File Access Manager Disaster Recovery

Version: 8.4

Revised: March 27, 2023

Copyright and Trademark Notices

Copyright © 2023 SailPoint Technologies, Inc. All Rights Reserved.

All logos, text, content, including underlying HTML code, designs, and graphics used and/or depicted on these written materials or in this Internet website are protected under United States and international copyright and trademark laws and treaties, and may not be used or reproduced without the prior express written permission of SailPoint Technologies, Inc.

"SailPoint," "SailPoint & Design," "SailPoint Technologies & Design," "Identity Cube," "Identity IQ," "IdentityAI," "IdentityNow," "SailPoint Predictive Identity" and "SecurityIQ" are registered trademarks of SailPoint Technologies, Inc. None of the foregoing marks may be used without the prior express written permission of SailPoint Technologies, Inc. All other trademarks shown herein are owned by the respective companies or persons indicated.

SailPoint Technologies, Inc. makes no warranty of any kind with regard to this manual or the information included therein, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. SailPoint Technologies shall not be liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Patents Notice. https://www.sailpoint.com/patents

Restricted Rights Legend. All rights are reserved. No part of this document may be published, distributed, reproduced, publicly displayed, used to create derivative works, or translated to another language, without the prior written consent of SailPoint Technologies. The information contained in this document is subject to change without notice.

Use, duplication or disclosure by the U.S. 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 252.227-7013 for DOD agencies, and subparagraphs (c)(1) and (c)(2) of the Commercial Computer Software Restricted Rights clause at FAR 52.227-19 for other agencies.

Regulatory/Export Compliance. The export and re-export of this software is controlled for export purposes by the U.S. Government. By accepting this software and/or documentation, licensee agrees to comply with all U.S. and foreign export laws and regulations as they relate to software and related documentation. Licensee will not export or re-export outside the United States software or documentation, whether directly or indirectly, to any Prohibited Party and will not cause, approve or otherwise intentionally facilitate others in so doing. A Prohibited Party includes: a party in a U.S. embargoed country or country the United States has named as a supporter of international terrorism; a party involved in proliferation; a party identified by the U.S. Government as a Denied Party; a party named on the U.S. Department of Commerce’s Entity List in Supplement No. 4 to 15 C.F.R. § 744; a party prohibited from participation in export or re-export transactions by a U.S. Government General Order; a party listed by the U.S. Government’s Office of Foreign Assets Control as ineligible to participate in transactions subject to U.S. jurisdiction; or any party that licensee knows or has reason to know has violated or plans to violate U.S. or foreign export laws or regulations. Licensee shall ensure that each of its software users complies with U.S. and foreign export laws and regulations as they relate to software and related documentation.

Contents

Overview

Installation Flow

- Configure all the prerequisites.
- Add a new application to the File Access Manager Administrative Client.
- Install the Activity Monitor / Permissions Collector / Data Classification services.

> Installing the permissions collector and data classification services is optional and should only be installed by someone with a full understanding of File Access Manager deployment architecture. The File Access Manager Administrator Guide has additional information on the architecture.

Configuration

Server Designation and Configuration

The server designation as Production or Disaster Recovery, as well as the configuration of the services in these servers is performed using the File Access Manager Server Installer.

Initial Configuration

When configuring the servers in the initial setup, mark Disaster Recovery (DR) servers, by ticking the Disaster Recovery tick box.

Disaster Recovery Configuration: Setting the Active Servers

You can set the servers to be active or inactive in the Server Installer. If a server is set to be "Toggle UP," the twin DR Server will be in "Sleep State" (idle). if the Production server is in "Toggle Down" Mode the Twin DR server will be in running state.

The user can set server up and down from any machine.

To set a production server Active \ Inactive

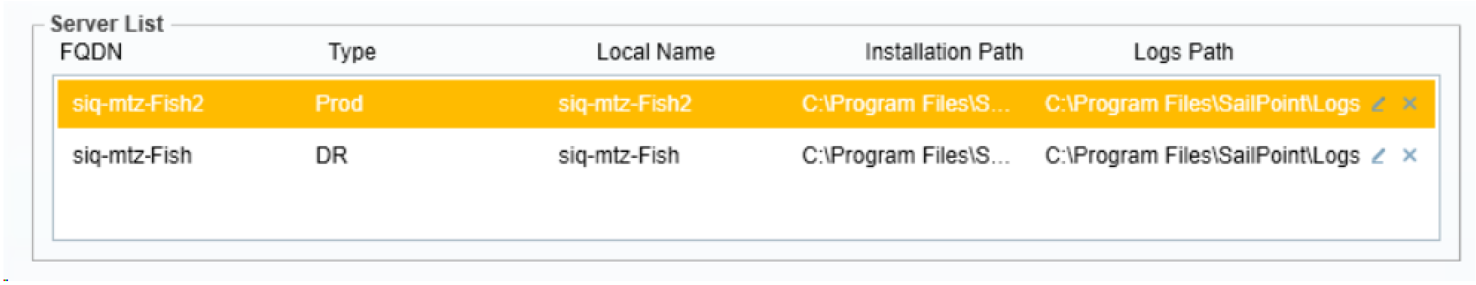

Server List

FQDN	Type	Local Name	Installation Path	Logs Path
siq-mtz-Fish2	Prod	siq-mtz-Fish2	C:\Program Files\S...	C:\Program Files\SailPoint\Logs
siq-mtz-Fish	DR	siq-mtz-Fish	C:\Program Files\S...	C:\Program Files\SailPoint\Logs

1. Open the Server Installer and connect to an existing DB.
2. Select the Production server which will be activated / deactivated and click the **Edit** button.

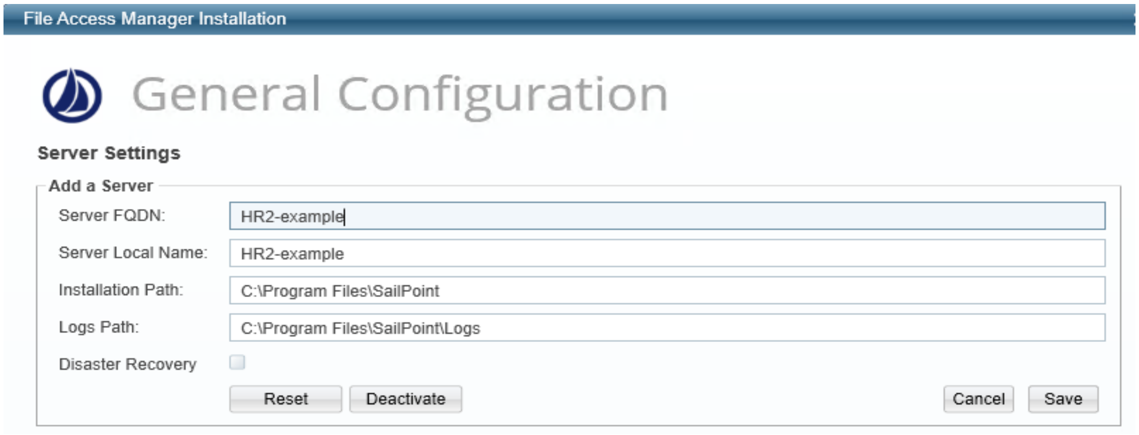

File Access Manager Installation

General Configuration

Server Settings

Add a Server

Server FQDN: HR2-example

Server Local Name: HR2-example

Installation Path: C:\Program Files\SailPoint

Logs Path: C:\Program Files\SailPoint\Logs

Disaster Recovery □

3. This will enable the server status buttons:

   Deactivate / Activate

   Toggle the server status. This will automatically toggle the corresponding DR server as well.

   Reset

   Mark the services running on this server as uninstalled. (This is necessary when a server should be reinstalled or replaced.)

4. Click **activate / deactivate** to change the server state.

5. Click **save**.

Disaster Recovery Flow

Switching on Disaster Recovery Mode

Select the production servers and set them back to Inactive (see Configuration).

This will set the corresponding DR servers to Active.

> Note: Only the Production servers have the Active / Inactive button enabled.

If the Web Site and / or API Servers are Down

If the web site and / or API servers are down, a manual action will be required for the clients to connect.

There are two possibilities to recover:

- In the DNS server, change the target URL of the web site / API to point to the disaster recovery environment, this is the recommended action and will be the simplest.
- Browse directly to the disaster recovery environment address.

High Availability Configuration Considerations

In addition to changing the designation of servers as active / inactive, you have to ensure that the relevant servers are listed in the high availability load balancer.

This depends on the types of servers listed in the load balancer:

Services listed in the load balancer	Configuration changes required when changing the disaster mode
all Prod services and dr services	Nothing to change however if you are monitoring the load balancer health checks you will receive many false positives of broken servers.
Production services only	Update the load balancer to point to the newly activated DR services
Different load balancer for Production & DR	Change the load balancer address inside the server installer load balancer screen.

For the website load balancer configuration only the second two options are relevant.

Disaster Recovery Fallback to Production Environment

Select the production servers and set them back to Active (see Toggle above).

This will set the corresponding DR servers to Inactive.

If the Return to Production Requires Installing the Services

1. Select **>Reset** on the server configuration screen.
2. Reinstall the services.
3. After completing the installation wait at least 2 minutes to allow all File Access Manager services to update their configuration.
4. Set the Production servers to Active.

If the Production Server Has to be Replaced

If the production server cannot be restored, and a new one is required:

1. In the Server Installer add a new server in the server list.
2. Continue until the end and click **>save configuration**.
3. Restart the server installer.
4. Select the new server in the Server Configuration list, click **edit**, then deactivate (the new server is not yet ready to be live).
5. Configure the required services on the new server and install them.
6. Wait for at least 2 minutes after the installation has completed.
7. Turn the new server to active.

If the Fallback Includes the Production Server of the Web Site/API Server

In case the fallback includes the production server of the Web site/API server, a manual action will be required for the clients to connect.

The action will be according to the recovery action taken in section “Switching on Disaster Recovery mode."

- If the DNS server was changed, it should be changed back to point to the Production environment.
- If the DNS server was not changed, you should browse directly to the Production environment address.

Elasticsearch Restoration

For more information, go to https://www.elastic.co/guide/en/elasticsearch/reference/current/snapshots-restore-snapshot.html#restore-different-cluster

> Note: Refer to the Data Backup guide for more information.

1. Stop the following FAM services in the DR environment:
   - File Access Manager Event Manager
   - File Access Manager Scheduled Task Handler
   - File Access Manager Central Data Classification
   - File Access Manager Activity Analytics
2. For each node in the disaster recovery Elasticsearch, find the “elasticsearch.yml” config file and change the “path.repo:” value to the Production backup base path.
3. Restart Elasticsearch.
4. Register to the Production continuous repository as Read Only for the disaster recovery Elasticsearch cluster:

   PUT _snapshot/continuous_backup

   ```
   {
    "type": "fs",
    "settings": {
     "location": "continuous_backup",
     "readonly": "true"
    }
   }
   ```
5. For a disaster recovery cluster, temporarily stop indexing and turn off the following features:

   GeoIP database downloader

   ```
   PUT _cluster/settings
   {
    "persistent": {
      "ingest.geoip.downloader.enabled": false
   ```

```
 }
}
```

ILM

```
POST _ilm/stop
```

Monitoring

```
PUT _cluster/settings
{
 "persistent": {
  "xpack.monitoring.collection.enabled": false
 }
}
```

Machine Learning

```
POST _ml/set_upgrade_mode?enabled=true
```

Watcher

```
POST _watcher/_stop
```

6. Use the cluster update settings API to set **action.destructive_requires_name to false**. This allows you delete data streams and indices using wildcards.

```
PUT _cluster/settings
{
 "persistent": {
  "action.destructive_requires_name": false
 }
}
```

7. Delete all existing data streams on the cluster.

```
DELETE _data_stream/*?expand_wildcards=all
```

8. Delete all existing indices on the cluster.

   ```
   DELETE *?expand_wildcards=all
   ```

9. Copy the name of the snapshot that you want and restore from the Production repository to the disaster recovery Elasticsearch (see step 5 in full cluster instructions).

   - Can get list of available snapshots:

     GET _snapshot/continuous_backup/*?order=desc

     Look for the first snapshot with "state": "SUCCESS".

10. When the restore operation is complete, resume indexing and restart any features you stopped:

GeoIP database downloader

```
PUT _cluster/settings
{
 "persistent": {
  "ingest.geoip.downloader.enabled": true
 }
}
```

ILM

```
POST _ilm/start
```

Machine Learning

```
POST _ml/set_upgrade_mode?enabled=false
```

Monitoring

```
PUT _cluster/settings
{
 "persistent": {
  "xpack.monitoring.collection.enabled": true
 }
}
```

Watcher

```
POST _watcher/_start
```

11. Reset the action.destructive_requires_name cluster setting.

```
PUT _cluster/settings
{
 "persistent": {
  "action.destructive_requires_name": null
 }
}
```

12. Unregister the production repository from the disaster recovery cluster:

    DELETE _snapshot/continuous_backup

13. For each node in the disaster recovery Elasticsearch, find the “elasticsearch.yml” config file and change the “path.repo:” value back to the disaster recovery backup base path and restart Elasticsearch.

14. Register the disaster recovery continuous repository for the disaster recovery Elasticsearch cluster:

    PUT _snapshot/continuous_backup

```
{
 "type": "fs",
 "settings": {
  "location": "continuous_backup"
 }
}
```

15. When returning back to Production environment, follow the above instructions. However, replace Production with disaster recovery and vice versa.

Troubleshooting

The Reindex Task Marked as Canceled Following a Dr Transfer

In the following scenario, the Reindex task might fail. In this case, run the Reindex manually.

1. Prod Elasticsearch server goes down.
2. The user marks the server Inactive in the server installer. This will toggle the parallel server in the DR environment to Active.
3. Reindex activities task starts automatically (which is run by the Scheduled Task Handler service).
4. The server where the Scheduled Task Handler service goes down.
5. The reindex task automatically gets canceled.

In this specific scenario and only if the reindex task was cancelled, create a new reindex task manually.

This is done from the Administartive client Health Center screen

Unable to see events

Symptom

The following error displays in a log file while you attempt to install the monitoring connector:

```
ERROR, WBX.whiteOPS.Agents.FilesMiniFilterActivity Mon-
itor.FileMiniFilterActivity MonitorManager,connect, An unexpected error
occurred while you attempt to start the mini-filter:
System.DllNotFoundException:Unable to load DLL 'wbapi.dll': The specified mod-
ule could not be found.
(Exception from HRESULT: 0x8007007E)—at
 WBX.whiteOPS.Agents.FilesMiniFilterActivity Monitor.SafeNativeMethods64.start
(UInt32 bufferSizeInBytes, UInt32 trustedProcessId)—at WBX.whiteOPS.A-
gents.FilesMiniFilterActivity Monitor.FileMiniFilterActivity Mon-
itorManager.connect()
```

Reason

Visual C++ 2010 redistributable package was not installed as part of the Activity Monitor service installation.

Solution Steps

Perform Step 3 of Activity Monitor Installation, Windows Server Core.

The application is not in the list of Collector Installation Managers

Symptom

The application does not appear in dropdown list of the Collector Installation Managers in the Activity Monitoring

Reason

Either the application was not defined or the Host Name (defined in section 5.11) does not match the server's short name on which the Collector Installation Manager was opened on.

Solution Steps

Create the application in case it does not exist.

In case it exists, make sure the Host Name is correct.