Integrating Linux with File Access Manager

Version: 8.2 Revised: July 01, 2021

Copyright and Trademark Notices.

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

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

This connector enables you to use IdentityIQ File Access Manager to access and analyze data stored in Linux 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.
- The Linux connector supports collection of Unix permissions as well as ACL permissions.

See the IdentityIQ File Access Manager documentation for a full description.

Linux Permission Types

The supported permission types are:

- Read
- Write
- Execute
- None

None permission

The “None” permission is used when a user or group has no permissions.

“Others” group

The “Others” group, which is part of the Unix permissions, is represented in File Access Manager as a calculated Everyone group.

When a resource has permissions for the “Others” group, this group contains all users except the users and groups that have explicit permissions.

Supported Linux Distributions

- Ubuntu versions 18.04 and 20.04
- Red Hat Enterprise Linux versions 7 and 8
- CentOS versions 7 and 8

Prerequisites

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

Software Requirements

IdentityIQ File Access Manager requires the latest ASP.NET Core 3.1.x Hosting Bundle. This bundle consists of .NET Runtime and ASP .NET Core Runtime. You can download the latest 3.1.x Hosting Bundle version from here .

- SSH and SFTP must be available on the Linux server.

Permissions

IdentityIQ File Access Manager connects to the Linux server using SFTP and SSH.

A user with elevated permissions is required in order to read directories with restricted permissions. We recommend granting the required permissions as described in Granting Read Permissions. This method grants File Access Manager the minimal required permissions to read any file or directory. Alternatively, it is possible to skip Granting Read Permissions and allow File Access Manager to use root instead.

Using a user other than root and not granting the permission as described in Recommended: Granting Read Permissions is not recommended. The fetched information will be limited to the permissions that the given user possesses. For instance, if the given user is not allowed to read the permissions of a directory, then the information of that directory will not be collected.

Mandatory Permissions

Permissions to run the commands:

`cat`

`getent` - Only if you plan to use Active Directory as an Identity Collector

`ypcat` - Only if you plan on using NIS as an Identity Collector

Permissions to read:

`/etc/passwd`

`/etc/group`.

In order to verify that a user has the required permissions, run the following commands with the desired user and make sure they succeed:

```
cat /etc/passwd
cat /etc/group
getent passwd 0 (Only if you plan on using Active Directory as an Identity Collector)
ypcat passwd (Only if you plan on using NIS as an Identity Collector)
```

Recommended: Granting Read Permissions

The SailPoint method of acquiring the required permissions is to use the “**cap_dac_read_search**” capability.

This capability allows us to bypass file read permission checks and directory read and execute permission checks.

Since Linux capabilities can be applied to files, but not to users, we will create dedicated executables that will only be used by IdentityIQ File Access Manager.

> If the SSH, SFTP or ACL packages are updated after following these steps, then the duplicated executables should be recreated, and the steps below should be repeated (except for creating the user for File Access Manager).

Using root, perform the following operations in the Linux server:

1. Create a user for File Access Manager

   a. Create the user famuser

   ```
   adduser  famuser
   ```

   b. Set password for the new user

   ```
   passwd famuser
   ```

   c. Make sure that famuser has the permissions as described in Mandatory Permissions.

2. Create a variable that contains the path of the sftp server executable:

   For RHEL or CentOS distributions

   ```
       sftpsrv=/usr/libexec/openssh/sftp-server
   ```

   For Ubuntu

   ```
       sftpsrv=/usr/lib/openssh/sftp-server
   ```

   > The sftp-server location could be different depending on the OS

3. Copy the sftp executable:

   ```
   cp -a ${sftpsrv} ${sftpsrv}-fam
   ```

4. Make File Access Manager’s user the only user that can read and execute it.

   ```
   chmod 500 ${sftpsrv}-fam
   
   chown famuser ${sftpsrv}-fam
   ```

5. Grant capability to bypass file read permission checks and directory read and execute permission checks

   ```
   /sbin/setcap cap_dac_read_search+ep ${sftpsrv}-fam
   ```

6. Next, we will create a new SSH Subsystem.

   Open your SSH configuration, For OpenSSH, use the following:

   ```
   nano /etc/ssh/sshd_config
   ```

7. Add the following line to the file. Make sure the path of the sftp executable matches the path described above, according to the distribution type.

> There will probably be a section for subsytems, look for a line that begins with “Subsystem” near the end of the file. it is best to add the line after the other subsystems.

For RHEL or CentOS distributions

```
Subsystem sftp-fam /usr/libexec/openssh/sftp-server-fam
```

For Ubuntu

```
Subsystem sftp-fam /usr/lib/openssh/sftp-server-fam
```

8. Restart the ssh service:

```
systemctl restart sshd
```

Optional - Grant Read Permissions for ACLs

This section should only be followed if you wish to read ACL permissions.

1. Copy the getfacl executeable:

```
cp -a /bin/getfacl /bin/getfacl-fam
```

2. Make File Access Manager’s user the only user that can read and execute it.

```
chmod 500 /bin/getfacl-fam

chown famuser /bin/getfacl-fam
```

3. 12. Grant the executable the capability to bypass file read permission checks and directory read and execute permission checks

```
/sbin/setcap cap_dac_read_search+ep /bin/getfacl-fam
```

Communications Requirements

Requirement	Source	Destination	Port
File Access Manager Internal Access	Application	File Access Manager Servers	8000-8008
File Access Manager Message Broker	Permissions Collector	RabbitMQ	5671
Permissions Collection	Permissions Collection service	Target Linux server	Configurable SSH port

Configuration Requirements

File Access Manager supports reading permissions of users from Active Directory only if the display format of Active Directory users in the Linux machine is user@domain (which is the default format).

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.

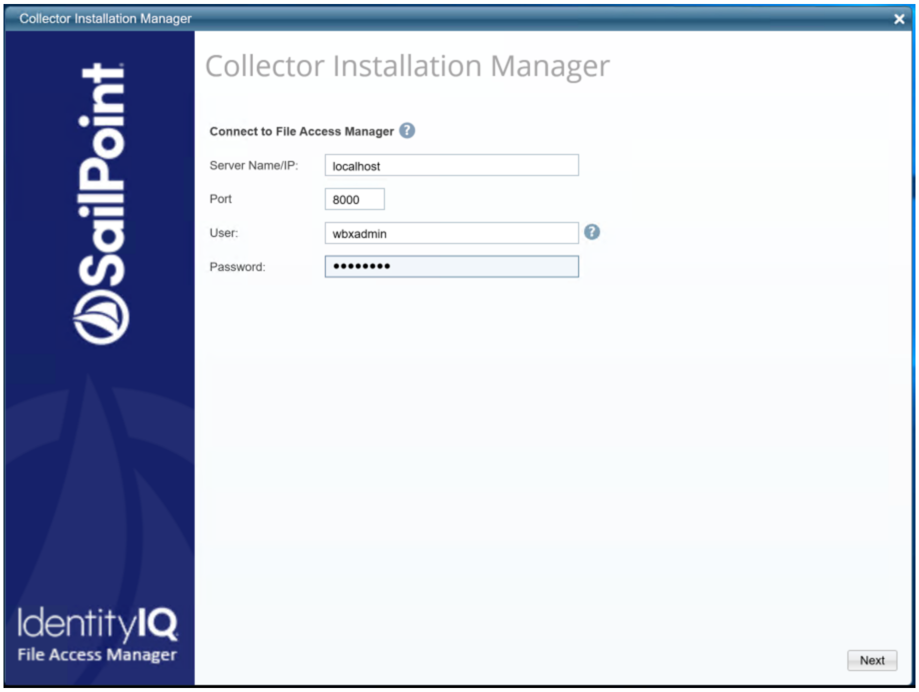

Collector Installation Manager

Connect to File Access Manager

Server Name/IP: localhost

Port: 8000

User: wbxadmin

Password: ••••••••

2. Enter the credentials to connect to IdentityIQ File Access Manager.

   a. ServerName/IP should be pointed to the Agent Configuration Manager service server.

   b. An IdentityIQ File Access Manager user with Collector Manager permission (permission to install collectors). For Active Directory authentication, use the format domain\username.

3. Click **Next**.

   The Service Configuration window displays.

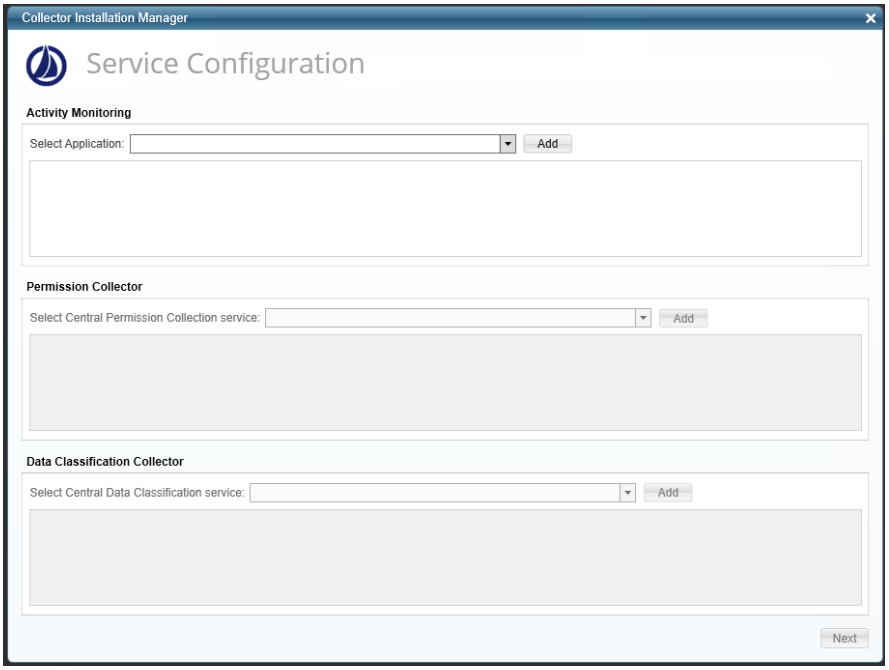
Collector Installation Manager
Service Configuration
Activity Monitoring
Select Application:
Add
Permission Collector
Select Central Permission Collection service:
Add
Data Classification Collector
Select Central Data Classification service:
Add
Next

4. > In some applications, additional credentials may be required to allow granting elevated permission for activity monitoring collection.

5. If you are installing the Permission Collector, select the Central Permission Collector to which to connect this service, and click **Add**
6. If you are installing the Data Classification, select the Central Classification Collector to which to connect this service, and click **Add**
7. 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, in which all future collectors will also be installed.

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

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

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

Adding a Linux Application

In order to integrate with Linux, we must first create an application entry in IdentityIQ 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

Linux

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

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

Connection Details

Enter the login details and credentials

- Server Address
- Shell Port

Shell Username

If you followed the section Recommended: Granting Read Permissions in this guide, enter “famuser”. Otherwise, enter another username as described in the Permissions section in this guide.

Select the login method, and either enter the user password, or upload a private key and passphrase

- Use Shell Password
- User Private Key

Use Dedicated Executables

Check this option if you followed the section Recommended: Granting Read Permissions and you have created dedicated executables for File Access Manager.

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 IdentityIQ File Access Manager Administrator Guide for further details.

Analyze ACL Permissions

Click to fetch and analyze ACL-type Permissions.

This option is checked by default

> If ACL is not supported by your server, make sure this field is unchecked.

Skip Identities Sync during Permission Collection

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

> 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 Resources' 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 for Linux

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. See the example below in section Business Resource Structure to better understand the business resource full path structure.

Crawler Regex Exclusion Examples

Exclude a path

Example: The path /root

```
^\/root($|\\.*)
```

Exclude multiple paths

Example: /root and /media

```
^(\/root|\/media)($|\\.*)
```

Include only a path (example: /home)

Please note that the parent directories must also be added, in this example we added the path '/'

```
^(?!(\/|\/home)($|\/.*)).*
```

Include multiple paths

Example: /home and /boot

Please note that their parent directories must also be added, in this example we added the path '/'

```
^(?!(\/|\/home|\/boot)($|\/.*)).*
```

> To write a slash or a Dollar sign, add a backslash before it as an escape character.

> To add a condition in a single command, use a pipe character "|" .

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 yo u 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.

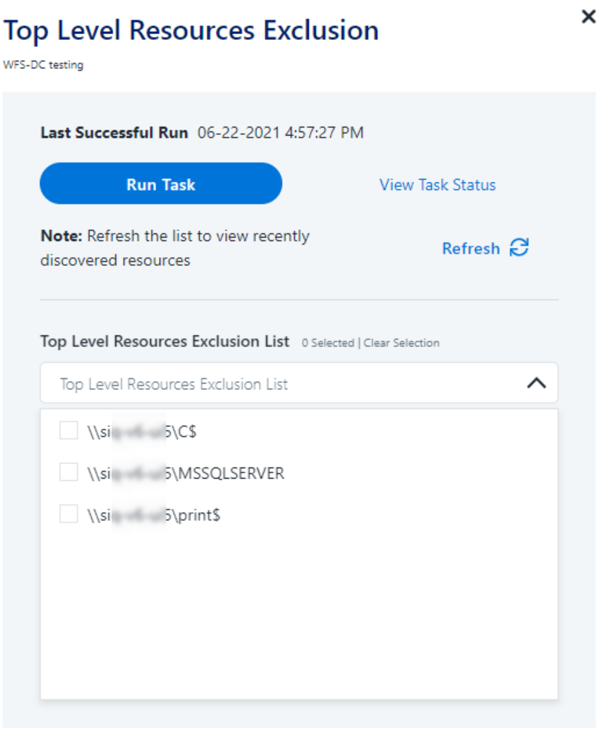

Top Level Resources Exclusion

WFS-DC testing

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

Note: Refresh the list to view recently discovered resources

Top Level Resources Exclusion List 0 Selected | Clear Selection

□ \\si[redacted]5\C$

□ \\si[redacted]5\MSSQLSERVER

□ \\si[redacted]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 ealier

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 IdentityIQ File Access Manager Administrator Guide for more information

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

Verifying the Linux 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 IdentityIQ File Access Manager services, and see that they are running.

for example:

- File Access Manager Central Permissions Collection - [Service 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*)

Troubleshooting

Check the issues below for common problems and suggested ways of handling them.

Dedicated SFTP System Not Found

Case

Tasks fail with the following error:

“Dedicated sftp system was not found, and 'use dedicated executables' is enabled, Please make sure you created the subsystem as instructed in the Linux connector installation guide”

Resolution / Suggestion

The cause of this error is that “**Use Dedicated Executables**” is selected in the connection details page in the application wizard, but the dedicated executables cannot be used. Please make sure you follow the Prerequisites section in this guide.

Get ACLs Command Not Found

Case

Tasks fail with the following error:

“get acls command not found”.

Resolution / Suggestion

“**Anaylze ACL Permissions**” is turned on in the connection details section of the application wizard, but the Permissions collector was unable to find the “getfacl” command.

it is possible that ACLs are not enabled in the Linux server.

Get ACLs Command Not Found and 'Use Dedicated Executeables' is Enabled

Case

Tasks fail with the following error:

“get acls command not found. 'use dedicated executeables' is enabled, make sure you followed the installation guide and created the dedicated executables for File Access Manager in the Linux server”

Resolution / Suggestion

The dedicated executable for getting ACLs cannot be found. Please make sure you followed the section Optional - Grant Read Permissions for ACLs