

Quest® InTrust 11.3.1

Auditing Guide



Contents

Task-Based Auditing Overview	5
Understanding Jobs and Tasks	6
Creating Tasks and Jobs	7
Organizing Tasks	8
Running or Stopping Tasks and Jobs	8
Creating a Job Sequence	8
Understanding Policies	9
Creating or Editing a Policy	9
Configuring Data Sources	10
Handling Descriptions of Windows Events	10
Handling Applications and Services Event Logs	11
Custom Text Log	11
External Events	11
Agent Side	11
InTrust Server Side	13
Keeping Event Data on the InTrust Agent Side	13
Tuning Jobs	15
Using Deadlines	15
Job Types	16
Gathering Job	16
Using Agents to Gather Data	17
Consolidation Job	18
Consolidation over a Firewall	18
Import Job	18
Reporting Job	19
Avoiding Report Timeouts	21
Reporting Job Security	22
Report-Driven Data Import	23
Credentials Settings for Report-Driven Data Import	24
Notification Job	26
Notification Templates	26
Creating Variable Messages	26
Specifying Evaluation and Notification Queries	27
Audit Database Cleanup Job and Repository Cleanup Job	27
Deleting a Repository	28
Windows Scheduled Task Job	28
Application Job	28
Alert Database Cleanup Job	28

Creating Your Task-Based Gathering Workflow	29
Using the Quick Start Wizard	29
Using the Configuration Wizard	29
Manual Workflow Configuration	30
Prerequisites	30
Example 1: A Simple Workflow	31
Example 2: A More Efficient Workflow with Two Reports	31
Example 3: Consolidating and Importing to Create a Single Report	32
Auditing Recommendations	33
Data Sources	33
What to Collect?	33
How Often?	34
How to Improve Repository Viewer Experience?	34
What Else to Consider?	35
Data Stores	36
What Will You Need?	36
How Many Repositories and Databases?	36
How Long to Keep?	37
Where to Locate?	37
Tasks and Sessions	38
About us	39
Contacting Quest	39
Technical support resources	39

© 2017 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.



Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are property of their respective owners.

Legend

-  **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.
-  **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

InTrust Auditing Guide

Updated - September 2017

Version - 11.3.1

Task-Based Auditing Overview

InTrust provides two toolsets for auditing:

- The newer real-time gathering, which is configured in the InTrust Deployment Manager console
- The older scheduled task-based workflow with specialized jobs, which is configured in the InTrust Manager console

This set of topics describes the scheduled task-based method. For details on the newer toolset, see the [Collecting Events in Real Time](#) topic.

Note the following important specifics of the two gathering methods:

- The repositories you use for real-time gathering should never be used by gathering jobs, and the other way around.
- Real-time gathering always uses InTrust agents.
- Real-time gathering uses only indexed repositories.
- The workflows differ only in the gathering part. All operations that follow gathering (such as repository consolidation, data import, reporting and event analysis in Repository Viewer) are similar no matter which of the two gathering methods you use.

For more details about tasks and jobs, see [Understanding Jobs and Tasks](#).

Understanding Jobs and Tasks

InTrust offers you a smooth, automated audit data gathering workflow based on scheduled *tasks*, which are sequences of *jobs*. Jobs process audit data, perform notification, or run applications. A task not only provides a container for jobs, but also ensures automated job execution—on schedule and in the specified order. InTrust's predefined tasks are available in InTrust Manager under **Workflow | Tasks**.

To collect event data from site computers into a repository and/or database using a gathering job, do the following:

- Assign a gathering policy to a site.
A gathering policy specifies what audit trails to collect from site computers (for example, Windows Security Log), and what filters to apply to this data (for example, you can include only the events from domain controllers, and cut off events older than one month).
- Specify the repository and/or audit database where the collected data will be stored.
Repositories are long-term file system-based storages, and use a special compressed format to store audit data.
Audit databases are used to store data for reporting purposes, because reporting cannot use repository data. It is recommended you keep databases lean and ensure that they store current and useful audit data.
- Specify whether to use agents for data collection.
At the scheduled time, if gathering is configured to use agents, an agent starts collecting data on the target computer and performs the following operations locally:
 - Analyzes what portion of the audit trail should be processed
 - Extracts the necessary data
 - Screens out irrelevant events (when filters are applied)
 - Compresses the data
 - Encrypts the data, if necessary
 - Sends the data to the server

When a data collection session is over, the collection stops automatically.

Generally, audit data is collected into a repository, and then an *import job* places it into the audit database, applying import policies which prescribe what portions of data must be imported. Then a *reporting job* is launched to generate the necessary reports on collected data.

i **NOTE:** InTrust gathers audit trails from heterogeneous environments while keeping the workflow uniform. For example, gathering the Solaris Syslog is no different from gathering Windows event logs in InTrust. However, in the case of Solaris, you need to take a few extra steps to prepare for auditing. See the [Gathering Job](#) topic for details about using agents.

Jobs that work with audit data (meaning gathering, consolidation, and import jobs) involve the following:

1. InTrust sites to determine from what computers audit data is gathered by gathering jobs
2. Gathering, consolidation and import policies to determine what audit data must be processed

Sites are discussed in the [InTrust Sites](#) topic.

The recommended auditing and reporting workflow is to gather audit data to a repository for storage, import only the necessary bits to the audit database for reporting, and clean up the audit database when the data is no longer needed.

To create, delete and modify audit databases and repositories, use the **Configuration | Data Stores** node in the InTrust Manager treeview.

i | **NOTES:**

- For details about using repositories, see the [Understanding Repositories](#) document.
- You can have as many audit databases as you like in addition to the default audit database.

Creating Tasks and Jobs

A task consists of one or more jobs. Jobs in a task can be configured to run simultaneously or one after another.

To create a task

1. In InTrust Manager, right-click **Workflow | Tasks** and select **New Task** to start the New Task Wizard.
2. Follow the wizard, supplying the task settings:
 - a. Enter the task name and description
 - b. On the next step, click **Modify** and specify the task schedule. Then select the **Schedule enabled** check box.
 - c. By default, the task will run under the server account; to specify another account, click **Set Account**.
3. On the last step, you can select the **Create a job in this task** check box—the New Job Wizard will be started automatically (see the procedure below for details about creating jobs)

To create a job within a task

1. Right-click the task and select **New Job** to start the New Job Wizard. Alternatively, select the task, then click and drag the cursor on the right pane.
2. After you finish the wizard, you can specify the account to be used when running the job. In the job's properties dialog box, you have two options as to what account to use:
 - The task account
 - Another account that you specify

Note that the job account must have access permissions on the repository and/or audit database.

If an account that will be used to work with the repository and/or database is specified in the repository/database properties, make sure it has been granted sufficient permissions for connection.

For details on required permissions, refer to [System Requirements](#).

i | **NOTE:** Before a newly-created (or modified) task or job can start, you must commit the changes.

Organizing Tasks

To reduce clutter in the treeview and group your tasks logically, you can put them in containers like files in folders. Task folders are created under the **Workflow | Tasks** node and can be nested.

To create a task folder, right-click **Workflow | Tasks** or an existing task folder and select **New Folder**. You can paste cut or copied tasks from other task folders if necessary.

Running or Stopping Tasks and Jobs

To run a task or a job, right-click this task or job and select **Run**.

To stop a running task or job, double-click **Workflow | Sessions**, right-click the task or job session and select **Stop**.

Creating a Job Sequence

To instruct InTrust task execute one job after another, you need to create a job sequence by linking these jobs. You can do it graphically or using the **Dependencies** dialog box:

- In graphical mode, do one of the following:
 - Click a job in the task map (on the right pane of InTrust Manager) and drag the cursor to the job you want to run next.
 - Click a job and drag the cursor to a blank area to create a new successor job.
- Use the **Dependencies** dialog box to select predecessor and successor jobs for existing jobs. To open the dialog box, right-click a job and select **Dependencies**.

You can select a job from the list of available ones; to run it after your job, put it in the **Child jobs** list; to run it before your job, put it in **Parents**. If necessary, specify a deadline (see the [Tuning Jobs](#) topic for details).

Understanding Policies

In terms of InTrust, gathering audit data means applying a policy to InTrust sites. You bind a policy to a site by creating and scheduling a task that contains corresponding jobs.

A policy is a collection of settings that defines what audit data to process, specifying data sources and filters:

- Data sources represent the audit trails from which audit data is retrieved (for example, Windows System log)
- Data filters applied to data sources specify events to be included or filtered out
- Object filters applied to entire gathering policies specify the site objects to be processed (for example, domain controllers)

So, you can configure a policy intended for collecting System logs from all domain controllers, or for collecting Security logs from the IIS servers, or any other policy you need.

Using policies spares you the effort of specifying logs and events every time you prepare for an audit data collection. Policies enable you to create a selection for a particular purpose once and for all, and to use it with all related jobs. Policies are easily editable, so you can copy any of them and make any necessary changes to the duplicate.

As described above, the gathering mechanism is used when InTrust does the following:

- Gathers data from a live network to a repository and/or audit database (gathering job)
- Consolidates data between repositories (consolidation job)
- Imports data from a repository to an audit database (import job)

For each of these operations there is a separate type of policy:

- Gathering policy
- Consolidation policy
- Import policy

These types are very similar functionally, and they are configured uniformly. The examples below involve gathering policies; you can work with other policies in a similar way.

Creating or Editing a Policy

Policies are created separately for each network environment (for example, Microsoft Windows Network).

To create a policy

1. Double-click **InTrust Manager | Gathering | Gathering Policies** in the left-pane tree view.
2. Right-click the environment for which you want to create the policy (for example, **Microsoft Windows Network**).
3. Select **New Policy** to start the New Policy Wizard.
4. Specify the name and optional description for the policy.
5. Add the data sources this policy will prescribe to process.
6. For each data source, specify gathering settings: what events to ignore, whether to clear log after gathering, and others. For details on data source configuration, see the [Configuring Data Sources](#) topic.
7. Configure a filter for the objects to which the policy is applied. For detailed instructions on using object filters, see [Using Filters](#).

To edit a policy

1. Right-click that policy and select **Properties**.
2. To edit policy name and description, open the **General** tab.
3. To configure a filter for the objects (computers) to which the policy is applied, use the **Filter** tab.
4. To add a new data source to the policy, right-click the policy and select **Add Data Source** from the shortcut menu to start the Add Data Source Wizard.

i | **NOTE:** After you create or modify a policy, commit the changes.

Configuring Data Sources

Data sources enable you to precisely select events that you need to process. Some data sources are used by gathering and monitoring processes, some provide data only for gathering, and others are monitoring-only. To see all available data sources, select **Configuration | Data Sources**. Right-click a data source to view its properties, including:

- Type of the data source, for example, Microsoft Windows Events, or Database Events, etc.
- Options specific to that particular data source

You can create a new data source by selecting the corresponding command from the **Data Sources** node's shortcut menu. Follow the wizard to specify the data source properties.

Handling Descriptions of Windows Events

There is an important set of options specific to the data sources of the Microsoft Windows Events type: you can select what libraries to use when retrieving standard descriptions for Windows events. The descriptions can be taken from libraries that exist locally on processed computers or from remote computers.

! | **CAUTION:** This setting works only if you gather events without agents. If you gather with agents, the local libraries of the processed computer are always used.

To select which libraries to use

1. In InTrust Manager, select **Configuration | Data Sources**, right-click the Microsoft Windows event log you need, and open its properties from the shortcut menu.
2. On the Microsoft Windows Events tab, specify order for using the libraries:
 - Select **Only local** to retrieve descriptions from libraries that exist on the InTrust server.
 - Select **Local, then remote** to get descriptions from the InTrust server as long as they are available.
 - Select **Remote, then local** to retrieve the necessary descriptions from libraries on remote computers as long as those descriptions are available.

Handling Applications and Services Event Logs

Windows Vista and later Windows versions support event logs with a hierarchical structure. In Event Viewer, these logs are available in the Application and Services Logs container. The Microsoft Windows Events data source type in InTrust works with event logs located at any level of the Event Viewer hierarchy.

! CAUTION: If you want to gather such logs using a Windows Server 2003 or 2003 R2-based InTrust server, you need to gather with InTrust agents. Agentless gathering will not work in this case, because Windows Server 2003 and 2003 R2 predate this type of log.
If the InTrust server that does the gathering is running Windows Server 2008 or later, then both agent-based and agentless gathering of such logs will work.

To enable InTrust to work with such a log, create a data source of the Microsoft Windows Events type. On the Windows Events Settings page of the New Data Source Wizard, specify the exact full name as the log name.

To look up and copy the full name of the log, run Event Viewer on a computer where the log is available, locate the log you need, and open its properties. Look in the **Full Name** text box.

Custom Text Log

For details about working with custom text logs, see [Auditing Custom Logs with InTrust](#).

External Events

The External Events data source type is not represented by any predefined data sources. It is different from other data source types in that it generates event records with fields that you define and hands them over to the InTrust agent to process.

Data sources of this type are represented by a command-line utility on the agent side and an InTrust data source object on the InTrust server side.

For example, you can use the utility if your application does not have its own log. The External Events data source simulates event records that can be processed by InTrust agents.

Agent Side

The External Events data source type relies on the cross-platform **ExtEvtProviderCmd** command-line utility that forces special events on the InTrust agent running on the same computer. The agent stores the events in its

backup cache. From there, the events can be captured by the gathering or real-time monitoring engine.

On each platform, the utility is implemented as two files: an executable program and a shared library. On Windows, these files are **ExtEvtProviderCmd.exe** and **ExtEvtProviderApi.dll**. On Unix-like systems, the files are **ExtEvtProviderCmd** and **ExtEvtProviderApi.so**.

Binary files for all supported platforms are located in **<InTrust_installation_folder>\Server\ADC\SupportTools**.

You can use the utility in either of two ways: copy the binaries to the target computers manually or deploy them as distributable files.

To deploy the utility as distributable files

1. In the InTrust Manager treeview, right-click **Configuration | Advanced | Distributable Files** and select **New File**.
2. Specify the appropriate executable file and complete the New Distributable File wizard.
3. Likewise, use the New Distributable File wizard to add the utility's shared library.
4. In the properties of your data source of the External Events type, on the **Distributable Modules** tab in the **Files** list, specify the files you have added.
5. Include the data source in a gathering policy.
6. Make sure that the data source is used in a gathering job and that the task with that gathering job has a schedule. This is a requirement of the External Events data source type.

By default, the utility components are automatically installed to the following locations:

- On Windows: **<agent_installation_folder>\Data\DDA**
The default path is **%WINDIR%\ADC\Agent\Data\DDA**
- On Unix-like systems: **/var/InTrust/{2A5211B3-98D5-4850-9458-29B411FBD1B6}/DDA**
You must copy the two files from there to the agent installation directory.

This utility is designed to be automated with scripts. Such scripts should launch the utility in situations when intervention is necessary. The scripts must construct and execute a command with the following syntax:

- **Windows:**
`ExtEvtProviderCmd -e`
-OR-
`ExtEvtProviderCmd <field_name> <field_value> [<field_name_2> <field_value_2>...<field_name_N> <field_value_N>]`
- **Unix (in the agent installation directory):**
`./adcrun <agent_installation_folder>/ExtEvtProviderCmd -e`
-OR-
`./adcrun <agent_installation_folder>/ExtEvtProviderCmd <field_name> <field_value> [<field_name_2> <field_value_2>...<field_name_N> <field_value_N>]`

The **-e** parameter means that an empty event record is sent.

The **<field_name>** parameter specifies the name of a field in the EventsStrings table of the InTrust audit database. The **<field_value>** parameter specifies the value that is written to this field.

The utility can fill in only a subset of the fields in audit database tables. The following table shows which fields of which tables can be used by a data source of this type:

Target Table Name in audit database	Value Name used as command parameter	Target Field Name in table
Events	Computer	Computer
	UserName	UserName
	UserDomain	UserDomain
	EventType	EventType
	Source	Source
	EventID	EventID
	StringCategory	Category
EventsDescriptions	Description	Description
EventsStrings	String<N> (You can use fields that have the format String1 String2 ... String<N>)	StringValue (the StringIndex field is also filled in; the StringIndex value equals the insertion string index (<N>))
	Any name not listed elsewhere in the table	StringValue (the StringIndex field is also filled in; the StringIndex value is greater than the maximum predefined insertion string index)

InTrust Server Side

To make InTrust aware of external events, create a data source of the External Events type in InTrust Manager.

To create an External Events data source

1. Right-click the **Configuration | Data Sources** node and select **New Data Source**.
2. In the New Data Source Wizard, select the **External Events** data source type.
3. Complete the remaining steps.

Now, if you include this data source in a policy and use the policy in a gathering job, InTrust will listen for external events from the site that the gathering job spans.

i **IMPORTANT:** The processing of a single event generated by a data source of this type can take several seconds. Therefore, you should not generate frequent external events.

Keeping Event Data on the InTrust Agent Side

To ensure the integrity of event data from the specified data source, you can create agent-side log backup. This will help you to protect data from losses that may occur due to accidental or malicious log cleanup on the target

machine. Log backup can be created for the most frequently used data sources (for example, Windows Event logs).

Agent-side log backup uses a compression method similar to that used in InTrust repositories. On average, the contents of the event cache are compressed to 1/15th their original size.

Agent-side cache is always used to process data from monitoring-only data sources. For the data sources used in both gathering and monitoring processes, you can select whether to create agent-side log backup.

Agent-side log backup is unavailable for gathering-only data sources, for example, MS ISA Server logs and MS Proxy Server logs.

i **TIP:** By default, the agent-side log backup feature is disabled but it is recommended to enable this option. This mismatch between the default and recommended setting value arises due to specificity of the InTrust task execution.

To configure agent-side log backup

1. Schedule the InTrust task that will process the selected data source.
2. Open the properties of the data source under the gathering policy.
3. Select the **Enable log backup and use it to gather events** option:
 - The **Clear the backup after gathering** check box is automatically selected together with the **Enable log backup and use it to gather events** option.
 - If a data source is used by more than one task and the agent-log backup feature is enabled for the data source, deselect the **Clear the backup after gathering** check box to avoid data loss.

! **CAUTION:** After you enable agent-side log backup, the log will be cleared the next time it is gathered. Subsequent gathering sessions do not clear the log.

4. Click **OK** to save the settings and close the dialog; commit the changes.

i **IMPORTANT:** Agent-side log backup will be created only if both of the following are true:

1. At least one gathering policy processing this log uses it with this option selected.
2. The schedule is enabled for at least one task involving this policy.

To set the log backup retention period

1. Select **Configuration | Data Sources**, and select the necessary data source.
2. From its shortcut menu, select **Properties**.
3. On the **General** tab, specify agent-side log backup retention period.

Tuning Jobs

Each type of InTrust job has a number of settings, both general and job type-specific. They general settings are available on the General tab of the job's properties, and specify the following:

1. Job name and optional description.
2. Account that is used to run a job. Jobs use accounts as follows:
 - To access objects in a site, a job can use the task account (by default) or a custom account that you specify. However, if a site account is specified in the site's properties, the job will use the site account.
 - To access data stores—that is, repositories and databases—the job uses the task account or a custom specified account. However, if an account is specified for a data store, the job uses it to access that data store.
 - Make sure that accounts used by jobs have the necessary permissions on the site objects and data stores for the job to be performed correctly.
3. Whether the job is enabled or disabled. If disabled, the job appears in the task workflow, but it is not executed.

So, to make your job available for execution, you must supply the job name and enable the job.

Using Deadlines

On the **Deadline** tab of the job's properties, you can provide an optional deadline value for a job. The deadline is the period of time that will pass before the job is started.

The countdown starts in the following situations:

- The moment the task is launched (if the Task option is selected).
- At the beginning of another job (if the Job option is selected and a preceding job is specified).

When the specified time runs out, this setting launches the next job or jobs without terminating the job in progress.

A deadline can be specified for any job in the task, unless it is the job that starts the task or one of such task-starting jobs running simultaneously.

i | **NOTE:** Turning on the deadline and specifying zero values for days, hours and minutes is the same as turning the deadline off.

Job Types

See the following topics for details about the specific job types:

- [Gathering job](#)
- [Consolidation job](#)
- [Import job](#)
- [Reporting job](#)
- [Notification job](#)
- [Repository or audit database cleanup job](#)
- [Windows scheduled task job](#)
- [Application job](#)
- [Alert database cleanup job](#)

Gathering Job

A gathering job collects audit data to a repository and/or audit database.

To configure a gathering job

1. In the job's properties dialog box, open the **Gathering** tab.
2. Assign a gathering policy to a site.
3. Select the repository and/or database where the collected data will be stored.

! **CAUTION:** Do not configure gathering jobs to collect events to repositories that are used for real-time gathering in InTrust Deployment Manager.

4. Specify whether to use agents for data collection.

Gathering to a database always starts with events that follow the last gathered event. Thus, if you have gathered data for a certain period of time using a certain filter, then you cannot gather data for the same period of time using a different filter.

When gathering to a repository, you always store events as specified by the filters you use. Duplicate events are also stored.

Events from the data sources of Microsoft Windows Events type (such as Windows Application log) have standard descriptions. If you are gathering these events to a repository, event descriptions are collected automatically. If you are gathering events to an audit database and you need to store these descriptions, take the following steps:

1. Locate the corresponding data source by selecting **Configuration | Data Sources**.
2. Open its properties dialog box. Click the **Microsoft Windows Events** tab and select **Store event descriptions to database**.

Using Agents to Gather Data

Data can be gathered with or without agents. To decide whether or not to use agents, consider the following:

- If collected with agents, data is communicated to the server in compressed form, which greatly decreases network traffic. Moreover, agent-side encryption ensures that data cannot be intercepted and decoded when being transferred over the network.
- There are few reasons, if any, for not using agents. For example, on some critical computers you may want no extra services, even if they start and stop automatically. You can collect data from such computers without agents.
- Finally, in some cases you cannot avoid using agents: for example, when you need to collect audit data from behind a firewall or when you want to collect Unix events.

To automatically install agents to all site computers in bulk, from site's shortcut menu, select **Install Agents**. Note that this is possible only in the Windows environment, on computers that are on the same side of a firewall.

For detailed description of manual agent installation, see [Installing Agents Manually](#).

You can automate the installation of agents using Group Policy settings. InTrust is shipped with a Windows Installer file containing the agent package.

To automatically install agents on specific computers using Group Policy

1. Copy the agent package from the **Agent** folder in your InTrust distribution to a share available to all those computers.
2. In the Active Directory Users and Computers MMC snap-in, create an OU that includes all of the required computers and add a Group Policy object for this OU.
3. Using the Group Policy Object Editor MMC snap-in, in **Computer Settings**, assign the agent package to the Group Policy object you added earlier.
4. To make InTrust process these computers with agents, make sure the computers are included in InTrust sites.

To prohibit automatic agent installation on site computers

1. Right-click the site and select **Properties**.
2. On the **General** tab, select the **Prohibit automatic agent deployment on site computers** option.

To gather audit data with agents

1. Open the gathering job properties.
2. Select **Use agents to execute this job on target computers**.
3. When the job starts, agents are installed automatically to site computers (if not deployed yet) and perform audit data gathering locally.

You can enable agent-server authentication using SRP and agent-side data encryption using 3DES.

To enable agent-server authentication and agent-side data encryption

1. Select **Configuration | InTrust Servers**, and right-click the InTrust server the agents belong to.
2. On the **Agent** tab of the properties dialog box, select the corresponding options.

Consolidation Job

A consolidation job copies audit data from one repository to another.

When you create or modify a consolidation job, you need to select the following:

- The server where to run the job
- The source and target repositories
- The policy according to which data will be picked for consolidation
- The repository objects from which the audit data is taken

Consolidation over a Firewall

You can consolidate audit data from a repository that is located on an InTrust server behind a firewall. To do it, first find out the local repository path on the InTrust server behind the firewall and the password of the InTrust organization behind the firewall. Then take the following steps:

1. Create a new repository. For that, right-click **Configuration | Data Stores | Repositories** and select **New Repository**.
Consider giving the new repository a name that indicates it is located behind the firewall.
2. On the Repository Location step of the New Repository Wizard, supply the local repository path so that it is identical to the local path of the repository behind the firewall. Complete the wizard.

CAUTION: The path you specify is not verified. The repository object you created on step 1 is just a representation of the source repository. The actual repository will be found as long as the path is correct.

3. Right-click the necessary task and select **New Job**; start creating a consolidation job.
4. On the Select Repositories step, do the following:
 - Select the source repository
 - Select **Use this server to manage source repository** and specify the InTrust server that hosts the repository.
 - Specify the port over which your firewall allows communication. By default, port 900 is used.
 - Specify and confirm the password of the InTrust organization that the InTrust server behind the firewall belongs to.
 - Select the destination repository.
5. Complete the wizard.

Now, you can configure job dependencies in the task and use the task as necessary.

Import Job

An import job copies audit data from an InTrust repository to an audit database; then InTrust reporting will be able to use this data for reports generation.

When you create or modify an import job, you need to select the following:

- The server where to run the job
- The source repository
- The destination audit database
- The policy according to which data will be picked for import
- The objects from which the data is taken

You can select whether events imported during previous gathering sessions will be imported again during a session that is going to take place. For that:

1. Select the import policy the job uses, and open the properties of the necessary data source
2. Click the **General** tab:
 - To import only those events that follow the last imported event, select the **Import incrementally, starting with last imported event** option.
 - To import events gathered at any time, clear this option.

When you are importing events from the data source of Microsoft Windows Events type, you may need to retain their standard descriptions. For that, select this data source from **Configuration | Data Sources**, open its properties, click the **Microsoft Windows Events** tab, and select **Store event descriptions to database**.

Reporting Job

A reporting job adds reporting capabilities to the InTrust workflow by using the Reporting Services feature of Microsoft SQL Server.

Reporting jobs are normally run after import or gathering jobs and prepare reports based on the newly-gathered data.

To generate a report, InTrust connects to the Reporting Web service on the SQL server. The actual report generation process takes place on the reporting server.

To configure the reporting job, specify the following:

- The URL of the reporting server's Web service
- The database to be used as the data source for the reports; the database you specify must exist and have the structure of an InTrust database
- Optionally, the repository to import the necessary data from (see the Report-Driven Data Import topic)
- Optionally, the credentials for creating the reports
- The reports and filters you need
- Where to deliver the ready reports—email address, network share or a Reporting Server snapshot that you can view using Knowledge Portal.
- Optionally, the repository from which to import data that is missing from the database.
- Optionally, settings for notification about job completion by email
- The InTrust server where the job runs

To modify the default URLs related to reporting jobs

1. Open the properties of the root node in the InTrust Manager treeview.
2. On the Parameters tab, select the Show all option.
3. Edit the following parameters:
 - Reporting_default_SRS—URL of the reporting server's Web service
 - Reporting_default_report_share—location of the folder where completed InTrust reports are stored
 - Reporting_default_QRS—URL of the Reporting Services user interface application's virtual directory

Web Service URL

The reporting server's Web service URL is first specified during InTrust setup. Although setup verifies it, you can select to proceed without a valid URL. If InTrust was installed like that, no default value is suggested in InTrust Manager, and you must supply the URL explicitly, for example, when creating a reporting job with New Job Wizard:

By default, the Web service URL is formed as follows:

```
http://<SQL_server_name>/reportserver
```

If you want to connect to a SQL server instance, then the URL may be formed differently. In a default configuration, it is constructed as follows:

```
http://<SQL_server_name>/reportserver$<SQL_server_instance_name>
```

i NOTES:

- If the HTTPS protocol is used in your SSRS deployment, the URLs you specify should begin with "https://".
- It is recommended that you specify the reporting server's DNS alias instead of its actual name. This will help avoid situations where changing the reporting server clears the list of reports selected for the reporting job. If you use an alias, the server switch involves changing the server that the alias points to, leaving the reporting job intact.

Contact your SQL Server administrator for the correct Web service URL.

Data Source Selection

When you create a reporting job, on the Reporting Server and InTrust Database step of the wizard, you can select the source from which data for reports should be taken:

The following options are available:

- **Use SRS data source associated with each report**—each report will be filled in with data from the data source associated with it. This data source can be found in report properties. Remember that each data source should point to a certain InTrust database—this is configured automatically during setup, or manually in Knowledge Portal or Reporting Services Report Manager.

! **CAUTION:** For this option to take effect properly, you should check the corresponding report properties (Data Sources property) and verify that the data source is properly associated with the desired InTrust database.

- **Select InTrust database for reporting**—use this option to specify the InTrust database you want to take data from. Click Credentials to specify the authentication method and database access credentials (for

details, see the Reporting Job Security section below).

Source Repository

If you need to report on data that is currently stored in an InTrust repository but not in the database, you should instruct InTrust to import missing data from the repository. For example, when creating a new job, on the Import Missing Data step, select **Import objects from the following repository** option, and select the source repository.

Filtering

You can instruct InTrust to cut off unnecessary events during import by configuring filters. For that, on the Reports step of the wizard, select a report from the list and click **Filters**. During data import, the following two filters can be applied: DateRange and Select Computers (if applicable). Select the filter and edit the filter value.

i | **NOTE:** Other filters configured for the report will be applied during report generation.

Report Storage Location

The default location for compiled reports is initially specified during InTrust installation. You can specify new defaults or use individual settings for each job.

Avoiding Report Timeouts

Reporting Services configuration includes the following timeout settings for reports that take too long to generate:

Option	Default Value	Configured Where
Report timeout	1800 seconds (30 minutes)	In the administration page for a Reporting Services site, on the General tab. You can set a custom value or disable the timeout altogether.
HTTP timeout	9000 seconds (2.5 hours)	In the Web.config file on the report server. This option has no UI representation. For details about changing it, see the procedure below.

If report generation times out for the reports you configure in your reporting job, consider changing the timeout settings.

To change the HTTP timeout

1. On the report server, locate the **Web.config** file. It should be in the `<installation_folder>\Reporting Services\ReportServer` folder.
2. In the file, find the `<httpRuntime>` tag, and change the value of the `executionTimeout` parameter (the value is in seconds). If `<httpRuntime>` doesn't exist, create it within the section enclosed in the `<system.web>` tag pair. For example:

```
<system.web>
...
  <httpRuntime executionTimeout = "18000" />
...
</system.web>
```

Reporting Job Security

Like any other job, a reporting job runs under the account it inherits from the task or the account that is set specifically for the job. However, to function properly, reporting requires more security settings than that.

Report Selection

To successfully create a reporting job, use an account that can read report definitions on the reporting server. Otherwise, you will not get the list of reports to select from.

The account you use to run InTrust Manager must have a role that enables read access to report definitions on the reporting server. The “Browser” role, which is a standard role in Reporting Services, has sufficient privileges.

Database Connection

When you create a reporting job using the New Job Wizard, you specify the location of Reporting Services and the database to be used as the Reporting Services data source.

The **Credentials** button lets you set the credentials that Reporting Services will use to connect to the database. You have the following choices as regards the credentials:

Option	Meaning and use
Windows authentication (using job account)	<p>Specifies that the Reporting Services will connect to the specified database using the credentials of the account that the job is running under.</p> <p>This authentication method is always used if you select to Import objects from the repository (that is, use the report-driven data import feature).</p> <p>This option is the best choice if Reporting Services and SQL Server with the specified database are deployed on the same computer.</p> <p>If they are deployed on different computers but you still want to use this option, enable delegation for the computer that runs Reporting Services. For that, take the following steps:</p> <ol style="list-style-type: none">1. Open the Active Directory Users and Computers MMC snap-in.2. Open the properties of the account that the reporting job runs under.3. Make sure that the Account is sensitive and cannot be delegated option on the Account tab is disabled.4. Open the properties of the computer that runs Reporting Services.5. On the General tab, select Trust computer for delegation.
Windows authentication	<p>Lets you explicitly specify the credentials. Use this option if Reporting Services and the database reside on different different computers. For secure transfer of these credentials, make sure Reporting Services communicate through the HTTPS protocol.</p> <p>An alternative to this option is to use the first option combined with delegation, as described above.</p>
SQL Server authentication	<p>Specifies that SQL Server-specific credentials are used. For secure transfer of these credentials, make sure Reporting Services communicate through the HTTPS protocol.</p>

Report-Driven Data Import

InTrust reporting uses the audit trails stored in the audit database. Typically, this database keeps information for the last 2–4 weeks (recommended retention period). However, an InTrust administrator may want to create a certain report, for example, on suspicious logons over 3 months. Data for this period is usually kept in the repository and has to be imported into the audit database for analysis and reporting. However, to report on the events you need, you do not necessarily have to create a chain of import and reporting jobs but configure the reporting job to import the necessary data from the specified repository right before report generation. To use this feature, you can do the following:

- While creating a new reporting job, on the Import Missing Data step of the New Job Wizard, select the **Import objects from the following repository** option.
- If you need to modify an existing reporting job so that it imports the necessary data from the repository, select the job, and on its **Properties** page, click the **Reporting** tab and select the same option.

So, whenever you need to report on events logged 3 months or a year ago, configure your reporting job like this, and all data required to generate the reports will be imported automatically.

i **NOTE:** When you specify a value for time period when configuring filters for a job that uses report-driven data import, time will be always treated as Local time (even if Use GMT time option was selected in the reporting job properties).

If you need to run such a reporting job periodically, you can schedule the task that contains this job. If you need to run it once, disable the job once the task session is complete. Importing and reporting operation details are written to the corresponding tasks' session logs and can be examined under the Workflow | Sessions node in InTrust Manager.

Access Rights

The following accounts are used during the reporting job that has data import enabled:

- Reporting job account—the one under which the reporting job is launched. Reporting job account is specified in the job properties on the **General** tab.
- Import job account—the one under which data is imported from the source repository to the audit database.
- Database connection account—the one under which the audit database is accessed to import data and report on it.

Access credentials and the authentication method for database access during import and reporting are specified on the **Reporting** tab of the job properties, where you should click **Credentials** to open the **Credentials Settings** dialog box.

Requirements for each account are listed in the table below. Some of these accounts may coincide depending on the authentication method you select, so refer to the next section to assign sufficient access rights to proper accounts.

Account	Requirements	Notes
Reporting job account	<ol style="list-style-type: none">1. Log on as a batch job on the InTrust Server2. Read permission on %WinDir%3. Full control permission on the InTrust Server installation folder	<ol style="list-style-type: none">1. This account must belong to the domain where SRS hosting Knowledge Portal is installed, Otherwise, membership in the Authenticated Users group (for the SRS server's domain) is required.

Account	Requirements	Notes
	<ol style="list-style-type: none"> Content Manager SRS role for the InTrust\SharedDatasources folder and for the folder where the report is located (under the InTrust folder) in SQL Reporting Services Browser SRS role for the Home folder in SQL Reporting Services Read permission on configuration objects used by the job 	<ol style="list-style-type: none"> To minimize access rights, the following item-level rights in SQL Server Reporting Services can be granted to the reporting job account (on the General tab): <ul style="list-style-type: none"> View Data Sources permission for the InTrust\SharedDatasources folder, View Folders permissions on the Home folder under InTrust Manage Reports, View Folders and View Reports permissions on the necessary subfolders (where the reports are stored) of the InTrust folder To access sub-reports, View Data Sources permission on the folder where sub-reports are stored. View Resources permission on the InTrust\SharedResources folder.
Import job account	<ol style="list-style-type: none"> Read permission on the source repository InTrust Gathering database role for the audit database (this role is created by setup) Read permission on configuration objects used by the job 	If a specific account for repository access is specified in repository properties, then the import job account can be assigned local administrative rights on the computer where the repository is located (instead of Read permission).
Database connection account	<ol style="list-style-type: none"> InTrust Gathering database role for the audit database Reporting Console User database role for the audit database 	

Credentials Settings for Report-Driven Data Import

After you click **Credentials** on the **Reporting** tab of the job properties, three authentication options are available to you:

- Windows authentication (using job account)
- Windows authentication
- SQL Server authentication

If you are using report-driven data import in your reporting job, the available authentication methods will depend on the database you select to get data from:

- If you use the **Select InTrust database for reporting** option and specify the database explicitly, then you can click **Credentials** and select any authentication method you need.
- If you select **Use SRS data source associated with each report** option, then Windows authentication (using job account) will always be used. In this case, make sure the job account has sufficient access rights to connect to the databases configured as data sources for the reports that will be processed.

Integrated Windows Authentication

If you select Windows authentication (using job account), then the job will use a single account for all operations. That means the database and repository will be accessed using the account that the reporting job runs under.

1. Open the **General** tab in the job properties.
2. Make sure the specified account meets all the requirements listed in the table.

! **CAUTION:** In case SQL Server and Knowledge Portal are installed separately from InTrust Server, take the steps described below to make Integrated Windows Authentication work properly.

To make Integrated Windows authentication work properly

1. In the Active Directory Users and Computers MMC snap-in, select the user account under which the reporting job will connect to the data source.
2. Select **Properties** and click the **Account** tab.
3. Make sure the **Account is sensitive and cannot be delegated** option is cleared.
4. Select **Account is trusted for delegation**.
5. Select the computer where Reporting Services and Knowledge Portal are installed.
6. Select **Properties** and click the **General** tab.
7. Select **Trust computer for delegation**.

Basic Windows Authentication

If you select Windows authentication, then you should specify credentials explicitly. They will be used to access the repository and the database.

1. Make sure the account you specified here meets the requirements for the import job account (middle row).
2. Open the **General** tab and make sure the reporting job account meets the requirements in the top row of the table.

SQL Server Authentication

If you select SQL Server authentication, credentials must be also specified explicitly in the Credentials Settings dialog box. This account will be used to connect to the audit database.

1. Make sure the account you specified here meets the requirements for the database connection account (bottom row in the table).

2. Open the **General** tab. The account specified there will be used as the reporting job account and import job account (that means, to launch the reporting job and to access the source repository to pick up required events). Make sure it meets corresponding requirements (top and middle rows) except the **InTrust Gathering database** role—the audit database will be accessed under the SQL Server account specified in the Credentials Settings dialog box.

Notification Job

A notification job sends net send or email messages to selected recipients, notifying them of the results of the task.

Before configuring a notification job that uses email notification, make sure the selected InTrust server is associated with an SMTP server. Open the job processing server's properties dialog box, click the **Notification Parameters** tab, and specify the SMTP server.

To configure a notification job, select the following:

- The server where to run the job
- The type of notification to be used
- The template for the message
- The database from which to get data to be included in the message (select the configuration database)
- The recipient or recipients of the message

Notification Templates

Messages are based on notification templates. Use notification templates to make InTrust notification messages informative by including data gathered from the network. Such messages are a faster means of notification than reports.

To create a notification template

1. In InTrust Manager, select Configuration | Advanced.
2. Right-click Notification Templates and select New Notification Template to start the New Notification Template Wizard.

Creating Variable Messages

To insert data in the message subject or body, you should use variable names delimited by two “%” signs. These variable names are substituted with values retrieved from a database. The rest of the message text that you specify is left unchanged.

The text between the delimiting “%” signs must match the name of a column in the record set returned by the SQL server when a database is queried. For example:

The event from %Source% occurred at %Time%.

would be resolved like this in the message:

The event from IISLog occurred at 13:51:00.

i | **NOTE:** To be able to send net send messages, make sure that the Messenger service is running on the InTrust server and the target operator's computer. By default, this service is disabled.

Specifying Evaluation and Notification Queries

During template configuration, you can provide the following two SQL queries:

- **Evaluation query.** This query performs a check against a database to determine whether to send the message. Along with an evaluation query, you must specify a conditional expression. This expression is compared to the number of rows that the query has retrieved from the database. If the expression is true, notification is performed.
- **Notification query.** Specify this query to include data from a database in a notification message.

Type these queries after specifying the subject and body of the template.

The two queries described are executed separately, and do not analyze the results of one another. However, notification depends greatly on what queries are specified. The following four situations are possible:

What is specified	What happens
Both queries	Notification takes place if the condition provided with the evaluation query is true. Data for the notification message is retrieved from the record set returned by the notification query.
Neither query	Notification takes place unconditionally. The notification message cannot contain any data from any database. Such a message is a fixed body of text.
Only the evaluation query	Notification takes place if the condition provided with the evaluation query is true. The notification message is a fixed body of text and cannot contain any data from any database.
Only the notification query	Notification takes place unconditionally. Data for the notification message is retrieved from the record set returned by the notification query.

Audit Database Cleanup Job and Repository Cleanup Job

There are two kinds of cleanup job that involve clearing audit data:

- An audit database cleanup job clears obsolete audit data from an audit database.
- A repository cleanup job clears obsolete audit data from an InTrust repository.

When configuring any of them, you need to select:

- The server where to run the job
- The database or repository

If necessary, you can also provide a date and time range filter for obsolete audit data.

You can schedule a cleanup job in a separate task rather than perform it each time you gather audit data. For example, a job that clears data older than a month should be scheduled to run monthly.

Note that though the gathered data is cleared, information about the gathering session is still kept. The next time a gathering job is started, InTrust collects data that has been written to audit trails since the last gathering session.

Deleting a Repository

The repository may contain too long directory or file names. Make sure that your operating system supports long file names. Otherwise, use the special utilities to work with these names or delete a repository from disk. To delete a repository, use the **ITRepositoryRemover.exe** command-line utility, as described in the [Removing Repositories](#) topic.

Windows Scheduled Task Job

A Windows scheduled task job binds Windows scheduled tasks and InTrust gathering workflow. This job requires that you specify a pre-configured scheduled task and provide the path to it. InTrust overrides the existing schedule settings for the task.

Configuration options for this job include synchronous operation. If this option is selected, InTrust will only consider the job completed once the scheduled task finishes. Otherwise, InTrust launches the scheduled task and regards the job as completed.

Application Job

An application job can launch an application, execute a command or a script, and so on. While creating an application job, the wizard requests the path to an executable file (*.exe, *.com, or any other type of file the operating system can execute). You can also set parameters for the application. The progress of the job depends on the file it opens.

Alert Database Cleanup Job

An alert database cleanup job clears obsolete alerts from an alert database. This job relates to the InTrust real-time monitoring process. Alert databases are used by the InTrust real-time monitoring service to store real-time alerts.

When configuring this job, you need to specify the following:

- The server where to run the job
- The alert database to be cleared
- A filter for the alerts to be cleared, based on such parameters as alert severity, age and status

Creating Your Task-Based Gathering Workflow

InTrust provides three ways to create a contiguous workflow for gathering audit data and reporting on it, as follows:

- Quick Start Wizard
- Configuration Wizard
- Manual configuration of InTrust objects

These three methods serve the same purpose, that is, to define workflow settings for InTrust to act on. However, these methods are progressively more complex and more flexible.

Using the Quick Start Wizard

The interactive Quick Start Wizard is the easiest gathering workflow configuration tool. It is mainly useful for quickly getting acquainted with InTrust's auditing and reporting functionality.

For the sake of simplicity, the Quick Start Wizard provides only the following subset of InTrust's features:

- Gathers audit data from the Microsoft Windows Security, System and Application event log
- Gathers with agents
- Due to rudimentary security configuration, requires that the audit database and the SQL Server Reporting Services are located on the same SQL server
- Uses the default data stores
- Shows only 16 most common reports from the Windows Report Pack
- Shows events only for the last 24 hours in reports.

i | **NOTE:** Links to sub-reports do not work in reports produced by the Quick Start Wizard.

The Quick Start Wizard takes you through five steps that introduce basic InTrust concepts and show what is required for successful auditing and reporting.

Using the Configuration Wizard

The Configuration Wizard can be successfully used to perform typical everyday auditing, reporting and real-time monitoring activity. This wizard suits most such needs. You can also use the Configuration Wizard to rough out a

basic workflow that you can manually tweak later.

To create a simple workflow with the Configuration Wizard, select Getting Started | Configuration Wizard node and proceed with the steps in the right pane.

This wizard has five documented steps, as follows:

1. Creating a site
2. Creating a policy
3. Creating a task and a gathering job in it
4. Creating a reporting job in the task
5. Setting up real-time monitoring of the site

The wizard shows where to look in the treeview for the objects you create. If you need to reconfigure these objects, do this by editing their properties.

i | **NOTE:** The Configuration Wizard configures auditing and real-time monitoring only for the Microsoft Windows environment. To audit and monitor Unix-like systems, create the necessary InTrust objects manually.

Manual Workflow Configuration

For full functionality and control, access the properties of specific elements that participate in workflow: jobs, tasks, policies and sites.

InTrust provides several predefined tasks, such as Windows and AD Security Daily collection and reporting, SQL Server logs daily collection, and so on. These are some of the most common tasks with typical settings.

To activate a predefined task

1. In InTrust Manager, expand **Workflow | Tasks**.
2. Right-click the necessary task and select **Properties** to open the properties dialog box.
3. Select **Schedule enabled**. If this option is unavailable, click **Modify** to provide a schedule and then enable it.
4. Right-click the root node and commit the changes.

The examples below explain how you can organize gathering workflow to obtain a report (or reports) on the network segment you are interested in.

Prerequisites

A typical workflow requires the following:

- At least one InTrust site from where the audit data will be processed
- At least one InTrust server responsible for data processing
- At least one policy
- That an SMTP server is associated with the InTrust server or servers if you need to perform email notification

- A notification group with at least one operator in it, or several such groups
- At least one notification template

Example 1: A Simple Workflow

The simplest scenario, which can be replicated using the Configuration Wizard, includes the following:

1. Gathering audit data from a site into a repository and audit database (gathering job)
2. Generating a report containing the data (reporting job)
3. Notify an operator or operators when the task is completed (notification job)

To implement this scenario, you will need:

1. A scheduled task that includes the required jobs.
2. A gathering job that uses a site spanning your network section. Configure the job to store gathered data in a repository and an audit database.
3. A reporting job that will use the gathered data to prepare the necessary report in the necessary format. If you need a notification as soon as the reports are ready, use the **Notify by email** option in the job.

This kind of workflow is easy to create but not always acceptable due to the following reasons:

- A large and/or distributed network or network section makes data processing too slow for such workflow organization.
- If there is a slow link within an InTrust site, the traffic rate is a consideration. It is not efficient to transfer large amounts of data uncompressed (meaning, gather directly to the database over a slow link), because this would be error-prone and would result in a congested network.

To boost efficiency, gather data from multiple sites that have fast reliable links within them. The following is an example of workflow management.

Example 2: A More Efficient Workflow with Two Reports

Assume that you have an InTrust organization with at least two InTrust sites in it. A separate server processes each of these two sites.

Now, the objectives are as follows:

1. Gather audit data from the two sites
2. Generate reports based on the result of the gathering
3. Notify an operator or operators of task completion

A possible way to split this course of action into several stages is as follows:

1. Simultaneously gather data from both sites, and put each site's data into a separate repository and a separate audit database.
2. Draw up two reports.

3. Send a notification message to an operator or operators.
4. Consolidate data from two repositories for archival and backup.

Follow these instructions to organize workflow based on the described model:

1. Create a task that will include the jobs required to achieve the goals.
2. To gather the audit data from two sites, configure two gathering jobs to run simultaneously. Specify a separate repository and a separate audit database as data stores for each job.
3. Create a reporting job that will use the gathered data to generate the necessary report or reports. Configure this job to have the first gathering job created earlier as its predecessor job.
4. Create another reporting job. Configure the job to have the second gathering job as its predecessor job. When configuring each of the gathering and reporting jobs, select the corresponding InTrust servers (assigned to the sites from which the gathering jobs collect audit data).
5. 5 Create a notification job to inform an operator or operators that the task is completed. Configure the job to have both reporting jobs as its predecessor jobs. Using a notification task instead of built-in reporting job notification means that the operators will get the message after both reporting jobs have completed.
6. 6 Create a consolidation job that runs simultaneously with the notification job. Configure the job to consolidate the repository used by the second gathering job into the one used by the first gathering job.

Example 3: Consolidating and Importing to Create a Single Report

Another way to split the operation into stages in this case is as follows:

- Simultaneously gather audit data from both sites and put it in separate repositories
- Consolidate repository data into one repository
- Import the result to an audit database
- Prepare a report based on the gathered data
- Send a notification message to the operator or operators

Follow these instructions to organize workflow based on the model:

1. Create a task which will include the jobs required to achieve the goals.
2. Configure two gathering jobs to run simultaneously. Specify a separate repository for each job.
3. Create a consolidation job that copies the data from one of the used repositories to the other. Configure this job to have both of the jobs created earlier as its predecessor jobs. This will ensure that the schedule job is not started until the gathering is finished.
4. Create an import job that imports the gathered audit data from the consolidated repository to an audit database.
5. Create a reporting job that will use the gathered data to generate the necessary report or reports. Enable job completion notification if necessary.

Auditing Recommendations

Generally, InTrust gathering workflow is organized as follows: data from the audit trails is gathered from across the network with or without agents, processed by the InTrust server, and placed to the specified repository in a compressed form.

i **NOTE:** If necessary, you can collect data into an audit database, but in this case you must precisely select the events you need to store, because database size grows rapidly.

A repository is a file-based structure intended for long-term data storage. For further analysis and reporting, the necessary data is imported to an audit database on a SQL server.

Usually, reports are generated automatically on schedule and saved to network shares, from where they become available to users. The reports can also be sent by email. Reports can also be processed interactively using Knowledge Portal.

To deploy InTrust gathering in the way that best fits your auditing needs and network environment, you need to identify:

- Which events need to be archived
- Where they are to be stored
- How often they must be gathered
- What reports you need, and how they must be scheduled and distributed

The related topics provide you with several recommendations on deploying and configuring the InTrust gathering workflow, including:

- Data sources to be processed, and data collection scheduling—see [Data Sources](#)
- Required data stores, their location, capacity assessment and recommended data retention periods—see [Data Stores](#)
- Optimizing tasks—see [Tasks and Sessions](#)

Data Sources

What to Collect?

First, decide what data you want to collect from which sources. For example, in order to be HIPAA-compliant, organizations need to archive logon attempts, so they need to collect and store events from server security logs. On the other hand, to comply with external regulations, usually there is no need to process logs on a large number of workstations (perhaps with the exception of the application log to troubleshoot new software).

However, domain controllers, servers, computers with business-critical resources and computers storing confidential information need to be audited extensively.

To simplify planning and adjustment of the audit data gathering process, InTrust offers predefined gathering policies which specify audit trails to be collected from specific sources. Also, you can create gathering policies of your own.

i | **NOTE:** Detailed recommendations on setting up your auditing can be found in [Windows Auditing References](#).

How Often?

Next, you need to determine how often data needs to be gathered. Gathering is best scheduled for off-peak hours, such as at night, when user activity is at a minimum. You will need to consider time differences if you have computers in different time zones, and log overwriting timeframe (especially on domain controllers).

How to Improve Repository Viewer Experience?

Use a Short-Term Repository and an Archive Repository

The entire auditing workflow should not use a single repository. The recommended minimal setup is to have one repository for recent audit data (short-term repository) and another for data that is older (archive repository). First, temporarily disable indexing of your existing repository, and then clone this repository. Do not use conventional file copying or regular file managers. These methods may fail, because the hierarchical file structure in InTrust repositories uses very long names. Instead, use specialized replication software such as Microsoft Robocopy, which has been shipped with Windows since Vista and was available as part of the Windows Resource Kit before Vista.

! | **CAUTION:** Disabling indexing for the duration of the replication procedure is required so that the index data is copied correctly. Indexing is enabled and disabled in InTrust Manager in the repository properties, on the Indexing tab.

After you have made a copy of the repository in the location of your choice, do the following:

1. Enable indexing of the original repository again.
2. Create a new repository in InTrust. In the New Repository wizard, specify the location of the cloned repository.

The resulting repository is ready to be used for audit data archival, and the original repository is ready for regular data extraction and cleanup.

To organize archival of repository data, first decide on the retention period in the short-term repository—for example, 90 days. To make the decision, take into account how far back the events you view usually date. Then create a task with the following jobs:

1. A repository cleanup job that clears from the short-term repository all data older than your preferred retention period
2. A consolidation job that copies the entire Windows network-related contents of the short-term repository to the archive repository; make this job the successor of the repository cleanup job
3. If applicable, a consolidation job that copies the entire Unix network-related contents of the short-term repository to the archive repository; make this job the successor of the repository cleanup job

Schedule the task to run at intervals equal to your retention period.

If your current repository setup is different from this configuration, consider converting your current production repository to a short-term repository and setting up a separate archive repository.

Use Specialized Audit Databases

In a default configuration, all best-practice scenarios use the same audit database. If you implement two or more scenarios at once, the data in the audit database becomes less specialized, and your reports may show data that is not strictly relevant to a given scenario, because it was gathered for a different scenario.

If you want to avoid this, take the following steps:

1. Create a separate audit database for each scenario that you use.
2. For each scenario, switch all the import and reporting jobs in the daily, weekly and ad-hoc reporting tasks to the correct database.
3. In the **Best Practices | Daily Audit Database Cleanup** task, make a copy of the existing job, and switch the new job to the correct database.

Enable Agent-Side Log Backup for Gathering

Use agent-side log backup to speed up gathering. For details about agent-side log backup, see the [Keeping Event Data on the InTrust Agent Side](#) topic.

Turn on the **Enable log backup and use it to gather events** option for all data sources included in the following gathering policies:

- Auditing Domain Controllers: Events from DCs
- Auditing Exchange Servers: Events from Exchange Servers
- Auditing File Servers: Events from File Servers
- Auditing Workstations: Events from Workstations

You should not enable the related **Clear the backup after gathering** option for these data sources, because gathering occurs on multiple schedules, and one gathering job clearing the backup may interfere with another job's operation.

What Else to Consider?

- Using agents helps to minimize network impact when communicating data from target computer to InTrust server.
- If any network segments are located behind a firewall, InTrust agents in these segments have to be installed manually, for example, to collect data from the Web farm and monitor suspicious activities in the DMZ.
- Data from untrusted domains can be collected with or without agents (if you specify appropriate credentials in the site properties).
- Repositories can be consolidated through the firewall.

Data Stores

What Will You Need?

InTrust uses the following two types of data storages:

- Repositories, which are used to store data for long periods in a compressed form
- SQL server databases, which are used for analysis and reporting, and for storing real-time monitoring alerts and configuration data

You need to plan for the following special databases:

- An InTrust configuration database
- One or more InTrust audit databases for analyzing and reporting on the audit data collected by InTrust
- An InTrust alert database where the alerts generated by InTrust real-time monitoring will be stored (one alert database is recommended for an InTrust server)

In accordance with your company's data retention policy and administrative needs, you need to plan for the following:

- The number and locations of your repositories and databases
- Data retention periods for repositories and databases

Typically, a SQL server hosting the audit database features a quad-core 3GHz processor with plenty of RAM and hard disk space. In particular, it is recommended that database size should be kept under 100GB; data retention period depends on your reporting needs.

How Many Repositories and Databases?

To plan how many repositories and databases you need to create, first estimate how much data you need to gather, even though your initial estimates may not be accurate. If you know approximately how many events per day are generated on your servers, you can plan the storage volumes. When calculating, consider the following:

- In a repository, data is stored in a compressed form with a compression ratio near 1:15 compared to EVT format. If stored directly in the database, each event takes about 5 times more space than the original format.
- While you can keep practically unlimited amount of data in a repository, it is recommended that you store only about 2–4 weeks' worth of the recent logs in the database for reporting, and import older data from repository when needed, as described later in this document.

You may require more than one repository, for example, for the following reasons:

- Company departments need to collect data separately.
- Security boundaries make centralized collection difficult.
- Network bandwidth, particularly slow links between sites, imposes limitations on centralized collection.

You can have as many repositories as you need. For centralized data analysis, reporting and backup on external media, data from several repositories can be consolidated into a central location.

For local reporting, data needs to be imported from the local repository into the local database. Consider this recommendation when estimating the number of databases required.

InTrust supports for both file-based and Centera-based repositories. File-based repositories are the structures of folders and compressed files secured by traditional NTFS permissions (or not secured if repositories are located on the FAT32 file system). Centera-based repositories use EMC Centera™ devices for storage and offer content-addressed storage system where data cannot be modified if once recorded. This native feature helps you to secure the archived data stored on EMC Centera. Note that if you have EMC Centera device operating in your environment, you can organize corresponding InTrust repository as described in [Understanding InTrust Repositories](#).

Generally, you can have more than one audit database—for example, one for ongoing reporting tasks, keeping the latest data (for example, for the last month), and other—for historical data analysis and reporting when necessary.

How Long to Keep?

There are no strong limitations on how long you can keep data in your repository. In accordance with your organization's data retention policy, you can periodically back up your repositories to magnetic tape or other external media. Also, you can automatically clean up your repositories and databases to get rid of obsolete data by configuring a special cleanup job and scheduling it to run as frequently as you need. Typically, a repository stores data collected during the last 6–12 months.

If you need to generate reports from data in a repository, for example, in case of internal investigation, then you can import the necessary data from the repository to the database, and generate the reports. Since the importing process copies information from repository to database, the events records will still be kept in the repository, allowing you to clean up the copies from the database.

Where to Locate?

It is recommended that you locate your repositories within one network area with an InTrust server to minimize network impact when communicating data between repository and server.

Audit databases must be located on a Microsoft SQL server with a reliable broadband connection (100Mbps+) with the InTrust server.

To provide for data integrity and availability, it is recommended to locate the default repository on a dedicated, well-protected file server (in case you plan to use a file-based repository).

Alternatively, you can locate an InTrust server, its configuration database, repository, and the audit and alert databases on one computer. Pros and cons of the solution are as follows:

- You ensure fast and reliable access to InTrust configuration data
- Data import takes less time than in any other case
- However, the hosting computer must be a high-powered server with fine-tuned resource allocation (considering that InTrust server requires 4 GB of RAM). Generally, to co-locate an InTrust server, repository and all databases, a hosting computer should have a multiprocessor (for example, 4GHz processor), plenty of RAM (8GB+ recommended) and hard disk space (100GB+ recommended).

Tasks and Sessions

Generally, it is recommended that you use copies of predefined InTrust objects, in particular, tasks, modifying them to fit your auditing needs. It is also recommended that when configuring such a task (copy), you set the **Keep sessions** option to the **Last <number_of_sessions>** or **During <time_interval>**, rather than **All**. This will prevent from performance loss that may occur if too much session data is stored to the database.

i | **NOTE:** If you create a task of your own, the **Keep sessions** value is set to **Last 10 sessions** by default.

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece — you — to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit www.quest.com/contact.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to-videos
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product