



Integrating NFS with File Access Manager

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

- Capabilities** 4
 - Supported Versions 4
- Connector Overview** 5
- Prerequisites** 6
 - Software Requirements 6
 - Required Permissions 6
 - Communications Requirements 7
- NFS Installation Flow Overview** 8
- Collecting Data Stored in an External Application** 9
- Adding an NFS Application** 11
 - Select Wizard Type 11
 - General Details 11
 - Connection Details 12
 - Configuring and Scheduling the Permissions Collection 14
 - Selecting and Scheduling the Data Classification Settings 22
 - Data Privacy 23
- Installing Services: Collector Installation** 24
- Verifying the NFS Connector Installation** 27
 - Installed Services 27
 - Log Files 27
 - Permissions Collection 27

Capabilities

This connector enables you to use File Access Manager to access and analyze data stored in NFS and do the following:

- Analyze the structure of your stored data.
- Classify the data being stored.
- Verify user permissions on the resources, and compare them against requirements.

See the File Access Manager documentation for a full description.

Supported Versions

- NFS v3
- NFS v4.1 (including Integrity and Privacy export security types)

Connector Overview

- File Access Manager NFS Connector supports Permissions Collection and Data Classification for Linux/Unix servers.

- **NFS protocol versions**

The NFS agent supports the following NFS protocol versions:

NFSv3

The agent uses the standard NFSv3 protocol to crawl NFS exports and directory structure.

The agent retrieves and analyzes UNIX-style object permissions.

NFSv4.1

The agent uses the standard NFSv4.1 protocol to crawl the NFS pseudo-filesystem.

- **Identity Collection schemes**

The NFS agent supports the following Identity Collection schemes:

Unix/Linux local users and groups are retrieved through SSH (or Telnet, if SSH is not available).

For environments with UNIX-style permissions, identities can be gathered from:

NIS server

List of local users and groups

For environments with NFSv4-style ACL permissions, identities can be gathered from:

Active Directory domain

NIS server

List of local users

Prerequisites

Make sure your system fits the descriptions below before starting the installation.

Software Requirements

File Access Manager requires the latest ASP.NET Core 6.0.x Hosting Bundle. This bundle consists of .NET Runtime and ASP .NET Core Runtime. You can download the latest 6.0.x Hosting Bundle version from [here](#) .

Required Permissions

- Local users and groups are gathered within the context of a proprietary Unix/Linux user, which requires the following permissions:
 - SSH (or Telnet) access to the NFS file server
 - Permissions to read the `/etc/passwd` and `/etc/group` files (using the `cat` shell command)
- Crawling and Permission Analysis is within the context of a proprietary NFS user, which requires permissions that vary in accordance with the file-system permission type being used:

Unix-style permissions

Read

Execute

NFSv4-style ACL permissions

r - read-data (files) / list-directory (directories)

x - execute (files) / change-directory (directories)

t - read-attributes - read the attributes of the file/directory.

n - read-named-attributes - read the named attributes of the file/directory.

c - read-ACL - read the file/directory NFSv4 ACL

y - synchronize - allow clients to use synchronous I/O with the server.

Communications Requirements

Requirement	Source	Target	Ports Protocol
File Access Manager Message Broker	Permissions Collector / Data Classification Collector	RabbitMQ	5671
NFSv3	Permissions Collector / Data Classification service	NFS file server	Port 111 - port-mapper Port 2049 - NFS
NFSv4	Permissions Collector / Data Classification service	NFS file server	Port 2049 - NFS
Permissions Collection – get local users and groups	Permissions Collector	NFS file server	SSH/Telnet

NFS Installation Flow Overview

To install the NFS connector:

1. Configure all the prerequisites.
2. Add a new NFS application in the Business Website.
3. Install the relevant services:
 - Permissions Collector
 - Data Classification Collector

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.

Collecting Data Stored in an External Application

Terminology:

Connector

The collection of features, components and capabilities that comprise File Access Manager support for an endpoint.

Collector

The “Agent” component or service in a Data Classification and or Permission Collection architecture.

Engine

The core service counterpart of this architecture.

Identity Collector

A logical component used to fetch identities from an identity store and holds the configuration, settings for that identity store, and the relations between these identities.

The identity collector has no “physical” manifest.

- The actual work is done by the Collector Synchronizer.

The list below describes the high level installation process required to collect and analyze data from an external application. Most of these should already be set up in your File Access Manager installation. See the server Installation guide for further details.

Install a Data Classification central engine

One or more central engines, installed using the server installer

Install a Permission Collection central engine

One or more central engines, installed using the server installer

Create an Application in File Access Manager

From the Business Website. The application is linked to central engines listed above.

Install Permission Collectors and / or Data Classification Collector (optional)

Optionally, you can install collectors that will run on a separate server and take some of the work from the central PC and DC engines (Where supported). When installing a collector, you attach it to an engine. If no collectors are installed, the central services act as both the engine and the collector.

To install a collector, you must have the **RabbitMQ** service installed for communication between the central engines and the collectors. RabbitMQ is installed

For further details, see section **Application > Central Service > Collector Relations** in the File Access Manager Administrator Guide

Adding an NFS Application

In order to integrate with NFS, we must first create an application entry in File Access Manager. This entry includes the identification, connection details, and other parameters necessary to create the link.

To add an application, use the **New Application Wizard**.

1. Navigate to *Admin > Applications*
2. Click **Add New** to open the wizard.

Select Wizard Type

1. Click **Standard Application**
2. Click **Next** to open the **General Details** page.

General Details

Application Type

NFS

Application Name

Logical name of the application

Description

Description of the application

Tags

Select tags for the application from the dropdown menu, and / or type a new name, and press **Enter** to create a new tag. The dropdown list of tags filters out matching tags as you type and displays up to 50 tags.

The **tags** replace the **Logical container** field that was used when creating applications in releases before 8.2

Event Manager Server

This option is available if there are more than one event manager servers configured in the system.

Select an event manager from the drop down menu.

NIS Identity Collector

You can choose to assign an identity collector that's configured over a Linux NIS server. Select an identity collector from the dropdown list.

AD Identity Collector

You can choose to assign an identity collector that's configured over a Windows Active Directory service. Select an identity collector from the dropdown list.

You can create identity collectors in the business website. Press the **Refresh** button to update the Identity Collector dropdown list after adding a new identity collector.

Click **Next**. to open the Connection Details page.

Connection Details

Server Address

The network address of the NFS server – any network-reachable address is valid - IP Address / WINS / DNS

Shell Protocol

Select whether to use SSH (default) or Telnet when connecting to the server to gather local users & groups

Shell Port

The target remote port for SSH/Telnet connection (Default: 22)

Shell Username

The username for interactive shell login.

This field can contain:

- NIS username
- Local user's UID
- "root" (not recommended for security reasons)

Shell Password

The password to use when connecting to SSH/Telnet

NFS Version

Select from the dropdown list. (default: v3)

Authentication Type (for NFS version v4)

When you configure an NFSv4.1 server, this dropdown allows configuring the scheme of NFS authentication.

Select **Unix** for classic NFS authentication (using UID/GID) or **Kerberos** for NFSv4-style authentication (using RPCSEC_GSS over NFSv4.1).

Default value: Unix

NFSv3 supports only Unix authentication. NFSv4.1 supports both Unix and Kerberos authentication.

For Unix authentication, fill in the following fields:

Username

When connecting to the NFS file server. This field can contain

- NIS user name
- Local user's UID
- "root" (not recommended for security reasons)

Group Name

When connecting to the NFS file server. This field can contain:

- NIS group name
- Local group's GID
- "root" (not recommended for security reasons)

For Kerberos authentication, fill in the following fields:

Server SPN

The NFS server's *service principle name* as defined in your Active Directory (or other KDC). This value can be found in:

- NFS server's Computer Account in the Active Directory – under the *servicePrincipleName* attribute.
- On the NFS server's Unix/Linux machine, run the following commands:

- – ktutil
- – read_kt /etc/krb5.keytab
- – list

/etc/krb5.keytab is the default location of the keytab file, but it may vary in your environment.

Domain Name

The NetBIOS domain name of the user when connecting to the NFS server

Username

When connecting to the NFS server

Password

The password when connecting to the NFS server

Click **Next**.

Configuring and Scheduling the Permissions Collection

Permissions can be analyzed to determine the application permissions of an out-of-the-box application, provided you have defined an identity store for File Access Manager to use in its analysis, and you have run a crawl for the application.

The permission collector is a software component responsible for analyzing the permissions in an application.

The Central Permission Collector Service is responsible for running the Permission Collector and Crawler tasks.

If the “IdentityIQ FAM Central Permission Collector” wasn’t installed during the installation of the server, this configuration setting will be disabled.

To configure the Permission Collection

- Open the edit screen of the required application.
 - a. Navigate to **Admin > Applications**.
 - b. Scroll through the list, or use the filter to find the application.

c. Click the edit icon  on the line of the application.

- Press **Next** till you reach the **Crawler & Permissions Collection** settings page.

The actual entry fields vary according to the application type.

When entering this page in edit mode, you can navigate between the various configuration windows using the **Next** and **Back** buttons.

Central Permissions Collection Service

Select a central permission collection service from the dropdown list. You can create permissions collection services as part of the service installation process. See section "Services Configuration" in the File Access Manager Administrator Guide for further details.

Skip Identities Sync during Permission Collection

Skip identity synchronization before running permission collection tasks when the identity collector is common to different connector.

This option is checked by default.

Scheduling a Task

Create a Schedule

Click on this option to view the schedule setting parameters.

Schedule Task Name

A name for this scheduling task

When creating a new schedule, the system generates a default name in the following format:

{appName} - {type} Scheduler

You can override or keep this name suggestion.

Schedule

Select a scheduling frequency from the dropdown menu.

- ***Schedule Types and Intervals***

Once

Single execution task runs.

Run After

Create dependency of tasks. The task starts running only upon successful completion of the first task.

Hourly

Set the start time.

Daily

Set the start date and time.

Weekly

Set the day(s) of the week on which to run.

Monthly

The start date defines the day of the month on which to run a task.

Quarterly

A monthly schedule with an interval of 3 months.

Half Yearly

A monthly schedule with an interval of 6 months.

Yearly

A monthly schedule with an interval of 12 months.

Date and time fields

Fill in the scheduling times. These fields differ, depending upon the scheduling frequency selected.


Active check box

Check this to activate the schedule.

Click **Next**.

Configuring and Scheduling the Crawler

To set or edit the Crawler configuration and scheduling

- Open the edit screen of the required application.
 - a. Navigate to **Admin > Applications**.
 - b. Scroll through the list, or use the filter to find the application.
 - c. Click the edit icon  on the line of the application.
- Press **Next** till you reach the **Crawler & Permissions Collection** settings page.

The actual entry fields vary according to the application type.

Calculate Resource Size

Determine when, or at what frequency, File Access Manager calculates the resources' size.

Select one of the following:

- Never
- Always
- Second crawl and on (This is the default)

Create a Schedule

Click to open the schedule panel. See [Scheduling a Task](#)


Setting the Crawl Scope

There are several options to set the crawl scope:

- Setting explicit list of resources to include and / or exclude from the scan.
- Creating a regex to define resources to exclude.

Including and Excluding Paths by List

To set the paths to include or exclude in the crawl process for an application

- Open the edit screen of the required application.
 - a. Navigate to **Admin > Applications**.
 - b. Scroll through the list, or use the filter to find the application.
 - c. Click the edit icon  on the line of the application.
- Press **Next** till you reach the **Crawler & Permissions Collection** settings page.


The actual entry fields vary according to the application type.

 1. Scroll down to the Crawl configuration settings.
 2. Click **Advanced Crawl Scope Configuration** to open the scope configuration panel.
 3. Click Include / Exclude Resources to open the input fields.
 4. To add a resource to a list, type in the full path to include / exclude in the top field and click **+** to add it to the list.
 5. To remove a resource from a list, find the resource from the list, and click the x icon on the resource row.

When creating exclusion lists, excludes take precedence over includes.

Excluding Paths by Regex

To set filters of paths to exclude in the crawl process for an application using regex.

- Open the edit screen of the required application.
 - a. Navigate to **Admin > Applications**.
 - b. Scroll through the list, or use the filter to find the application.
 - c. Click the edit icon  on the line of the application.
- Press **Next** till you reach the **Crawler & Permissions Collection** settings page.

The actual entry fields vary according to the application type.

1. Click **Exclude Paths by Regex** to open the configuration panel.
2. Type in the paths to exclude by Regex, See regex examples in the section below. Since the system does not collect BRs that match this Regex, it also does not analyze them for permissions.

Crawler Regex Exclusion Example

The following are examples of crawler Regex exclusions:

Exclude all shares which start with one or more shares names:

Example: Starting with /shareName

Regex: `^\\shareName$`

Real Example:

Path: /exports/recent/Automation

Regex: `^\\exports\\recent\\Automation$`

Example: Starting with /shareName or /OtherShareName

Regex: `^\\(shareName|OtherShareName)$`

Include ONLY shares which start with one or more shares names:

Example: Starting with /shareName

Regex: `^(?!\\shareName($|V.*)).*`

Example: Starting with /shareName or /OtherShareName

Regex: `^(?!\\(shareName|OtherShareName)($|V.*)).*`

Excluding Top Level Resources

Use the top level exclusion screen to select top level roots to exclude from the crawl. This setting is done per application.

To exclude top level resources from the crawl process

1. Open the application screen **Admin > Applications**.
2. Find the application to configure and click the drop down menu on the application line. Select **Exclude Top Level Resources** to open the configuration panel.

3. ***Run Task***

The Run Task button triggers a task that runs a short detection scan to detect the current top level resources.

Before running the task for the first time, the message above this button is:

"Note: Run task to detect the top-level resources"

If the top level resource list has changed in the application while you are on this screen, press this button to retrieve the updated structure.

Once triggered, you can see the task status in **Settings > Task Management > Tasks**.

This will only work if the user has access to the task page

When the task has completed, press **Refresh** to update the page with the list of top level resources.

4. Click the top level resource list, and select top level resources to exclude.
5. Click **Save** to save the change.
6. To refresh the list of top level resources, run the task again. Running the task will not clear the list of top level resources to exclude.

Top Level Resources Exclusion

WFS-DC testing

Last Successful Run 06-22-2021 4:57:27 PM

[Run Task](#) [View Task Status](#)

Note: Refresh the list to view recently discovered resources [Refresh](#)

Top Level Resources Exclusion List 0 Selected | [Clear Selection](#)

Top Level Resources Exclusion List

- \\si-...-5\C\$
- \\si-...-5\MSSQLSERVER
- \\si-...-5\print\$

Special Consideration for Long File Paths in Crawl

If you need to support long file paths above 4,000 characters for the crawl, set the flag

`excludeVeryLongResourcePaths`

in the Permission Collection Engine App.config file to true.

By default this value will be commented out and set to false.

This key ensures, when enabled, that paths longer than 4000 characters are excluded from the applications' resource discovery (Crawl), to avoid issues while storing them in the SQLServer database.

When enabled, business resources with full paths longer than 4000 characters, and everything included in the hierarchical structure below them, will be excluded from the crawl, and will not be collected by File Access Manager. This scenario is extremely rare.

You should not enable exclusion of long paths, unless you experience an issue.

Background

File Access Manager uses a hashing mechanism to create a unique identifier for each business resource stored in the File Access Manager database. The hashing mechanism in SQLServer versions 2014 and earlier, is unable to process (hash) values with 4,000 or more characters.

Though resources with paths of 4000 characters or longer are extremely rare, File Access Manager is designed to handle that limitation.

Identifying the Problem

When using an SQL Server database version 2014 and earlier

The following error message in the Permission Collection Engine log file:

```
System.Data.SqlClient.SqlException (0x80131904): String or binary data would be truncated.
```

In all other cases, this feature should not be enabled.

Setting the Long Resource Path Key


The Permission Collection Engine App.config file is `RoleAnalyticsServiceHost.exe.config`, and can be found in the folder

`%SailPoint_Home%\FileAccessManager\[Permission Collection instance]\`

Search for the key **`excludeVeryLongResourcePaths`** and correct it as described above.

Selecting and Scheduling the Data Classification Settings

To associate an application with a data classification service, and set the schedule

- Open the edit screen of the required application
 - a. Navigate to **Admin > Applications**
 - b. Scroll through the list, or use the filter to find the application
 - c. Click the edit icon  on the line of the application

- Press **Next** till you reach the **Data Classification** settings page.

The actual entry fields vary according to the application type

Central Data Classification Service

Associate the application with a Central Data Classification Service. This service is responsible for running the Data Classification tasks.

If the “Central Data Classification” wasn’t installed during the installation of the server, this field is disabled.

Disabling Data Classification

To disable data classification, delete the entry from the central data classification field.

Disabling data classification can also be achieved by setting the scheduler to be inactive (which is the default setting for data classification).

Create a Schedule

This option is enabled only if a central data classification service is selected.

See [Scheduling a Task](#)

See the chapter “Data Classification” in the File Access Manager Administrator Guide for more information

Click **Next** or **Finish**.

Data Privacy

A user can associate the application with a Central Data Classification Engine Service. This engine will be responsible for executed Data Privacy tasks.

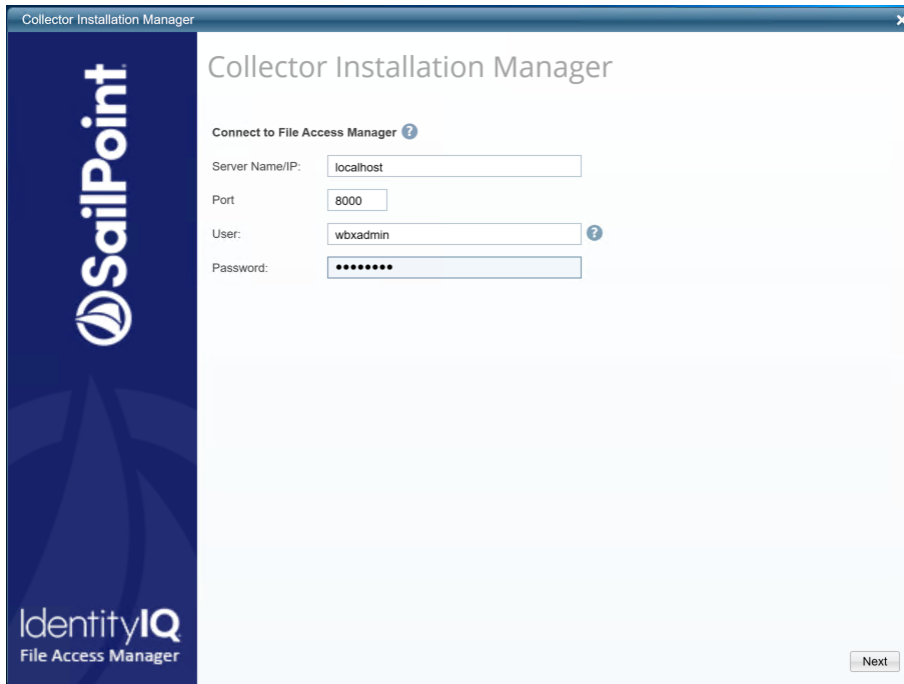
Though using different processes for each, the Data Classification engine service is in charge for both Data Privacy and Data Classification discovery tasks.

You may choose the same service for both, or use a different one for each, to run them in parallel.

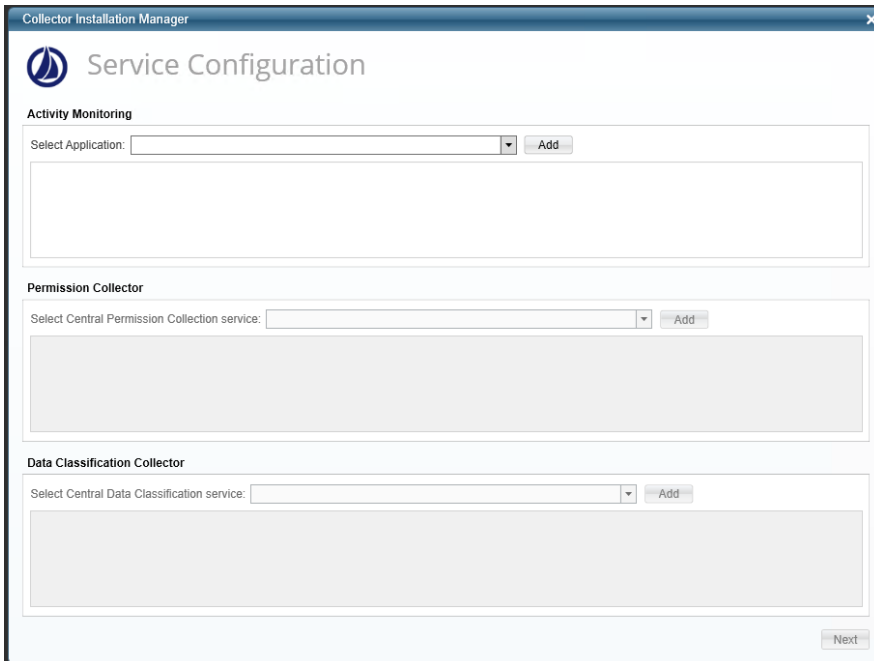
The fields on the Data Privacy step are the same as the Data Classification step.

Installing Services: Collector Installation

1. Run the **Collector Installation Manager** as an Administrator.
The installation files are in the installation package under the folder Collectors.
The Collector Installation Manager window displays.



2. Enter the credentials to connect to File Access Manager.
 - a. ServerName/IP should be pointed to the Agent Configuration Manager service server.
 - b. A File Access Manager user with Collector Manager permission (permission to install collectors). For Active Directory authentication, use the format domain\username.
3. Click **Next** to open the Service Configuration window.



4. If you are installing the Permission Collector, select the Central Permission Collector to which to connect this service. Click **Add**.
5. If you are installing the Data Classification, select the Central Classification Collector to which to connect this service. Click **Add**.
6. Click **Next**.

The Installation Folder window displays.

If this is the first time you are installing collectors on this machine, you will be prompted to select an installation folder. All future collectors will be installed in this folder.

7. Browse and select the location of the target folder for installation.
8. Browse and select the location of the folder for system logs.
9. Click **Next**.
10. The system begins installing the selected components.
11. Click **Finish**.

The Finish button is displayed after all the selected components have been installed.

The File Access Manager Administrator Guide provides more information on the collector services.

Verifying the NFS Connector Installation

Installed Services

Verify that the services installed for the connector are available and active. Using windows Service manager, or other tool, look for the File Access Manager services, and see that they are running.

for example:

- File Access Manager Central Permissions Collection - <Application_Name>
- File Access Manager Central Data Classification - <Application_Name>

Log Files

Check the log files listed below for errors

- "%SAILPOINT_HOME_LOGS%\PermissionCollection_<Service_Name>.log"
- "%SAILPOINT_HOME_LOGS%\DataClassification_<Service_Name>.log"

Permissions Collection

1. Run the Crawler and Permissions Collector tasks (*Settings > Task Management > Scheduled Tasks*)
2. Verify that:
 - The tasks completed successfully
 - Business resources were created in the resource explorer (*Admin > Applications > [application column] > Manage Resources*)
 - Permissions display in the Permission Forensics page (*Forensics > Permissions*)