EVIDEN

Identity and Access Management

DirX Audit

Best Practices

Version 7.2, Edition June 2025

Table of Contents

Preface

This document provides administrators with guidelines and tips intended to help avoid common mistakes and improve the experience of a DirX Audit installation.

DirX Audit Documentation Set

DirX Audit provides a powerful set of documentation that helps you configure your audit server and its applications.

The DirX Audit document set consists of the following manuals:

- *DirX Audit Introduction*. Use this book to obtain a description of DirX Audit architecture and components.
- *DirX Audit Tutorial*. Use this book to get familiar quickly with your DirX Audit installation.
- *DirX Audit Administration Guide*. Use this book to understand the basic tasks of DirX Audit connectivity administration.
- *DirX Audit Manager Classic Guide*. Use this book to obtain a description of the DirX Audit Manager Classic user interface provided with DirX Audit.
- *DirX Audit Manager Guide*. Use this book to obtain a description of the DirX Audit Manager user interface provided with DirX Audit.
- *DirX Audit Command Line Interface Guide*. Use this book to obtain a description of the command line interface provided with DirX Audit.
- *DirX Audit Customization Guide*. Use this book to customize your DirX Audit environment.
- *DirX Audit History Synchronization Guide*. Use this book to obtain information about the DirX Audit History synchronization jobs.
- *DirX Audit Best Practices*. Use this book to obtain guidelines and tips for avoiding common mistakes and improving the experience of a DirX Audit installation.
- *DirX Audit Installation Guide*. Use this book to install DirX Audit.
- *DirX Audit Migration Guide*. Use this book to migrate DirX Audit from previous versions.
- *DirX Audit Release Notes*. Use this book to understand the features and limitations of the current release.
- *DirX Audit History of Changes*. Use this book to understand the features of previous releases.

Notation Conventions

Boldface type
In command syntax, bold words and characters represent commands or keywords that must be entered exactly as shown.

In examples, bold words and characters represent user input.

Italic type
In command syntax, italic words and characters represent placeholders for information that you must supply.

[ ]
In command syntax, square braces enclose optional items.

{ }
In command syntax, braces enclose a list from which you must choose one item.

In Tcl syntax, you must actually type in the braces, which will appear in boldface type.

|
In command syntax, the vertical bar separates items in a list of choices.

...
In command syntax, ellipses indicate that the previous item can be repeated.

install_path
The exact name of the root of the directory where DirX Audit programs and files are installed. The default installation directory is *userID_home_directory*/**DirX Audit** on UNIX systems and **C:\Program Files\Atos\DirX Audit** on Windows systems. During installation the installation directory can be specified. In this manual, the installation-specific portion of pathnames is represented by the notation *install_path*.

install_media
The exact path where the DirX Audit installation media is located.

1. Best Practices

This guide gives administrators guidelines and tips intended to improve the experience of a DirX Audit installation and avoid common mistakes. To fully benefit from this guide, any administrator needs to have a solid knowledge of the management and administration of the underlying database.

The guide is organized into the following chapters:

- "Before Deployment" describes best practices to be performed before DirX Audit is installed and deployed, mostly in the audited products.
- "Initial Deployment" describes items to be considered during DirX Audit configuration.
- "Day-to-day Management" gives checklists and provides hints for running in a production environment.
- "Troubleshooting" describes how to track down and solve problems.

2. Before Deployment

This chapter describes best practices that you perform before DirX Audit is installed and deployed, mostly in the audited products.

2.1. Controlling the Number and Size of Audit Events

The number of audit events and their content are the most important factors that affect the size of the database for audit events. Minimizing their number and size supports performance and allows you to keep the events in the DirX Audit Database for a longer time.

2.1.1. Controlling Audit Events in DirX Identity

You use the audit policies of a DirX Identity domain to control the audit events and their content. Use Audit Policies in the Audit Trail folder of the Auditing view in DirX Identity's Provisioning view group to manage these policies.

An audit policy must be active for each entry type to be audited. The entries are identified either by their object description name (for service layer components) or by LDAP object class and optionally dxrType. Each object requires an **audit object type**. Please do not change the default types! They are evaluated in subsequent processes; for example, to calculate the digest, which is a business summary of the audit event.

The list of **audited attributes** controls when an audit event is produced: always when one of these attributes has been changed. If none of these attributes is changed, no event is generated.

When an event is produced, the list of **identifying attributes** is always entered into the event. Identifying attributes help to:

- Display a recognizable name for the entry in DirX Audit Manager's Audit analysis and in reports. For example, the surname and given name of a user or account. On the other hand, you should pay attention to personal data protection and consider an appropriate pseudonym like a corporate user identifier as an option.
- Categorize the entry; for example, that it is an organizational unit.

Fewer audited entries, fewer audited attributes and fewer identifying attributes reduce the total size of the DirX Audit Database.

The above audit policies apply when entries in the DirX Identity Store are changed. They do not apply to changes in connected systems by the Provisioning workflows. They write audit events only when the flag "Write Audit Log" is set for the controller of the join activity of the workflow. As lot of the attributes in the connected systems are different, only a standard set of identifying attributes can be inserted into an audit event and only if its attributes contain these standard names.

Please also pay attention to the channel names of the Provisioning workflows. They are used as the audit object type in the audit event, so they appear, for example, as dimension values in charts next to the standard types like user, role, account, and group. As mentioned before, they also affect the digest generation. This is especially important for accounts and groups. The channel names in the Provisioning workflow configuration should at least contain the string “account” or “group” respectively.

2.1.2. Controlling Audit Events in DirX Access

Use DirX Access Manager’s **Audit Service** configuration section to control the number of audit events produced by DirX Access. You can configure the auditing subsystem to push a substantial number of audit events to auditing sinks. Carefully consider individual auditing levels based on the planned server characteristics and the expected audit messages usage.

To access the configuration, select **Servers** in the main navigation bar and then choose the **Cluster** page and its **Audit Service** section. Here you can configure various auditing levels (**info, warning, error, none**) for the following DirX Access features and subsystems: server lifecycle, application repository service, authentication service, authorization service (decision making), authorization service (subject attribute finder), authorization service (policy finder), configuration service, federation service, policy service, SSO service and user service. For more information, see “Managing DirX Access Cluster Properties” in the *DirX Access Administration Guide*.

2.2. Managing History Entries

DirX Identity entries and their changes are synchronized together with validity timestamps to the DirX Audit History Database by History Synchronization jobs that run as scheduled jobs on the DirX Audit Server. There are three aspects relevant to managing these jobs:

- Managing the number of history entries, history entry types, and attributes stored in the DirX Audit History Database
- Managing the mapping of synchronized attributes to DirX Audit History Database tables
- Managing the scheduling of the jobs

2.2.1. Controlling Database Size

You must decide which entry types and which of their attribute types are stored in the DirX Audit History Database. Note that invalidated entries and invalidated attribute values remain in the DirX Audit History Database even after they have been deleted from the DirX Identity domain; they receive a VALID_TO timestamp to indicate the end of their lifetime.

DirX Identity entry types are represented in DirX Audit History Database by entry type definitions. Each type definition in a DirX Audit tenant properties file contains an entry type name in the DirX Audit History Database and LDAP search base and a search filter that defines a set of synchronized entries from DirX Identity. There are predefined types for all standard entry types in DirX Identity. LDAP History Synchronization jobs are scheduled for one or more entry type.

The first thing to consider is which entry types you need to have synchronized in the DirX Audit History Database. Important entries are users, target systems with accounts and groups, roles, and permissions, if they are used. Also important are request (approval) workflow instances, but what about their definitions? Do you work with tickets, which are change orders to be executed at a future due date? Do you use certification campaigns? Is it important that you have delegations or configuration objects in the DirX Audit History Database?

Once you select the entry types to be stored, you must select which of their attributes are stored. Some attributes that should not be stored in the DirX Audit History Database are already excluded; for example, dxrHistory or passwords. You can also exclude attributes that are not important for you. For details, see the section “Excluding Attributes from the Synchronization“ in the *DirX Audit History Synchronization Guide*.

Over time, the DirX Audit History Database size will grow as new entries and attributes are added, but nothing is deleted automatically. Typically, entries and attributes invalidated more than perhaps one year ago are no longer relevant. Consider removing them with the DirX Audit **dxthistdbtool** tool. See the section "Remove Old Data" in the chapter “Day-to-Day Management” in this guide for details. Note that this tool can also archive data. In a subsequent version, it will allow you to re-import the archived data later when necessary.

2.2.2. Controlling Attributes Mapping

DirX Identity entry attributes are stored in small, link, and large attributes database tables. The configuration of attribute mapping to tables is in the DirX Audit tenant properties file. See the section “Configuring and Customizing DirX Audit History Synchronization Jobs“ in the *DirX Audit History Synchronization Guide* for details. The table for small attributes can only take values with a maximum size of 850 characters. If History Synchronization jobs try to insert an attribute value longer than this, DirX Audit Server logs a warning and only the truncated value is stored in the DirX Audit History Database.

Attributes with a value longer than the limit should be stored in the large attributes table. Note that large attributes are not currently displayed in DirX Audit Manager, so storing them in the DirX Audit History Database is not of much value and thus they are excluded from History Synchronization jobs by default. See the section “Customizing Attribute Mapping“ in the *DirX Audit History Synchronization Guide* for information on modifying attribute mapping.

Attributes that are links to other entries, for example, DN and dxrGroupMember* are synchronized to the link attributes table. This table shares maximum attribute size restrictions with the small attributes table. Attribute values exceeding the limit are stored truncated.

2.2.3. Scheduling History Synchronization Jobs

There are two types of History Synchronization jobs:

- Modify job – Synchronizes newly created and modified DirX Identity entries and their attributes to the DirX Audit History Database. Modify job can be scheduled for one or multiple entry types. Multiple types are synchronized sequentially one by one.
- Delete job – Detects deleted DirX Identity entries and sets the “VALID TO” timestamp for them and their attributes in the DirX Audit History Database.

Normally it is sufficient to run Modify job for all entry types that you need to have in your DirX Audit History Database once a day before running other DirX Audit History Database jobs. You should schedule one Modify job for all synchronized entry types. If you have one entry type in DirX Identity that is frequently modified, you should consider running Modify job for the type more often.

Generally, DirX Identity entries are not frequently deleted. It should be sufficient to run Delete History Synchronization job for all your synchronized entry types once a week.

You must consider the History Synchronization jobs schedule and adjust it to your needs. Keep in mind that it is not possible to run more than one Modify or Delete job at a time. This situation is indicated in the DirX Audit Server log. If you try to start a DirX Audit History Synchronization job while another History Synchronization job is running, the second job will not start.

2.3. Establishing Secure Communication

We recommend using secure communication over SSL/TLS channels between the DirX Audit components and to and from external components, especially from DirX Identity / DirX Access audit plug-ins and to the DirX Audit Database(s). You can find detailed information on how to set up secure communication in the *DirX Audit Administration Guide* and the *DirX Audit Installation Guide*.

For SSL/TLS communication, a trust-store and often for client authentication a key-store is required. We recommend storing them in the folder *install_path*/**conf/crypto/stores**. For more information, see the section “Managing Cryptographic Material” in the *DirX Audit Administration Guide*.

2.4. Securing Apache Tomcat

Tomcat is configured to be reasonably secure for most use cases by default. However, we recommend to secure Tomcat beyond the default installation and especially consider the advice from the following OWASP and Apache Tomcat resources:

https://tomcat.apache.org/tomcat-9.0-doc/security-howto.html

https://wiki.owasp.org/index.php/Securing_tomcat

Here a few recommendations that can easily be realized:

2.4.1. Webapps Folder

The folder *install_path*/**webapps** contains default web applications that should not publicly be available in productive deployments. Especially remove the **documentation**, **examples** and **root** applications. If you decide to allow remote administration and keep the **Manager** or **Host Manager** applications, secure them appropriately.

2.4.2. Version Info

Some default components like the Error Report Valve include Tomcat server version info in their responses. Eliminate that by setting appropriate information in the Server Info properties:

Create the file **ServerInfo.properties** in the folder *install_path*/**lib/org/apache/catalina/util** with the following content:

```
server.info=Apache Tomcat
```

2.5. LDAP Authentication

This section describes aspects of LDAP authentication that need your attention before deploying DirX Audit.

2.5.1. Managing Group Search

When an auditor logs in to DirX Audit Manager, the Manager reads the user's groups from LDAP to find the right DirX Audit application role for the auditor. If there are a lot of groups that match the group search criteria, this operation can take a very long time.

We recommend refining the search criteria and if possible, moving the auditor groups to a separate subfolder in LDAP. Adapt the search base and search filter accordingly using the Authentication Configuration in the Configuration Wizard for the Tenant Configuration.

2.5.2. Disabling Endpoint Identification

As of Java 1.8.181, the endpoint identification for LDAPS (secure LDAP over TLS) connections is enabled by default. If the certificate doesn't contain the right host (hostname or IP address) of the LDAP server (in SAN = subject alternative name), the Java runtime cannot set up the TLS connection.

The best solution is to add the host to the certificate. If this is not possible, you can disable endpoint identification by setting the following Java system property when starting the JVM:

```
-Dcom.sun.jndi.ldap.object.disableEndpointIdentification=true
```

3. Initial Deployment

This chapter describes topics that you should consider when installing DirX Audit and running the initial configuration.

3.1. Use the JMS Audit Event Collector

Several event collectors for importing audit events into the DirX Audit Database are provided. For DirX Identity, the LDAP collector is mandatory: it reads the audit events from the **dxrHistory** attribute of the LDAP entries.

For DirX Access and the additional DirX Identity audit messages, you must choose between the JMS collector and the File collector. For both cases, we recommend selecting the JMS collector. The disadvantages to using the File collector are:

- Audit files must be moved to the import folder of the DirX Audit File collector. Make sure that they are moved only after the provider (DirX Identity / DirX Access) has closed them.
- There is a risk of inadvertently copying files repeatedly to the import folder and thereby provoking error logs of duplicate import.
- The amount of disk space required to store the files is typically greater than what the Message Broker requires; with the Message Broker, the messages are processed immediately and are then automatically removed from the Message Broker repository.

3.2. Exclude Unused Fact Tables

The DirX Audit Server's fact population job regularly computes a lot of facts and then stores them in fact tables whose names start with **FCT_**. Your company may not use some of these tables. If you exclude these unused tables from processing, the fact population job is faster and the database requires less disk space.

The default fact tables along with the configuration for dimensions and facts are configured in the folder *install_path***/conf/fact_configuration**. The configuration of the fact tables themselves is in the file **confFactTables.xml**. If you want to customize them, copy these files to the tenant-specific folder:

*install_path***/conf/tenants/***tenantID***/fact_configuration**

and then adapt them. For more detailed information about this task, see the section "Managing Fact and Dimension Tables" in the *DirX Audit Administration Guide*.

If you do not have DirX Access installed, consider dropping the fact table for authorizations:

- FCT_AUTHORIZATIONS

If you do not have the license for the History feature installed, consider dropping the fact tables based on history entries; the names of these fact tables start with **FCT_HST_**.

You must also remove references to the SQL scripts modifying the excluded fact tables in the following configuration files:

*install_path***/conf/sql/common/factpopulation/factpopulation_scripts_datadb_list.txt**

*install_path***/conf/sql/common/factpopulation/factpopulation_scripts_historydb_list.txt**

3.3. Schedule Post-Processing Jobs in Logical Order

The DirX Audit Server hosts several data post-processing jobs that extend audit messages and history entries. History entries are imported into the DirX Audit History Database by History Synchronization jobs running on DirX Audit Server. All these jobs have a logical order of execution as some jobs process data that have been updated by other jobs.

We recommend scheduling them in the following order:

1. DirX Audit History Synchronization jobs – imports the history entries from the DirX Identity domain into the DirX Audit Database. You can schedule Modify or Delete job for multiple entry types. They execute sequentially one by one.
2. History DB update – adds virtual attributes to history entries that support subsequent dimension and fact calculation. Especially it adds the database foreign keys to link table attributes and thereby establishes the references between history entries.
3. Fact population – fills the dimension tables (DIM_) and then the fact tables (FCT_). Lots of fact tables are based on the history entries and especially on their references and virtual attributes.

Schedule the following DirX Audit jobs to be repeated short term:

- Context records calculation – calculates relations between audit messages that were caused by the same core event; they are considered to be in the same context. The job processes only audit messages not yet associated to a context. If it is scheduled every 5 minutes (the default), there is only a small number of events and it finishes quickly.
- Scheduled reports – produces and sends the requested reports. It checks not only regularly scheduled reports, for example daily or weekly, but also those that are started in DirX Audit Manager to be generated “as soon as possible”. Therefore, the job should run frequently; the default is every 30 seconds.

Reports are scheduled individually. In order to get up-to-date data, they should be scheduled after the respective jobs have run. For example, if they display history entries, they should be scheduled after the history update job. If they use facts from fact tables, they should be scheduled after the fact population job.

Additional scheduling considerations include the following:

- The jobs generate high load both in the DirX Audit Server and in the database server. If your configuration has multiple tenants, the jobs for the tenants should not all run at the same time.
- The jobs update database tables and during this time block them for other updates and sometimes also for parallel reading.

- Indexes should be reorganized or rebuilt from time to time as the content of the table changes. During an index update, the database server locks the corresponding table. The need for index update depends on the number of changed records, namely those in the tables for audit messages and history entries. As a rule of thumb, update indexes and statistics once per week and outside of business hours. For more information on indexes, see the section “Maintain Database Indexes” in chapter “Day-to-Day Management” of this guide.

4. Day-to-Day Management

This chapter suggests some tasks you should perform in day-to-day management of a DirX Audit deployment to maintain its performance.

4.1. Check the Error Logs

All DirX Audit services write log files. Check them daily for any warnings and errors. For the location of the log files see the section “Log Files” in the Troubleshooting chapter of this guide.

4.2. Check for Audit Message Import Errors

When viewing the DirX Audit Server logs, you may encounter error logs about “unique key constraint violation”. These logs occur when the server repeatedly tries to import the same audit message.

Most commonly, it occurs when a file with audit messages has been put repeatedly into the file collector’s input folder. Rarely, it occurs when the import of audit messages from the DirX Identity domain via LDAP has been interrupted so that the LDAP collector can insert the message into the DirX Audit Database but not remove the record from the **dxrHistory** attribute in DirX Identity domain.

The collectors identify such issues and just store the not-imported audit message into a separate error folder. You can safely remove these files.

4.3. Monitoring Missing Links between History Entries

In DirX Audit Manager’s History view or in reports on history entries, you might miss a link from one entry to another; for example, from an account to a group where it is a member. In such a case, the link information is not available in the database. Reasons might be:

- The target entry (in the account / group example: the group) is not available in the DirX Identity domain and is therefore also not in the DirX Audit History Database.
- The target entry has not yet been imported into the DirX Audit History Database.
- The target entry and the link value were available in the DirX Identity domain only for a short time between two consecutive runs of the History Synchronization jobs and were therefore never imported into the DirX Audit History Database.
- The referenced target attribute is not mapped to the HST_LINK_ATTRS_IN_TIME table of the DirX Audit History Database.
- The lifetime of the link value and the value in the referenced entry do not overlap.
- The reference value in the DirX Identity domain is wrong.

Some background details:

The link value in DirX Identity domain is either an LDAP DN or, in the case of account-group memberships, a value referring to a configurable attribute in the target entry. In the default case for accounts and groups, the value in, for example, **dxrGroupMemberOK** should refer to the **dxrPrimaryKey** of the target entry with a value relevant only in the target system; for example, a DN in Active Directory or a primary key in a relational database.

For performance reasons, DirX Audit Server calculates a foreign key from this original information and stores it in the column LNK_ENTRIES_ID of the table HST_LINK_ATTRS_IN_TIME. This calculation is part of the regularly scheduled History DB update job. If for some reason a foreign key cannot be calculated, the column remains empty and DirX Audit Manager and reports cannot navigate to the referenced entry.

To generate an overview of these cases and allow administrators to remediate them, DirX Audit provides some reports on missing links. The reports can be identified by the tag **Monitoring**. A link is considered missing if the foreign key in the LNK_ENTRIES_ID column of the HST_LINK_ATTRS_IN_TIME table is empty.

- The report **Missing entry links** shows the number of missing links for all entry types. For each entry type (such as Account or User), it lists the link attributes and, for two selectable time points, the number of missing links for all link values. The last column **Difference** gives the difference between the two time points. A summary line after each entry type shows the trend as a graph and the aggregated numbers for all link attributes.
- The report **Missing entry links by target system** provides the same information for each target system. As missing links most often occur for account-group memberships, this report is considered the most important.
- The report **Missing entry links with details (CSV format)** provides the most detail for all entry types in a format that can easily be post-processed by tools like Microsoft Excel. For each link attribute and value, it contains the **dxrUid**, **dirxEntryUUID** and entry name (normally the **cn**) along with the lifetime of the value (VALID_FROM - VALID_TO).

For more details on these reports, see the chapter “Monitoring DirX Audit Data with Reports” in *DirX Audit Administration Guide*.

We recommend running at least one of these reports regularly (weekly or monthly). If you observe higher numbers of missing links that do not eventually disappear, investigate these cases in more detail. Values that have been created on the current or last day might be explained by the schedules of the DirX Audit History Synchronization jobs and the DirX Audit History DB update job: the target entry might not have been updated in time.

For the other cases, check the potential reasons listed at the beginning of this section. Maybe you need to run the History Synchronization Modify jobs multiple times per day. Or if a referenced entry is never imported into the DirX Audit History Database, consider scheduling History Synchronization jobs for the type of the referenced history entry or exclude such the link attribute from the synchronization.

4.4. Minimize Directory Server Restores

Try to keep directory server restores to a minimum to avoid inconsistency in your DirX Audit History Database. Directory server restores may reimport already deleted DirX Identity entry attributes. These attributes are then synchronized to the DirX Audit History Database, where they affect the consistency of the history snapshot of your DirX Identity domain.

4.5. Check the Audit Database Size

The DirX Audit Database increases every day: new messages are inserted, history entries and their attributes are added and changed, facts on all that are added. For large queries or index updates, the database engine temporarily needs considerable space.

Make sure that the database has enough disk space and consider updating that monthly. See also the section “Remove Old Data” in this guide.

4.6. Maintain Database Indexes

Database indexes are a very important factor for the performance of database queries. This section provides hints on maintaining database indexes for best performance.

4.6.1. Check for Missing Indexes

DirX Audit comes with default indexes out-of-the-box. These default indexes may not be enough for the queries or the data of every customer. If some reports or some queries performed in DirX Audit Manager take unnecessarily long, a missing index can be the reason. In this case, you may need to analyze the query.

The *DirX Audit Administration Guide* in section “Tuning Database Performance” gives hints on how to analyze queries and tune database performance.

The following SQL queries provide a list of current indexes in the database. Choose one according to your platform.

*install_path***/conf/sql/adm/get_indexes_MSSQL.sql**
*install_path***/conf/sql/adm/get_indexes_Oracle.sql**

4.6.2. Update the Database Index

When database tables are changed (records created, updated, and deleted), the index also needs to be updated; otherwise, the index quality degrades over time and at some point, when building its execution plan, the database engine may decide to stop using the index and resort to full table scans.

Modern database engines support regular index updates automatically by default. However, the way they implement this function differs depending on the engine; in some cases, the automatic processes may not be enough.

For example, the Microsoft SQL Server database engine continuously modifies the index. This process scatters the information in the index, fragmenting it over time. Index fragmentation decreases the query performance until at some point, the database engine decides to perform full table scans. As a result, Microsoft recommends regularly and automatically rebuilding or reorganizing the indexes.

What is better: reorganizing or rebuilding? A rebuilding process locks the affected tables and views and thus blocks the clients working on them. However, rebuilt indexes are more effective. Microsoft recommends reorganizing the index if the percentage of fragmentation is between 10% and 30% and rebuilding the index if the percentage is higher than 30%.

The Oracle database engine, on the other hand, considers table statistics when calculating its execution plans. By default, it schedules a statistics-gathering job in the maintenance window at night to regularly update statistics.

Even automatic periodic index and/or statistics updates may not be adequate when, for example, a bulk load inserts or deletes a high portion of table rows and columns. Query performance may be significantly impacted until the next index rebuild. In these cases, administrators should start the index rebuilding process manually; for example, by running a regular job after the bulk load or before follow-up jobs.

For determining appropriate time ranges, it helps to build logical groups of tables and schedule the index rebuild for each group separately. Natural logical groups are: tables with history entries (HST*, HDB*), tables with audit events (DAT*), and fact with dimension and risk tables (FCT*, DIM*, TAG*, RSK*). For sample scripts, see the folder:

*install_path***/conf/sql/adm**

4.7. Remove Old Data

From time to time, you should remove old data from the DirX Audit Database. Otherwise, the disks may become full and cause performance to degrade.

4.7.1. Audit Messages

For audit messages, export and delete old messages by running the DirX Audit DB tool. See the chapter “Using the DirX Audit Tools” in the *DirX Audit User Interface Guide*.

Use a function for the parameter **to** to specify up to which date events should be removed. For example value **`$bday(-365)`** tells the tool to export – and optionally delete – events older than a year.

Note that the fact and dimension tables remain untouched. That allows you to see charts on aggregated data even for events that have been removed. Of course, a drill-through to details of such events is then not possible anymore.

If necessary, you can re-import the archived audit events using the import function of the DirX Audit DB tool.

4.7.2. History Entries

For history entries, delete old history data (mainly attribute values) by running the DirX Audit DB tool. See the chapter “Using the DirX Audit Tools” in the *DirX Audit User Interface Guide*.

Use a function for the parameter **to** to specify up to which date data should be removed. For example value **`$bday(-365)`** tells the tool to delete history data older than a year.

Our experience shows that approximately half a million history entries can be completely exported into files having the total size of about 1 GB.

Note that the fact and dimension tables remain untouched. This allows you to see charts on aggregated data even for history entries from which data have been removed.

There is currently no option to re-import the deleted history data using the DirX Audit DB tool.

4.7.3. Schedule Jobs

DirX Audit Configuration Wizard allows you to schedule the Purge job to remove history entries, audit messages, and original messages on a regular basis. For example, you can remove audit messages as well as history entries and attributes older than a year.

5. Troubleshooting

This chapter provides tips to track down and solve problems.

5.1. Checklist

Check that:

- All services are running and are connected to the correct target system.
- All required features share the same and correct configuration (distributed installation).
- All required port connections are allowed in the firewall.
- Target systems, components and applications are correctly configured and available.
- All passwords are correct.
- All services have administrative access to required repositories and disk folders.
- Truststores and keystores are in place and have the correct certificates and passwords.
- Connection / configuration has been performed for the correct tenant.

5.2. Log Files

Location of log files:

- For DirX Audit Manager (Apache Tomcat):
  *tomcat_install_path***/logs/dirxaudit-manager.log**
- For DirX Audit Server (Spring Boot):
  in the folder *install_path***/server_container/tenants/***tenant_name***/logs**:

  dirxaudit-server.log
  dirxaudit-server-errors.log
  security.log
- For DirX Audit Message Broker (Apache ActiveMQ):
  *install_path***/message_broker/data/activemq.log**

To manage logging, see the chapter “Configuring Logging” in the *DirX Audit Administration Guide*. If you need more detailed information in the logs, set the log level temporarily to DEBUG/TRACE (Log4j) or FINE/FINEST (Java Logging) depending on the logging framework in use. Don’t forget to reduce the log level to WARNING or INFO after you’ve finished debugging so that performance and disk space are not affected more than necessary.

5.3. DirX Audit Server

When you check the logs of DirX Audit Server, here is a list of error, warning, or information logs to which you should pay attention:

5.3.1. Slow Authentication Due to Many Groups

The authentication process checks that the authenticated user is in one of the configured auditor groups. This process may take too long if there are too many groups in the LDAP tree.

Try to refine the LDAP search for groups so that there are fewer matching entries. For example, move the auditor groups into a separate LDAP subfolder and take this folder as search base.

5.3.2. Reports Are Not Sent

When reports are not sent, check that:

- Sending mails is enabled in the Common Configuration dialog in the DirX Audit Core Configuration Wizard.
- The DirX Audit Server service is running.
- The default email address and the SMTP server address are correctly entered.
- The user who created the report has a correct mail address LDAP attribute value.
- The attachment size does not exceed the target mail server limit.
- The SMTP server is configured to allow relaying from the machine where DirX Audit is installed.

5.3.3. Unique Key Constraint Violation

A “unique key constraint violation” in audit collector logs indicates duplicate events. A collector tried to import an audit message that is already stored in the DirX Audit Database. That can occur, for example, if:

- A file is put to the import folder that has already been imported previously
- A **dxrHistory** LDAP attribute is read a second time after an LDAP server restore
- An imported **dxrHistory** attribute value could not be deleted after a shutdown of the LDAP or the DirX Audit Server

The collector’s error handling stores such messages in a special subfolder of the error folder. You can ignore them. Just make sure that the folders are cleared regularly.

5.3.4. Resetting the DirX Audit Server

When the source for errors cannot be identified, sometimes it helps to restart the server:

- Stop the DirX Audit Server - Windows service or UNIX daemon.
- Start the DirX Audit Server. No additional resetting steps are required, unlike previous versions.

5.3.5. Truncating Long Values in the DirX Audit History Database

A “checkAttributeLength: Value 'attributeValue' too long for database table 'table' and column 'column'.“ in the DirX Audit Server History Synchronization job logs indicates that History Synchronization job detected a value that exceeded the database column size. Such a value is truncated to fit the DirX Audit History Database table and column and synchronized. You can consider changing the attribute mapping. See the section “Customizing Attribute Mapping“ in the *DirX Audit History Synchronization Guide* for details.

5.4. SQL Knowhow

This section gives some advice for working with the DirX Audit Database and lists some SQL query scripts that help to find more detailed information in the database when DirX Audit Manager is not sufficient. The queries can be found in the folder:

*install_path***/conf/sql/common/adm**.

5.4.1. Refreshing Materialized Views (Oracle)

To enhance performance, some materialized (or indexed) views are defined in the DirX Audit History Database. If reports or charts do not seem to contain up-to-date information, one of the reasons may be an inadequate refresh of these views.

In Oracle Database, a schedule must be configured for the view refresh. These views typically are about relationships with users, roles, groups, etc. They depend on foreign keys that are regularly calculated by the History DB job. Hence, the view refresh should run after the History DB job finished.

To configure a schedule with Oracle SQL Developer, select Scheduler and navigate to DBMS Jobs >> Run Job.

The materialized views can be also refreshed manually on request.

5.4.2. Generating a List of Indexes

The queries **get_indexes_MSSQL.sql** and **get_indexes_Oracle.sql** lists all the indexes with the table and column to which they apply.

5.4.3. Searching a User in the History Database

If you want to see details on a user, you can first find the user's primary key with a search based, for example, on the family name and then query on all or specific attributes.

A sample query for finding a user based on the family name is in **hst_user_search.sql**; just replace the default family name string with the appropriate one.

To view all the link attributes of that user entry later than a certain day, use the following query and replace the values for the HST_ENTRIES_ID (primary key) and the VALID_FROM:

```
select
 *
from
 HST_LINK_ATTRS_IN_TIME
where
 HST_ENTRIES_ID = 15298
    and VALID_FROM >= '2019-01-01'
```

If you are only interested in a particular attribute, add the following clause with the right attribute name:

```
and ATTRIBUTE_NAME = 'dxrassignedaccounts'
```

If you are interested in other entry types, for example, accounts, replace the type in the **hst_user_search.sql** with the appropriate one; for example, Account. Find all the type names in the table **HST_ENTRY_TYPES**; they are the names that are configured in the corresponding History Synchronization job.

5.4.4. Listing Occurrences of dxrUids

The query **list_dxruid.sql** lists dxrUid values with their type and count. A count that is not equal to 1 indicates entries with duplicate dxrUid values.

5.4.5. Listing Duplicate dxrUids

The query **get_duplicated_dxruid.sql** lists duplicate entries with multiple unfinished records in the HST_ENTRIES_IN_TIME table and therefor duplicate dxrUID values for joining. Besides the dxrUid, it returns the entry type, DN, the number of occurrences and the time from when they are valid (VALID_FROM).

DirX Product Suite

The DirX product suite provides the basis for fully integrated identity and access management; it includes the following products, which can be ordered separately.

DirX Identity

DirX Identity provides a comprehensive, process-driven, customizable, cloud-enabled, scalable, and highly available identity management solution for businesses and organizations. It provides overarching, risk-based identity and access governance functionality seamlessly integrated with automated provisioning. Functionality includes lifecycle management for users and roles, cross-platform and rule-based real-time provisioning, web-based self-service functions for users, delegated administration, request workflows, access certification, password management, metadirectory as well as auditing and reporting functionality.

DirX Directory

DirX Directory provides a standards-compliant, high-performance, highly available, highly reliable, highly scalable, and secure LDAP and X.500 Directory Server and LDAP Proxy with very high linear scalability. DirX Directory can serve as an identity store for employees, customers, partners, subscribers, and other IoT entities. It can also serve as a provisioning, access management and metadirectory repository, to provide a single point of access to the information within disparate and heterogeneous directories available in an enterprise network or cloud environment for user management and provisioning.

DirX Access

DirX Access is a comprehensive, cloud-ready, scalable, and highly available access management solution providing policy- and risk-based authentication, authorization based on XACML and federation for Web applications and services. DirX Access delivers single sign-on, versatile authentication including FIDO, identity federation based on SAML, OAuth and OpenID Connect, just-in-time provisioning, entitlement management and policy enforcement for applications and services in the cloud or on-premises.

DirX Audit

DirX Audit provides auditors, security compliance officers and audit administrators with analytical insight and transparency for identity and access. Based on historical identity data and recorded events from the identity and access management processes, DirX Audit allows answering the “what, when, where, who and why” questions of user access and entitlements. DirX Audit features historical views and reports on identity data, a graphical dashboard with drill-down into individual events, an analysis view for filtering, evaluating, correlating, and reviewing of identity-related events and job management for report generation.

For more information: support.dirx.solutions/about

Legal remarks

On the account of certain regional limitations of sales rights and service availability, we cannot guarantee that all products included in this document are available through the Eviden sales organization worldwide. Availability and packaging may vary by country and is subject to change without prior notice. Some/All of the features and products described herein may not be available locally. The information in this document contains general technical descriptions of specifications and options as well as standard and optional features which do not always have to be present in individual cases. Eviden reserves the right to modify the design, packaging, specifications and options described herein without prior notice. Please contact your local Eviden sales representative for the most current information. Note: Any technical data contained in this document may vary within defined tolerances. Original images always lose a certain amount of detail when reproduced.