers file, press `Enter`. After three attempts, you receive an error message, but the software is installed.

```

Sorry, try again.
[sudo] password for intrafind:
sudo: 3 incorrect password attempts

```

Otherwise, the service is now registered as available in `/etc/init.d/if-sv`.

3. If you do not sudo, manually install all services with the root or a sudo user by calling `/home/intrafind/services/if-sv-.../bin/installDaemonNoPriv.sh`.

```

=====
Installation was SUCCESSFUL.
Consider changing firewall configurations
Check install.log for more information
=====
intrafind@localhost:~/install$

```

Note: Be sure to change the firewall rules (at least port 8080 for the front-end user) if the service must be accessible from outside the firewall.

4. You can also edit `if-sv-search/config.cfg` and extend the security settings for the services (ports 9600 through 9620) with the host IP addresses or range, which can access the services.

```

## network security
net.server.pipeline.subnet-restriction.subnet : localhost
es.security.subnet                          : localhost

```

Installer Packages 12 and 13: Tomcat and Search Front End

```

cd ~/install
unzip apache-tomcat-8.0.14.zip -d /home/intrafind/
cd ..
mv apache-tomcat-8.0.14 tomcat8

```

```

cd /home/intrafind/tomcat8/bin
tar xzf commons-daemon-native.tar.gz
cd commons-daemon-1.0.15-native-src/unix
./configure --with-java=$JAVA_HOME
make

```

Note: You can typically ignore the following warnings:

```

cp jsvc /home/intrafind/tomcat8/bin/jsvc
nano /home/intrafind/tomcat8/bin/daemon.sh

```

The configuration file opens. Change the content by completing the following steps:

1. Change `$2`. Insert `intrafind` if you have only local users on your installation (for example, `DEV-Installation`) or insert a domain user if you have an AD environment.

```

--tomcat-user )
TOMCAT_USER="$2"
shift; shift;
continue
;;
test ".$TOMCAT_USER" = . && TOMCAT_USER=tomcat

```

2. Change `tomcat` at the end of the line. Insert `intrafind` if you have only local users on your installation (for example, `DEV-Installation`) or insert a domain user if you have an AD environment.

```

cd /etc/init.d
sudo ln -s /home/intrafind/tomcat8/bin/daemon.sh tomcat
cd /home/intrafind/bin/
chmod +x *.sh

```

4.4 Configure IntraFind Software

This section provides detailed information about configuring IntraFind software.

If-elasticsearch and linguistic

To add languages to your installation, edit the following file:

```
/home/intrafind/services/if-elasticsearch/elasticsearch/config/elasticsearch.yml
```

Look for the following line and add more languages:

Note: If the preceding configured languages are not licensed in the license file, the service does not start.

```
intrafind.search_plugin.languages: de, en
```

If-sv-access

If you are installing in an AD domain, start this service with a domain user, which can read all files. This approach is because the delivery of a document over the search front end is routed through this access service, and the document is read by this service.

If-sv-converter

If you are installing in an AD domain, start this service with a domain user, which can read all files. This approach is required so that the converter service can open and read the documents for the thumbnail and preview creation. The converter service creates images of the documents to provide the thumbnails and previews.

If you are using an e-indexer and want to index files on the local server, you must change the configuration of the converter service. If the e-indexer runs on a machine other than the converter service, do not change the setting.

Edit the following file:

```
/home/intrafind/services/if-sv-converter/config.cfg
```

Change the following line:

```
converter.allow.local.files: false
```

Note: By changing `false` to `true`, you can convert local files.

Broken Thumbnails and Previews for Webpages

If you receive the following error in the log file, install the font configs:

```
sudo apt-get install libfontconfig:
```

```
INFO|1812/0|Service if-sv-converter|17-05-02 15:19:21|2017-05-02 15:19:21,587 {convert 01812}
[ERROR] <PhantomRunner > |@2| error executing external command: ./phantomjs/phantomjs: error
while loading shared libraries: libfontconfig.so.1: cannot open shared object file: No such file
or directory
```

If-sv-search

1. Edit the following file:

```
/home/intrafind/services/if-sv-search/config.cfg
```

2. Scroll to the end of this file and uncomment four of the six iFinder lines. The result should look like this example:

```
## the following keys must be installed in the admin-ui for secure search to work. If not, please
set them here and uncomment them.
ifinder.ldap.user :
```

```
ifinder.ldap.user.password :
ifinder.ldap.server       :
ifinder.ldap.server.domain :
#ifinder.ui.roleinfo      :
#ifinder.ui.features      :
```

Note: If you are in a local environment (not an AD environment), leave the settings blank. If you are in an AD environment, provide the specific information. The password must be encrypted with the admin UI.

If you do not have an AD, complete the following step:

1. Look for `permissionsearch.permission-check` and make the following changes:
 - a. Comment the `aclRetriever` and the first `permCheck`.
 - b. Uncomment the second `permCheck`.

```
permissionsearch.permission-check:
var namespace = com.intrafind.ifinder.secure;

# sample configuration for a LDAP based secure search
#var aclRetriever = new namespace.ACLRetrieverWindowsLDAP()
#var permCheck    = new namespace.ACLBasedPermissionCheck(aclRetriever)

# allows all documents which have a field "_raw.aclallow" containing a value "S-1-1-0"
var permCheck    = new namespace.AllowPublicPermissionCheck()

# var termFilter  = new namespace.TermFilter("_facet.indexname", "Test")
var termFilter  = new namespace.MatchAllDocsFilter()
```

iFinder5

1. Copy the `if-app-search-ui... war` file to `/home/intrafind/tomcat8/webapps/iFinder5.war`.
2. Start and stop the Tomcat one time.

```
sudo /etc/init.d tomcat start
sudo /etc/init.d tomcat stop
```

3. Change the directory to the following:

```
/home/intrafind/tomcat8/webapps/iFinder5/WEB-INF/classes
```

4. Edit the `config.cfg` file.
5. Scroll to the bottom and uncomment the two lines. For example:

```
_config_store: com.intrafind.common.beans.Beans.of(com.intrafind.api.config.ConfigStore.class,
"http://localhost:9600/hessian/configstore")
common.init.config-incubating.refresh-secs: 30
```

Note: There must no space at the beginning of the lines.

Note: If you are not in an AD environment, activate the local users.

6. Look for the following path:

```
ifinder.loginfilter.enabled
```

7. Change the value to `true`.

```
ifinder.loginfilter.enabled      : true
```

if-app-sitemapcrawler

1. Edit the `/home/intrafind/services/if-sv-search/config.cfg` file and add the following lines before the `_config` line:

```

es.index-stats-connector:
  var client = Beans.of("es.client")
  new IndexES(client, "index-stats-connector", "es.template.default", "schema.stats")

es.search-stats-connector:
  var client = Beans.of("es.client")
  new SearchES(client, "index-stats-connector")

```

2. Edit the `/home/intrafind/services/if-sv-search/conf/wrapper.conf` file. Scroll to the bottom and uncomment or create the following entries:

Note: The numbers must be ascending and must not be duplicated.

```

#index-stats-connector
wrapper.app.parameter.21 = index-stats-connector
wrapper.app.parameter.22 = es.index-stats-connector

#search-stats-connector
wrapper.app.parameter.23 = search-stats-connector
wrapper.app.parameter.24 = es.search-stats-connector

```

3. Open the `/home/intrafind/apps/if-sv-sitemapcrawler/conf/config.cfg` config file of the sitemap crawler.
4. Change the directory of the temp directory to a directory to which the `intrafind` user has access.

```
sitemapindexer.tmpdir:/home/intrafind/apps/if-app-sitemapcrawler/tmp
```

5. Scroll down to:

```
sitemap.sitemapseeds: {
```

6. Change the following settings:

```

sitemapId -> Has to be unique in your installation
sitemapUrls -> Sitemap-URLs in double quotes and brackets
tenant -> change this to a valid tenant ai, if you are using tenants
indexName -> Change this to the desired Index Name which is used in the Frontend
baseurl -> internal data, has to be the Domain
rootPath -> internal data, has to be the Domain

```

For example:

```

"sitemapId"      : "INTRAFIND",
"sitemapUrls"   : ["https://www.intrafind.de/sitemap.xml"],
"ignoreRobots"  : "true",
"tenant"        : "public",
"#ignoreLastMod": "true",
"indexName"     : "Web",
"sitemapType"   : "xml",
"prio"          : "true",
"inactive"      : "false",
"excludes"      : [],
"crawlDelay"    : "1",
"baseurl"       : "https://www.intrafind.de",
"parentFolderField": "_facet.canonical",
"rootPath"      : "https://www.intrafind.de",

```

7. Scroll down to the service URLs.

```
sitemapindexer.converter.url:
```

8. For all five lines, verify that the ports are valid: port 9602 for the first line (converter) and port 9605 for the other four lines.
9. Save the config file.

4.5 Complete IntraFind Software Installation

To complete the installation, complete the following steps:

1. Restart the system.
2. Review the `elasticsearch` log at `/home/intrafind/services/if-elasticsearch/logs/if-cluster.log` and the search service log at `/home/intrafind/services/if-sv-search/logs/wrapper.log`.

4.6 Additional Information

To verify which services are running, run the following commands:

```
ps -ef|grep wrapper.jar
systemctl | grep if-
```

To test a service and view the log file directly, run the `./runConsole.sh` command in the bin directory of a service; this mode is the interactive one.

5 Configuring IntraFind FPolicy Server

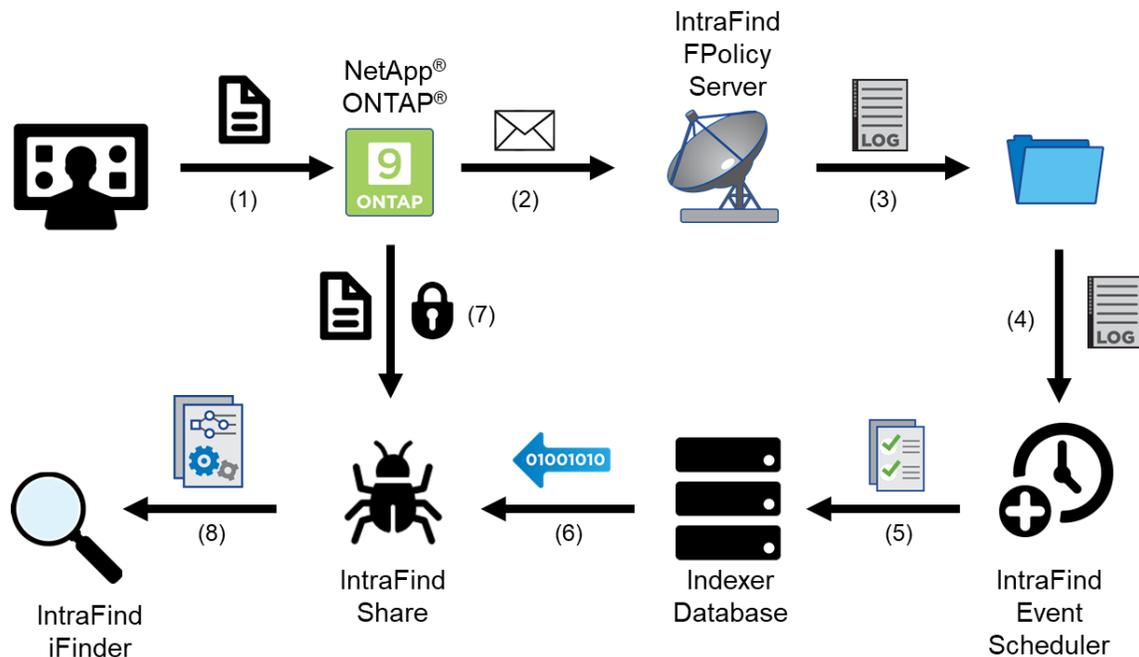
This section describes how to configure an IntraFind FPolicy server.

5.1 Service Description

The FPolicy service with the technical name `if-sv-fpolicy` implements the FPolicy interface for ONTAP. This service is used by IntraFind to index changes in a file system much faster than a pull mechanism (typical time of a push request is approximately 30 seconds) into the index. A standard pull mechanism would not allow this kind of performance.

Figure 3 demonstrates the procedures for processing FPolicy events.

Figure 3) FPolicy events.



1. A user modifies a file on the ONTAP file share.
2. ONTAP informs the IntraFind FPolicy server by issuing an event informing about the changes. This process takes 1ms.
3. The FPolicy server persists this information into the file system (event log). This process takes 1ms.
4. A second service permanently analyzes the event log. This process takes less than 2 seconds.
5. Insert the tasks based on the crawl strategy into the connector task list. This process takes 1 to 30 seconds, depending on the scheduling strategy and the workload.
6. The share indexer worker processes the tasks. This process takes approximately 10 seconds, depending on the workload.
7. The new version of the file, including its access rights from the share, is read. This process takes approximately 500ms.
8. The index is updated. This process takes approximately 5 seconds.

5.2 Prerequisites

In addition to the IntraFind FPolicy server, an installation of the share indexers (`if-app-indexer-share`) is required.

Between the ONTAP SVMs and the FPolicy servers, socket connections must be established (the default port is 9000). If required, firewall rules must be adapted accordingly.

Note: The software was tested and certified with clustered Data ONTAP 8.3 and ONTAP 9.0.

5.3 Update Behavior

The FPolicy server receives notifications of changes in the file system and transforms those into respective tasks for the share indexer. Therefore, latencies between changes and resulting index updates are drastically reduced (often from many hours to just a couple of seconds). The overall duration depends on the workload of the involved systems and varies depending on the workload.

While combining distributed and complex components, a loss of events cannot be guaranteed (for example, due to network failure), even if setup, planning, and sizing of the involved components are as exact and redundant as possible. The IntraFind architecture recognizes this issue and introduces additional pull events for the involved pull events. As a result, changes in a share are recognized by the index, even if a notification is missed. This approach might involve a significant higher latency.

Therefore, even after failure of the services, no complex operations are required to offer synchronism between share and index, because the parties automatically synchronize again.

Note: If, for any reason, the recovery of those changes is not indexed within 24 hours, contact your IntraFind representative.

5.4 Dependencies with Other IntraFind Components

For complete functionality, the following IntraFind service must be installed and configured:

- Share indexer

5.5 Install IntraFind FPolicy Infrastructure

Execute the installation package for `if-sv-fpolicy`.

The installed IntraFind FPolicy infrastructure consists of the following components:

- FPolicy server ONTAP
- Event scheduler

Note: To close the processing gap between ONTAP and iFinder5, you must install the `if-sv-indexer-share` share indexers.

Install FPolicy Server for ONTAP

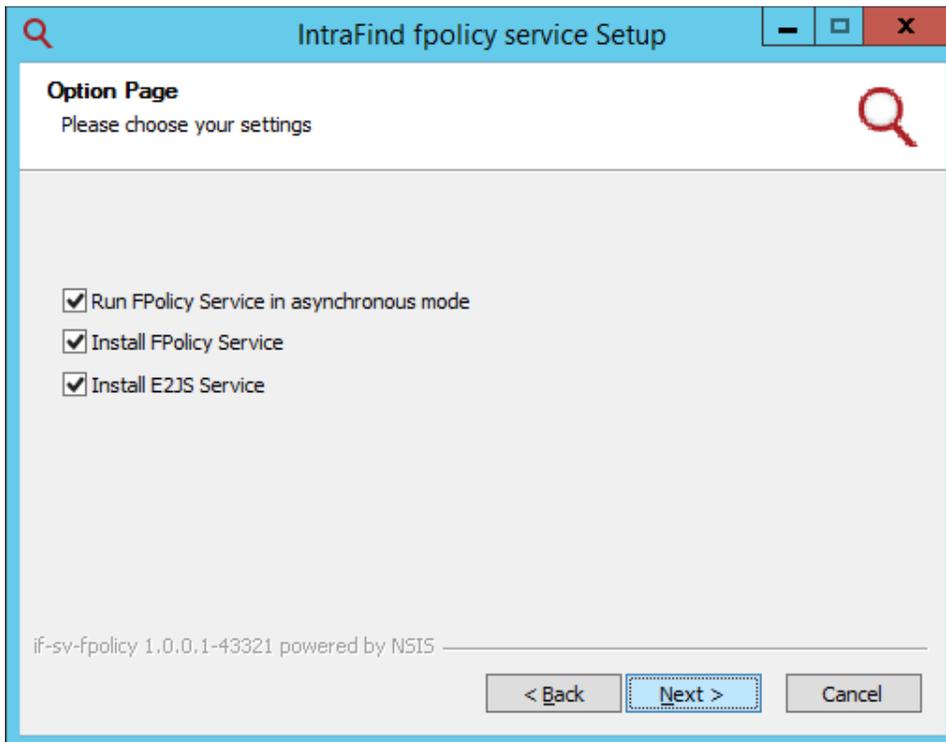
During the FPolicy server installation process, two important configurations should be acknowledged: FPolicy mode and FPolicy port.

Note: You can change these configurations later.

FPolicy Mode

An important configuration is the selection of the FPolicy mode: synchronous or asynchronous. Figure 4 shows the FPolicy mode options page.

Figure 4) FPolicy mode configuration.



Note: In synchronous mode, ONTAP waits on a confirmation of events. This wait time slightly increases the overall duration of the process, but the confirmation assures that all operations are processed completely.

Note: The recommended setting for FPolicy is asynchronous mode. For this mode, the wait time for confirmation of the event is almost eliminated, and the initiated process on the file share (for example, copy or move) is immediately completed. But, due to technical issues (for example, network failure), the event can get lost.

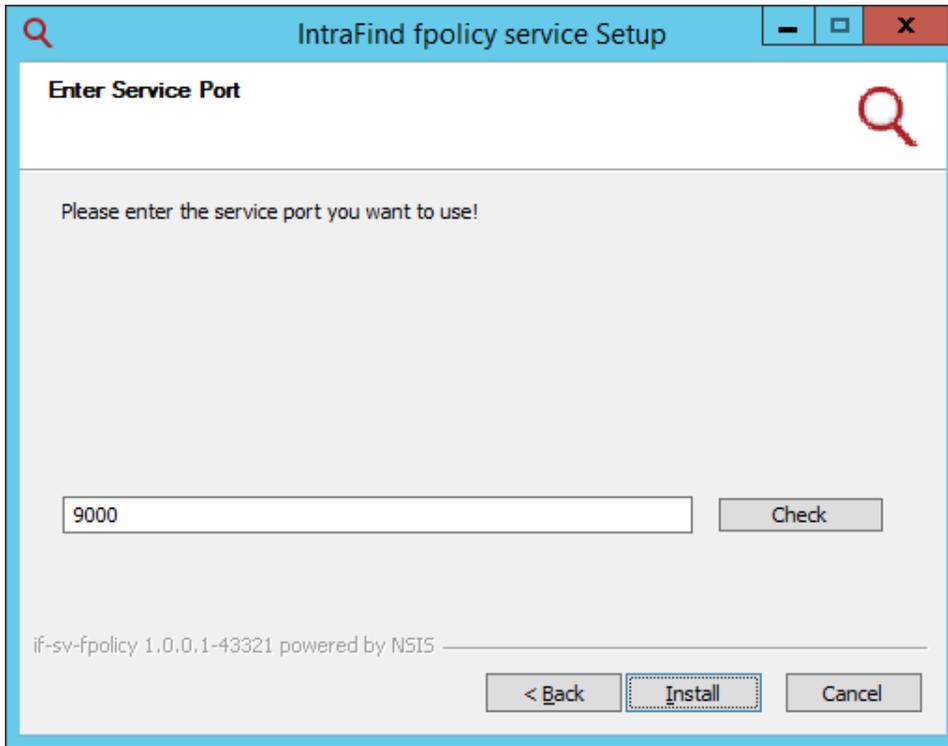
During the FPolicy configuration in ONTAP, the appropriate setting must be selected.

Note: If the settings for this feature differ between ONTAP SVM and the IntraFind FPolicy servers, a communication can be established first, but then disconnects.

FPolicy Port

The FPolicy port setting must match the configuration, done later, in ONTAP.

Figure 5) FPolicy port configuration.



Note: If the communication settings differ, ONTAP and FPolicy server cannot communicate with each other.

After installation, the following steps are required:

1. Configure the FPolicy server.
2. Configure FPolicy in ONTAP.

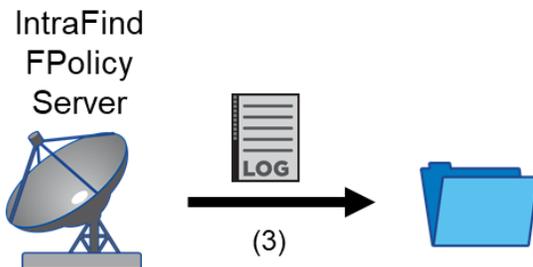
Detailed information about these additional steps is provided in the appendix.

6 Configuring IntraFind Policy Servers

This section describes the implementation by IntraFind software of the ONTAP FPolicy protocol for ONTAP.

To increase robustness and to provide a fast response time of the FPolicy servers, the FPolicy server only writes the events into an event log and does not conduct any further processing of the events.

Figure 6) IntraFind policy servers for ONTAP configuration.



6.1 Configure FPolicy Server Mode

The most important decision when configuring the FPolicy server for ONTAP is to decide between synchronous and asynchronous mode. In asynchronous mode, the system works without confirmations; for example, notifications are not confirmed. This mode increases the risk of losing notifications, but offers a significantly lower latency and is, therefore, recommended.

The mode and port can be changed in the `conf-FPOL/wrapper.conf` file. This change requires a restart of the service.

```
#service_base parameter
wrapper.app.parameter.1 = -p
wrapper.app.parameter.2 = 9000
wrapper.app.parameter.3 = --async
```

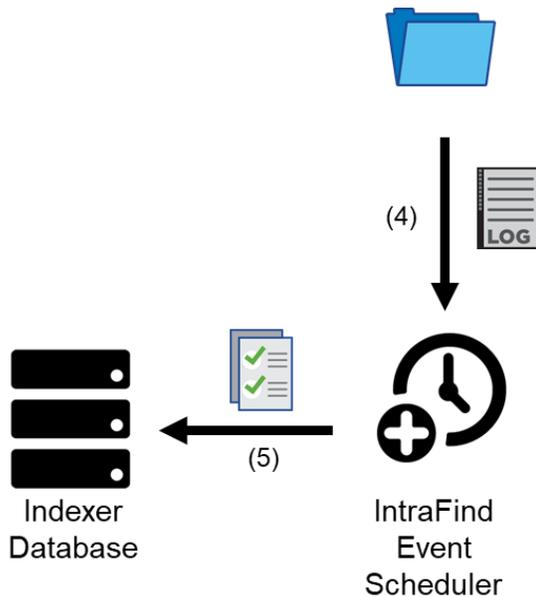
Note: If the parameter 3 is deleted, the system works in synchronous mode.

Note: ONTAP must also be set to synchronous mode.

7 Configuring Event Processing

The event processing permanently reads the event logs generated by the FPolicy server and transforms the included notifications in tasks for the share indexer: the database (JobStore) in which they are inserted.

Figure 7) Event processing configuration.



The processing of push events is possible only after these three components work successfully together:

- FPolicy server
- Event handling (event scheduler)
- Share indexer

The event logs are realized as text files and stored in the directory `fpolicy-events`, which is rotated hourly. For example, the `2017-02-17_13.xx.txt` file contains all entries from February 17, 2017, between the 13th and 14th hours.

The event logs contain a notification per line and consist of the event source, date, type (FILE or DIR), operation, SVM-UUID, volume-MSID, and a JSON-coded relative path.

```
2016-07-27 11:17:34,089 CDOT FILE SMB_WR df2189cb-f73c-11e5-a6bf-000c294fa338
2164258650 "/files/file.txt"
```

Note: This format is subject to change.

7.1 Configure Multiple Event Logs

It is necessary to process additional event logs in only selected situations: for example, if multiple SVMs are being connected and one is asynchronous and the other is asynchronous.

In this example, an additional source must be added to the configuration `wrapper.conf` from the event scheduler. The recommended configuration for this example is:

```
#service_base parameter
wrapper.app.parameter.1 = -p
wrapper.app.parameter.2 = fpolicy.events2jobstore
wrapper.app.parameter.3 = -s
wrapper.app.parameter.4 = -7200
wrapper.app.parameter.5 = -S
wrapper.app.parameter.6 = 2040-01-01
wrapper.app.parameter.7 = fpolicy-events/%.14sxx.txt##era=hour
wrapper.app.parameter.8 = fpolicy-events/sync-%.14sxx.txt##era=hour
```

7.2 Configure in the config.cfg File

In the `config.cfg` file, the settings are managed for the IntraFind services. The configuration file is in the `if-sv-fpolicy` folder.

```
indexer.job-store.index: http://localhost:9705/hessian/index-jobstore
indexer.job-store.search: http://localhost:9705/hessian/search-jobstore

fpolicy.table:
    1 , , share_7m , , NULL,
    05-95-3e, 3218, share_cdots, //filesrv, Predicates.find("/logs(/|$)")

fpolicy.processor.reconnect-bat:
    Beans.of("FPolicy.processor.factory-process")(Lists.of("cmd.exe", "/K", "reconnect.bat"))
```

Note: For a description of these options, see Table 1.

Note: Changes to the `config.cfg` file require appropriate knowledge. Make sure that while changing the configuration file, all existing formatting is observed; for example, make sure that indents are not changed.

Table 1) Options for configuring in the `config.cfg` file.

Key Value	Description
<code>indexer.job-store.index</code>	The key configures the index for the JobStore. In that database, information about the tasks is stored. In this case, events are pushed. This value needs to be configured only if the JobStore is not available locally.
<code>indexer.job-store.search</code>	This value sets the endpoint for searching of the JobStore. This value needs to be configured only if the JobStore is not available locally.
<code>fpolicy.table</code>	This value is described in section "Configure <code>fpolicy.table</code> ."
<code>fpolicy.processor.reconnect-bat</code>	The ONTAP FPolicy server compares the incoming connections with the FPolicy table. It is treated as an error if not all configured connections can be established.

Key Value	Description
	<p>In this example, the FPolicy server can be configured in such way that an external batch file is processed (configured in this example: <code>reconnect.bat</code>). For example, this batch file can send an email to the administrator or even restart the connections through the SSH command.</p> <p>For each case, a log entry is created. If not, every connection could be established.</p> <p>For example:</p> <pre>2017-02-16 16:24:37,944 {08112} [WARN] <CheckUuids> sl missing vserver uuid(s) '[df2189cb-f73c-11e5-a6bf-000c294fa338]'</pre>

Configure `fpolicy.table`

The `fpolicy.table` is the essential configuration. It creates the link between the ONTAP SVM UUID, the volume MSID for the share configuration, and the associated connector context.

Note: Table 2 cannot be automatically determined; it must be set manually.

These entries can be retrieved through commands in the NetApp console. Alternatively, the values can be retrieved from the event log and, including the full path, be added to the share.

Note: The SVM is referred to as Vserver in the GUI and CLI.

Table 2) Options for configuring in the `fpolicy.table` file.

Key Value	Description
<code>fpolicy.table</code>	<p>This value is a comma-separated table, which is created based on the following scheme:</p> <pre><VServer-UUID>, <Volume-MSID>, <Context-Id>, <networkpath>, <Blacklist Predicate></pre> <p>This table can contain any number of those tuples:</p> <ul style="list-style-type: none"> • FPolicy events for ONTAP contain, if changing files, the UUID of the SVMs, the MSID of the respective volumes, and a relative path. With this table, a relative path can be converted to an absolute path. • For each event, the context ID assigns a connector context. • With the blacklist predicate, push events can be ignored. For example, some directories might show frequent changes, but do not need to be processed simultaneously (for example, processing of log files). • The configuration of such expressions is part of the documentation of the share indexer (<code>if-app-indexer-share</code>). • If you need to switch off the filtering completely, use the reserved key word <code>NULL</code>.

Configure `fpolicy.table` for ONTAP

Essential settings for the functioning of the system are performed through the `fpolicy.table` parameter.

For a simple exploratory configuration of the `fpolicy.table`, complete the following steps:

1. Start the IntraFind FPolicy server.
2. Configure FPolicy in ONTAP on an SVM.
3. On the share, a temporary file with a specific and easy to remember file name is created: for example, `\\ml-if-vs01\bigfiles\files\simple-configuration.temp.txt`.
4. An entry like the following is in the event log:

```
2016-07-27 11:17:34,089 CCDOT FILE SMB_WR 925eb706-5d5e-11e6-a4c8-000c294636bb
2158102436 "/files/simple-configuration.temp.txt"
```

5. Column 4 contains the UUID (925eb...36bb), and column 5 contains the MSID (2158102436).
6. The prefix for changing the relative path into an absolute path is //ml-if-vs01/bigfiles.

The parameters also can be determined systematically up front. This concept is explained in more detail in the following sections.

Determine SVM UUID

Determine the SVM UUID by running the `vserver show` command, as shown in the following example:

Note: Replace `ml-if-vs01` with the name of the appropriate SVM.

Note: The SVM is referred to as Vserver in the GUI and CLI.

```
cluster1::> vserver show -vserver ml-if-vs01

          Vserver: ml-if-vs01
          Vserver Type: data
          Vserver Subtype: default
          Vserver UUID: 925eb706-5d5e-11e6-a4c8-000c294636bb
          Root Volume: ml_vs01_root
          Aggregate: clu1_aggr1
          NIS Domain: -
          Root Volume Security Style: ntfs
          LDAP Client: -
          Default Volume Language Code: C.UTF-8
```

Determine Volume MSID

Determine the volume MSID by running the `volume show` command:

Note: The SVM and the volume name must be specified based on your settings.

Note: The SVM is referred to as Vserver in the GUI and CLI.

```
cluster1::> volume show -vserver ml-if-vs01 -volume data01

          Vserver Name: ml-if-vs01
          Volume Name: data01
          Aggregate Name: clu1_aggr1
          Volume Size: 5GB
          Volume Data Set ID: 1028
          Volume Master Data Set ID: 2158102436
          Volume State: online
          Volume Type: RW
```

Configure Context ID

If you are connecting one share, the settings should be `share`. If you are connecting multiple shares, the settings for the context description must be correct.

Validate the configuration of the contexts of the share indexer (`if-app-indexer-share`) and determine the name of the context for the respective share. It is assumed that two shares will be indexed with the following configuration; the setting should be `share_b`.

```
indexer.context.share_a.parent: indexer.context.share
indexer.context.share_a.seeds: //share/root1, //share-a/root2
indexer.context.share_a.root: //share

indexer.context.share_b.parent: indexer.context.share
indexer.context.share_b.seeds: //ml-if-vs01/bigfiles
indexer.context.share_b.root: //ml-if-vs01
```

Determining the absolute Path

The absolute path for the share must be determined. In this example, the path is `//ml-if-vs01/bigfiles`. Proceed as follows for reviewing the absolute path:

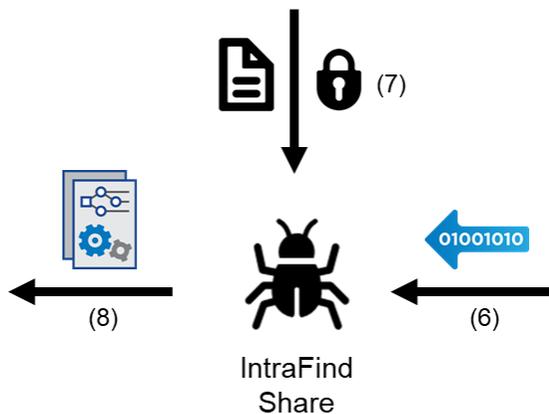
```
cluster1::> cifs share show
```

Vserver	Share	Path	Properties	Comment	ACL
ml-if-vs01	admin\$	/	browsable	-	-
ml-if-vs01	bigfiles	/data01	oplocks browsable showsnapshot changenotify	-	Everyone / Full Control

8 Reconfiguring Share Indexer

The share indexer is responsible for the actual retrieval of the files and their access control list (ACL). It is also responsible for the indexing into the iFinder index.

Figure 8) IntraFind share indexer.



In the standard configuration, the share indexer is not configured to process push events. Instead, it just processes the regular pull crawl. To change the worker configuration, complete the following steps:

1. Open the `if-app-indexer-share-?/conf/wrapper.conf` file.
2. Change the current configuration from this example:

```
wrapper.app.parameter.1 = --context  
wrapper.app.parameter.2 = share  
wrapper.app.parameter.3 = --profile  
wrapper.app.parameter.4 = share
```

To this example:

```
wrapper.app.parameter.1 = --context  
wrapper.app.parameter.2 = share  
wrapper.app.parameter.3 = --profile  
wrapper.app.parameter.4 = share-push
```

With this changed profile, the crawler (share indexer) can now process push events (in addition to the standard file crawl operations).

3. Restart the services.

9 Miscellaneous

9.1 Scaling and Failover

Many SVMs can communicate with the IntraFind ONTAP FPolicy server. Also, several hundred events per second can be processed.

For failover, multiple FPolicy servers operate in parallel. To accomplish this parallel operation, install additional instances with the same configuration settings. During the ONTAP configuration, specify the IP address of the additional servers through the parameter `secondary servers`.

Note: For more information about scaling, contact your IntraFind representative.

9.2 Memory Settings

To change the assigned memory for the service, use the parameters `-XMS` (initial memory allocation) and `-XMX` (maximum memory allocation). These settings are performed in `wrapper.conf` in the `conf` subdirectory.

Note: For more information about memory settings, contact your IntraFind representative.

9.3 Licensing

The connector requires a license file provided by IntraFind. For more information, contact your IntraFind representative.

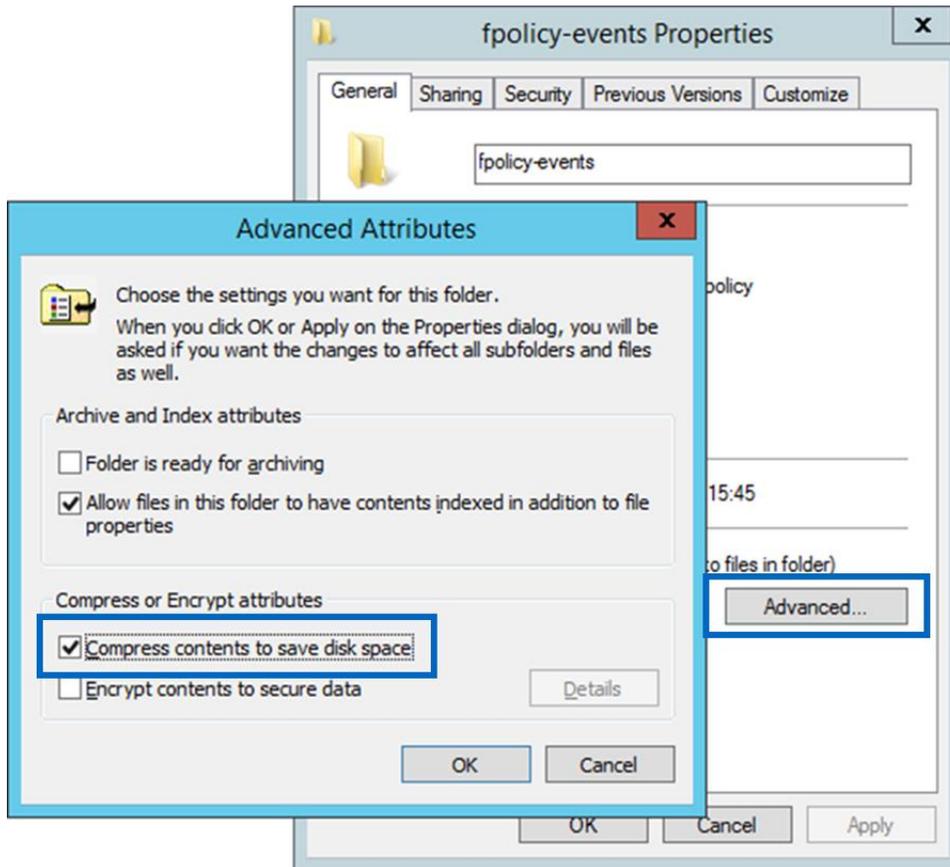
9.4 Logging

By default, the log files are stored in the subdirectory `logs` of the installation directory.

9.5 Compression of Event Logs

For Windows, it might be beneficial to compress the contents of the `fpolicy-events` directory, as shown in Figure 9.

Figure 9) Compression of event logs.



10 Configuring FPolicy on ONTAP

This section provides instructions for configuring FPolicy policies for NetApp file servers operating in cluster mode.

The FPolicy structure is defined as follows:

- **Event.** Defines which operations and protocol types the FPolicy audits.
- **External engine.** Defines the endpoint to which the FPolicy sends notification information.
- **Policy.** The aggregation of events policy, external engine, and scope.
- **Scope.** Defines the volumes, shares, export policies, and file extensions to which the FPolicy policy applies. You can also include and exclude all relevant filters.

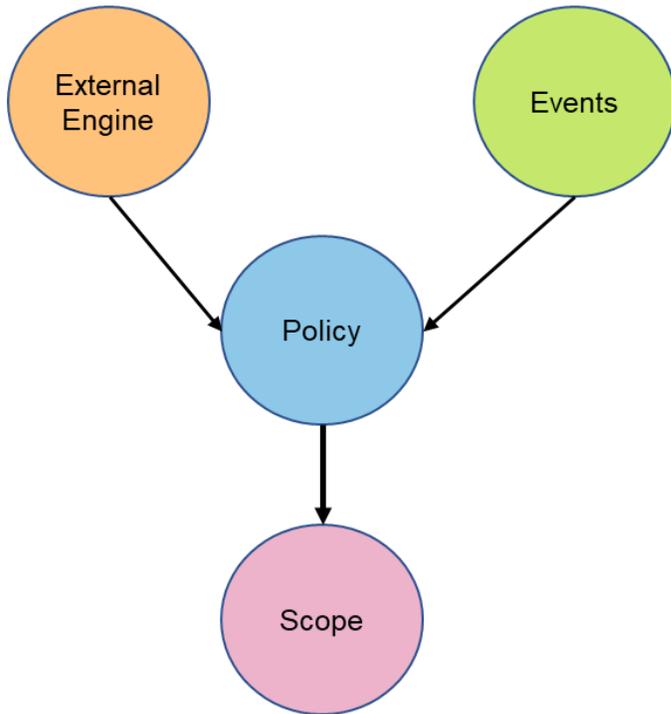
The FPolicy configuration requirements include:

- The shares must reside on the volume monitored for CIFS events.
- The export policy must be created on and applied to the volume monitored for NFS events.

10.1 FPolicy Configuration Workflow

Figure 10 shows a workflow that specifies the high-level steps that you must perform to configure and manage FPolicy.

Figure 10) FPolicy configuration workflow.



The workflow for creating a resident policy is as shown Figure 10. You should create the external engine and event before you create a policy. After the policy is defined, you must associate it with a scope.

After the scope is created, you must enable a policy with the sequence number. The sequence number helps to define the priority of the policy in a multipolicy environment. The sequence of 1 has the highest priority, and 10 has the lowest.

10.2 Create an FPolicy Event

To connect to a NetApp file server operating in cluster mode, you must configure an FPolicy event. You must be a user with the `vsadmin` role and have a user name that is associated with the `ontapi` application. The order in which you create an FPolicy event is important.

To create an FPolicy event using TCP, complete the following steps:

1. Connect to NetApp Data ONTAP management.
2. To create an FPolicy event object, run the following command:

```
fpolicy policy event create -vserver <vserver name> -event-name intrafind -protocol cifs -file-operations delete_dir,create_dir,rename_dir,delete,setattr,open,close -filters close-with_modification,open-with-delete-intent
```

Where:

- `-vserver` is the name of the SVM (formerly known as Vserver) on which you want to create an FPolicy external engine.
- `-event-name` is the name of the FPolicy event that you want to create.
- `-file-operations` is the file operations for the FPolicy event.

The values are:

- Create: `create_dir`

- Delete: delete_dir
- Read
- Close
- Rename: rename_dir
- -protocol is the name of the protocol for which the event is created. The protocol value is cifs.
- -filters specifies the filters used with a given file operation for the protocol specified in the -protocol parameter: for example, first-read, close-with-modification.

3. To view and verify an FPolicy event object, run the following command:

```
fpolicy policy event show <event name> -instance
```

10.3 Create an FPolicy External Engine

To create an FPolicy external engine, run the following command:

Notes:

- Make sure that the firewall does not prohibit communications.
- Replace the placeholders accordingly.
- Usually, the setting selects the asynchronous mode, which reduces the delay for receiving a confirmation to a minimum.
- Set the FPolicy server according to the FPolicy either synchronous or asynchronous. For more information, see section 5 "Configuring IntraFind FPolicy Server."

```
fpolicy policy external-engine create -vserver <Vserver Name> -engine-name intrafind -primary servers < IP address of FPolicy server> -port <port no> -extern-engine-type asynchronous -ssl-option no-auth
```

Where:

- -vserver is the name of the SVM on which you want to create an FPolicy external engine.
- -engine-name is the name of the external engine that you want to create.
- -primary-servers is the IP addresses for the primary FPolicy servers.
- -port is the port number for the FPolicy service.
- -extern-engine-type is the type of external engine. Only asynchronous type is supported.
- -ssl-option is the SSL option for external communication with the FPolicy server.

Possible values include:

- server-auth provides for server authentication.
- mutual-auth provides both server and NetApp authentication.

To view and verify the external engine, run the following command:

```
FPolicy policy external-engine show
```

10.4 Create an FPolicy Policy

To create an FPolicy policy, run the following command:

Notes:

- NFS events are not monitored.
- If time-sensitive processing is required, set the -is-mandatory configuration to true.

- If `-is-mandatory` is set, file operations are blocked. If the FPolicy server cannot be reached, working with the share is not possible.

```
fpolicy policy create -vserver <vserver> -policy-name intrafind -events intrafind -engine intrafind -is-mandatory false
```

Where:

- `-vserver` is the name of the SVM on which you want to create an FPolicy external engine.
- `-policy-name` is the name of the FPolicy policy that you want to create.
- `-events` is a list of events to monitor for the FPolicy policy.
- `-engine` is the name of the external engine that you want to create.
- `-is-mandatory` determines whether the FPolicy object is mandatory.

To view the policy, run the following command:

```
fpolicy policy show
```

10.5 Create an FPolicy Scope

To create an FPolicy scope, run the following commands:

Note: The rules should be valid for all volumes and FPolicy policies.

```
fpolicy policy scope create -vserver <Vserver Name>
-policy-name intrafind -volumes-to-include "*" -
export-policies-to-include "*"

```

Where:

- `-vserver` is the name of the SVM on which you want to create an FPolicy external engine.
- `-policy-name` is the name of the FPolicy policy that you want to create.
- `-volumes-to-include` is a comma-separated list of volumes to be monitored.
- `-export-policies-to-include` is a comma-separated list of export policies for monitoring file access. Wildcards are supported.

To view the FPolicy scope, run the following command:

```
fpolicy policy scope show -vserver <Vserver Name> - policy-name <Policy name>
```

10.6 Enable an FPolicy Policy

Starting the probe service enables the new FPolicy policy. The following command is for reference only:

Note: The FPolicy server must be live before you can activate the FPolicy policy.

```
fpolicy enable -vserver <Vserver Name> -policy-name intrafind -sequence-number <seq no>
```

11 Configuring Security Login for FPolicy Server

The configuration of the FPolicy service is performed through the ONTAP admin shell.

12 ONTAP Best Practices

NetApp recommends the FPolicy best practices described in this section for server hardware, operating systems, patches, and so on.

12.1 Policy Configuration

FPolicy External Engine for SVM

Providing additional security comes with a performance cost. Enabling SSL communication has a performance effect on CIFS.

FPolicy Events for SVM

Monitoring file operations has an effect on the overall user experience. In fact, filtering unwanted file operations on the storage side improves the overall user experience. NetApp recommends monitoring the minimum number of file operations and enabling the maximum number of filters without breaking the use case. The CIFS home directory environment has a high percentage of `getattr`, `read`, `write`, `open`, and `close` operations. NetApp recommends using filters for these operations. For recommended filters, refer to the section "Create an FPolicy Event."

FPolicy Scope for SVM

Restrain the scope of the policies to relevant storage objects, such as shares, volumes, and exports, rather than enabling them throughout the SVM. NetApp recommends checking directory extensions. If `is-file-extension-check-on-directories-enabled` is set to true, directory objects are subjected to the same extension checks as regular files.

12.2 Network Configuration

Network connectivity between the FPolicy server and the controller should be of low latency. NetApp recommends separating FPolicy traffic from client traffic by using a private network.

Note: In a scenario where the LIF for FPolicy traffic is configured on a different port than the LIF for client traffic, the FPolicy LIF might fail over to another node due to a port failure. This situation can make the FPolicy server not reachable from the node and can also make the FPolicy notifications for the file operations on the node fail.

Make sure that the FPolicy server is reachable through at least one LIF on the node to process FPolicy requests for the file operations performed on that node.

12.3 Hardware Configuration

The FPolicy server can be on either a physical server or a virtual server. If the FPolicy server is in a virtual environment, make sure to allocate dedicated resources (CPU, network, and memory) to the virtual server.

12.4 Multiple Policy Configuration

The FPolicy policy for native blocking has the highest priority, respective of the sequence number. Decision-altering policies have a higher priority than others. Policy priority depends on use cases. To determine the appropriate priority, NetApp recommends working with partners.

12.5 Managing FPolicy Workflow and Dependency on Other Technologies

NetApp recommends disabling an FPolicy policy before making any configuration changes. For example, if you want to add or modify an IP address in the external engine configured for the enabled policy, then first disable the policy.

If you configure FPolicy to monitor NetApp FlexCache® volumes, NetApp recommends that you do not configure FPolicy to monitor `read` and `getattr` file operations. Monitoring these operations in ONTAP requires the retrieval of inode-to-path (I2P) data. Because I2P data cannot be retrieved from FlexCache

volumes, it must be retrieved from the origin volume. Therefore, monitoring these operations eliminates the performance benefits that FlexCache can provide.

When both FPolicy and an offbox antivirus (AV) solution are deployed, the AV solution receives notifications first. FPolicy processing starts only after AV scanning is complete. A slow AV scanner could affect overall performance, so AV solutions must be sized properly.

12.6 Sizing Considerations

FPolicy performs inline monitoring of CIFS operations, sends notifications to the external server, and waits for a response, depending on the mode of external engine communication (synchronous or asynchronous). This process affects the performance of CIFS access and CPU resources. To mitigate any issues, NetApp recommends assessing and sizing the environment before enabling FPolicy. Performance is affected by the number of users, workload characteristics such as operations per user, data size, and network latency.

13 iFinder5 Elastic Edition for NetApp Best Practices

In every enterprise search project there is low-hanging fruit. All projects start in the sales cycle, determining what the customer wants or needs to achieve. Key questions to ask at this starting point include:

- Secure search: does the search software need to reflect the authorization information?
- Which sources need to be connected?
- Which metadata or entities are supposed to be indexed?

With iFinder5 elastic edition for NetApp, the file share source is a given. Indexing a file share provides quick success in a project.

14 Troubleshooting

14.1 Problem 1: The FPolicy Server Is Disconnected

Potential solution: If the server is not connected, try to connect it by running the `engine-connect` command. Look for the reason for FPolicy server disconnection by running the `show-engine -instance` command and take appropriate action.

Example command:

```
1. fpolicy show-engine
2. fpolicy engine-connect -node <node name> -vserver <vserver name> -policy intrafind -server <ip
address of FPolicy server>
3. fpolicy show-engine -instance
```

14.2 Problem 2: The FPolicy Server Does Not Connect

Precheck: Verify that the SVM has a data LIF through which the FPolicy server is reachable.

Example command:

```
network interface show
network ping -lif <vserver_data_lif> -destination <fpolicy server IP address> -lif- owner
<vserver_name>.
```

Potential cause number 1: There are issues with routing.

Potential solution: Check the routing table entries by running the `routing-groups route show` command to verify whether a route is available for the SVM. If not, add a route by running the `routing-groups route create` command.

Example command:

```
routing-groups route create -vserver <vserver name> -routing-group d10.X.0.0/18 -destination 0.0.0.0/0 -gateway 10.X.X.X
```

Potential cause number 2: The FPolicy server is not listening on the port specified.

Potential solution: Look for the log entry `connect failed. errno = 61 Establish TCP connection returned error` in the FPolicy user space log file (`fpolicy.log`). Then, check the port on which the FPolicy server is listening and modify the external engine configuration to use the same port.

Example command:

```
fpolicy policy external-engine modify -vserver <vserver name> -engine-name intrafind -port <tcp port no>
```

Potential cause number 3: The security options for the external engine are not the same as for the FPolicy server.

Potential solution: Run the `fpolicy policy external-engine show -instance` command. If the FPolicy server is using SSL, then the field `SSL Option for External Communication` is either `mutual-auth` or `server-auth`.

Also, check the fields `FQDN` or `Custom Common Name`, `Serial Number of Certificate`, and `Certificate Authority` to verify that the certificates are properly configured.

To correct this problem, modify `ssl-auth` to `no-auth` if the FPolicy server is not using SSL. Otherwise, use `mutual-auth/server-auth`, depending upon the level of security needed.

Example command:

```
fpolicy policy external-engine modify -vserver <vserver name> -engine-name intrafind -primary-servers <ip address> -port <tcp port no> -ssl-option no-auth
```

Potential cause number 4: The dedicated LIF for the FPolicy traffic failed over to a different node.

Potential solution: Make sure that the FPolicy server is reachable through at least one LIF for that SVM on the node to process FPolicy requests for the file operations performed on that node.

Example command:

```
network interface show
fpolicy show-engine
```

14.3 Problem 3: The External Engine Is Not Native for the Policy

Potential solution: Run the `fpolicy policy show` command to verify whether the `Engine` field is set to `Native`. Then create an external engine for the FPolicy server and attach it to the policy.

Example command:

```
fpolicy policy external-engine create
fpolicy policy modify
```

14.4 Problem 4: Notifications Are Not Being Received for the File Operations on Volume, Share, and Export

Potential cause: The FPolicy policy scope is not set properly.

Potential solution: Run the `fpolicy policy scope show` command to verify whether the scope contains the `vol/share` on which the `ops` are performed. Then, create or modify the scope for the policy to add the necessary volume, share, or export.

Example command:

```
fpolicy policy scope create/modify
```

15 Performance Monitoring

FPolicy is a notification-based system. Notifications are sent to an external server for processing, and a response is then sent back to the ONTAP software. This roundtrip process adds latency to client access.

Monitoring the performance counters on the FPolicy server and ONTAP helps to identify bottlenecks in the solution and allows you to tune the parameters necessary for an optimal solution. For example, an increase in FPolicy latency has a cascading effect on CIFS latency. Therefore, you should monitor both workload (CIFS) and FPolicy latency. Also, you can use quality-of-service policies in ONTAP to set up a workload for each volume or SVM that is enabled for FPolicy.

NetApp recommends displaying workload statistics by running the `statistics show -object workload` command. NetApp also recommends that you monitor the average read and write latencies, the total number of operations, and the read and write counters. You can also use the ONTAP FPolicy counters described in this section to monitor the performance of FPolicy subsystems.

Note: To collect statistics related to FPolicy, you must be in diagnostic mode.

15.1 Collect and Display FPolicy Counters

To collect FPolicy counters, run the following commands:

```
statistics start -object fpolicy -instance <instance name> -sample-id <id>
statistics start -object fpolicy_policy -instance <instance name> -sample-id <id>
```

To display FPolicy counters, run the following commands:

```
statistics show -object fpolicy -instance <instance_name> -sample-id <id>
statistics show -object fpolicy_server -instance <instance_name> -sample-id <id>
```

15.2 Counters to Monitor

Table 3 lists the FPolicy counters that can be monitored.

Table 3) FPolicy counters.

Counters	Description
<code>max_request_latency</code>	Maximum screen requests latency
<code>outstanding_requests</code>	Total number of screen requests in process
<code>request_latency_hist</code>	Histogram of latency for screen requests
<code>requests_dispatched_rate</code>	Number of screen requests dispatched per second
<code>requests_received_rate</code>	Number of screen requests received per second
<code>max_request_latency</code>	Maximum latency for a screen request
<code>outstanding_requests</code>	Total number of screen requests waiting for response

Counters	Description
request_latency	Average latency for screen request
request_latency_hist	Histogram of latency for screen requests
request_sent_rate	Number of screen requests sent to FPolicy server per second
response_received_rate	Number of screen responses received from FPolicy server per second

15.3 Performance Monitoring for iFinder5 Elastic Edition for NetApp

As described in the appendix, various log files are maintained. These log files also contain statistical data (number of items processed over time) that can be used to compare real-life performance against planned performance.

In future versions of the IntraFind administration, review of the worker queue of the underlying processing engine will be possible.

Appendix: FPolicy Server Cheat Sheets

FPolicy Server Installation Cheat Sheet

wrapper.conf (if-sv-indexer-share)

```
wrapper.app.parameter.1 = --context
wrapper.app.parameter.2 = share
wrapper.app.parameter.3 = --profile
wrapper.app.parameter.4 = share-push
```

config.cfg

```
fpolicy.table:
  <VServer-UUID>, <Volume-MSID>, share, //<host>/<share>, NULL, # cDOT
  1, , share, , NULL # 7mode
```

FPolicy for ONTAP Configuration

```
fpolicy policy event create -vserver <vserver> -event-name intrafind -protocol cifs -file-
operations delete_dir, rename_dir, create_dir, rename, delete, setattr, open, close -filters
close_with_modification, open_with_delete_intent
```

```
fpolicy policy external-engine create -vserver <vserver> -engine-name intrafind -primary-servers
<ip-fpolicy-server> -port 9000 -extern-engine-type asynchronous -ssl-option no-auth
```

```
fpolicy policy create -vserver <vserver> -policy-name intrafind -events intrafind -engine
intrafind -is-mandatory false
```

```
fpolicy policy scope create -vserver <vserver> -policy-name intrafind -volumes-to-include "*" -
export-policies-to-include "*"

```

```
fpolicy enable -policy-name intrafind -sequence-number 1 -vserver <vserver>
```

FPolicy for 7-Mode Configuration

```
fpolicy create intrafind screen
```

```
fpolicy options intrafind cifs_setattr on
```

```
fpolicy enable intrafind
```

Event2LogStore: Additional Event Logs for 7-Mode

```
wrapper.app.parameter.7 = fpolicy-events/%.14sxx.txt##era=hour  
wrapper.app.parameter.8 = fpolicy-events-7mode/%.14sxx.txt##era=hour  
wrapper.app.parameter.9 = fpolicy-events-7mode/2-%.14sxx.txt##era=hour
```

FPolicy Server Troubleshooting Cheat Sheet

Table 4) FPolicy server troubleshooting.

Symptom	Explanation/Solution
Error 0x5 reported while trying	FPolicy server (7-Mode) does not run with a user that has backup operator rights.
There is no share configured for <UUID>:<MSID>	There is a missing configuration in <code>fpolicy.table</code> .
Cannot create context <code>share_d</code>	The context <code>share_d</code> is not correctly configured. There is a possible missing class or parent key.
Missing SVM uuid(s) '[<UUID>]	The SVM does not connect with the FPolicy server. Solution: Disable FPolicy: <pre>fpolicy disable -policy-name intrafind -vserver <vserver></pre> Enable it again (see preceding).
No 7-Mode events are recorded	Was the <code>fpolicy-events-7mode</code> directory already created?
Events are generated (see event log), but the files are not crawled	Are the key seeds in context correctly configured? ONTAP: <ul style="list-style-type: none">Were <code>seeds</code> and <code>fpolicy.table</code> consistently used for the same share for the IP or the host name? 7-Mode: <ul style="list-style-type: none">Do the paths of the events (host name) match to <code>seeds</code>?Do the share indexer instances run under a user with sufficient rights?Were all event directories (or patterns) also configured for the event scheduler?

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- ONTAP 9 Documentation Center
<http://docs.netapp.com/ontap-9/index.jsp>

Version History

As an option, use the NetApp Table style to create a Version History table. Do not add a table number or caption.

Version	Date	Document Version History
Version 1.0	April 2018	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2018 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.

TR-4670-0418