



IdentityIQ Tasks

Version: 8.4

Revised: September 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 Technologies,” (design and word mark), “SailPoint,” (design and word mark), “Identity IQ,” “IdentityNow,” “SecurityIQ,” “Identity AI,” “Identity Cube,” and “SailPoint Predictive Identity” are registered trademarks of SailPoint Technologies, Inc. “Identity is Everything,” “The Power of Identity,” and “Identity University” are 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 regarding these materials 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

- Tasks Overview** 1
 - Template-Style Tasks and Special-Purpose Tasks 1
 - How Tasks Are Run 1
 - Process Overview for Creating Tasks 2
- Working with Tasks** 3
 - How to Create a New Task 3
 - How to Edit a Task 5
 - How to Schedule a Task 8
- Working with Schedules** 9
 - How to Edit a Schedule 9
- Task Administration** 10
 - Running Tasks on Specific Hosts 10
 - Tracking Task Run Length for Troubleshooting 10
 - Administrator Console: Tasks Page 11
- Tasks Page** 12
 - Predefined Tasks 13
- Scheduled Tasks Page** 16
- Task Results Page** 17
- Task Types** 19
 - Access History Tasks 22
 - Account Aggregation 22
 - Account Group Aggregation 27
 - Activity Aggregation 28
 - Alert Aggregation 29
 - Alert Processor 29
 - Application Builder 30
 - ArcSight Data Export 32
 - Data Export 35
 - Effective Access Indexing 36

Encrypted Data Synchronization Task	37
Entitlement Role Generator	37
File Access Manager Classification	38
ITIM Application Creator	39
IdentityIQ Cloud Gateway Synchronization	40
Identity Refresh	40
Identity Request Maintenance	49
Missing Managed Entitlements Scan	50
OIM Application Creator	51
Policy Scan	51
Propagate Role Changes	51
Refresh Logical Accounts	54
Reset Failed NativeIdentityChangeEvents	55
Role Index Refresh	56
Role-Entitlement Associations	56
Run Rule	56
Sequential Task Launcher	57
System Maintenance	57
Target Aggregation	61
How to Complete Task Work Items	63

Tasks Overview

Caution: When working with tasks, do not open multiple tabs or browsers. Opening multiple tabs might cause a change in one tab to overwrite changes made in another.

Tasks are used to automate the processes which build, update, and maintain the information in IdentityIQ. Tasks perform periodic operations such as aggregating data from applications, refreshing Identity Cubes to update entitlements and roles, running rules, performing system maintenance, and more.

IdentityIQ offers many predefined tasks to handle typical IdentityIQ operations. Tasks are categorized by type and function, to help you find the right task to accomplish what you need. In most cases, you can set specific parameters for a task to control its behavior and results. You can also create custom tasks to meet specialized needs of your organization.

Template-Style Tasks and Special-Purpose Tasks

Most of the tasks IdentityIQ provides as standard types are designed to be used as templates. This means that you can create multiple instances of a particular type of task, and use parameters to determine what the task does and how it behaves. For example, you can set up a unique instance of an Account Aggregation task for each of your applications – such as one task for your LDAP application, a different one for your procurement system, etc. As another example, you could set up one instance of an Identity Refresh task that runs daily to update entitlements and roles, and a different instance of Identity Refresh that runs weekly to do analysis of policy violations.

Other tasks have a specialized purpose and do not support parameters or multiple instances. An example of this type of task is the Check Expired Work Items task.

How Tasks Are Run

Tasks can be run on demand from the Task UI, by right-clicking on a task in the Task tabs and choosing Execute in Background, or by choosing Save and Execute after you have configured a particular task.

Tasks can be scheduled by right-clicking on a configured task in the Tasks tab and choosing Schedule. For more information, see [How to Schedule a Task](#) and [Working with Schedules](#).

There is a specialized task called the Sequential Task Launcher that lets you run a series of tasks sequentially, without having to schedule each of them individually. For more information, see [Sequential Task Launcher](#).

Note: Access to the Tasks feature, and the ability to run tasks or only view tasks, are controlled by IdentityIQ Capabilities and scopes. Talk to your system administrator if you do not have access to the Task features you need.

Process Overview for Creating Tasks

The basic process for creating a new instance of task is as follows. Specific details about parameters and options for each type of task is in [Task Types](#).

- Choose the type of task you want to run from the New Task dropdown.
- Give the task instance and name, and configure the task parameters to define the behavior you want.
- Save the task. If you want to run the task immediately, choose **Save and Execute**. If you want to save the task without running it yet, choose **Save**. When you save a task, an entry is created for it on the Tasks tab of the [Tasks Page](#).
- If you want to schedule the task to run at a specific time, or to repeat on a specific schedule, right-click the task in the Tasks tab of the Tasks Page and choose **Schedule**. Scheduled tasks are listed on the [Scheduled Tasks Page](#). For more information on how to schedule tasks, see [How to Schedule a Task](#) and [Working with Schedules](#).
- When a task has completed, you can see task results from the [Task Results Page](#).

Working with Tasks

To run or execute a task, right-click on the task name and select **Execute** or **Execute in background**. **Execute** displays a pop-up progress window and opens the Task Result page when it is complete. **Execute in background** launches the task in the background and you must go to the Tasks Results page to track progress or view the finished task. See [Task Results Page](#).

Tasks that require sign off generate work items and email notifications that are assigned to the designated signers. Sign off decisions are retained with the task results for tracking purposes. See [How to Complete Task Work Items](#).

To create a new task, use the **Create new task** dropdown list to select a task type and display the New Task page.

Note: The predefined tasks are not templates that can be used to create new tasks. Changes made to these tasks overwrite existing information. To create new task you must use the **Create New Task** dropdown menu at the bottom of the page.

See [How to Create a New Task](#).

To edit an existing task, click a task or right-click and select **Edit** to display the Edit Task page. See [How to Edit a Task](#).

To schedule a task, right-click and select **Schedule** from the dropdown list to display the New Schedule dialog. You can schedule task to run once, hourly, daily, weekly, monthly, quarterly or annually to meet the requirements of your enterprise and auditors. Go to the Scheduled Tasks tab to view or edit existing schedules. See [How to Schedule a Task](#) and [Scheduled Tasks Page](#).

To terminate a currently running task, access the Task Results page, right-click on the task to terminate and select **Terminate** from the dropdown menu. You are asked to confirm the termination request. Task that are currently running are flagged as pending in the Date Complete column of the Task Results table.

To delete a task, right-click the task and select **Delete** from the dropdown menu. Click **Yes** on the confirmation pop-up to delete the task. When you delete a task from the Tasks table, all associated task results are deleted as well.

How to Create a New Task

Use the New Task page to create a task based on the task types provided. Tasks can be as general or specific as required.

See [Task Types](#) for the complete list of tasks types available in IdentityIQ.

1. Click or mouse over the Setup tab and select **Tasks** to open the Tasks page.
2. Select a task type from the **New Task** dropdown list to open the New Task page.

3. Enter a **Name** and brief **Description** for the new task. This information is displayed on the Tasks table when the new task is saved.
4. Select a **Previous Result Action** from the dropdown list. **Delete** is selected by default. Previous result actions determine how subsequent runs of this tasks react to existing task results.
 - a. **Delete** – overwrite the previous task results with the new information.
 - b. **Rename Old** – append a numeral to the name of the old task result and preserve both.
 - c. **Rename New** – append a numeral to the name of the new task result and preserve both.
 - d. **Cancel** – cancel the new run of the task.
 - e. *Optional: Allow concurrency.* Select **Allow Concurrency** to enable two identical tasks to run at the same time. If enabled, allow concurrency appends a numeric value to the name of the task that started second. If disabled, the second task is canceled and an exception sent to the requestor.
5. *Optional: Require sign off.*
 - a. Select **Required sign off** to expand the Signoff Properties section.
 - b. Select an email notification template from the Initial Notification Email dropdown list. For example, the Task Result Signoff template.
Templates are created and defined when the application is configured.
 - c. Specify the escalation criteria for the sign off request. Use the options displayed to set your escalation parameters.
None – no reminder emails are sent and no escalation is performed for this work item.
Send Reminders – email reminders are sent at the configured interval.
Reminders then Escalation – the configured number of reminders are sent and then the work item is escalated to the signers manager.
Escalation only – this work item is escalated after the configured interval with no reminders being sent.
 - d. Specify the required signers.
Enter the first letter, or letters, of an identity or workgroup to display a selection list of valid identities or workgroups containing that letter string or click the arrow to the right of the field to display all identities and workgroups and select a signer.
You can add as many signers as required.
6. *Optional: Host.* If you want to choose a specific host or set of hosts to run the task on, add a comma separated list of host names. If multiple hosts are specified, the task manager selects the first active host If there are no

active hosts, or if an incorrect host name is given, the task terminates, and an error message is left in the result.

7. **Email Task Alerts.** Choose an **Email Notification** type, if you want to send email notifications upon task completion, to the recipients whose email is being registered to receive the task status. These settings overwrite the email notification configured at the IdentityIQ Configuration level setting. Use these options to set your notification:

Caution: If your installation has enabled the task reanimator service on the host where this task is run, be aware that when a task hangs and is restarted by the reanimator service, email task alert notifications may not be sent. For more information on the task reanimator service, see the IdentityIQ **System Administration** documentation.

- **Disabled** – no email notification will be sent.
 - **Warning** – email notification will be sent in case of any warning after completion of task.
 - **Failure** – email notification will be sent in case of task Failure.
 - **Always** – email notification will be sent at completion of task irrespective of the task status.
 - **Email Notification Template** – (*Applicable only if **Disabled** is not selected*) Select **Task Status** template to send emails on task completion. Templates are customizable.
 - **Email Recipients** – (*Applicable only if **Disabled** is not selected*) Select the identity to register them to receive task status notification on emails associated with it.
8. Specify the task options required for the task you are creating. Each task type displays unique task options. See the [Tasks Page](#) for details on each type.
 9. Click **Save** to save the new task to the Tasks table, or Click **Save and Execute** to save the task to the Tasks table and run it immediately. The Tasks Results page displays when the task completes.

See [Task Results Page](#).

How to Edit a Task

Use the Edit Task page to make changes to an existing task.

Note: There is no **Save As** function on the Edit Task page. Any changes made to an existing task overwrite the task you are editing. You must use the **Create New Task** dropdown menu to create a new task.

Procedure

1. Click or mouse over the Monitor tab and select **Tasks** to open the Tasks page.
2. Click on a task, or right-click on a task and select **Edit** from the dropdown list to open the Edit Task page.
3. Edit the **Name** and **Description** section as needed.
4. Note: Changing the name does not save this as a new task and preserve the task being edited. Anything entered here overwrites the existing information.
5. Select a **Previous Result Action** from the dropdown list. **Delete** is select by default.

Previous result actions determine how subsequent runs of this tasks react to existing task results.

Delete – overwrite the previous task results with the new information.

Rename Old – append a numeral to the name of the old task result and preserve both.

Rename New – append a numeral to the name of the new task result and preserve both.

Cancel – cancel the new run of the task.

6. **Optional:** Allow concurrency. Select **Allow Concurrency** to enable two identical tasks to run at the same time.

If enabled, allow concurrency appends a numeric value to the name of the task that started second.

If disabled, the second task is canceled and an exception sent to the requestor.

7. **Optional:** Require sign off.
 - a. Select **Required sign off** to expand the Signoff Properties section.
 - b. Select an email notification template from the Initial Notification Email drop-down list. For example, the Task Result Signoff template.
Templates are created and defined when the application is configured.
 - c. Specify the escalation criteria for the sign off request. Use the options displayed to set your escalation parameters.
 - None** – no reminder emails are sent and no escalation is performed for this work item.
 - Send Reminders** – email reminders are sent at the configured interval.
 - Reminders then Escalation** – the configured number of reminders are sent and then the work item is escalated to the signers manager.
 - Escalation only** – this work item is escalated after the configured interval with no reminders being sent.

- d. Specify the required signers.

Enter the first letter, or letters, of an identity or workgroup to display a selection list of valid identities or workgroups containing that letter string or click the arrow to the right of the field to display all identities and workgroups and select a signer.

You can add as many signers as required.

8. **Optional:** Host.

If you want to choose a specific host or set of hosts to run the task on, add a comma separated list of host names. If multiple hosts are specified, the task manager selects the first active host. If there are no active hosts, or if an incorrect host name is given, the task terminates, and an error message is left in the result.

9. **Email Task Alerts.** Choose an **Email Notification** type, if you want to send email notifications upon task completion, to the recipients whose email is being registered to receive the task status. These settings overwrite the email notification configured at the IdentityIQ Configuration level setting. Use these options to set your notification:

Important: If your installation has enabled the task reanimator service on the host where this task is run, be aware that when a task hangs and is restarted by the reanimator service, email task alert notifications may not be sent. For more information on the task reanimator service, see the IdentityIQ **System Administration** documentation.

- **Disabled** – no email notification will be sent.
 - **Warning** – email notification will be sent in case of any warning after completion of task.
 - **Failure** – email notification will be sent in case of task Failure.
 - **Always** – email notification will be sent at completion of task irrespective of the task status.
 - **Email Notification Template** – (*Applicable only if **Disabled** is not selected*) Select **Task Status** template to send emails on task completion. Templates are customizable.
 - **Email Recipients** – (*Applicable only if **Disabled** is not selected*) Select the identity to register them to receive task status notification on emails associated with it.
10. Specify the task options required for the task you are creating. Each task type displays unique task options. See the [Tasks Page](#) for details on each type.
 11. Click **Save** to save the new task to the Tasks table, or Click **Save and Execute** to save the task to the Tasks table and run it immediately. The Tasks Results page displays when the task completes.

See [Task Types](#) for details on each type.

12. Click **Save** to save the new task to the Tasks table.

— OR —

Click **Save and Execute** to save the task to the Tasks table and run it immediately. The Tasks Results page displays when the task completes.

For more information, see [Task Results Page](#).

How to Schedule a Task

Use the New Schedule dialog to schedule tasks to run during times of low business activity. Schedule recurring tasks as needed to maintain routine compliance within your enterprise.

The New Schedule dialog enables you to assign a unique name and description to the task schedule. This information is stored on the Scheduled Tasks page and displays in the Task Results table.

See [Scheduled Tasks Page](#) and [Task Results Page](#).

Procedure

1. Click or mouse over the Monitor tab and select **Tasks** to open the Tasks page.
2. Right-click on a task name and select **Schedule** from the dropdown list to open the New Schedule dialog.
3. Enter a unique name and description for this schedule task.

Note: Task that run across time zones run at the time scheduled, relative to the time zone in which they are scheduled. For example, a task scheduled to run at 4:00 PDT runs at 1:00 EDT.

4. Enter the date and time to launch the first execution of this task.

You can enter the date manually, or click the [...] icon to select a date from the calendar.

— OR —

Select the **Run Now** field to schedule the task to run immediately after clicking **Schedule**. For recurring tasks, the task runs at the current time at the specified **Execution Frequency**.

5. Specify how often this task should run with the **Execution Frequency** drop-down list. Subsequent executions of this task occur at the time specified in the **First Execution** fields.
6. Click **Schedule** to save this scheduled task. You can go to the Schedule Tasks page to view a list of all scheduled tasks. See [Scheduled Tasks Page](#).

Working with Schedules

To edit an existing schedule, click a schedule name or right-click and select **Edit** to display the Edit Schedule dialog.

See [How to Edit a Schedule](#).

To delete a schedule, right-click the schedule name and select **Delete** from the drop-down menu. Click **Yes** on the confirmation pop-up to delete the schedule.

How to Edit a Schedule

Use the Edit Schedule dialog to edit task schedules:

1. Click or mouse over the **Monitor** tab and select **Tasks** to open the Tasks page.
2. Click on the **Schedule Tasks** tab to display the list of scheduled tasks.
3. Click a schedule name or right-click and select **Edit** to display the **Edit Schedule** dialog.
4. Edit the name or description for this scheduled task.

Note: Task that run across time zones run at the time scheduled, relative to the time zone in which they are scheduled. For example, a task scheduled to run at 4:00 PDT runs at 1:00 EDT.

5. Change the date and time to launch the first execution of this task.

You can enter the date manually, or click the [...] icon to select a date from the calendar.

— OR —

Select the **Run Now** field to execute the task immediately after clicking **Schedule**. For recurring tasks, the task runs at the current time at the specified **Execution Frequency**.

6. Specify how often this task should run with the **Execution Frequency** dropdown list. Subsequent executions of this task occur at the time specified in the **First Execution** fields.
7. Click **Save** to save this scheduled task and return to the Schedule Tasks page.

See [Scheduled Tasks Page](#).

Task Administration

Task administrators and system administrators have some tools both within the Task feature itself and in IdentityIQ's Administrator Console, to help them manage where and how tasks are run, and to troubleshoot common task issues.

Running Tasks on Specific Hosts

If your IdentityIQ instance is configured to use multiple hosts, you can configure tasks to run on specific host machines. This is done in the Task definition.

1. Navigate to **Setup > Tasks** and either create your new task, or select an existing task you want to edit.
2. In the **Host** field, enter the name of the host to use. You can specify more than one host in a comma-delimited list, to provide alternative options if the first-choice server is unavailable. Note that IdentityIQ does not validate the host names you enter; if you enter an invalid or inactive host, the task will fail.
3. If you specify more than one host, the task service will run the task on the first host in the list which responds (that is, which has an active heartbeat).
4. When the task is run, either manually or based on a task schedule, the host configuration you have specified will be used.

The task results include information about which host the task ran on.

Tracking Task Run Length for Troubleshooting

IdentityIQ keeps track of how long a task takes to execute, both per execution and on average. This is useful for evaluating when a task is running longer than normal, so that administrators can investigate and troubleshoot system or network problems.

IdentityIQ reports the number of runs and the average run time in the **task definition** (that is, the UI where you configure parameters and other settings for the task) and in the **task results**. The task definition shows the number of times the task has been run since the last reset of run statistics, along with the average run time for the task across those task executions.

The task result shows both the current run time and the average run time, as well as the percent change in run time of this execution vs. the average. The **Average Run Time** value indicates the average not including the current task run. The **Run Time** and **Run Time Change** values are not updated until the task completes (they are valued at zero during task execution). See [Task Results Page](#).

The **Tasks** page in the **Administrator Console** also shows this information for completed tasks. **Average Runtime** reflects the average time of runs prior to this execution. Tasks will only have average and difference values once they have been run multiple times. For more information, see the **System Administration** documentation.

Restarting Run Time Average Calculations

When system configuration changes occur which should substantially impact a task's run time, you may want to reset the statistics to ensure the average run time value represents clean and meaningful data. Within the task definition (**Setup > Tasks >** select the relevant task), click **Reset Run Statistics** to clear out the saved run count and average run time.

It is recommended that you reset run statistics any time you change the options selected for a task, since the selected options will directly impact the set of activities the task is performing and therefore its execution time. Likewise, changes to the designated Host for the task (if you have chosen to select specific host(s) for running the task) should also prompt resetting run statistics, as should any hardware changes on the host(s) which would affect task (and overall system) performance.

Administrator Console: Tasks Page

The Administrator Console page for managing tasks is accessible through the **gear menu > Administrator Console > Tasks**. On this page, administrative users can postpone a scheduled task, terminate a running task, or dump a stack trace of a running task to see what is currently happening in it (typically used when the task is running long and the administrator wants to diagnose problems). For more information, see the **System Administration** documentation.

Tasks Page

The Tasks page contains a list of all of the tasks that have been created. The first time you access the Task page you see the predefined tasks provided by SailPoint. The tasks are grouped into categories based on the task type. You can expand or contract the categories on the grid using the plus (+) or minus (-) icon next to the category name.

Note: Task category headings are only displayed if a task exists in that category.

The task categories are:

- Account Aggregation
- Account Group Aggregation
- Activity Aggregation
- Activity Alerts
- Certification Refresh
- Generic
- Identity
- Scoring
- System
- Target Aggregation

See [Predefined Tasks](#).

Use the search options to limit the number of tasks displayed in the table. Entering a letter, or partial name, in the **Search** field displays any tasks with names containing that letter pattern.

Use this page to create, edit, run, schedule or delete task.

See [Working with Tasks](#).

The Tasks page contains the following information:

Field Name	Description
Name	The name of the task as defined when the it was created.
Description	A brief description of the specific task.

Predefined Tasks

SailPoint provides a number of predefined tasks that can be run to aggregate, correlate and refresh information within your enterprise.

Note: The predefined tasks are not templates that can be used to create new tasks. Changes made to these tasks overwrite existing information. To create new task you must use the **New Task** dropdown menu at the bottom of the page.

Caution: These tasks are defined to perform specific functions within your enterprise. Deleting or altering these tasks might have negative affects on the performance of IdentityIQ.

SailPoint provides the following tasks:

Generic Tasks:

- Refresh Role Indexes – update all role information and create the indexes needed to perform role searches. You must run this task before performing any role searching.

Identity Tasks:

- Check Active Policies – scan all users for policy violations and update Identity Risks Scores. Edit this task to specify how policy violations are handled when detected.
- Prune Identity Cubes – delete identities that have no account links and have no important references. Identities in any of the following states are protected:
 - Marked protected
 - Is a manager (managerStatus flag true)
 - Has capabilities
 - Bundle, Application, Workitem, or TaskResult owners
 - Work item requestor
 - Application secondary owner
 - Application remediator
 - Creator of a MitigationExpiration

If the **protectIfCertifying** option is on, identities are protected if they are in an active certification. There is also an option to run the scan for analysis but not delete any identities.

- Refresh Entitlement Correlation – scan all user entitlements and applications to update role assignments.
- Refresh Groups – scan all users and update the group indexes for all identity groups.
- Refresh Identity Cube – perform a full refresh of the identity cubes for all users. Edit this task to specify which portions of the identity cubes are refreshed by this task.
- Refresh Risk Scores – scan all users and update the Identity Risk Scores for each.

Scoring Tasks:

- Refresh Application Scores – runs the scoring algorithms against all specified applications and updates the Application Risk Scores page.
- Refresh Role Scorecard – analyzes each role in the system and collects statistics about them.

System Tasks:

- Check Expired Mitigations – scans all users for temporary exceptions allowed in a certification that have now expired. The original certifier can optionally be notified when allowed exceptions expire.
- Check Expired Work Items – scans all work items looking for those that need to be canceled or escalated to a different user.
- Check Sunset Requests – send notifications when a Role with a sunset date is approaching expiration.
- Complete Orphaned Identity Requests – removes completed requests for roles that exist in your system.
- Dispatch Access History – transforms IdentityIQ objects to create a JSON message and dispatches it to the access history service.
- Effective Access Index Refresh – refreshes or rebuilds the effective access index.
- Full Text Index Refresh – builds and refreshes the index files used for full text searches on defined fields on the access request pages of the Lifecycle Manager. The index files are rebuilt each time this task is run.
- Perform Identity Request Maintenance – prunes old identity request objects and scans unverified access requests to check for provisioning completeness.
- Perform Maintenance – prunes identity snapshots, task results, and certifications, escalates orphaned work items, and performs other background maintenance tasks.

Note: Electronically signed objects are not affected by this task.

- Refresh Role-Entitlement Associations – deletes all existing role-entitlement associations and re-creates them for all roles.

Remove Orphan Role Requests – stops and removes requests for roles that no longer exists in your system. For example, if the sunset date for a role passes before the request is processed, this task removes that request.

- Role Overlap Analysis – performs impact analysis on a specified role. The task result name is annotated with the name of the selected role so you can tell multiple analysis results apart.
- Synchronize Roles – synchronizes IdentityIQ roles with the roles on the identity management systems that are configured to work through a provisioning provider.

Scheduled Tasks Page

The Scheduled Tasks page contains a list of all scheduled tasks, whether recurring or one time only. One-time tasks are removed from the list after they are executed.

Tasks that are scheduled, but do not execute due to malformed task definitions are displayed in the Scheduled Task table with an error icon (!) in the **Last Executed** column. Tasks that fail in this way never execute and, therefore, never display results on the Tasks Results page. To see details of the execution error, click on the task to display the Edit Schedule dialog. The error information is displayed in the **Last Launch Error** field. Errors of this type should only occur for custom task definitions, not for any of the tasks supplied with the product. To correct the error, delete the task schedule, correct the task definition, and recreate the schedule.

Use the Scheduled Tasks page to edit or delete schedules.

Use the search options to limit the number of tasks displayed in the table. Entering a letter, or partial name, in the **Search** field displays any tasks with names containing that letter pattern. Click **Result** to search by task results. See [Working with Schedules](#).

To create a scheduled task see [How to Schedule a Task](#).

The Scheduled Tasks page contains the following information:

Field Name	Description
Name	The name of the schedule as defined on the New Schedule page.
Next Execution	The date and time at which the task is next scheduled to execute.
Last Execution	The date and time at which the task most recently executed. This field displays an error icon (!) for tasks that do not execute due to malformed tasks definitions.
Last Result	The result of the last run of this task, for example Success or Failed.
Owner	The creator of the schedule.

Task Results Page

The Task Results page contains a list of all of the tasks that have run or are currently running.

Use the search options to limit the number of tasks displayed in the table. Entering a letter, or partial name, in the **Search** field displays any tasks with names containing that letter pattern. Click **Advanced Search** to filter by start date, end date, or results.

Column	Description
Name	The name of the task.
Date Complete	The date and time stamp of when the task completed running.
Result	The result status, Pending , Success , or Failed . A result of Success with an exclamation point (!) indicates that there are warnings associated with the results.
Signoff	The status of the sign off request for the task results. None – no sign off required Waiting – sign off request not complete Signed – a sign off decision has been made
Owner	The name of the user who launched this task.

Click on a task name in the Tasks Results table to display the Task Results page. Each task type returns information specific to the options that were selected. Tasks that executed with partitioning enabled also display the partitioned results, broken down by the host name of the partitions on which they ran.

Several statistics related to task run length are maintained to help identify tasks that are running longer or shorter than expected. Each time a task is run, we save the start time. When the task is complete we calculate the run time in seconds.

These statistics are not set until the task complete. Until then they are zero. The run time change is a positive or negative integer representing the percent change in run length for this task relative to the average at the time was started. A value of 25 means the task ran 25% longer than average, and a value of -10 means the task was 10% faster.

See [Task Types](#) for details on the information that might be on the Task Results page.

To terminate a currently running task, a task flagged as pending in the Date Complete column, right-click on the task and select **Terminate** from the drop-down menu. You are asked to confirm the termination.

To delete task results, right-click on a result and select **Delete**. Tasks that require a sign off can only be deleted by a user with the Signoff Administrator capability.

If a task was scheduled to run but no results were returned, go to the Scheduled Task tab to ensure that errors did not occur during the task execution.

Task Types

The task types are:

- **Access History** – gathers objects and places them in a queue to create history events related to identities. See [Access History Tasks](#).
- **Account Aggregation** – scan all applications, discover users and entitlements on those applications, and then correlate those users and entitlements with roles. See [Account Aggregation](#).
- **Account Group Aggregation** – scans applications and aggregates account groups and application object types. These are then used for group certification (either permissions or membership) or for displaying group information in identity certifications. See [Account Group Aggregation](#).
- **Activity Aggregation** – scan all applications, discover activity on the applications, and then correlate that activity with identity cubes. This enables you to track and monitor all activity for possible policy violations. See [Activity Aggregation](#).
- **Alert Aggregation** – scan applications and aggregates alerts from a set of Alert Collectors. These are then used to generate alert actions. See [Alert Aggregation](#)
- **Alert Processor** – process the aggregated alerts against the alert definitions and launch the appropriate action. See [Alert Processor](#)
- **Application Builder** – create multiple IdentityIQ applications or update the attribute map of an existing IdentityIQ application. See [Application Builder](#)
- **ArcSight Data Export** – export data for HP ArcSight Database Connector to an external database table. See [ArcSight Data Export](#)
- **Data Export** – generate a denormalized data report to export to an external database table. See [Data Export](#)
- **Effective Access Indexing** – generate an index of any indirect access that was granted through another object. For example a nested group, an unstructured target, or another role. See [Effective Access Indexing](#)
- **Encrypted Data Synchronization Task** – re-encrypt data with user-generated encryption key. See [Encrypted Data Synchronization Task](#)
- **Entitlement Role Generator** – scans the entitlements in the system and automatically generates a simple role and appropriates a profile for each one that it finds. See [Entitlement Role Generator](#)

- File Access Manager Classification – retrieve classification data from File Access Manager and assigns it to entitlements according to correlation logic defined in the applications that aggregate relevant account and group data or in the File Access Manager global configuration settings. See [File Access Manager Classification](#).
- ITIM Application Creator – inspect the IBM Tivoli Identity Manager (ITIM) and retrieve information about the ITIM services (applications). This task auto-generates an application for each service defined in ITIM. See [ITIM Application Creator](#)
- Identity IQ Cloud Gateway Synchronization – synchronize the specified objects to the Cloud Gateway. See [IdentityIQ Cloud Gateway Synchronization](#)
- Identity Refresh – scan all applications, including the IdentityIQ application, to ensure that all identity information is up-to-date and accurate. Refresh identity scans are also used to detect and report on policy violations and trigger event certifications. See [Identity Refresh](#).
- Identity Request Maintenance – scan for completed Lifecycle Manager access requests. See [Identity Request Maintenance](#).
- Missing Managed Entitlements Scan – scan the selected application to create entitlement objects for items added after the application was last aggregated. See [Missing Managed Entitlements Scan](#)
- OIM Application Creator – inspect the OIM application and retrieve information about all connected applications. See [OIM Application Creator](#).
- Policy Scan – runs policies against identity cubes and update identity score cards with any policy violations discovered. See [Policy Scan](#).
- Propagate Role Changes – refreshes identities who have an assigned role whose associated entitlements have changed. See [Propagate Role Changes](#).
- Refresh Logical Accounts – is used to refresh composite accounts for all identities that could, potentially, have a composite account on the composite applications selected. See [Refresh Logical Accounts](#).
- Role Index Refresh – updates all role information and creates the indexes needed to perform role searches. You must run this task before performing any role searching. See [Role Index Refresh](#)
- Run Rule – runs the specified rule with name / value pairs. See [Run Rule](#)
- Role-Entitlement Associations – deletes existing role-entitlement associations then analyzes each role in the system and creates associations between the role and any granted entitlements. See [Role-Entitlement Associations](#).

- Sequential Task Launcher – launches the specified tasks in the order defined. This enables you to launch tasks that must be run sequentially in the proper order without having to schedule each separately based on estimated run times. See [Sequential Task Launcher](#)
- [System Maintenance](#) – tasks designed to run in the background.
- Target Aggregation – scan selected applications for activity targets. See [Target Aggregation](#).

See [Tasks Page](#) for information on working with these task types.

All task types contain the following standard properties:

Field	Description
Name	The name of the task as defined when the task was created
Description	Brief description of the task.
Previous Result Action	<p>Previous result actions determine how subsequent runs of this task react to existing task results.</p> <p>Delete – overwrite the previous task results with the new information.</p> <p>Rename Old – append a numeral to the name of the old task result.</p> <p>Rename New – append a numeral to the name of the new task result.</p> <p>Cancel – cancel the new run of the task if a task result with the same name exists.</p>
Allow Concurrency	<p>Enable two identical tasks to run at the same time.</p> <p>If enabled, allow concurrency appends a numeric value to the name of the task that started second.</p> <p>If disabled, the second task is cancelled and an exception sent to the requestor.</p>
Require Signoff	<p>Require sign off on the results of this task.</p> <p>Tasks that require sign off generate work items and email notifications that are assigned to the designated signers. Sign off decisions are retained with the task results for tracking purposes.</p>
Host	<p>A comma separated list of host names on which to run this task. If multiple hosts are specified, the task manager selects the first active host</p> <p>If there are no active hosts, or if an incorrect host name is given, the task terminates, and an error message is left in the result.</p>
Number of Runs	The number of times this task has been run.
Average Run Time	The average time it takes to run this task, based on prior runs.

Field	Description
Reset Run Statistics	<p>Reset the statistic if you reconfigure the task and expect the run times to change.</p> <p>When you reconfigure complex tasks like aggregation or refresh, you should consider resetting run statistics. For example, enabling provisioning in the refresh task can profoundly influence run time so statistics should not be diluted by the previous average before provisioning was enabled.</p>
Email Task Alerts	
Email Notification	<p>Select a frequency for email notification to be sent upon task completion.</p> <p>Disable – no email notification sent on task completion</p> <p>Warning – send an email notification if the task results in a warning</p> <p>Failure – send an email notification if the task fails</p> <p>Always – always send an email notification upon task completion</p>
Email Notification Template	Select a notification email template from the dropdown list.
Email Recipients	<p>The list of users to receive the task completion notification.</p> <p>Use the dropdown arrow to display all identities, or type the first few letters of a name. select names from the list.</p>

Access History Tasks

Access History tasks should be run periodically to gather objects and place them in a queue to create history events related to identities. When Access History runs, it processes the objects from the queue and writes them into database tables to populate the access history store.

Before you can use the Access History task, Access History must be enabled.

Under Access History Options, select an Access History YAMLConfig from the dropdown menu.

Account Aggregation

Account Aggregation tasks scan all applications, discover users and entitlements on those applications, and, optionally correlates those users and entitlements with roles.

Identities that have changed since the last aggregation performed on an application are marked as needing refresh to increase the performance of identity refresh tasks. You can disable this function.

You can perform the correlation functions as part of this task or run account aggregation on all of the applications in your enterprise and then correlate the Identity Cubes with all of the aggregated information using an identity refresh task.

To perform aggregation on a composite application you must include the composite application and all of the applications that have accounts with which it is associated in the task definition.

Partitioning

Partitioning is available to speed the processing time for account aggregations and level the load on the machines running these tasks. Partitioning is used to break operations into multiple pieces, or partitions. Each partition is then placed in a global queue, and machines, or hosts, in a cluster compete to execute the partitions in the queue.

Machines are added or removed from the cluster dynamically with automatic balancing. If a machine fails or is taken down while processing a partition, the partition is placed back into the queue and reassigned to a different machine. A single result object is shared by all partitions and is continually updated so you can monitor the overall progress of the partitioned operation. When all partitions have finished executing the result is marked complete. See **the System Administration** documentation.

Note: You must run the Target Aggregation task after this task is complete if you have activity targets set. This task removes all targets when it is run. See [Target Aggregation](#).

Handling Duplicate Distinguished Names in Active Directory

During aggregation, if IdentityIQ detects two Active Directory accounts or account groups with the same Distinguished Name but different UUIDs, it will update the UUID to the most recent value, and treat the two accounts or account groups as the same. This handles the case where an account or group is accidentally deleted and readded. Consequently, it is not advisable to reuse the same DN with a different meaning. IdentityIQ will not detect this as an account or account group change to any attribute but UUID.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

Option	Description
Select an application to scan	The dropdown list of all applications.
Optionally select a rule to assign capabilities or perform other	<p>If accounts are discovered that do not have matching identities in the IdentityIQ application, the rule specified here is used to create a new Identity Cube.</p> <p>These rules are created during configuration and deployment.</p> <div style="background-color: #e6f2ff; padding: 5px; margin-top: 10px;"> <p>Note: Click the [...] icon to launch the Rule Editor to make</p> </div>

Option	Description
processing on new identities	changes to your rules if needed.
Refresh assigned and detected roles	Scan for newly assigned roles and update Identity Cubes. IMPORTANT: it is not recommended to use this setting if you are also using the Identity Refresh task to manage entitlement correlation or provisioning, as this can result in unintended or incomplete provisioning.
Check active policies	Scan for policy violations and update Identity Cubes.
Only create links if they can be correlated to an existing identity.	Check to updated existing identities, but not to create new identities if a match is not found
Refresh the identity risk scorecards	Scan for risk score information and update identity risk score cards.
Maintain identity histories	Compare current Identity Cubes to existing Identity Cube history, snapshots, and create new snapshots if any changes are discovered.
Enable Delta Aggregation	Enable the connector to aggregate only those accounts that have changed since the last aggregation. This requires support by the connector.
Detect deleted accounts	Compare current aggregated accounts with the accounts previously aggregated and report any deleted accounts. Maximum deleted accounts: This is the maximum number of accounts that can be flagged for deletion after an account aggregation. If this number is passed, no accounts are deleted from the application.
Refresh assigned scope	Refresh assigned scope based on changes discovered during the aggregation and correlation process.
Disable auto creation of scopes	Do not automatically assign scope to identities as part of this task.

Option	Description
Disable optimization of unchanged accounts	Use this option to force the aggregation of all accounts, changed or unchanged since the last aggregation.
Promote managed attributes	When enabled, any values for entitlement or permissions encountered while running the task automatically get promoted as managed attributes.
Enable rename detection on managed attributes	<p>This option affects aggregation from Active Directory. It enables IdentityIQ to detect when an account group DN has changed due to being renamed. IdentityIQ determines whether a DN is new or is a rename of an existing DN, by examining the relevant account group's GUID or UUID. Enabling this option can prevent unintended changes to access that is based on assignment rules which use DN as assignment criteria. For more information on Active Directory moves and renames, see the IdentityIQ Application Management documentation.</p> <p>Note that when a change is made in Active Directory to an OU which contains accounts or groups (such as renaming or moving it), a delta aggregation does not pick up the changes. This is due to a limitation in Microsoft DirSync Control.</p> <p>To avoid this issue, perform a full aggregation to capture the changes and update the child objects. You might have to do this regularly to ensure the data is up to date.</p>
Disable auto-creation of applications	Do not automatically create application objects for multiplexed accounts.
Disable marking the identity as needing refresh	<p>Disable marking only identities on which change was detected as need to be refreshed.</p> <p>All identities are included in subsequent identity refresh task.</p> <p>For more information on using this option to optimize performance, see Refreshing Changed Identities Only (Delta Identity Refresh).</p>
Enable Partitioning	<p>Enable partitioning of this task across multiple hosts.</p> <p>Partitioning is not supported for PE2 based connectors.</p> <p>Partitioning has to be configured on the applications and connectors before this option is valid.</p>
Objects per partition	If the connector(s) for the selected application(s) do not support partitioning, use this field to specify the number of objects per partition. The default value is 1000.

Option	Description
Loss Limit	The loss limit sets the maximum number of accounts that will be reprocessed in case of a sudden termination of a partitioned refresh. This option is used only when partitioning is enabled. See the System Administration documentation.
Terminate when maximum number of errors is exceeded	<p>Terminate after the specified number of errors occurs.</p> <p>If the database is available, the task result contains a message indicating that the task was terminated due to excessive errors. If the database is down, the task result cannot be persisted and the task might appear to remain in the pending state.</p> <p>Maximum errors before termination Number of errors to tolerate before terminating the task.</p>
Sequential Execution - Terminate on Error	Force applications to aggregate in the listed order and stop the aggregation task if an error is encountered.
Actions to include in the task result	<p>Select the actions performed as part of the aggregation task for which detailed information should be included in the task results.</p> <p>This task performs a number of individual actions on accounts and Identity Cubes during the aggregation and configuration processes. By default only the final results of the task are included in the task results report.</p> <p>To include detailed information on the actions performed as part of the task, select those actions from the list.</p> <p>Correlate Manual – identities with accounts that were manually correlated. These are not changed by the task.</p> <p>Correlate Maintain – correlation information has not changed since the last time this task ran.</p> <p>Correlate New Account – a new account was discovered for an existing identity and assigned.</p> <p>Correlate Reassign – an existing account was reassigned from one identity to another as part of the correlation process.</p> <p>Create New Identity – an account was discovered for an identity that did not exist in IdentityIQ. An Identity Cube was created for the new identity.</p> <p>Ignore – an account for a new identity was discovered, but a new Identity Cube was not created. This might occur if this task is configured to perform correlation only.</p> <p>Remove Account – an account discovered as part of a previous aggregation</p>

Option	Description
	was not found during this aggregation. These accounts are removed from IdentityIQ.

Account Group Aggregation

An Account Group Aggregation task scans applications and aggregates account groups and application object types. These results are then used for group certification (either permissions or membership), for displaying group information in certifications, and for performing identity searches.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

Option	Description
Select applications to scan	The dropdown list of all applications configured to work with IdentityIQ.
Filter object types to scan	<p>This option is only available for applications on which multiple application objects can exist.</p> <p>This option is not available if you select to scan more than one application.</p> <p>The list of all object types or account groups associated with the selected application. If nothing is selected, all object types and account groups are included.</p> <p>It might become important to scan object types separately if they share attributes.</p>
Enable Delta Aggregation	Enable the connector to aggregate only those account groups or application objects that have changed since the last aggregation. This requires support by the connector.
Detect deleted account groups	Detect and delete any account group or application object that was deleted on the native application since the last aggregation task was run.
Automatically promote descriptions to this locale	The default locale for the description attribute of the account group or application object. This option is used if an existing description locale is not found.
Description attribute (default description)	<p>The Description Attribute defined in the Application Group Schema overwrites any value set here.</p> <p>The attribute that stores the description. This value defaults to the value description if this option is not set.</p>

Option	Description
Group Aggregation Refresh Rule	The rule used to set the owner or modify the account group when it is created or refreshed. Click the [...] icon to launch the Rule Editor to modify the rule if needed.
Promote Classifications	Promote classification from the ResourceObject classification to the ManagedAttribute.
Enable partitioning	Enable partitioning of this task across multiple hosts. Partitioning must be configured globally before this option can be used. See the System Administration documentation.
Number of partitions	Specify a number of partitions. If no number is specified, IdentityIQ calculates an optimal number based on available request servers.
Loss Limit	The loss limit sets the maximum number of identities that will be reprocessed in case of a sudden termination of a partitioned refresh. This option is used only when partitioning is enabled. See the System Administration documentation.
Terminate when maximum number of errors is exceeded	Terminate after the specified number of errors occurs. If the database is available, the task result contains a message indicating that the task was terminated due to excessive errors. If the database is down, the task result cannot be persisted and the task might appear to remain in the pending state. Maximum errors before termination Number of errors to tolerate before terminating the task.

Activity Aggregation

Activity Aggregation tasks scan all applications, discover activity on the applications, and then correlate that activity with identity cubes. Using these tasks enables you to track and monitor activity within your enterprise.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

Option	Description
Select an activity data source	The dropdown list of all activity data sources discovered by IdentityIQ. If no data source is selected, all available data sources are scanned as part of the task. Your applications must be configured to support activity tracking.

Option	Description
Enable storage of the last activity position scanned on the data source	<p>If enabled, the task marks the last activity scanned on the data source so that subsequent runs of this task begin scans at that mark instead of rescanning information.</p> <p>If this option is not enabled, each run of the task scans the entire data source.</p>
Store uncorrelated activities so they can be re-scanned and correlated at a later time	If enabled, all activity discovered on the data source is stored, even if that activity does not correlate to a user in the application. Storing this activity enables you to add users and update their identity cubes without having to rescan your data sources.

Alert Aggregation

Alert Aggregation tasks scan applications and aggregates alerts from a set of Alert Collectors. These are then used to generate alert actions.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

Option	Description
Select sources to scan	<p>The dropdown list of all alert data sources discovered by IdentityIQ.</p> <p>If no data source is selected, all available data sources are scanned as part of the task.</p> <p>You applications must be configured to support alert tracking.</p>
Enable Delta Aggregation	<p>If enabled, the task only aggregates alerts that have occurred since the last run of this task.</p> <p>This option requires support from the connectors being scanned.</p>
Process Alerts	<p>If enabled, an Alert Processor task will launch as soon as this Alert Aggregation task is complete.</p> <p>If you select this option, you can limit the alert types processed in a comma separated list, or process all of the alerts collected.</p>

Alert Processor

Alert Processor tasks process the aggregated alerts against the alert definitions and launch the appropriate action.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

Option	Description
Optional filter string to constrain the alerts processed	If not set, all are processed.
Exclude alerts previously processed	Enable to exclude Alerts that were previously processed.
Optional filter string to constrain the alert definitions to match against alerts	If not set, all Alert Definitions are evaluated.
Enable Partitioning	Enable the task to split into partitions and run across multiple threads and hosts, if available.

Application Builder

The Application Builder task lets you create multiple IdentityIQ applications, and update existing applications in bulk. The task also includes the ability to perform account and group aggregation for a host using the associated application. It can also export essential data about your existing applications.

The task accepts the inputs required to create or update applications from a .csv file. Sample.csv files for Linux-Direct and Windows-Local are provided with this task as examples of how input data can be defined. The sample files are located in the WEB-INF/config directory of your IdentityIQ installation. You can also use the task's **Read** option to create .csv files from your existing applications, to use as models for creating .csv files that support the create and update options.

By default, before creating or updating an application on IdentityIQ, a test connection is performed to ensure that the connector is performing correctly. To skip the Test Connection operation, use **Skip Test Connection** in the Application Builder options.

To enable logging for the Application Builder task, add this entry to the log4j2.properties file:

```
logger.ApplicationBuilderExecutor.name=sailpoint.task.ApplicationBuilderExecutor
logger.ApplicationBuilderExecutor.level=debug
```

Before using the task to update an existing application, it is recommended that you use the iiq console to export the application definition, in case you need to restore them to their original state.

When you use this task to **Update** an existing application, the update is partial; that is, the update operation can add new attribute definitions to an existing schema, as well as adding a new schema.

Use the account or group aggregation options to trigger a background aggregation task.

Working with Flexible Schemas and Provisioning Forms

The Application Builder task supports including XML definitions in your csv files if you need to create or update flexible account schemas, or provisioning forms. Refer to the sample.csv files provided with this task for examples of how a

schema definition can be included in the.csv file. Sample files are provided in the WEB-INF/config directory for Linux-Direct and Windows-Local.

If your input file includes an XML definition of a Provisioning Form, be aware that importing a Provisioning Form definition in a create or update operation will replace all existing Provisioning Forms with the new form as defined in the .csv

Option	Description
Application Type	Select an application type from the drop-down list. This is type of application you want to bulk-process. A single application builder task can only process applications of the same IdentityIQ-supported type, such as JDBC, Active Directory, or LDAP
Operation	<p>Select an operation from the drop-down list.</p> <p>Create – create multiple applications by providing parameters in the .csv file in the specified format</p> <p>Update – update existing applications by providing parameters in the .csv file in the specified format</p> <p>Read – export existing applications to the .csv file format. Any existing exported files will be overwritten if the task is run again using the same filename.</p> <p>The Read operation reads the attribute map, account schema, and provisioning policy of an existing application present in IdentityIQ and exports it to the file path provided in CSV format. You must provide the application type and file path to which the file is to be exported before running the operation.</p>
File Path	<p>The file path, including file name, for the.csv file. For the Read option, this is the path to the location and name of the file the task will create. For Create and Update options, this is the path to the file containing the data for creating or updating your applications; these files must be present on the application server or accessible within the network.</p> <p>Sample .csv files are provided in the WEB-INF/config directory for Linux-Direct and Windows-Local:</p> <p>Application-builder_linux.csv</p> <p>Application-builder-windows-local.csv</p>
Account Aggregation	<p>Executes the account aggregation task. The account aggregation task is triggered sequentially.</p> <p>The aggregation task will use the following format; the UID (unique identifier) is generated automatically:</p> <p><Application type> + <Account Aggregation> + <Current time stamp> + <UID></p>

Option	Description
Group Aggregation	Executes the group aggregation task. The group aggregation task is triggered sequentially. The aggregation task will use the following format; the UID (unique identifier) is generated automatically: <Application type> + <Group Aggregation> + <Current time stamp> + <UID>
Number of Applications per Aggregation Task	The number of application included in each aggregation task. Default: 10
Skip Test Connection	Skip the default test connection operation.

ArcSight Data Export

Export data for HP ArcSight Database Connector to an external database table.

The ArcSight data export task enables you to export IdentityIQ data to external tables.

Before you can use the ArcSight data export task, you must create the export databases on your destination data source.

The task schedule user interface includes a button that generates a customized DDL which you can hand off to a database administrator for execution. Once the data source parameters are entered, click **Generate Table Creation SQL**. The task adds the following tables in database:

Tables	Description
sptr_arcsight_export	Table to maintain the task execution history.
sptr_arcsight_identity	Table contains exported data of Identity.
sptr_arcsight_audit_event	Table contains Audit Events information.

Option	Description
Datasource Parameters	
Database	Select a database type from the dropdown list.
User Name	Enter the user name parameter of the database table.
Password	Enter the password of the database table.
Driver	Enter the driver class used for database.

Option	Description
Class	
URL	Enter the URL of the database.
Object Export Options	
Export Identities	Export Identity related data in ArcSight tables. It provides the following options: Full: Exports all the records irrespective if they were exported earlier. Incremental: Exports only records that are updated since last run of this task. This option can even be selected when running the task for first time. When the task is running for first time, this option exports all records similar to the Full option.
Export Audits	Export Audit Events in ArcSight table. It provides the following options: Full: Exports all the records irrespective if they were exported earlier. Incremental: Exports only records that are updated since last run of this task. This option can even be selected when running the task for first time. When the task is running for first time, this option exports all records similar to the Full option.

After you complete customizing your task options, click **Save** for later use or **Save and Execute** to save the task and run it immediately.

Configuring HP ArcSight Task to populate host name or IP

The value of column `application_host` can be populated by adding a map with the value as `arcsightAppNameHostMap` as shown in the following example. The field `This` is read from the map as explained below:

It is difficult to determine the host name or IP address of the account as the field is not constant in Application definition in IdentityIQ. Hence, customer can define a map in TaskDefinition and select the task added to export data in ArcSight table. The key in the map should be name of the application defined in IdentityIQ and value should be `hostname`, `IP`, or any string that ArcSight administrator understands.

To add the map:

1. Go to debug page, navigate to TaskDefinition and open the ArcSight task configured above.
2. Add the entry as key = Name of Application defined in IdentityIQ and value as the string to identify host of Account like `Hostname` or `IP`.
3. Save the task definition. For example:

```
<entry key="arcsightAppNameHostMap">
  <value>
```

```

    <Map>
      <entry key="LinuxApp1" value="linux01.iiq.com"/>
      <entry key="LinuxApp2" value="127.15.19.21"/>
      <entry key="ADDirectApp" value="AD.iiq.com"/>
      <entry key="ServiceNowApp" value="https://iiq.service-now.com"/>
      <entry key="ACF2App" value="ACF2-Mainframe"/>
    </Map>
  </value>
</entry>

```

Note: If the application name is not defined in the map the host field is blank.

The following fields are added in export table:

Fields	Description
linkid	Primary key for Link table in IdentityIQ database. This field is copied from spt_link table id field and is the primary key for export table.
identityid	Primary key in Identity table. This field is copied from spt_Identity table.
modified_dt	Populates timestamp when the record is exported in export table. The field can be referred while configuring time based ArcSight database connector.
identity_display_name	Represents Display Name of Identity which is copied from spt_identity table field (display_name).
identity_first_name	Represents first name of Identity which is copied from spt_identity table field (first-name).
identity_last_name	Represents last name of Identity which is copied from spt_identity table field (last-name).
application_type	Populates the type of Account which is connected to the Identity like ActiveDirectory – Direct, ACF2 – Full, Box, Cloud Gateway, ServiceNow and so on.
application_host	The host name, IP, or any string which can be used by ArcSight administrator to identify the host of link/account uniquely. Customer can enter any string which can be sent to ArcSight to identify the host of link. This field can be populated as explained in ArcSight Data Export .
application_name	Populates the name of Application of the Account connected to the Identity.
link_display_name	The account connected to the identity which is copied from spt_link table, field display_name.
entitlements	Represents comma separated list of entitlements to the link of Identity.
risk_score	Represents the composite risk score of Identity.

Fields	Description
auditid	The audit ID which is primary key for the export Audit table. The field is copied from spt_audit_event table id field.
created_dt	Populates timestamp when the record is exported in export table. The field can be referred while configuring time based ArcSight database connector.
owner	Describes the Owner of the audit generated.
source	Provides more details to help ArcSight administrator determine the source of audit.
action	Describes the action taken on entity.
target	Provides target details.
application	Describes the name of application the target belongs to.
account_name	The name of Account is populated in this field.
attribute_name	The name of attribute modified.
attribute_value	The value provided to the attribute.

Data Export

The Data Export task enables you to export IdentityIQ data to an external database. You can select to export any combination of identity, account, and certification data.

Before you can use the Data Export task, you must create the export database tables on your destination data source.

The task schedule user interface includes a button that generates a customized DDL which you can hand off to a database administrator for execution. Once the data source parameters are entered, click **Generate Table Creation SQL**.

Option	Description
Datasource Parameters	
Database	Select a database type from the dropdown list.
User Name	Enter the user name parameter of the database table.
Password	Enter the password of the database table.
Driver Class	Enter the driver class used for database.
URL	Enter the URL of the database.
Object Export Options	
Export Iden-	Export identity related data. You can perform a full or incremental export. Use

Option	Description
titles.	the Export Filter field to apply any database filters.
Export Accounts.	Export account related data. You can perform a full or incremental export. Use the Export Filter field to apply any database filters.
Export Certifications.	Export certification related data. You can perform a full or incremental export. Use the Export Filter field to apply any database filters.

After you complete customizing your task options, click **Save** for later use or **Save and Execute** to save the task and run it immediately.

Effective Access Indexing

Effective Access is any indirect access that was granted through another object, such as a nested group, an unstructured target, or another role.

Option	Description
Index Entitlement Targets	Include any effective entitlements associated with application that support effective access searching.
Index Role Targets	Include any effective roles associated with application that support effective access searching.
Index direct role permissions	Include any effective direct role permissions associated with application that support effective access searching.
Index direct entitlement permissions	Include any effective direct entitlement permissions associated with application that support effective access searching.
Index unstructured targets	Include any unstructured targets.
Refresh Fulltext Index	Refresh the Fulltext index as part of this task.
Index classifications	Add an entitlement's classifications to the target association that is created when the entitlement target is indexed; in the UI, this means that an entitlement's classifications will be displayed whenever that entitlement occurs as Effective Access.
Promote classifications	Promote classifications up the effective access "chain" to the entitlement that grants the effective access. For example, if EntitlementA grants you effective access to EntitlementB, and EntitlementB has a classification assigned to it, then with the Promote Classifications option enabled, the classification

Option	Description
	assigned to EntitlementB will also be displayed in the UI for EntitlementA.
Index elevated access	This will make the elevated access for a role or entitlement show as effective access if it is associated with other objects
Promote elevated access	Marks the role as Elevated if it has any target associations. They must be indexed first.
Clean elevated access	Removal of any items marked as Promote Elevate Access if the association is broken or removed.
Delete all current targets before indexing	Clear an existing Effective Access Index before running this task.

After you complete customizing your task options, click **Save** for later use or **Save and Execute** to save the task and run it immediately.

Encrypted Data Synchronization Task

The Encrypted Data Synchronization Task is used to re-encrypt IdentityIQ data when a new custom encryption key is generated.

Option	Description
Disable Application Synchronization	Select this option to ignore encryption key synchronization against applications.
Disable Identity Synchronization	Select this option to ignore encryption key synchronization against identities.
Disable IntegrationConfig Synchronization	Select this option to ignore encryption key synchronization against IntegrationConfig objects.
Disable Attachment Synchronization	Select this option to ignore encryption key synchronization against attachments.
Convert Encrypted Identity Secrets to Hashing	Select this option to convert any encryption keys to use hashing.

Once you have completed customizing your task options, click **Save** for later use or **Save and Execute** to save the task and run it immediately.

Entitlement Role Generator

The Entitlement Role Generator creates an Entitlement Role for every entitlement found in a specified application. Recommended role types are Entitlement or IT.

You can further refine creation by specifying an entitlement name or permission target so that only entitlements matching the specified criteria are used.

It is recommended to specify a template to be used to name the created roles. IdentityIQ uses Velocity templates. If no template is used, a generic name based on either the entitlement or role is created.

Option	Description
Applications	Select one or more applications from the dropdown list.
Type of Role to Create	Input the name of the role based on the specifications for your enterprise.
Enter the locale to check for descriptions. (If left blank the default Locale is used)	Enter the location of the role description.
Generate entitlements from attributes whose name starts with	Enter letters in the attribute name to filter the scan.
Generate entitlements from permissions whose target starts with	Enter letters in the permission name to filter the scan
Velocity template from which to generate entitlement role names. The template is always passed the applicationName parameter. The description, attributeName, attributeValue, permissionTarget, and/or permissionRights parameters are set when available.	Enter the Velocity template string.

File Access Manager Classification

The File Access Manager Classification task is used when you are integrating with File Access Manager, to use File Access Manager's classification to flag and categorize entitlements within IdentityIQ. This task retrieves classification data from File Access Manager and assigns it to entitlements according to the correlation logic that is defined in the applications that aggregate relevant account and group data, or in IdentityIQs File Access Manager global configuration settings.

Option	Description
Classification Customization Rule	You can use a rule to customize your classification object, for example to add or modify attributes in the object. Rules must be of the type "ClassificationCustomization" to appear in this selection list.

Option	Description
Automatically promote descriptions to this locale	The locale that any description that is included in the File Access Manager objects will be promoted to, by default. This is used if an existing description locale is not found.

Classification Filter Rule

The Classification Filter Rule feature on the File Access Manager Classification task is exposed as of version 8.4. It is used to narrow the scope of Classifications for the task. Like most rules, it accepts a SailPointContext and a Logger.

The return value is a QueryOptions that contains a Filter that is used to generate the SCIM query. The Filter supports "and" operations but not "or." If expression values include reserved URL characters, they need to be encoded.

Page size is now configurable on the task. This argument specifies the number of records fetched with each SCIM call to File Access Manager.

Several arguments added to the File Access Manager Classification task allow users to adjust its tolerance. They are Retry Limit, Retry Gap, and Max Errors.

Option	Description
Retry Limit	The number of times to retry a failed query before giving up and moving on.
Retry Gap	The number of milliseconds to wait before retrying a failed query
Max Errors	The number of times that the retry limit can be exceeded before giving up on the task.

After you complete customizing your task options, click **Save** for later use, or **Save and Execute** to save the task and run it immediately.

ITIM Application Creator

Run the ITIM Application Creator task to inspect IBM Tivoli Identity Manager (ITIM) and retrieve information about the ITIM services (applications). This task auto-generates an application for each service defined in ITIM. Each ITIM application contains a list of services that are roughly equivalent to the list of applications maintained in IdentityIQ. The applications generated by this task are added to the list of applications in IdentityIQ.

Option	Description
ITIM Applications	Select applications to inspect and from which applications should be generated based on the services found.
Generated application	Specify a prefix to append to any applications created by this task.

Option	Description
name prefix	
Generated application name suffix	Specify a suffix to append to any applications created by this task.

IdentityIQ Cloud Gateway Synchronization

IdentityIQ Cloud Gateway Synchronization tasks scan selected IdentityIQ applications for specified objects and synchronizes them with IdentityIQ Cloud applications. It is intended for use when IdentityIQ is not able to communicate directly with the managed system.

Option	Description
IdentityIQ Cloud Gateway application name	Select the name of the application to synchronize.
Applications hosted on the IdentityIQ Cloud Gateway	Select the name of the hosted cloud gateway application with which to synchronize the IdentityIQ application.
Rules to be executed on the IdentityIQ Cloud Gateway	Select which rules to execute against selected applications.

Identity Refresh

Refresh identity tasks scan all identities to ensure that all identity information is up-to-date and accurate. Refresh identity scans are also used to detect and report on policy violations and launch event certifications.

Incremental identity refresh can be configured to only refresh those identities on which information has changed since the last refresh was performed, to increase performance.

Note: Partitioning is disabled if you enable **Mark dormant scopes after refresh** or **Refresh the group scorecards** options.

Note: The **Number of Refresh Threads** option is not supported when partitioning is enabled.

Partitioning is available to speed the processing time for identity refresh tasks and level the load on the machines running these tasks. Partitioning is used to break operations into multiple pieces, or partitions. Each partition is then placed in a global queue, and machines, or hosts, in a cluster compete to execute the partitions in the queue. Machines are added or removed from the cluster dynamically with automatic balancing. If a machine fails or is taken down while processing a partition, the partition is placed back into the queue and reassigned to a different machine. A single result object is shared by all partitions and is continually updated so you can monitor the overall progress of the partitioned operation. When all partitions have finished executing the result is marked complete. See the **System Administration** documentation.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

To reduce or eliminate the possibility of getting an `ObjectAlreadyLocked` exception, there are additional parameters available on the IdentityIQ debug pages. `enableTriggerIdentityQueue`, set to `true` to enable the queuing feature, and `triggerIdentityQueueSize`, to specify the number of triggers to queue prior to processing, without this setting, the default is 10.

Option	Description
Optional filter string to constrain the identities refreshed	A filtering string used to limit the number of identity cubes updated by this task. For example you can limit the refresh to one department within your enterprise, such as Finance, by entering: <code>department == "Finance"</code>
Optional list of group or population names to constrain the identities refreshed	A filtering string used to limit the number of identity cubes updated by this task. For example you can limit the refresh to one group or population within your enterprise.
Refresh identities whose last refresh date is before this date	<p>Refresh any identities not refreshed since the date entered.</p> <p>Enter and date manually or click the [...] icon to display the calendar view.</p> <p>Use this to recover from a refresh that ended abnormally. For example, you start a refresh task and it runs for a day before stopping abnormally. After resolving the issue with the task, instead of repeating the refresh of all the identities that completed before the task stopped, you can only refresh the ones that were missed on the last refresh. Enter the approximate date the last refresh stopped and only refresh the remainder.</p>
Refresh identities whose last refresh date is at least this number of hours ago	<p>Enter the number of hours manually.</p> <p>Use this option to refresh identities that have not been refreshed recently. The time in this option is relative rather than absolute. Instead of remembering a specific task launch date and typing that in each time you run the refresh task you can have just one task and run that repeatedly. For example you can run it for every thing more than an hour old.</p>
Refresh identities whose last refresh date is within this number of hours	<p>Enter the number of hours manually.</p> <p>Use this option to refresh identities that were refreshed recently. The primary use case for this is to refresh things that were recently</p>

Option	Description
	<p>touched by aggregation.</p> <p>For example, if you have several aggregation sources but those sources tend to touch different subsets of all identities, and you would rather not refresh the identities that were not touch be the last aggregation.</p>
Include modified identities in the refresh window	<p>Refresh any identities modified within the specified time frames.</p> <p>There are two dates stored on each Identity, the date of last refresh and the date of last modification.</p> <p>The last refresh date is set whenever you run the refresh or aggregation tasks and the identity is changed in some way.</p> <p>The last modification date is set whenever you edit the identity in some way outside of a refresh or aggregation task, for example from a Lifecycle Manager workflow or a custom task.</p> <p>Use this option to refresh identities that were edited within a period of time, but not necessarily by the refresh task. For example, you might do a full refresh once a week but during the week people were adding or removing roles, changing extended identity attributes, doing manual correlation, or changing identities in some other way. Most of those cases have options to do a targeted refresh immediately after the change happens but this is not always the case and sometimes it is better to batch up a number of refreshes rather than have hundreds of individual refreshes occurring concurrently. If you ran the refresh task with one of the date-based options you would not necessarily pick up identities that were manually edited. If you want to include those select this option.</p>
Refresh only identities marked as needing refresh during aggregation	<p>Only refresh identities marked as needing refresh during the most recent aggregation task.</p> <p>For more information on using this option to optimize performance, see Refreshing Changed Identities Only (Delta Identity Refresh).</p>
Do not reset the needing refresh marker after refresh	<p>Do not clear the needing refresh marker set during aggregation.</p> <p>Use this option if you have multiple refresh tasks scheduled, such as entitlement and risk refresh. Then you can set the final refresh to clear the markers.</p>
Exclude identities	Exclude inactive identities from the refresh.

Option	Description
marked inactive	
Refresh identity attributes	Update Identity Cubes with any changes made to the attributes used to define identities.
Refresh Identity Entitlements for all links	<p>Refresh any account attribute mark as an entitlement in the application schema.</p> <p>This process is resource intensive as it refreshes all entitlement values for all links.</p>
Refresh manager status	Update all Identity Cubes in which the manager status has changed. For example, if a user was promoted to manager in their department, their Identity Cube would be updated by this task.
Refresh assigned and detected roles and promote additional entitlements	Update any assigned or detected role assignments that have change since the last time this task was run. Any additional entitlements found in this refresh are promoted during this task.
Provision assignments	Provision any assigned roles and entitlements detected since the last time this task was run.
Disable deprovisioning of deassigned roles	Prevents assigned roles from being deprovisioned after they have been deassigned.
Refresh role metadata for each identity	<p>Update information about the identity's relationship to their role. For example, information regarding whether or not an identity has all the roles required by the given role.</p> <div data-bbox="565 1226 1273 1339" style="border-left: 2px solid #0056b3; padding-left: 10px; background-color: #e6f2ff;"> <p>Note: This option must be selected in order to generate Role Statistics.</p> </div>
Enable manual account selection	Sent Account Selection Notification emails to users with more than one account on any application where the system cannot determine the provisioning account. By default, no provisioning is done in this case.
Synchronize Attributes	Provision identity mapping targets if their value has changed.
Refresh the identity risk scorecards	Update Identity Risk Scores with any information discovered by the scan performed by this task.
Maintain identity histories	Update the identity history by creating a snapshot of any identities with information that has changed since the last refresh.
Refresh the group scorecards	Update Group Risk Scores with any information discovered by the scan performed by this task.

Option	Description
	Partitioning is disabled if you select this option.
Clean up groups definitions that are no longer referenced	Delete unreferenced group definitions. This option is only supported if it is selected in conjunction with the Refresh the group score card option and they run in the same task.
Check active policies	Scan for active policies and apply those policies to the identities included in the task.
Keep previous violations	Maintain a history of violations that are no longer active.
A comma separated list of specific policy names. When set this overrides the default policies	Scan for and apply only those policies included in this list to the identities included in the task.
Refresh assigned scope	Refresh assigned scope based on changes discovered.
Disable auto creation of scopes	Do not automatically assign scope to identities as part of this task.
Mark dormant scopes after refresh	Mark scopes that are not assigned to any identities as dormant. Partitioning is disabled if you select this option.
Process Events	Enable event certifications. Use the snapshots created during aggregation to approximate the previous state of the identities at the beginning of the refresh. This copied identity is compared to the updated identity to determine if event certifications are launched.
Disable identity processing threshold	Identity processing thresholds let you stop lifecycle events before they are fully processed to prevent any dangerous workflows from accidentally being triggered. They can ben enabled in Rapid Setup events and in Lifecycle and Certification events. If identity processing thresholds are enabled, use this field to disable the identity processing threshold for this task. See the Rapid Setup documentation for more details.
Refresh logical application links	Scan for changes to composite applications and refresh the link information.
Promote managed attributes	When enabled, any values for entitlement or permissions encountered while running the task automatically get promoted as managed attributes

Option	Description
Number of Refresh Threads	<p>Specify the number of concurrent threads used during task processing.</p> <p>The number of threads should not exceed 10.</p> <p>This option is not supported with partitioning enabled.</p>
Always launch the workflow (even if the usual triggers do not apply)	<p>Launch a workflow for each identity even if no identity triggers or provisioning policy questions apply.</p>
Enable the generation of work items for unmanaged parts of the provisioning plan	<p>Create work items for role entitlements that are not managed by available connectors or provisioning integration modules so the appropriate action can be taken.</p>
Disable connector lookup of managers that do not correlate	<p>Disable the default <code>MANAGER_LOOKUP</code> feature and stop the automatic lookup/bootstrap of the manager account at the connector level.</p>
Enable partitioning	<p>Enable partitioning of this task across multiple hosts.</p> <p>Partitioning must be configured globally before this option can be used.</p> <p>See the System Administration documentation.</p>
Number of partitions	<p>Specify a number of partitions. If no number is specified, IdentityIQ calculates an optimal number based on available request servers.</p>
Loss Limit	<p>The loss limit sets the maximum number of identities that will be reprocessed in case of a sudden termination of a partitioned refresh. This option is used only when partitioning is enabled.</p> <p>See Loss Limits. See the System Administration documentation.</p>
Do not schedule retry requests during application maintenance windows	<p>Disables the scheduling of provisioning retry requests, when provisioning fails due to an application being within a maintenance window. Application maintenance windows can be set for each application.</p> <p>See the Application Configuration documentation.</p>

Refreshing Changed Identities Only (Delta Identity Refresh)

A "delta" identity refresh lets you update only those Identity Cubes that have changed since your last aggregation(s), rather than updating all identities. This can result in a significant reduction in refresh time, and can remove or reduce the need to partition your identities into subsets for efficient refresh processing.

In most cases, identities which have had no changes to their attributes or accounts ("link" objects) as a result of aggregations are not likely to have new policy violations or need new workflows launched to handle state changes. These identities can therefore be skipped or not processed by the Identity Refresh task. In contrast, identities that *have* undergone some kind of change, referred to as some kind of "delta," should be processed by the Identity Refresh task.

IdentityIQ lets you set up your tasks to refresh only the changed or delta identities; this is a two-step process:

1. Configure and run an aggregation task to **mark identities as changed** when attribute or account data on the identities has been modified.
2. Configure and run an Identity Refresh task to perform their functions **only on the marked identities**.

Marking Identities as Changed

During an aggregation, details of some identities are changed, while some others may not be. IdentityIQ's aggregation tasks include a setting that lets the task flag any identities updated by the aggregation as needing a refresh. This lets you single out only updated identities for a refresh when the Refresh Identity Cube task is run. The default behavior of aggregation task is to set this flag; if you don't want an aggregation task to flag identities that need a refresh, you can turn this option off.

During the aggregation task, IdentityIQ marks the identities that have changed by setting the attribute `needsRefresh` to `true` on the changed identities as they are updated. This is a default operation performed in all aggregation tasks, although it can be turned off with an option on each of the aggregation tasks if desired.

This `needsRefresh` flag can then be used by the Identity Refresh tasks to target only those identities with accounts that were modified in a recent aggregation. The refresh tasks can then reset that flag to `false` when they are done with the identities so that subsequent aggregations can set the flag anew, and subsequent refresh cycles will again only pick up changed identities.

If you want to use this delta identity refresh feature, you should carefully consider which attributes you choose to aggregate from your applications into IdentityIQ. Aggregating attributes such as last login date, for example, would likely cause IdentityIQ to reflect changes to identities more frequently than choosing to aggregate only more static data fields, and would therefore flag more identities for delta refresh.

Note that aggregation is the only process which automatically sets this `needsRefresh` flag on identities. If other processes (such as Lifecycle Manager requests) make attribute or account changes to an identity which would affect identity refresh functionality, a full refresh that does not rely on this flag would be required to process those other identities' changes. Alternatively, the Lifecycle Manager workflows also include an optional targeted identity refresh step which could refresh the single changed identity immediately, and could be configured either to clear or not clear that identity's current `needsRefresh` flag value at that time.

To set up an aggregation task that will mark identities that have changed:

1. Click **Setup > Tasks**.
2. Choose the aggregation task to edit.
3. Uncheck the **Disable marking the identity as needing a refresh** option.
4. **Save** the task.

Refreshing Only Identities Marked As Changed

In the Refresh Identity Cube task (or other refresh tasks), select the option to **Refresh only identities marked as needing refresh during aggregation**. It is important to note that this operation is disabled by default; that is, default behavior for Identity Refresh tasks is to ignore the `needsRefresh` flag that was set by the aggregation task. If you want to use the delta identity refresh feature, you have to explicitly set this option in your refresh task(s):

1. Click **Setup > Tasks**.
2. Choose the refresh task to edit.
3. Check the **Refresh only identities marked as needing refresh during aggregation** option.
4. **Save** the task.

When the refresh task runs, it resets the `needsRefresh` flag to `false` for every identity it processes. This way, IdentityIQ knows that the identity has been refreshed already and so will not refresh it again until the next aggregation. However, you can change this behavior if you want. Depending on how you run refresh tasks, you may or may not want to reset this flag.

For example, if you aggregate and refresh infrequently, it can be a good practice to have the refresh task clear the `needsRefresh` tag, to avoid needlessly repeating refreshes on an identity that has just been refreshed. However, if you segment the refresh task to, for example, split out the refreshing of entitlements, attributes, and policies, you would **not** want to clear the `needsRefresh` tag. Leaving the `needsRefresh` tag in place as you iterate through all the refresh segments lets you avoid a situation where an identity is updated only for one segment of the full refresh process, rather than all segments that might apply.

To configure the refresh task so that it does **not** clear the `needsRefresh` flag from an Identity Cube when it runs:

1. Click **Setup > Tasks**.
2. Choose the refresh task to edit.
3. Check the **Do not reset the needing refresh marker after refresh** option.
4. **Save** the task.

Best Practices for Delta Identity Refresh

Delta identity refresh offers flexibility for managing the important identity refresh functions more efficiently. Here are some recommended best practices for using this functionality.

Interspersing Delta and Full Refreshes

Delta identity refresh can be a helpful performance boost for IdentityIQ, as it streamlines time-intensive identity refresh processes. It is still strongly recommended, however, that customers regularly run a full refresh to ensure that all identities get updated and processed following other changes which may occur outside the data flows which would trigger identities to be marked as needing refresh. For example, Lifecycle Manager request workflows may or may not refresh identities after entitlement and role assignments, depending on options configured in the workflows, and they do not set the `needsRefresh` flag by default; a full refresh on a periodic basis would catch those identities and update them as needed.

Different customers run refreshes on different schedules, often depending on the volume of ongoing changes they anticipate or experience in their environment. For example, customers who run an hourly delta refresh should plan to run a daily full refresh, while customers with lower change volumes may run daily delta refreshes with weekly full refreshes.

Timed Triggers

It is common for installations to define workflow triggers (IdentityTrigger objects) which are set to run based on the calendar date when the trigger executes. For example, a "pre-offboarding" workflow could be configured to run the day before a user's termination date or a "pre-onboarding" workflow could run a day or a week before the user's start date, with those dates being specified through attributes in an HR or authoritative system feed. The Identity Refresh task that launches those workflow triggers (through its "Process Events" option) should be configured as a full refresh which processes all identities, not a delta refresh handling only identities that need refreshing; the Identity's attributes or accounts might not change in any aggregation run on the day these events need to execute, so this workflow needs to evaluate all identities, not just recently changed ones.

Delta Identity Refresh and Logical Applications

Logical applications in IdentityIQ are applications which are defined based on accounts, entitlements, and / or permissions granted through one or more other applications. For example, Application X, for which access is controlled by membership in the AppX group in LDAP, could be configured as a logical application in IdentityIQ.

Logical application "aggregation" is actually performed by processing information already stored in IdentityIQ for other applications (for example, to identify Application X accounts and entitlements, IdentityIQ looks at the LDAP application accounts with the AppX group), rather than by reading data from an external source like traditional application aggregations do. Standard applications, with a data source external to IdentityIQ, support an optimized aggregation process in which IdentityIQ skips the bulk of processing on any records it finds to be unchanged as it reads the source.

Because the data source for logical applications is existing identity data which must be dynamically examined and

assessed for logical accounts based on the logical application definition, there is no built-in option for optimizing aggregation from logical applications.

Delta identity refresh offers functionality which can be used to approximate an "optimized refresh," so it is possible to use this feature to manage a more optimized aggregation of logical applications. These are the steps to implement that option:

1. Run an optimized aggregation from the tier applications in the logical application definition. This aggregation will then mark only the changed identities from those applications with the `needsRefresh` flag. For best results, this should be run at a time when there are no identities previous marked as needing refresh, that is, a full or delta refresh process has reset the `needsRefresh` flag on all marked identities prior to this aggregation.)
2. Run a delta identity refresh task to recalculate the logical application accounts. Since it will only look at the identities whose accounts changed in the optimized aggregation, it will only recalculate the logical accounts for those identities. This refresh should have these options set:
 - **Refresh only identities marked as needing refresh during aggregation** – enable this option to have the task process only identities which were marked as changed in the optimized aggregations in step 1.
 - **Refresh logical application links** – enable this option to use this refresh for logical application account calculation.
 - Edit the task's definition in the Debug pages to add a **compositeApplications** key, set to a comma-separated list of logical application names which this refresh should process; this can be a single logical application at a time or a limited list so the task only processes the logical application(s) which are based on the tier applications just aggregated

Identity Request Maintenance

The Identity Request Maintenance task scans all Lifecycle Manager access requests to ensure that all identity change requests were provisioned.

Partitioning is available to speed the processing time for this task, and to level the load on the machines running these tasks. Partitioning is used to break operations into multiple pieces, or partitions. Each partition is then placed in a global queue, and machines, or hosts, in a cluster compete to execute the partitions in the queue. Machines are added or removed from the cluster dynamically with automatic balancing. If a machine fails or is taken down while processing a partition, the partition is placed back into the queue and reassigned to a different machine. A single result object is shared by all partitions and is continually updated so you can monitor the overall progress of the partitioned operation. When all partitions have finished executing the result is marked complete. See the **System Administration** documentation.

Option	Description
Max age (in days) for Identity Request objects	<p>The maximum number of days that an identity request object (AccessRequest) is stored in the IdentityIQ database before it is removed.</p> <p>Set this according to your policy on how long access request details are required.</p> <p>The default is zero (0), which indicates that they are stored forever.</p>
Verify provisioning for requests?	<p>Scan for provisioning requests which have been verified.</p>
Number of days to attempt to verify the request with the Identity model before failing.	<p>The number of days the task attempts to scan for verified access requests before reporting a failure.</p> <p>When a timeout occurs, any item not verified is left in the non-finished provisioning state, either Committed or Pending, and the overall request is marked Partially Complete if any item succeeded. If no item succeed the entire request is marked failed.</p> <p>Set this value based on the type of connectors and their expected provisioning times. The default setting is continuous checking forever.</p>
Enable partitioning	<p>Enable partitioning of this task across multiple hosts.</p> <p>Partitioning has to be configured before this option is valid.</p>
Number of partitions	<p>Specify a number of partitions. If no number is specified, IdentityIQ calculates an optimal number based on available request servers.</p>

Missing Managed Entitlements Scan

Missing Managed Entitlement Scan tasks scan the selected application and create any entitlement objects for items added after the application was last aggregated.

Select the applications to include in the scan. At least one application must be specified. Click the arrow to the right of the suggestion field to display a list of all applications, or enter a few letters in the field to display a list of applications that start with that letter string.

This task returns a list of entitlement names, values, and the application on which they were detected.

OIM Application Creator

Run the OIM Application Creator task to inspect Oracle Identity Manager applications and retrieve information about the applications to which they are connected. This task generates an IdentityIQ application for applications connected to the OIM application specified. The applications generated by this task are added to the list of applications in IdentityIQ.

Option	Description
OIM Application	Select an OIM application to inspect and from which applications should be generated.

Policy Scan

The Policy task type is used to run policies against identity cubes and update identity score cards with any policy violations discovered. IdentityIQ provides the Check Active Policies task as a global policy update task.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

Option	Description
Optional filter string to constrain the identities refreshed	A filtering string used to limit the number of identity cubes updated by this task. For example you can limit the refresh to one department within your enterprise, such as Finance, by entering: department == "Finance"
Optional list of group or population names to constrain the identities refreshed	Use this list to further limit the number of identities included in this policy scan.
Apply all active policies	Scan for active policies and apply those policies to the identities included in the task.
A comma separated list of specific policy names. When set, this overrides the default policies	Scan for and apply only those policies included in this list to the identities included in the task.

Propagate Role Changes

Note: IdentityIQ does not propagate role changes for entitlements on applications that do not support direct provisioning and would require the creation of multiple work items. If required, a business process can be enabled in the System Configuration settings to handle that situation.

Note: Entitlements that were detected are not removed from an identity during role propagation, unless they are also part of an assignment. Only those entitlements that were assigned, individually or as part of a role assignment, are removed during propagation.

The Propagate Role Changes task updates any identities that have assigned roles whose associated entitlements have changed. This is the only task that can propagate the removal of entitlements from an assigned role.

Option	Description
Number of minutes task should run	The number of minutes for the task to run. The task stops only after finishing current event processing.
Check active policies	Scan for active policies and apply those to the identities included in the task.
Keep previous violations	Mark old policies as inactive but do not delete them.
A comma separated list of specific policy names. When set, this overrides the default policies.	Scan for and apply only those policies included in this list to the identities included in this task.
Enable Partition	Allow the task to split into partitions and run across multiple threads and hosts, if available.
Maximum failures before event pruning	This parameter sets the number of times a role change event can fail to progress before it is pruned. A failure to progress is defined as zero successes on the event during the task. Events that are blocked by other pending events are not counted as failing to progress. If this value is left blank, the event will never be pruned until it has been fully processed.
Maximum failure threshold	This parameter limits how many identities can fail to be provisioned by a single role change event, expressed as a percentage of the total number of identities affected by the event. All partitions for a role change event are allowed to run to completion, and once finished, the transition request computes the actual failure percentage

Option	Description
	and compares it to the maximum failure threshold. If the percentage is exceeded, the propagation terminates. Note that this does not mean that a single role change event will stop as soon as it hits the maximum; it means that if an event exceeds the maximum, no more subsequent events will be processed.

Once you have completed customizing your task options, click **Save** for later use or **Save and Execute** to save the task and run it immediately.

After executing the task, the Task Result page displays the following output:

- **Number of Identity Updates** – displays the total number of Identity updates propagated. It is different than number of Identities updated, since multiple role events include some common identities and are counted multiple times, for each role event.
- **Number of Events Processed** – displays the total number of role events propagated. This is not the number of role modifications but the role change events in the queue. As single role modification results in multiple role change events in the queue.
- **Number of Events Pending** – displays the total number of pending role change events in the queue. If timeout is not defined, Role Propagation task completes only after propagating all the events. If timeout is defined, there could be pending events in the queue even after successful completion of this task.
- **Number of Events with No Impacted Identities** – displays the total number of events which are not impacting on connected identities. This event count is based on those bundles which are not directly assigned to the identities.
- Role change events are propagated sequentially and are not consolidated to cancel out redundant changes.
- If Refresh Identity Task is run before Role propagation task, and if it adds any entitlement as part of role changes, processing of role change event through Role Propagation Task would be redundant.
- In case of retry status, the transaction would be marked as failed and role propagation task would be stopped.
- While adding an entitlement, if account is missing, transaction would be marked as failed and role propagation task would be stopped. User has to run Refresh Identity task to resolve this.
- While processing an event, if the following exception is from target system, the task would remain blocked until the events are successfully processed.

```
mandatory group cannot be removed
```

This issue can be resolved by deleting the event from the database.

- When the Role Propagation task is running, if a user creates events in database, these events are not processed by the current task. These events will be processed the next time the task runs.

Refresh Logical Accounts

The Refresh Logical Accounts task type is used to refresh composite accounts for all identities that could, potentially, have a logical account on the applications selected. This refresh occurs without performing aggregation on the logical or tiered applications containing the links.

A logical account rule is run on each identity that has a logical link, or a link on the primary tier application of the logical application. If no primary tier has been defined, the rule is run on all identities that have an account on any of the tier applications.

The information scanned and updated is determined by the following criteria when the task is created or edited. You can use any combination of options to build a task.

Option	Description
Logical Applications	Select composite applications to refresh from the dropdown list.
Refresh identities whose last refresh date is before this date	Refresh any identities that have not been refreshed since the date entered. Enter and date manually or click the [...] icon to display the calendar view.
Refresh all application account attributes	Perform an aggregation of identity information on each application and update the account attributes on each identity as required. Selecting this attribute initiates a full application aggregation for each identity included in this task. This might impact the performance of IdentityIQ.
Refresh identity attributes	Update identity cubes with any changes made to the attributes used to define identities.
Refresh manager status	Update all identity cubes in which the manager status has changed. For example, if a user was promoted to manager in their department, their identity cube would be updated by this task.
Refresh the identity risk scorecards	Update Identity Risk Scores with any information discovered by the scan performed by this task.
Maintain identity histories	Update the identity history by creating a snapshot of any identities with information that has changed since the last refresh.
Refresh the group scorecards	Update Group Risk Scores with any information discovered by

Option	Description
	the scan performed by this task.
Apply all active policies	Scan for active policies and apply those policies to the identities included in the task.
A comma separated list of specific policy names. When set this overrides the default policies	Scan for and apply only those policies included in this list to the identities included in the task.
Number of Refresh Threads	Number of threads to use simultaneously while running this task.

Reset Failed NativeIdentityChangeEvents

This task is used with the Active Directory Native Move/Rename feature.

IdentityIQ uses the Active Directory GUID, a globally unique identifier, to determine when an account or group object's DN has changed. When a change is detected, the object is updated, and the change is propagated to all DN references throughout IdentityIQ.

When a changed DN is updated on aggregation, IdentityIQ creates an event to propagate the changes.

To process the event, a Native Identity Change Propagation Request is created, which propagates changes to the appropriate destinations based on the event type. After all propagations to the destinations have completed successfully, the event is marked as DONE. If the event fails for any reason, the event and request are both marked as FAILED.

Failed events can be restarted or pruned from your system using the Reset Failed NativeIdentityChangeEvents task.

Option	Description
Prune Failed Events whose old and new Native Identity differ only by case	Select this option to prune events that failed due to a mismatch in the case of the Native Identity.
Reset and process Failed Native Identity Change Events	Select this option to create a Native Identity Change Propagation request to process failed events, and reset event status from FAILED to PENDING.

For more information on the Active Directory Native Move/Rename feature, see the *IdentityIQ Application Management Guide*.

Role Index Refresh

A role index refresh task updates all role information and creates the indexes needed to perform role searches. You must run this task before performing any role searching.

Role-Entitlement Associations

This task deletes existing role-entitlement associations then analyzes each role in the system and creates associations between the role and any granted entitlements.

This task only needs to be run one time to establish role associations to entitlements and permissions; once it has been run, IdentityIQ automatically updates the relationship table any time changes are made to role profiles. This task is run by default when upgrading from an earlier version of IdentityIQ to the current version; in an upgrade scenario, you do not need to run the task independently of the upgrade process in order to establish these relationships.

Although there is no requirement to run the Role-Entitlement Associations task again after it is first run, you can choose to run it if you want to – for example, if you have onboarded many applications in a short timeframe and want to take extra care to ensure that your relationship table is up to date.

For more information about the associations between roles and the entitlements and permissions they grant, see [Understanding Relationships Between Roles and Entitlements/Permissions](#).

Option	Description
Enable Partitioning	<p>Enable partitioning of this task across multiple hosts.</p> <p>Partitioning must be configured globally before this option can be used.</p> <p>See the System Administration documentation.</p>
Number of Partitions	Specify a number of partitions. If no number is specified, IdentityIQ calculates an optimal number based on available request servers.

After you complete customizing your task options, click **Save** for later use, or **Save and Execute** to save the task and run it immediately.

Run Rule

A task used to run an arbitrary rule with a series of name / value pairs.

You must have to configure some return statement as string. From your code, you have to return some meaningful string to the task. In your task definition declare:

```
<Returns>
  <Argument name="tskSuccess" type="string">
    <Prompt>Task Result:</Prompt>
  </Argument>
</Returns>
```

And in your code:

```
String tskSuccess = "failed";
if ( Do some condition check here) {
//Do something;
String tskSuccess = "Success";
}
return tskSuccess;
```

The rule is expected to return a string value representing its status. Any string other than Success results in a failed task result.

For an overview of developing and using rules in IdentityIQ, see the *IdentityIQ System Administration* guide.

Sequential Task Launcher

A sequential task launcher initiates the specified tasks in the order defined. This enables you to run tasks sequentially without having to schedule each separately based on estimated run times.

Option	Description
Enter the list of tasks you would like to run. Tasks are run in the order that they are entered.	Select the tasks you would like to run and the order in which they should run.
Task execution timeout	Specify the timeout argument for the sequential task. This argument is applied to each task defined as part of the sequential task.
Print log statements to indicate which tasks have been completed.	Select to print log statements so the sequential tasks can be tracked.
Cease execution if one of the executing tasks encounters an error.	Select to stop the sequential task if one of the tasks in the list fails. If this option is not selected the task continues in order.

System Maintenance

SailPoint provides System Maintenance tasks with the IdentityIQ application, the Work Item Expiration Scanner, Mitigation Expiration Scanner, System Maintenance, System Maintenance Object Pruner, Role Overlap Analysis, and the Synchronize Roles task. These tasks are configured, by default, to run in the background of the application and update score card, application, and role information as needed.

In addition to the [Perform Maintenance](#) task and the System Tasks defined to perform specific functions within your enterprise as described in [Predefined Tasks](#), the following maintenance tasks may be defined via **Setup > Tasks > New Task**:

- The **System Maintenance Object Pruner** prunes objects in batches to improve performance. This task is not part of the System Maintenance task pruning operations and is run independently when necessary. This task is always run with partitioning enabled. This task is useful if you want to set up tasks specifically for pruning objects; pruning can also be accomplished using the System Maintenance/Perform Maintenance task, with partitioning enabled.
- The **Reset Orphaned WorkItem Events** task is designed to recover orphaned work item events. This is a user-driven task to determine the conditions in which a work item event is determined to be orphaned by way of workflow name, which should be restarted or discarded. The determination condition would be for work items that are of type Event, have expired locks, and are x-time beyond the date in which the Perform Maintenance task was to have run. You control x-time. A list of workflows can also be supplied used to declare for which workflows a work item would restart. Any work item that otherwise matches the search condition is purged. Details of the purged work items can be captured in the task result or Work Item Archive.

Perform Maintenance

The Perform Maintenance is a predefined system maintenance task that performs a variety of essential operational activities. It prunes identity snapshots, task results, and certifications, escalates orphaned work items, and performs other background maintenance tasks.

The predefined Perform Maintenance task is accessed from the **Tasks** tab in the **Setup > Tasks** UI. It is grouped with other **System** tasks. To create a new maintenance task, as an alternative to using IdentityIQ's preconfigured version, choose **New Task > System Maintenance**.

Object Pruning Options in the Perform Maintenance Task

The Perform Maintenance task's pruning options control which objects are pruned, and when. A good practice is to configure dedicated Perform Maintenance tasks that run on independent and separate cycles, to prune and archive objects. For example, you may choose to configure specific Perform Maintenance tasks for each type of object.

Global settings are used in conjunction with Perform Maintenance, to control what is pruned and when. Settings under the **gear menu > Global Settings > IdentityIQ Configuration > Miscellaneous** tab control the timing for object "expirations." The expiration timeframes for objects determine when these objects are eligible to be pruned by the Perform Maintenance task. For more information, see the **System Configuration** documentation.

IMPORTANT: Once objects are pruned, they are **unrecoverable** unless a backup has been made.

Perform Maintenance Task Options

Option	Description
Prune identity snapshots	<p>Identity snapshots are copies of Identity data that can be maintained for historical purposes. These snapshots are created during certification generation, and by using the option "Maintain identity histories" in the Identity Refresh task.</p> <p>Snapshots are deleted by this task according to the expiration days set in the "Days before snapshot deletion" option in the gear menu > Global Settings > IdentityIQ Configuration > Miscellaneous tab.</p>
Prune task results	<p>Deletes the results of any tasks that are complete and do not have pending sign-offs, per the expiration days set in the "Days before task result deletion" option in the gear menu > Global Settings > IdentityIQ Configuration > Miscellaneous tab.</p>
Prune requests	<p>Background requests are created internally by IdentityIQ, to handle future execution, like mitigation expirations, email requests, and sunset/sunrise.</p> <p>This option deletes background requests whose creation date has passed the <code>requestMaxAge</code> set in the System Configuration object, and uncompleted requests whose expiration has passed.</p> <p>The timeframe for pruning background requests can only be set in the System Configuration object, using the <code>requestMaxAge</code> parameter; it can not be set in the UI.</p>
Prune provisioning transactions	<p>Deletes any provisioning transactions older than the age set in the "Days before provisioning transaction event deletion" option of the Provisioning Transaction Log Settings in the gear menu > Global Settings > IdentityIQ Configuration > Miscellaneous tab.</p>
Archive and prune certifications	<p>Archives and / or deletes completed certifications based on the expiration days set in the gear menu > Global Settings > IdentityIQ Configuration > Miscellaneous tab.</p> <p>This option first archives completed certifications; when the archive expiration date is reached, the option deletes the archived certification.</p>
Automatically close certifications	<p>Finds and closes all certifications that have an automatic closing date earlier than right now and that are not yet marked signed.</p> <p>For more information on automatic closing of certifications, see the Certifications and Access Reviews documentation.</p>
Finish certifications	<p>Finishing is the final step of a certification after it has been completed /</p>

Option	Description
	<p>signed off.</p> <p>This option checks certifications for completion status and any other final validations. If the certification is ready to be finished, this option also generates any needed remediation work items.</p>
Number of finisher threads	<p>Set a number of concurrent threads to use during the task, to improve performance. This option is not supported if you are using partitioning for this task.</p> <p>The maximum allowable number of threads is five times the number of cores.</p>
Transition certifications phases	<p>This option finds any certifications that have passed their phase transition date, and advances them to the next phase (for example, from an active to a remediation phase).</p>
Scan for completed revocations	<p>Finds any certifications that contain items that have been revoked, but not yet marked complete. This option finds all entities requiring remediation, and marks whether entitlements requiring remediation have been remediated. By default, revocations for a certification are only scanned once per day</p>
Forward inactive user work items	<p>Escalate any inactive work items to the designated user or workgroup.</p> <p>Forwarding is determined first by rules configured in the Work Items Rules under the gear menu > Global Settings > IdentityIQ Configuration > Work Items tab. If no rule is specified the item is forwarded to the identity's manager. If the identity does not have a manager, the item is forwarded to the Administrator.</p>
Denormalize scopes	<p>Updates any object whose assigned scope has changed in the scope hierarchy.</p>
Prune batch requests	<p>Deletes any batch requests older than 30 days.</p>
Prune syslog events	<p>A syslog event is a capture of an error in the system, including the stack trace of the error.</p> <p>This option deletes any syslog events older than the age set in the Syslog Settings section of the gear menu > Global Settings > IdentityIQ Configuration > Miscellaneous tab</p>
Process background workflow events	<p>Processes workflow events that have moved to the background.</p>
Number of background workflow threads	<p>The number of threads that should be created to handle background workflow processes.</p>
Workflow thread	<p>The number of seconds to wait before aborting the background thread.</p>

Option	Description
timeout (seconds)	This variable can be overridden by specifying a variable within the workflow, but by default it is left blank and the thread never times out.
Prune Attachments	<p>Delete attachments older than 30 days that are not associated with an access request.</p> <p>For auditing purposes, there is an audit event called Prune Pending Attachments which can be triggered during the cleanup in the System Maintenance Task. To enable auditing for attachment pruning, enable the Prune Pending Attachments option in IdentityIQ's Audit Configuration (gear menu > Global Settings > Audit Configuration).</p>
Prune Pending Attachments	<p>Delete files attached to abandoned access requests that are older than 12 hours. The attachment is considered "pending" and eligible to be deleted when the attachment file has been uploaded but the request has not been submitted.</p> <p>This timeframe can be overridden by adding an entry to the System Configuration object called <code>pendingAttachmentPruneAge</code> with a value that represents a number of hours.</p>
Prune Native Identity Change Events	Deletes native identity change events that are older than the age specified in the System Configuration.
Prune Inoperable Intercepted Deletes	Deletes records that are unable to be processed
Prune Post Commit Notification Objects	<p>Deletes PostCommitNotificationObjects older than 30 days.</p> <p>For details about the PostCommitNotificationObject itself, see the IdentityIQ Object Model and Usage technical white paper on Compass.</p>
Enable Partitioning	<p>Enable partitioning of this task across multiple hosts.</p> <p>For more information, see the System Administration documentation.</p>

Target Aggregation

A target aggregation task scans applications or target sources, and aggregates activity targets from those applications. These targets are then used for activity monitoring and risk assessment.

Select the application or target source to include in the scan. If no applications are specified, all applications are included. Click the arrow to the right of the suggestion field to display a list of all target sources.

Select **Include empty targets** to aggregate targets with no associated users or groups. By default, empty containers will not be included.

Task Types

For Privileged Account Management (PAM), use this task to aggregate PAM container permissions; select your PAM Target Source (that is, the Unstructured Target Collector you configured when setting up your PAM application) as the target source to aggregate.

How to Complete Task Work Items

Task work items are generated by task that require sign off on the results they create, and those sign off request that are forwarded by a designated signer. Sign off request are displayed on your home page and you are notified through an email when the work item is created.

Sign-off decisions are retained with the task results for tracking purposes.

1. Select **My Work > Work Items** to view your current work items.
2. Select the **View** button to the right of a task sign-off work item to display the request.
3. Review the work item information in the Summary section.
4. Review the Comments section for any information associated with this work item.
5. Select **Click to View Task Results** in the Details sections to review the Task Results page. After reviewing the task results, select **Return to Work Item**.
6. Select an action button to open the associated comments dialog and conclude this work item. Action options are **Signoff, Reject, Forward, Save, or Cancel**.

Note: If you sign off on, or reject the sign off request, the status of the task results is updated to reflect that decision. If you forward the work item, you must specify a recipient.