

Data Management Framework

Microsoft Dynamics® AX

Intelligent Data Management Framework for
Microsoft Dynamics AX

Administration Guide

White paper

This document describes the administration and use of the
Intelligent Data Management Framework for
Microsoft Dynamics AX (Data Management Framework).

Date: October 2010

<http://www.microsoft.com/dynamics/ax/>



Table of Contents

- Introduction.....5**
- Audience6**
- Prerequisites.....6**
- System Architecture.....6**
- Deployment scenarios.....8**
 - Single-server deployment 8
 - Multi-server deployment 9
 - Distributed deployment 9
- Before you begin.....10**
- Using the Data Management Framework.....12**
 - Menu bar12
 - 1. Analysis menu12
 - 2. Configure menu12
 - 3. Schedule menu12
 - 4. Status menu12
 - 5. Administer menu12
 - 6. About menu.....12
 - Toolbar.....12
 - Status bar.....12
- Analysis13**
 - Navigation13
 - Analysis dashboard (Production database group)14
 - Navigation of the Analysis dashboard workspace15
 - Analysis details (Production database group).....16
 - Navigation of the Analysis details workspace17
 - Manage indexes.....19
 - Analysis Dashboard (Archive database group)20
 - Analysis Details for (Archive database group)20
 - Show system health20
 - Navigation of the Application health check workspace21
 - Performance.....22
 - Navigation of the Performance dashboard workspace22
 - Export to Excel27
- Configure28**
 - Overview of the purge and archive functionality.....28
 - Introduction to the discovery process32
 - Navigation33
 - Create Purge Objects.....34
 - Considerations for Purge Object, driver table, relation, and rules34

Create a new Purge Object.....	36
Navigation of the Create Purge Object workspace	39
Walkthrough: Create a Purge Object	40
Walkthrough: Disable a relation.....	42
Create Archive Object.....	43
Overview of the data archival process	43
Considerations for Archive Object, driver table, relation, and rules.....	43
Overview of the Related Archive Objects.....	45
Understanding the Archive Object exception list	47
Navigation of the Archive Object exception tables window.....	48
Create a new Archive Object	49
Navigation of the Create Archive Object workspace.....	52
Walkthrough: Work with an Archive Object	55
Purge templates/Purge Object	56
Walkthrough: Working with a purge template.....	56
Archive templates/Archive Object	57
Walkthrough: Working with an archive template.....	57
Add/Edit rules	58
Navigation of the Add/Edit rules window	58
Walkthrough: Add or edit a rule in a Purge Object	59
Walkthrough: Add or edit a rule in an Archive Object	60
Add relations.....	61
Navigation of the Add relations window	61
Walkthrough: Add a relation in a Purge Object	63
Walkthrough: Add a relation in an Archive Object	64
Export	64
Import.....	65
Save	65
Show versions	66
Validate all.....	67
Navigation of the Validate all templates and Objects window.....	67
Schedule	69
Navigation	70
Navigation of the Scheduled tasks window	71
Production database.....	75
Archive database	76
Ledger periods.....	76
System health check	76
Metadata	77
Master data.....	77
Purge	78
Archive	78
Restart of the failed or aborted archive schedule	79
Revert the failed or aborted archive schedule	80
Restore archive	80
Walkthrough: Create a restore archive schedule for the Archive Object	80
Restart the failed or aborted restore archive schedule	82
Export to Excel	82

Introduction

The Intelligent Data Management Framework for Microsoft Dynamics® AX (the Data Management Framework) enables system administrators to optimize the performance of Microsoft Dynamics AX installations. The Data Management Framework assesses the health of the Microsoft Dynamics AX application, analyzes current usage patterns, and assists in reducing database size.

The Data Management Framework allows you to analyze the database and maintain an optimal database size. To maintain the database size, the Data Management Framework provides the purge and archive functions. The purge function removes or deletes data from a set of related entities (tables) from the production database. The archive function moves data from all related tables from the production database to a standby database called archive database. Users can use the archive database for reporting purposes but cannot update it. The Data Management Framework and this document use the term *offlining* interchangeably with the term *archiving* and *purge* or *recycling* with *deleting*.

Both the purge and archive operations depend upon a carefully determined hierarchical relationship tree of related tables based on the Microsoft Dynamics AX metadata. To create a hierarchical relationship, you select a parent table and discover all related tables based on the Microsoft Dynamics AX metadata. The parent table in the relationship is placed at the root level of the tree. The discovery process creates a nested relationship tree from a parent entity to a child entity until there are no relationships left at the lowest level. The relationship tree that is created based on the driver table is called Purge Object or Archive Object.

A Purge Object is used to remove selected records from all tables in the relationship tree from Microsoft Dynamics AX database. Similarly, an Archive Object is used to move selected records from all tables in the relationship tree from Microsoft Dynamics AX database to the archive database. After creating a Purge Object or an Archive Object, you can apply business rules and selection criteria to entities and transactions to determine which records should be deleted or moved from the production database. Finally, you create a schedule to purge or archive the database. At run time, a purge schedule will select matching records from all the tables in the Purge Object and delete them from the production database. Similarly, an archive schedule will select matching records from all the tables in the Archive Object and move them from the production database to the archive database. With careful planning and system testing, the purge and archive schedules will delete or move records from the production database while maintaining the consistency and integrity of your database.

Note: The Data Management Framework provides archive templates for only versions 4.0 and 2009 of Microsoft Dynamics AX. In Microsoft Dynamics AX 3.0, you must create your own Archive Objects by using the discovery process to use the archive functionality.

Audience

This white paper is designed for database and system administrators who are responsible for administration and maintenance of the Microsoft Dynamics AX application and Microsoft® SQL Server® database.

Prerequisites

To benefit from this white paper, you must have knowledge in the following areas:

- Microsoft Dynamics AX application and system administration.
- Microsoft SQL Server database administration, backup, recovery, and performance tuning.
- Windows Server® administration, backup, recovery, and performance tuning.

Warning: To use this white paper, you must have experience in maintenance and administration of the Microsoft Dynamics AX application and database. The Data Management Framework allows you to create or modify Purge Objects and Archive Objects. A Purge Object or an Archive Object defines a hierarchical relationship tree among the Microsoft Dynamics AX application tables. You can then apply rules to select records based on specific criteria. A purge schedule selects records that match the criteria in the Purge Object and deletes them from all the tables in the relationship tree. An archive schedule selects and moves records from the production database to the archive database that matches the criteria in the Archive Object.

System Architecture

This section provides a high-level overview of the system architecture for the Data Management Framework.

The Data Management Framework was created using the Microsoft .NET development environment and provides a single document interface (SDI). The Data Management Framework uses a database, called the management database, for storage and retrieval of data, and communicates with the Application Object Server (AOS) through the COM Business Connector or the .NET Business Connector. The AOS processes all the business logic and database queries to access the Microsoft Dynamics AX application database. The Data Management Framework uses the COM Business Connector for Microsoft Dynamics AX version 3.0 and the .NET Business Connector for versions 4.0 and 2009. You must install the Business Connector on the computer where you install the Data Management Framework.

During the installation, the Data Management Framework installs a Windows service called Intelligent Data Management Framework for Microsoft Dynamics AX service. This service is used to run scheduled jobs and is referred to as the scheduler service for the Data Management Framework.

In the post-installation stage, the Data Management Framework uses the Microsoft Dynamics AX client to import and compile two X++ projects (XPOs). One XPO is used to create the application entities that are required by Data Management Framework, such as classes, tables, and a job. The Microsoft Dynamics AX Windows client is used to import and synchronize the metadata from Microsoft Dynamics AX with the Data Management Framework.

The Microsoft Dynamics AX documentation is available on the [TechNet](#).

The following diagram provides a high-level overview of the Data Management Framework system architecture.

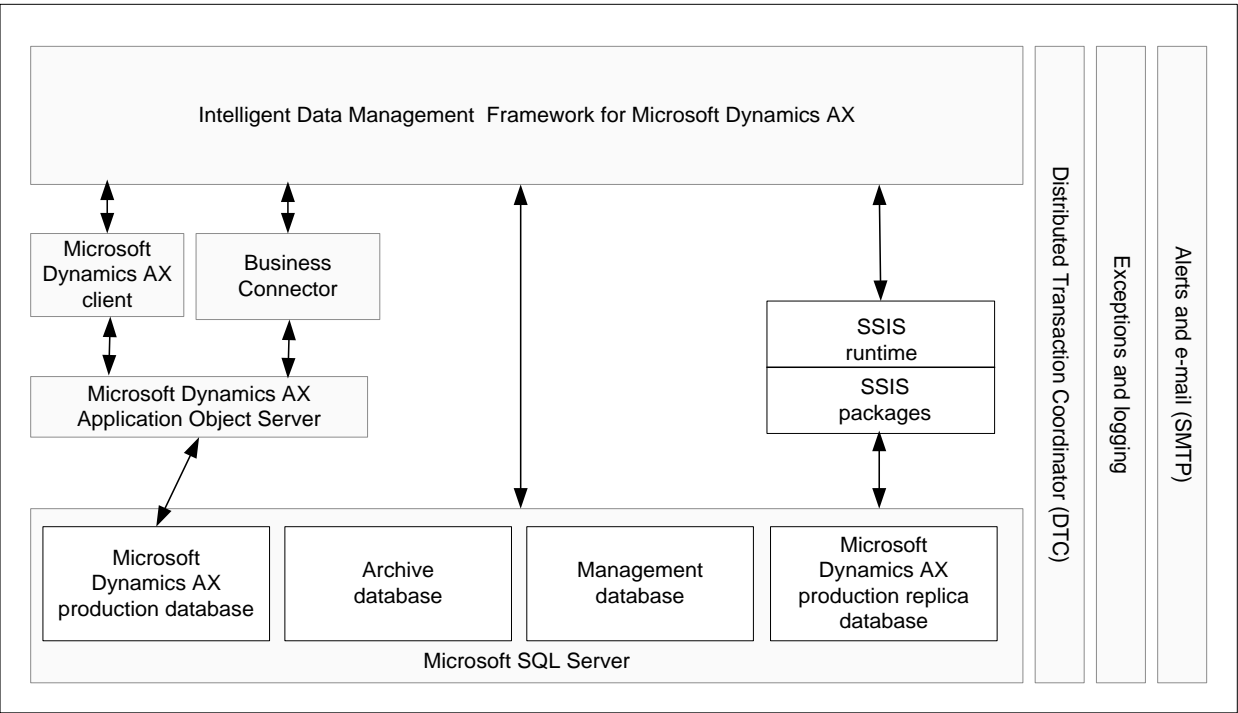


Figure 1. System architecture of the Data Management Framework

Note: The system architecture diagram shows a single SQL Server computer with different databases for the ease of illustration. You can deploy these databases on a single database server or on multiple database servers.

Deployment scenarios

This section describes the deployment scenarios for the Data Management Framework. As shown in the following sections, the system topology can range from a single-server to a multi-server deployment.

Single-server deployment

The following diagram shows the deployment of all components on a single server. This topology is not recommended for the production environment.

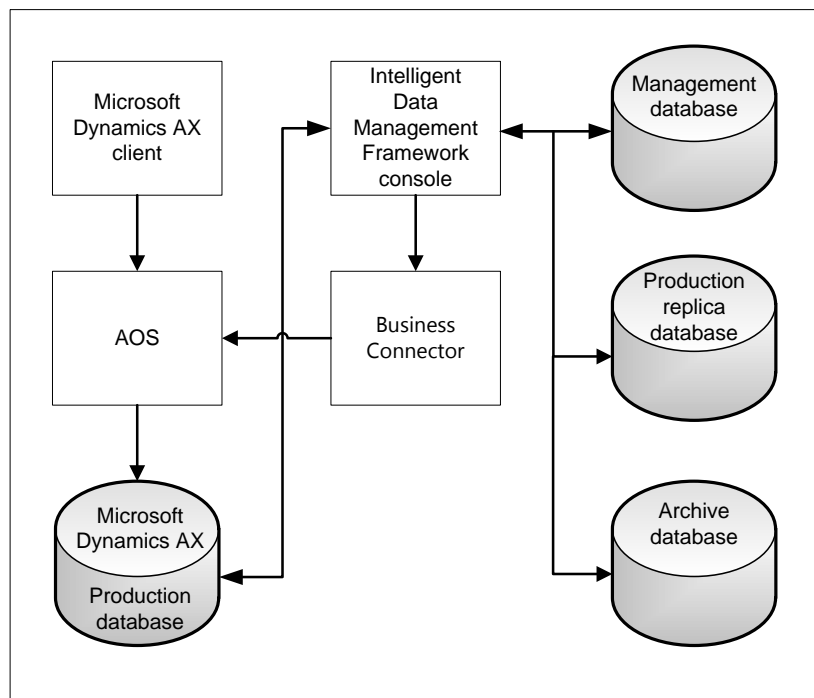


Figure 2. Single-server deployment

Multi-server deployment

The following diagram shows the multi-server deployment where the Application Object Server (AOS), the database server, and the Data Management Framework are deployed on dedicated servers.

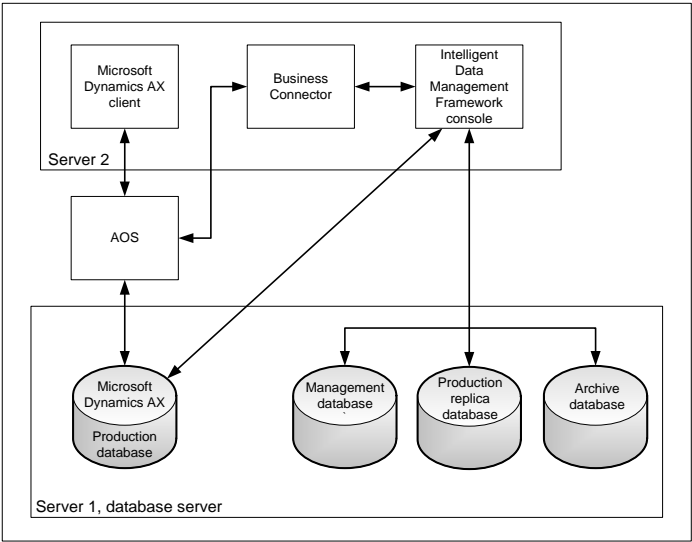


Figure 3. Multi-server deployment

Distributed deployment

The following diagram shows a distributed deployment that extends the multi-server topology by placing each database on a dedicated database server. Optionally, you can combine the management database and the production replica database on a single server.

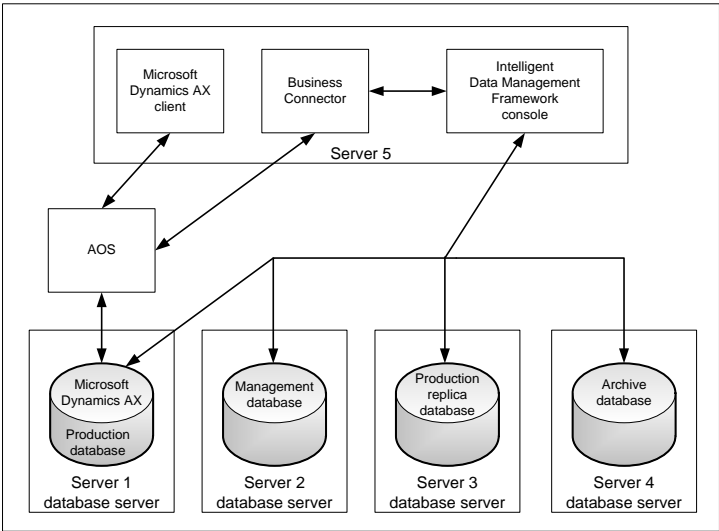


Figure 4. Distributed deployment

Before you begin

When you launch the Data Management Framework for the first time, you will be prompted to create a baseline snapshot of the database and the application. The baseline process takes a snapshot of the production database and determines the health of the Microsoft Dynamics AX application. You must create a baseline before you can use the Data Management Framework. For instructions to create the baseline schedules, see the [Data Management Framework Installation Guide](#).

Consider the following points for before working with the Data Management Framework:

- The archive and purge operations are resource intensive and will impact hardware resources such as memory, processor, disk drive (or storage), and network. Schedule these operations during the scheduled maintenance when users are offline.
- Work with the purge and archive schedules in the test environment to understand their impact on the hardware and software resources and the amount of time required for successful completion. Use the results of the tests performed in the test environment to determine your maintenance window in the production environment.
- Verify that the initial database size for the production and archive databases provide sufficient room for future growth with a large volume of data. Ensure that database files in the production and archive databases have sufficient room complete the purge and archive operations without resorting to the automatic growth. The automatic growth option should be considered a safety mechanism that allows the database files to grow when absolutely necessary and thus prevent errors. For more information about the database configuration checklist for SQL Server, see [Planning database configuration for Microsoft Dynamics AX](#).
- Carefully plan archive schedules for tables with the data types of **blob** and **memo** due to their large storage requirements.
- For each archive and purge schedule:
 - a. Estimate the number of records to be purged or archived.
 - b. Calculate the size of impacted records in megabytes.
 - c. For a purge operation, verify that the available space in the production database is at least twice as much as the size you calculated in the preceding step. For an archive operation, verify that the available space in the production and archive databases is at least twice as much as the size you calculated in the preceding step.
 - d. Determine the time it will take to successfully complete the purge or archive schedule. Run the schedule in a maintenance window to minimize the impact on online users.
- A purge operation makes database changes in the production database. An archive operation makes database changes in the production and archive database. Select the full recovery model for the production database. Carefully determine the recovery model for the archive database. The simple recovery model provides better performance for the archive operation but risks significant work-loss exposure if the database is damaged. Data is recoverable only to the most recent backup of the lost data. The full recovery model provides greater protection for data than the simple recovery model. The full recovery model relies on backing up the transaction log to provide full recoverability and to prevent work loss in the broadest range of failure scenarios. However, the archive schedule will take a longer time to complete with a full recover model than a simple recovery model. For more information, see [Choosing the Recovery Model for a Database](#).

-
- Perform a complete database backup of the production and archive databases immediately before running a purge or an archive schedule. Backup the transaction logs of the production database before and after each purge and archive operations. Backup the transaction log of the archive database immediately after each archive operation when in the full recovery mode. Perform a complete database backup of the archive database immediately after each archive operation when in the simple recovery mode.
 - You can configure an Application Object Server (AOS) to connect to archive database to view archived records. However, this connection must be read-only. Any attempt to update the data in the archive database or operations such as the database synchronization from the Application Object Tree (AOT) will result in data and application corruption.

Using the Data Management Framework

The Data Management Framework workspace is the main application area. You can use the workspace to open windows and dialogs that are used by the application.

The Data Management Framework workspace is divided into the following area.

Menu bar

The menu bar appears at the top of the Data Management Framework workspace and can be used to navigate within Data Management Framework.

The menu bar consists of the following menus.

1. Analysis menu

The [Analysis](#) menu allows you to work with the analysis dashboard, performance dashboard, analysis details, and application health check analysis.

2. Configure menu

The [Configure](#) menu allows you to work with the archive templates, purge templates, Archive Objects, and Purge Objects.

3. Schedule menu

The [Schedule](#) menu allows you to schedule various jobs for the database analysis, application health check, synchronization, purge, and archive functionality.

4. Status menu

The [Status](#) menu allows you to check the status of scheduled jobs.

5. Administer menu

The [Administer](#) menu allows you to configure application settings such as database connection, alerts, e-mail parameters, threshold values, master data tables, and the exception list.

6. About menu

Provides information about the Data Management Framework application version and the license holder.

Toolbar

The standard toolbar provides commands that you can use for the selected menu.

Status bar

The status bar appears below the toolbar and provides information about the database size. This information is captured when the Data Management Framework connects with the production or archive database.

Analysis

The **Analysis** menu allows you to analyze the production and archive databases and application health statistics from the production replica database. The Data Management Framework automatically starts in the **Analysis dashboard** view for the production database.

Navigation

This section describes the toolbar commands that are available from the **Analysis** menu. These commands are explained in greater detail later in this document.

1. The [Analysis dashboard](#) command in the **Production database** group provides a graphical view of the top 10 tables by row count, data size, and index size for the production database
2. The [Analysis details](#) command in the **Production database** group provides a detailed analysis of top 50, top 100, or all tables and their indexes for the production database.
3. The [Manage indexes](#) command allows you to view all indexes, including unused indexes, and create a schedule to defragment fragmented indexes.
4. The [Analysis dashboard](#) command in the **Archive database** group provides a graphical view of the top 10 tables by row count, data size, and index size for the archive database
5. The [Analysis details](#) command in the **Archive database** group provides a detailed analysis of top 50, top 100, or all tables and their indexes for the archive database.
6. The [Show system health](#) command allows you to work with application health measures. These measures are aggregated by company and by year.
7. The [Performance dashboard](#) command provides configuration details for all databases. This command also provides the table and query statistics for the production database and archive database.
8. The [Export to Excel](#) command allows you to export selected rows from the data grid from all workplaces where this command is available.

Analysis dashboard (Production database group)

The Analysis dashboard command shows a graphical view of the database that is similar to the following screen shot. At least two database analysis schedules must be completed before the **Database growth trend** chart will be displayed.

The controls and commands available for the analysis dashboard are similar for the production database and the archive database. The command you click becomes unavailable to visually distinguish the database in use. For example, when you click **Analysis dashboard** from the **Production database** group, the command becomes unavailable, as shown in the following screen shot diagram.

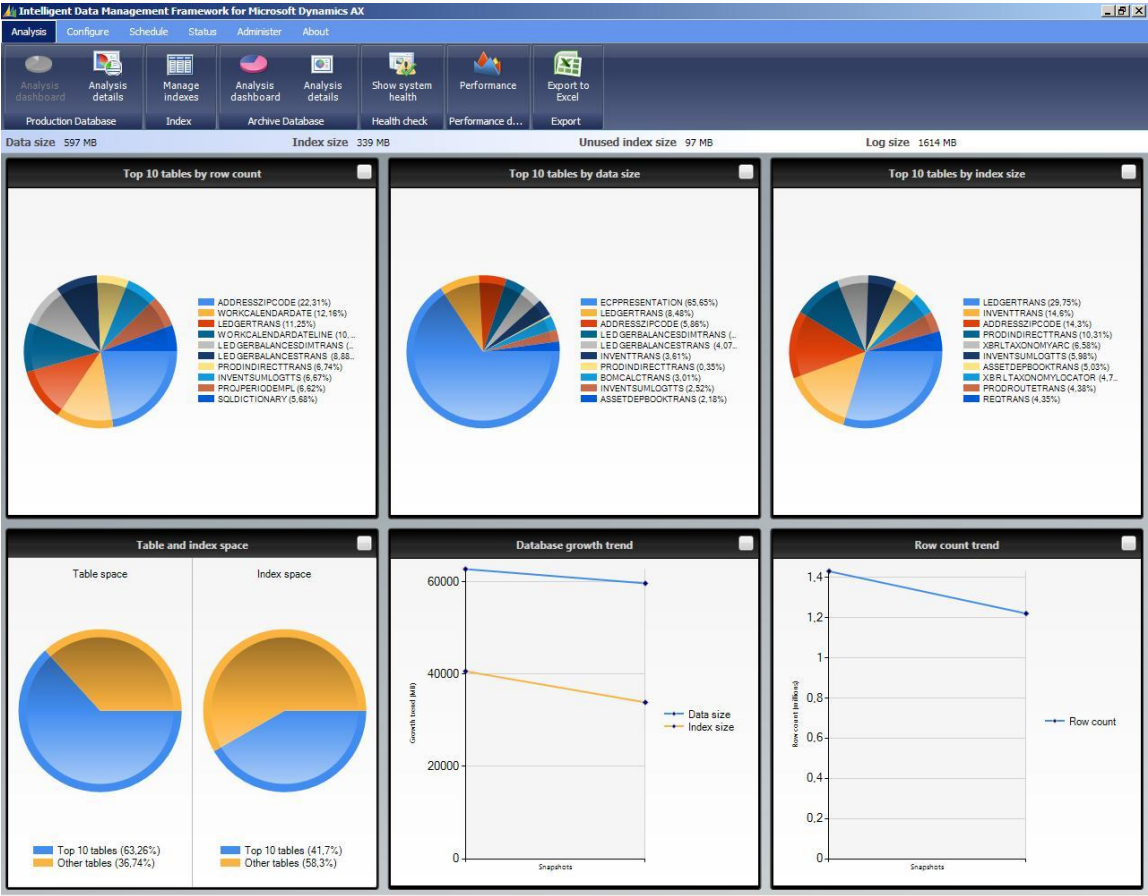


Figure 5. Analysis dashboard

Navigation of the Analysis dashboard workspace

The following table describes the panes in the **Analysis dashboard** workspace.

Pane	Description
Top 10 tables by row count	Provides a graphical view of top 10 tables by row count. The pie-chart shows the name of the table and number of rows when you hover over an area. The percentage shown for each table is calculated from the number of rows in the table versus the total rows in the top 10 tables.
Top 10 tables by data size	Provides a graphical view of top 10 tables by data size. The pie-chart shows the name and size of the table when you hover over an area. The percentage shown for each table is calculated from the size of the table versus the total size of the top 10 tables.
Top 10 tables by index size	Provides a graphical view of top 10 tables by index size. The pie-chart shows the name of the table and size of the index when you hover over an area. The percentage shown for each table is calculated from the index size of the table versus the total index size of the top 10 tables.
Table and index space	Provides two pie-charts. One chart provides a graphical view of table space used by top 10 tables versus other tables. The other chart provides a graphical view of the index space used by top 10 tables versus other tables.
Database growth trend	Provides a trend analysis of your production or archive database over multiple snapshots. You must have a minimum of two snapshots to see the trend analysis. The trend analysis processes the 10 most recent snapshots. Snapshots that are older than the most recent 10 snapshots are ignored by the trend analysis feature.
Row count trend	Provides the row count trend analysis of your production or archive database over multiple snapshots. You must have a minimum of two snapshots to see the trend analysis. The trend analysis processes the 10 most recent snapshots. Snapshots that are older than the most recent 10 snapshots are ignored by the trend analysis feature.

Analysis details (Production database group)

The **Analysis details** command provides detailed database analysis information that is similar to the following screen shot. At least two database analysis schedules must be completed before the table growth trend chart and the **Growth trend** button will be displayed. Use the **Analysis details** workspace to work with database and index analysis for the production or archive database, performance dashboard for the selected database, or manage indexes for the production database.

The controls and commands available for the analysis details are similar for the production database and the archive database. The command you click becomes unavailable to visually distinguish the database in use. For example, when you click **Analysis details** from the **Production database** group, the command becomes unavailable, as shown in the following screen shot diagram.

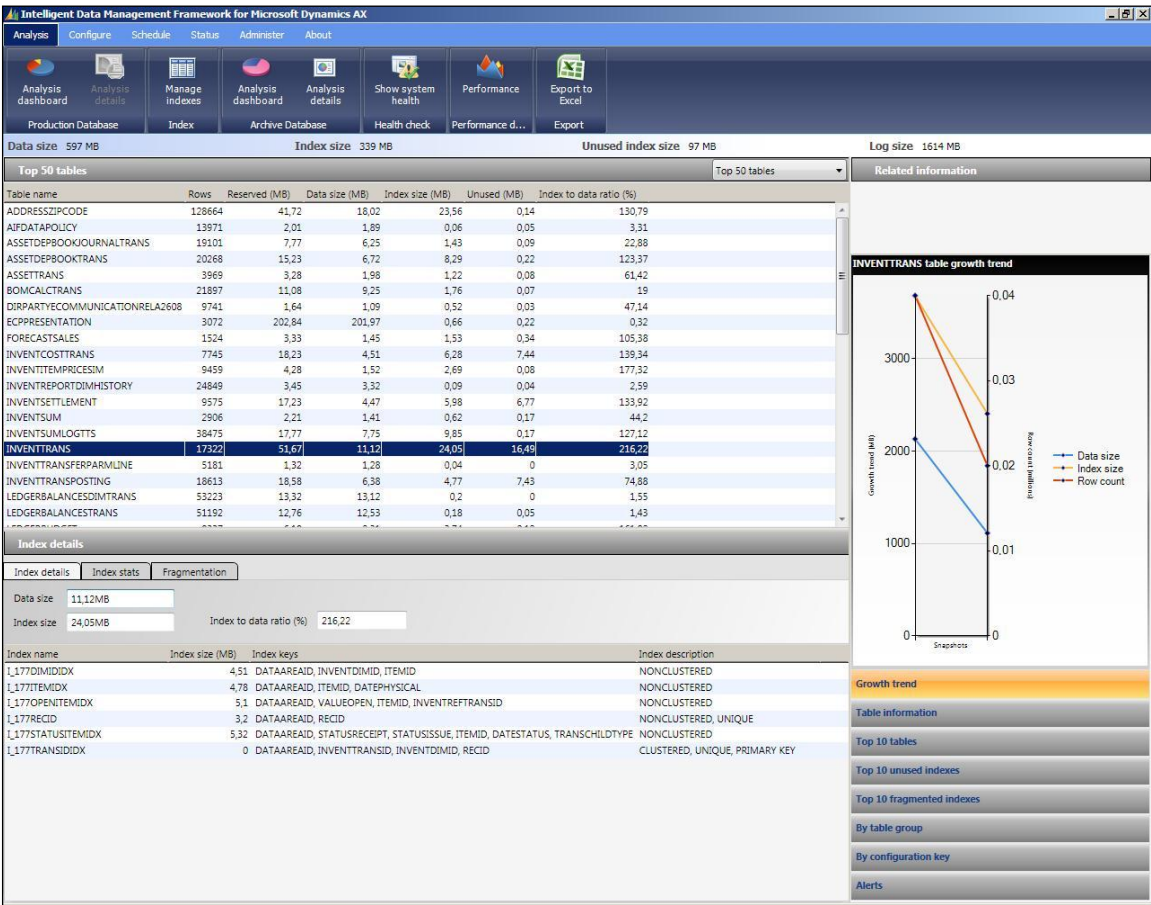


Figure 6. Analysis details

Navigation of the Analysis details workspace

The following tables provide descriptions for the controls in this workspace.

Panes

Pane	Description
Upper-left pane (defaults to Top 50 tables)	Displays sizing statistics for each table. The Data Management Framework displays the Top 50 tables by default. You can select Top 100 tables or All tables by using the filter on the title bar of the pane.
Index details	Displays index statistics for the selected table and allows you to defragment selected fragmented indexes.
Related information	Provides additional information about some tables. If the selected table has any additional information, it will appear in this pane. Read the related information carefully as it provides critical insights about the selected table.
Graphical view	Appears below the Related information pane and defaults to the Table information view. Change the graphical view by selecting a button that appears below the Graphical view pane.

Tabs (in the Index details pane)

Tab	Description
Index details	Provides information about the data size, index size, and index information for indexes of the selected table.
Index stats	Provides information about the index statistics such as the index reads, index writes, and system fragmentation percentage.
Fragmentation	Provides information for index fragmentation. Select one or more rows in the grid and click Defragment index to create a schedule to defragment selected indexes. Creating and managing schedules is covered later in this document.

Buttons

Button	Description
Defragment index (Fragmentation tab)	Creates a schedule to defragment selected indexes.
Growth trend	Provides a graphical view of the growth trend of the selected table. This button appears only after successful completion of two or more database analysis snapshot schedules.
Table information	Provides a graphical view of record count by company for the selected table.
Top 10 tables	Provides a graphical view of top 10 tables by data size in megabytes.
Top 10 unused indexes	Provides a graphical view of top 10 unused indexes by index size.
Top 10 fragmented indexes	Provides a graphical view of top 10 fragmented indexes.
By table group	Provides a graphical view of data size by table group.
By configuration key	Provides a graphical view of data size by configuration key.

Button	Description
Alerts	Provides a list of alerts. Additional details for the selected alert, if any, appear in the Related information pane.

Fields (Upper-left pane)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

Field	Description
Table name	Name of the table
Rows	Number of rows in the table
Reserved (MB)	The size that is reserved for the table in the database, in megabytes
Data size (MB)	The total data size of the table in megabytes
Index size (MB)	The total index size for the table in megabytes
Unused (MB)	The unused database space in this table in megabytes
Index to data ratio (%)	The index-to-data ratio for this table expressed as a percentage

Fields (Index details pane, Index details tab)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

Field	Description
Index name	Name of the index
Index size (MB)	The index size
Index keys	The columns that are used to create the index
Index description	Index properties such as clustered, non-clustered, unique, or primary key

Fields (Index details pane, Index stats tab)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation. The data grid displays an average of all snapshots by default (**All** in the **Stats time** column). Expand each index to see the statistics for each snapshot.

Field	Description
Index name	Name of the index. Expand this node to see statistics for each snapshot.
Stats time	The timestamp for each snapshot or an average of all snapshots (shown as All). The data grid defaults to All .
Index reads	Number of index reads
Index writes	Number of index writes
Avg fragmentation %	Index fragmentation percentage
Index description	Index properties such as clustered, non-clustered, unique, or primary key
Index keys	The fields that are used to create the index
User seeks	Number of user seeks

Field	Description
User scans	Number of user scans
User lookups	Number of user lookups
System seeks	Number of system seeks
System scans	Number of system scans
System lookups	Number of system lookups
Index size (MB)	Index size in megabytes

Fields (Index details pane, Fragmentation tab)

Field	Description
Index name	Name of the index
Average fragmentation (%)	The average of fragmentation percent for the index across all database snapshots. Select one or more indexes and click Defragment index to create a schedule to defragment selected indexes.

Manage indexes

This command allows you to create a schedule to defragment fragmented indexes. Use the following steps to create a schedule to defragment fragmented indexes:

1. Click **Analysis > Manage indexes** to open **Manage indexes** window. This window contains the **Unused indexes** and the **Fragmented indexes** tabs. Both tabs provide a data grid containing **Table** name and **Index name**. The **Fragmented indexes** tabs also provides fragmentation percentages for indexes. Use this window to create a schedule to defragment fragmented indexes.
2. In the **Manage indexes** window, select the **Fragmented indexes** tab. This tab provides a list of all the tables and their fragmented indexes. Use the **Advanced filter** control to search using table name or fragmentation percentage. Click column headings in the data grid to sort the column by ascending or descending sequence.
3. Select the table or index of interest. You can select multiple indexes from multiple tables. You must select at least one index to create a schedule. Click **Schedule**.
4. In the **Task details** pane, enter the required information and click **Save**.
5. Confirm that the schedule appears in the **Scheduled tasks** pane.

Note: Some indexes will retain a non-zero fragmentation percentage even after the defragmentation schedule completes successfully. This is a normal process. For more information, see [Troubleshoot administration of the Data Management Framework](#).

Analysis Dashboard (Archive database group)

This command provides similar information for the archive database as the **Analysis dashboard (Production database group)** provides for the production database. For detailed instructions, see [Analysis dashboard \(Production database group\)](#).

Analysis Details for (Archive database group)

This command provides similar information for the archive database as the Analysis details (Production database group) provides for the production database. For detailed instructions, see [Analysis details \(Production database group\)](#).

Show system health

This command provides **graphical views and details** for key measures from the Microsoft Dynamics AX application based on pre-determined queries. You can also create your own queries using the **Administer > Application health check** command.

A measure captures aggregated statistics for key business processes for each company in the Microsoft Dynamics AX application. This information is captured across calendar years based on the ledger periods you set in the application. For example, the measure “Number of inactive Sales Quotations” is calculated as “the total number of sales quotations which are cancelled, confirmed, or lost”. The total is grouped by company and by year.

The application health check provides key measures for the Inventory, Accounts receivable, Accounts payable, General ledger, and Administration modules as shown in the following screen shot.

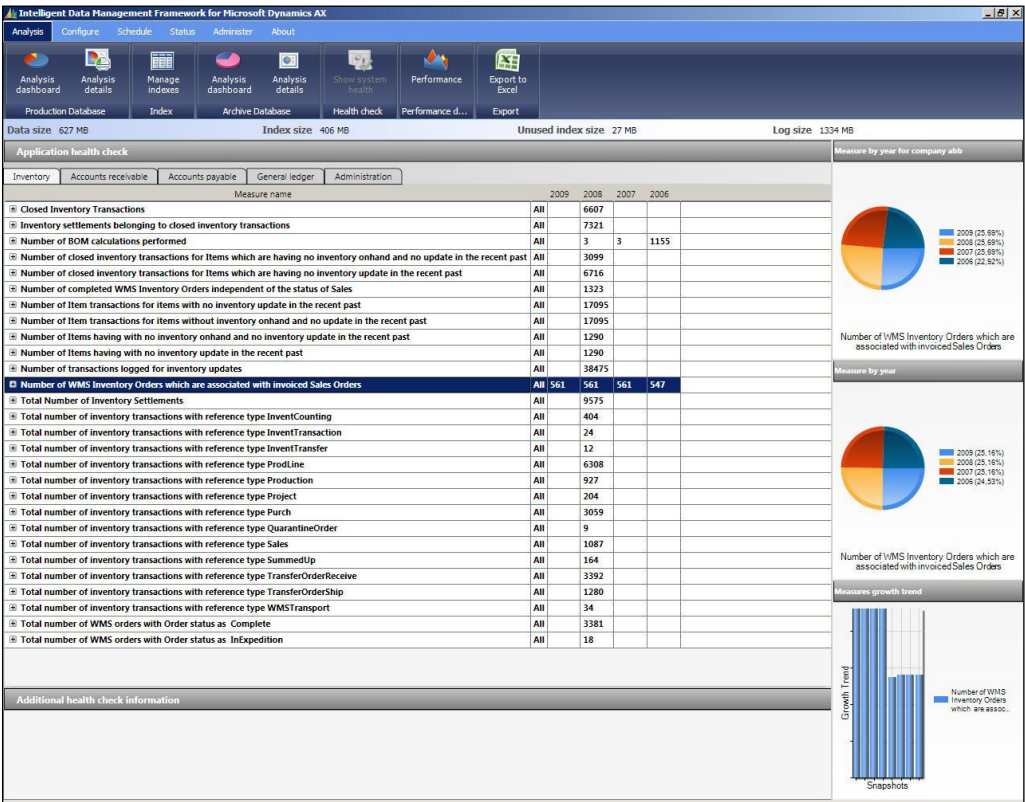


Figure 7. Show system health

Navigation of the Application health check workspace

The following tables provide descriptions for the controls in this workspace.

Panes

Pane	Description
Application health check	Provides aggregates of measures by company and by year for selected modules from the Microsoft Dynamics AX application. The default view is the aggregate of all companies by year. Expand the measure to select the measure statistics for individual companies.
Additional health check information	Provides additional information for the selected measure including a brief description, archival relevance, and the formula that was used to calculate the measure.
Measure by year for company <company name>	Shows a graphical view of the selected measure for the first company in the data grid. If there are multiple companies, expand the measure and select another company to change the graphical view. When you hover over an area, the chart will show the year and value for the selected measure.
Measure by year	Shows a graphical view of aggregate value for all companies for the selected measure by year. The chart will show the year and value for the selected measure when you hover over an area.
Measures growth trend	Shows the aggregate trend analysis from different snapshots. You need to complete at least two health check analysis schedules to see a graphical view in this pane. The trend analysis captures the 10 most recent snapshots and ignores older snapshots, if any.

Tabs (Application health check pane)

Tab	Description
Inventory	Provides selected measures from the Inventory management module.
Accounts receivable	Provides selected measures from the Accounts receivable module.
Accounts payable	Provides selected measures from the Accounts payable module.
General ledger	Provides selected measures from the General ledger module.
Administration	Provides selected measures from the Administration module.

Fields (Application health check pane)

Field	Description
Measure name	Name of the measure
Company	Company accounts in Microsoft Dynamics AX represent the organizational structure of a company. An aggregate of the measure for all the companies by year is listed as "All".
Year	The aggregate of the measure is placed in a column for each year based on the ledger periods configuration.

Performance

This command provides information about SQL Server configuration options, database options, and statistics for tables and indexes for production and archive database. This command displays the database configuration details for all databases used by the Data Management Framework and table statistics and query statistics for the production and archive databases.

Navigation of the Performance dashboard workspace

The following tables provide descriptions for the controls in this workspace.

Tabs (in the Index details pane)

Tab	Description
Database configuration	Provides information about the database properties and configuration options for the selected database such as information about the data and log files, database options, and other database properties. This information is available for the selected database.
Table stats	<p>This tab is only available when you select the production or archive database from the list in the Database Configuration tab.</p> <p>Provides information about table statistics for the selected database such as tables without clustered index, tables with missing indexes, tables with locks, and tables that can be enabled for caching in the Application Object Server (AOS).</p>
Query stats	<p>This tab is only available when you select the production or archive database from the list in the Database Configuration tab.</p> <p>Provides information about the query statistics such as the cached query plans, query statement and query plan for the selected cached query plan.</p>

Panes (Database configuration)

Pane	Description
Database properties	Provides database properties for the selected database such as the database name, database owner, number of active users, information about database and log files, and currently configured values for the database options.
Advanced SQL Server configuration options	Provides the minimum, maximum, configured, and currently used (run) value of advanced configuration options for SQL Server.

Panes (Table stats)

Pane	Description
Cacheable tables	You can enable these tables so that the Application Object Server (AOS) caches the entire table. To enable a table for caching, change the CacheLookup property of the table to EntireTable . For more information, refer to the Microsoft Dynamics AX developer documentation on the MSDN Library .
Tables without clustered index	Provides a list of tables without a clustered index.
Missing indexes	Provides a list of tables that do not have any index.
Tables with locks	Provides a list of locked tables with relevant details such as row or page lock count and row or page lock wait time.

Panes (Query stats)

Pane	Description
Query plan statistics	Provides statistics for the cached query plans such as usage count, logical and physical reads and writes.
Query text	This is the Transact-SQL query statement that was used to create the cached query plan. Click Export to export the query text as the SQL text file.
Query plan	This is the query plan for the query text that was used to create the selected cached query plan. Click Export to export the query plan as XML file.

Fields (Database configuration tab)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

Field	Description
Select database	Select a database from the list. The performance dashboard displays relevant information for the selected database.
Property name	Name of the database property such as database name, owner, status, creation date, collation, last backup date, and number of users.
Value	Value of the database property. For example, the property value owner displays the current owner of the database.
Name (Files list)	The name of the data or log file
File name	The filename of the data or log file
File group	The name of the file group that is used for the data file
Size	The size of the data or log file in kilobytes (KB)
Growth	The growth size for the data or log file in kilobytes (KB) or percentage
Usage	Whether the file is used for data or transaction log
Property name (Options list)	Name of the database option such as database name, compatibility level, recovery model, and auto close.
Value	Value of the database option.

Field	Description
Name (Advance SQL Server configuration options pane)	Name of the SQL Server configuration option.
Minimum	Minimum value of the configuration option.
Maximum	Maximum value of the configuration option.
Configuration value	Value to which the configuration option was set using sp_configure (value in sys.configurations.value). For more information about these options, see Setting Server Configuration Options and sys.configurations (Transact-SQL) .
Run value	Currently running value of the configuration option (value in sys.configurations.value_in_use). For more information, see sys.configurations (Transact-SQL) .

Fields (Table stats tab, Cacheable tables list)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

Field	Description
Table name	The table name
Row count	Number of rows in the table
User seeks	Number of user seeks
User scans	Number of user scans
User lookups	Number of user lookups
User updates	Number of user updates

Fields (Table stats tab, Tables without clustered index list)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

Field	Description
Table index	Index name
Row count	Number of rows in the table
Total index reads	The number of times this index was used for read operations

Fields (Table stats tab, Missing indexes list)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

Field	Description
Table name	Name of the table.
Equality columns	Comma-separated list of columns that contribute to equality predicates of the form: <i>table.column</i> = <i>constant_value</i> .
Inequality columns	Comma-separated list of columns that contribute to inequality predicates, for example, predicates of the form: <i>table.column</i> > <i>constant_value</i> .

Field	Description
Included columns	Comma-separated list of columns needed as covering columns for the query.
Avg total user cost	Average cost of the user queries that could be reduced by the index in the group.
Avg user impact	Average percentage benefit that user queries could experience if this missing index group was implemented. The value means that the query cost would on average drop by this percentage if this missing index group was implemented.
Last user seek	Date and time of last seek caused by user queries that the recommended index in the group could have been used for.
Unique compiles	Number of compilations and recompilations that would benefit from this missing index group. Compilations and recompilations of many different queries can contribute to this column value.

Fields (Tables with locks list)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

Field	Description
Table name	Name of the table.
Index name	Name of the index.
Index description	Index properties such as clustered, nonclustered, unique, primary key, and heap
Row lock waits (milliseconds)	Total number of milliseconds the database engine waited on a row lock.
Page lock waits (milliseconds)	Total number of milliseconds the database engine waited on a page lock.
Row lock wait count	Cumulative number of times (milliseconds) the database engine waited on a row lock.
Row lock count	Cumulative number of row locks requested.
Page lock count	Cumulative number of page locks requested.
Page lock wait count	Cumulative number of times (milliseconds) the Database Engine waited on a page lock.
Index lock promotion count	Cumulative number of times the database engine escalated locks.

Fields (Query stats tab, all panes)

For a detailed explanation of SQL Server terms, refer to the SQL Server documentation.

The **Select a snapshot time** list allows you to select query plan statistics, query text, and query plan for a specific analysis snapshot. The list defaults to the most recent time. The data grid filters the information based on your selection. The following table describes the fields in the data grid.

Field	Description
Usage count	Number of time the plan has been run since it was last compiled
Total elapsed time	Total elapsed time, in microseconds, for completed processing of this plan.
Total worker time	Total amount of CPU time, in microseconds, that was consumed by processing of this plan since it was compiled.
Avg elapsed time	Average amount of CPU time, in microseconds, based on Usage count and Total worker time .
Avg physical reads	Average number of physical reads performed by processing of this plan since it was compiled.
Avg logical reads	Average number of logical reads performed by processing of this plan since it was compiled.
Last elapsed time	Elapsed time, in microseconds, for the most recently completed processing of this plan.
Min elapsed time	Minimum elapsed time, in microseconds, for any completed processing of this plan.
Max elapsed time	Maximum elapsed time, in microseconds, for any completed processing of this plan.
Total physical reads	Total number of physical reads performed by processing of this plan since it was compiled.
Last physical reads	Number of physical reads performed the last time the plan was processed.
Min physical reads	Minimum number of physical reads that this plan has ever performed during a single processing.
Max physical reads	Maximum number of physical reads that this plan has ever performed during a single processing.
Total logical reads	Total number of logical reads performed by processing of this plan since it was compiled.
Last logical reads	Number of logical reads performed the last time the plan was processed.
Min logical reads	Minimum number of logical reads that this plan has ever performed during a single processing.
Max logical reads	Maximum number of logical reads that this plan has ever performed during a single processing.
Total logical writes	Total number of logical writes performed by processing of this plan since it was compiled.
Last logical writes	Number of logical writes performed the last time the plan was processed.

Field	Description
Min logical writes	Minimum number of logical writes that this plan has ever performed during a single processing.
Max logical writes	Maximum number of logical writes that this plan has ever performed during a single processing.
Last worker time	CPU time, in microseconds, that was consumed the last time the plan was processed.
Min worker time	Minimum CPU time, in microseconds, that this plan has ever consumed during a single processing.
Max worker time	Maximum CPU time, in microseconds, that this plan has ever consumed during a single processing.

Export to Excel

This command allows you to export selected rows from the data grid from all workplaces where this command is available.

Configure

The Configure menu allows you to work with the purge templates, Purge Objects, archive templates and Archive objects.

Overview of the purge and archive functionality

The Microsoft Dynamics AX application provides real-time updates for a transaction. For example, when users enter, modify, or delete orders, Microsoft Dynamics AX updates a main table and some of its related tables. Over time, the volume of data in the database will grow.

The Data Management Framework allows you to analyze the database and maintain an optimal database size. To maintain the database size, the Data Management Framework provides the purge and archive functions. The purge function removes or deletes data from a set of related entities (tables) from the production database. The archive function moves data from a set of related tables from the production database to a standby database called the archive database. Users can use the archive database for reporting but should not be allowed to update it. The Data Management Framework and this document use the term *offlining* interchangeably with the term *archiving* and *recycling* with *purge* or *delete* operation.

Both the purge and archive operations depend upon a carefully determined hierarchical relationship tree of related tables based on the Microsoft Dynamics AX metadata. For example, when a user creates updates or deletes a sales order, the SalesTable table and other tables that are related to SalesTable are updated. When you archive or purge data from a main table such as the SalesTable, you must archive or purge data from all the tables that are related to SalesTable. The archive and purge operations always work on a set of related tables. This set of related table forms a hierarchical relationship. Each hierarchical relationship starts with the main table, called a driver table, at the root level or level 0 of the hierarchical tree. All the child tables that are related to the driver table are at the next level or at level 1. All the child tables of tables in level 1 are at the next level or at level 2. The nested hierarchy or parent-child relationship continues until there are no more child tables for the lowest level of tables.

To help you understand the hierarchical relationships, the Data Management Framework includes some templates that show sample hierarchical relationships. Templates for the archive function are called archive templates. Templates for the purge function are called purge templates. Each template provides the sample hierarchy based on a standard installation without any customizations. You cannot use these templates directly. You must open these templates, review them to verify applicability in your environment, and then save them before you can use them for archive or purge function.

When you open a template, verify that the relationship hierarchy contains all the relevant tables from your implementation of Microsoft Dynamics AX. The hierarchical tree in the template may or may not match your implementation for a number of reasons. The following list provides some of the reasons for discrepancy between a template and relationships in your implementation of Microsoft Dynamics AX:

- Your implementation may not have certain tables that are in the template depending upon customization, configuration, and licensing of your implementation.
- Certain fields in the tables within the template might not be available in your implementation.
- You may have created custom tables that are not contained in the templates.

It is critical that your hierarchical relationship includes all relevant tables in your implementation. A purge or archive operation that is based on an improper hierarchy will result in data corruption. Therefore, you must create a relationship tree that is complete and accurate for your implementation. Because of the uniqueness of each implementation, you cannot use a template in the archive or purge function. Instead, you must create a Purge Object or an Archive Object that encompasses all relevant relationships for your implementation. There are two ways you can create a purge object or an Archive Object.

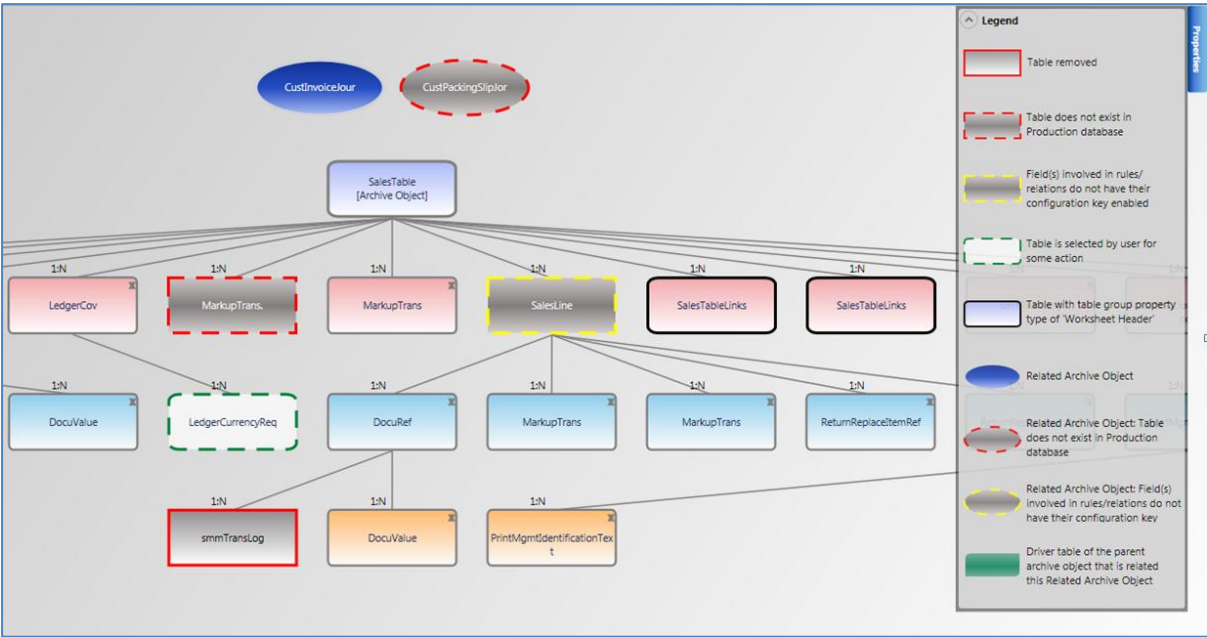
1. Work with a template. Open a template that is included with the Data Management Framework. Review the template and then Add or remove tables and rules to match your implementation. Verify that the relationship tree is complete for your implementation. Save the template. An archive template is saved as an Archive Object. A purge template is saved as a Purge Object.
2. Create an object on your own rather than using the template that is included with the Data Management Framework. Depending upon licensing, configuration, and customization, the hierarchical relationship in your implementation may be significantly different than the templates that are included with the Data Management Framework. In this case, you must use the discovery process which allows you to select a driver table that is the parent table that forms the root of the relationship tree. Based on the driver table that you select and the Microsoft Dynamics AX metadata, the Data Management Framework generates the hierarchical relationship tree. The Data Management Framework will check an exception list and will ignore tables contained in the exception list as detailed in the later sections of this document.

We recommend that you create a Purge Object or an Archive Object through the discovery process instead of opening and saving a template. Compare your object with the template, if a similar template exists, to see if your object is missing any tables that the template contains. Depending on the comparison and verification, add tables to or remove tables from your object, as appropriate.




Each Archive Object or Purge Object contains certain rules that govern the archival or purge of data. Rules enable you to determine the selection criteria to select the data that is qualified for purge or archival. You can add, remove, or modify these rules to match your archival and purge guidelines.







You must save an Archive Object or a Purge Object before you can use it in a purge or archive schedule. The Data Management Framework provides a graphical view of your Purge Object or Archive Object that is similar to the sample that is shown in the following diagram. Familiarize yourself with the legend so that you can easily understand and validate your Archive and Purge Objects.

Figure 8. A sample Archive Object with legend



The following table explains the legend that is shown in a relationship tree of an object or a template:

Legend	Text	Description
	Table removed	This table has been removed from the object you are working with. The Purge Object or Archive Object will not purge or archive data from this table.
	Table is missing from the production database	This table exists in the template or object you are working with but is not found in your production database. You must remove this table from the relationship before you can save this object.
	Fields with disabled configuration keys	<p>In Microsoft Dynamics AX, each field is governed by the security key and a configuration key. Depending upon the licensing, security keys, and the configuration you use, each field can be an enabled or disabled field. For the purpose of this document, a disabled field refers to a field that is disabled via the configuration key.</p> <p>A driver table with any rules that use a disable field or a child table with any relationship that use a disabled field will appear with a yellow border. You must resolve the rules and relationships with disabled fields before you can save the object. For more information, see Add/Edit rules and Add relations.</p>

Legend	Text	Description
	Selected table	You have selected this table. You can select multiple tables and then perform actions such as remove all occurrences of the selected tables or add selected tables to the data replicate feature.
	Table group property is worksheet header	This table has a table group property of type worksheet header. You can right click such tables and rediscover related child tables.
	Related Archive Object	This is a Related Archive Object. To avoid complexity, you can add a Related Archive Object to your Archive Object instead of adding a child table with many levels of nested relationships.
	Related Archive Object: Table is missing from the production database	This Related Archive Object does not exist in the production database. You must remove this table from the Related Archive Object and save it before you can save your current object.
	Related Archive Object: Fields with disabled configuration keys	This Related Archive Object contains relationships or rules with disabled configuration keys. You must enable the configuration keys or remove the relationships and rules based on disabled fields and save the Related Archive Object before you can use it in this object.
	Driver table of the parent Archive Object	Signifies the driver table of the Archive Object that is the parent object of this Related Archive Object.

Introduction to the discovery process

The discovery process creates a hierarchical relationship tree for a Purge Object or an Archive Object. When you create a new Purge Object or an Archive Object, you select a parent table, called the driver table. The Data Management Framework uses a process called discovery to create a hierarchical relationship tree based on the driver table you select.

The following steps provide a high-level walkthrough of the discovery process assuming you selected WMSOrder table as the driver table:

1. The discovery process starts when you select a table as the driver table that is the root parent in a parent-child hierarchy. However, the Data Management Framework will not display a table in the driver table list that is part of an exception list. WMSOrder table appears in the list as is not in the exception parameters list. For more information about the exception parameters list, see [Discovery](#).
2. The xRefTableRelation table in Microsoft Dynamics AX contains relationships between two tables, and the fields that are used by the relationship. The Data Management Framework queries the xRefTableRelation table to retrieve the names of the tables that are child tables of the parent WMSOrder. Explanation of the xRefTableRelation field is beyond the scope of this document. For more information, see [xRefTableRelation table](#).
3. Continue the preceding step for every table until there are no child tables left. In certain cases, a child table is excluded from the relationship tree, if:
 - a. A child table has the TableGroup value of Main, Parameter, or Group.
 - b. A child table is in the **Exception parameters** list with the **Purge discovery** field selected. In this case, the child table is excluded from the relationship tree of the Purge Object.
 - c. A child table is in the **Exception parameters** list with the **Archive discovery** field selected. In this case, the child table is excluded from the relationship tree of the Archive Object.

Note: Tables with the TableGroup value of WorksheetHeader are ideal candidates for driver tables and appear with a dark black border in the relationship tree when they are discovered as a child table. The discovery process shows these tables as child tables but does not show their child tables in the tree. Determine the best way to incorporate these tables in your purge and archival strategy from the following list:

1. Keep these tables in the relationship tree but do not add their child tables.
2. Keep these tables in the relationship tree and add their child tables. Right-click the table and then select **Discover** to discover, or add, the child tables that form the parent-child hierarchy for this table.
3. In case of an Archive Object, remove these tables from the relationship tree and add them as Related Archive Objects or create independent Archive Objects. For more information, see [Overview of the Related Archive Objects](#).
4. In case of a Purge Object, remove these tables and create new Purge Objects by using these tables as the driver tables.

Warning: You must understand the Microsoft Dynamics AX metadata and application before you start working with the exception parameters list, Archive Objects, or Purge Objects. You must functionally validate the Purge Object or Archive Object to maintain database and application integrity.

Navigation

This section describes the toolbar commands that are available for the **Configure** menu. These commands are explained in later sections of this document.

1. The [Create Purge Object](#) command allows you to create a new Purge Object.
2. The [Create Archive Object](#) command allows you to create a new Archive Object.
3. The [Purge templates/Purge Object](#) command allows you to work with Purge templates and save them as Purge Objects.
4. The [Archive templates /Archive Object](#) command allows you to work with Archive templates and save them as Archive Objects.
5. The [Add/Edit rules](#) command allows you to create rules. A rule allows you to define the selection criteria for a table. These rules are used to filter the records when a purge or archive schedule runs.
6. The [Add relations](#) command allows you to add new tables to the Purge Objects or Archive Objects by creating a relationship between an existing table from the selected Purge Object or Archive Object (parent) and the added table (child).
7. The [Import](#) command allows you to import a Purge Object or an Archive Object that was saved in XML format.
8. The [Export](#) command allows you to export a Purge Object or an Archive Object in XML format.
9. The [Save](#) command allows you to save the Purge Object, Purge template, Archive Object or archive template that you are working with. You must review and save a Purge template or an Archive template before you can use it in a purge schedule or an archive schedule.
10. The [Show versions](#) command provides the version history for Archive Objects.
11. The [Validate all](#) command validates all Purge Objects, purge templates, Archive Objects, and archive templates.

Create Purge Objects

This command allows you to define and create your own Purge Objects.

Warning: To use this white paper, you must have experience in maintenance and administration of the Microsoft Dynamics AX application and database. The Data Management Framework allows you to create a Purge Object that defines a hierarchical relationship tree among the Microsoft Dynamics AX application tables. You can then apply rules to select records based on specific criteria. Records matching the selection criteria are deleted from the entire relationship tree in the Purge Object when a purge schedule runs. Improper use of this framework can cause unexpected results, database corruption, and application downtime requiring full database and application recovery. Exercise extreme caution and thoroughly test your recycling strategy in a test environment before working in the production environment.

Considerations for Purge Object, driver table, relation, and rules

Consider the following points when working with a Purge Object.

- Tables that store transient or intermediate data are usually good candidates for a Purge Object. You must validate the data dependency and application functionality before creating a Purge Object with such tables.
- Verify that the table you select as the driver table for a Purge Object is a header table such as the SalesParmTable or SalesQuotationTable. A good indication, although not always accurate, is that the TableGroup value for such a table should be WorkSheetHeader, Miscellaneous, or Transaction.
- Tables with the TableGroup value of Main, Parameter, Group, and WorkSheetLine cannot be driver tables.
- The driver table becomes the root parent in the hierarchical relationship tree. The Data Management Framework queries the XRefTableRelation table to determine the parent-child relationships as detailed in the preceding section.
- Removing records from a Purge Object may result in data inconsistency in your application if:
 - The driver table has known parent tables.
 - The driver table has known child tables that are not added to the Purge Table as a relation.
- Always minimize the Purge Object and keep only necessary relations and rules for your implementations. Consider removing a table from the Purge Object if:
 - The table does not contain any data and it is not likely to contain data in the future.
 - The table stores only transient data such as CustTransOpen or SpecTrans.
 - The table stores open transactions.
 - The table creates nested relationships. If a table appears multiple times on different levels in the relationship tree with the same qualifying relationship condition or set of primary keys, then you should consider keeping the occurrence only in the lowest-numbered level in the relationship tree.

- There may be multiple relationships between two tables. In that case, evaluate the relationships and functionality carefully. We recommend that use only a single relationship in the Purge Object. However, you can use multiple relationships if necessary. Be sure to select a valid relationship using a unique index or primary key if one exists.
- A table is unrelated to the driver table and will make its own Purge Object. For example, if you use the ProdTable as a driver table, it will also discover ProdJournalTable which itself is a separate Purge Object. Another example is the PurchReqTable. If you use the PurchReqTable as a driver table, it will discover PurchParmLine table which itself is a part of the PurchParmTable Purge Object.
- Verify that the relationship tree that you have created in the Purge Object actually purges the intended data from your system. A Purge Object that you create must look at all related records. Missing relationships lead to orphaned data.
- Thoroughly test the impact of the purge on your database and application in a test environment that is similar to your production environment. Test all business processes and obtain user sign-off before attempting to implement the Purge Object in the production environment.
- When you create a Purge Object from a purge template that are included with the Data Management Framework, validate and thoroughly test the Purge Object in a test environment. These templates are created in a standard Microsoft Dynamics AX application and may or may not match your implementation. We recommend that you use the discovery process to create your own Purge Objects or Archive Objects and compare them with the templates, if a matching template is available. Carefully analyze any discrepancy between the object you created through the discovery process and the template to determine the cause of the discrepancy. Manually add relations and rules to or remove relations and rules from the Purge Objects and Archive Objects to fit your requirements.
- Be sure that the TableGroup property is set for all custom tables in Microsoft Dynamics AX, so that the discovery process can retrieve all related tables.
- Synchronize changes such as adding, deleting, or updating a table in Application Object Tree (AOT) or the data dictionary synchronization in the Microsoft Dynamics AX metadata with Data Management Framework. Run the post-installation tasks to synchronize the metadata. For information, see the **Post-installation tasks** section in the [Data Management Framework Installation Guide](#).

Warning: Verify that any tables with the table group of type **Transaction** are part of the Purge Object only if it is acceptable to purge data from these tables.

Create a new Purge Object

Use the following steps to create a new Purge Object.

1. On the toolbar, click **Create Purge Object**.
2. In the **Create Purge Object** dialog window, select a table from the **Driver table** list. (The list excludes tables from the exception list, which is covered later in this document.) Navigate to and select **PurchParmTable**. A driver table is the parent table that defines the relationship tree for a specific Purge Object. In a Purge Object, the driver table may have child entities but does not have a parent entity in the relationship tree.
3. To continue, select the most unique index for the table. This index will be used to create the parent child relationship in the Purge Object. In the **Unique keys** box, select **ParmId**. The **ParmTableRefIdx** index is created using the **ParmID** and **TableRefID** fields. Notice that by selecting **ParmID**, you are selecting a partial key. A unique index enables the discovery process to identify related tables within the same functional area such as finance. If you do not select a unique index, the discovery process may find related tables from all functional areas. You must verify and ensure that the discovery process identifies the correct functional relationships for your implementation of Microsoft Dynamics AX.
4. Click **Discover**. The Data Management Framework uses metadata that is imported from the Microsoft Dynamics AX database and the exception parameters list to generate a relationship tree. The relationship tree starts with the driver table and creates a hierarchy of parent-child relationships. Your Purge Object should look similar to Figure 9.

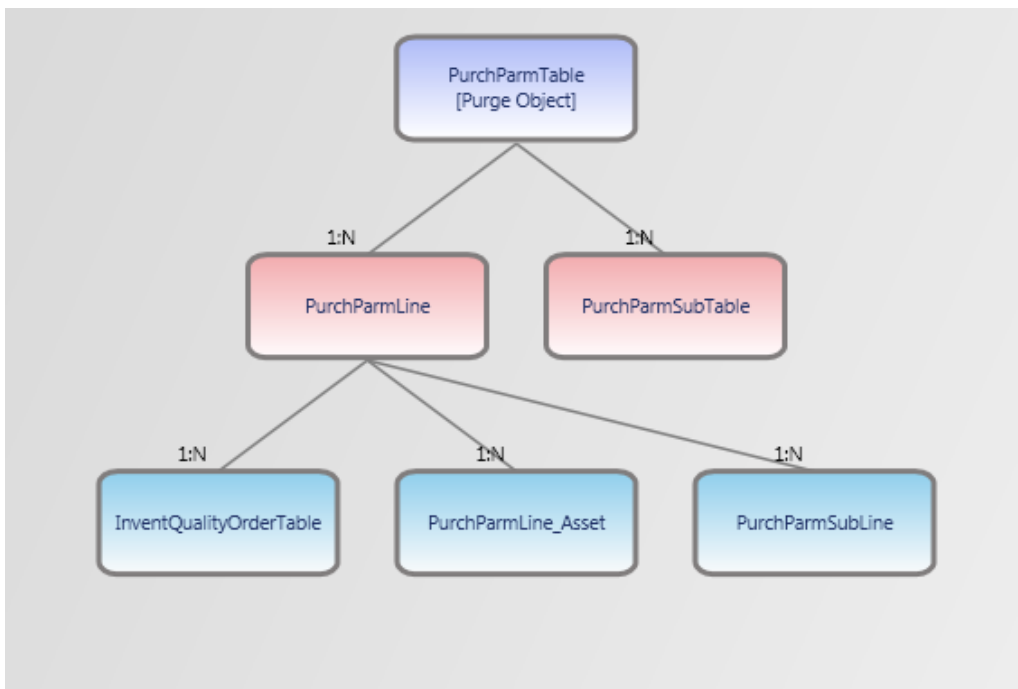


Figure 9. The Purge Object

The relationship tree in the preceding Purge Object is three levels deep. Level 0, the topmost level, contains the driver table, **PurchParmTable**. Level 1 contains the child entities of the driver table. The last level, level 2, contains the child entities of tables in Level 1.

This relationship tree is created based on:

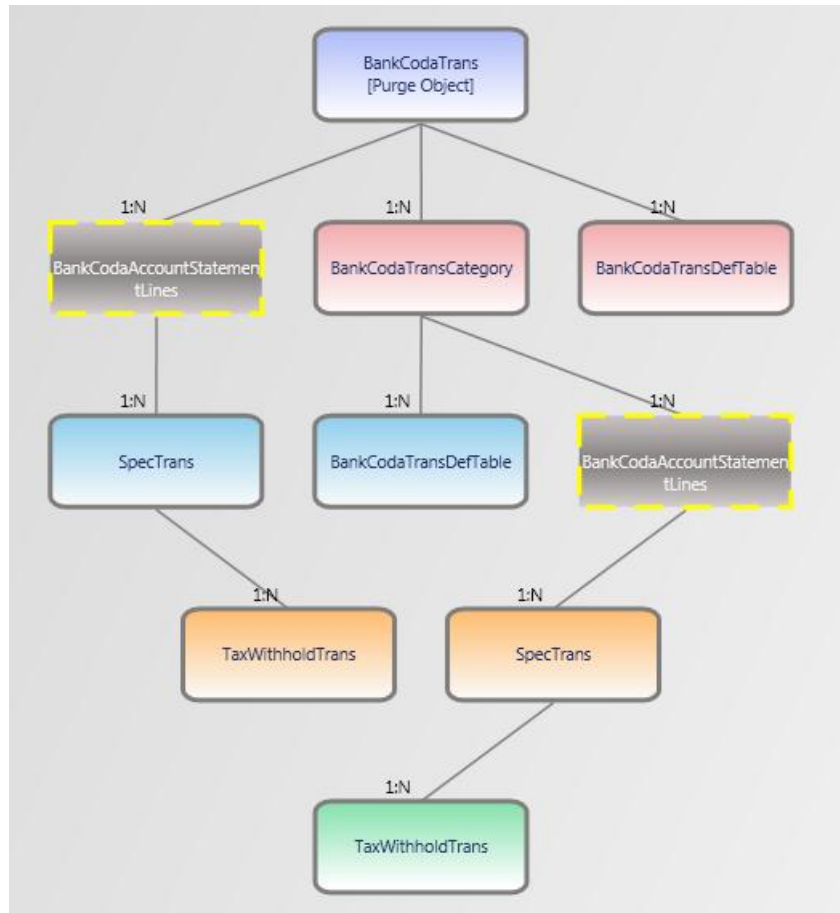
1. The Microsoft Dynamics AX metadata. The discovery process uses the metadata to create a list of related tables that form the parent-child hierarchy based on the driver table.
2. The exception parameters. Tables that belong to the exception parameters list are filtered out from the relationship tree. Use the **Administer** menu to configure the exception parameters.

Verify the relationship tree and assess the functional impact of the tree on your Microsoft Dynamics AX application. As a result of your assessment, you may choose to manually add or remove some relations and rules. For example, when you click **Restore** in the **Properties** pane, the Purge Object will be restored from the default template that ships with Data Management Framework. The restored relationship tree in this case will be different from the Purge Object you created using the PurchParmTable as the discovery table. This is because the default template in the Data Management Framework is modified based on the functionality assessment of a standard Microsoft Dynamics AX application. The Purge Object you create through discovery may or may not agree with the default template that is included with Data Management Framework, if one exists for the discovery table.

Be careful with your selection of the driver table. If you select a table such as the ProdTable table, the relationship tree can go many levels deep, spanning many tables on each level. Such a Purge Object creates a complex relationship tree and increases the potential for error.

The default template that is included with the Data Management Framework may not match with the metadata of your Microsoft Dynamics AX implementation. A table in the default template that is not found in the metadata from the production database is marked invalid. An invalid table appears with a dotted red border. You must remove all invalid tables from the Purge Object. The fields that you use to create relations and rules from the valid tables are also validated. A field with a disabled configuration key is considered invalid. A valid table with invalid fields is shown with a yellow dotted border as shown in Figure 10. You must remove any rules or relations in the Purge Object that are defined with the disabled fields. You can continue using the object after removing the invalid table or tables with invalid fields.

Figure 10 – Tables with invalid fields in the Purge Object



Navigation of the Create Purge Object workspace

The following tables provide descriptions for the controls in this workspace.

Panes

Pane	Description
Relationship tree pane	<p>Provides a graphical view of all the tables in the Purge Object.</p> <p>The following list provides controls and command that you can use to navigate the Relationship tree pane:</p> <ul style="list-style-type: none">• With the Level slider you can select a level and show only the tables at the selected and higher levels.• With the Zoom slider you can change the diagram size.• Right-click the driver table and then click Remove all invalid tables to remove all invalid tables with all related tables in the nested hierarchy.• Right-click a table to display a pop-up menu with the following options:<ul style="list-style-type: none">• Click Remove to remove the selected table and its related tables from the hierarchical tree. For example, if SpecTrans table appears in level 2 and level 3. If you right-click the table in level 2 and then select Remove. The table will be removed from level 2 along with its nested child relations. The occurrence of the table will remain intact in level 3.• Click Select all occurrences to display all occurrences of the selected table and their related tables in a dotted green border. In this case, the table SpecTrans will be selected in levels 2 and 3 along with all its nested child relationships.• Click Remove all occurrences to remove all occurrences of the selected table and their related tables. In this case, the table SpecTrans will be removed from levels 2 and 3 along with all its nested child relationships. The removed tables will appear in a grey colored shape with a solid red border.• Click Discover to regenerate the hierarchical parent-child relationship for the selected table. This command is available only when you select a table with the table group property of WorhSheetHeader.
Properties	<p>Shows properties for the selected table and provides commands that you can use for the selected tables.</p> <p>When a table in the Purge Object contains multiple relations, you can use the Relations node in the tree view to disable one or more relations, as detailed later in this section.</p>
Remove table	<p>Provides a data grid that you can use to select tables to be removed from the relationship tree. This pane also provides an advanced filter control that you can use to filter the data grid.</p>
Related Information	<p>Provides additional information for some of the tables. If related information is provided, read it carefully to understand the recycle relevance of the table.</p>

Buttons

Button	Description
Remove	Remove tables that are selected in the data grid in the Remove table pane. When you remove a table, all of the child tables for the selected tables are removed. The removal spans the entire relationship tree, and all nested parent-child relations for the selected table's child tables will be removed.
Revert	Reverse the modifications made to the Purge Object back to the last save. This is similar to an "undo" command.
Show all\Show selected	Show all tables or selected tables and toggles the label of the command button. Removed tables show with a broken red line around them.
Restore	Revert the Purge Object to the original source that you used to create it. In case of a purge template, the Purge Object reverts to the template that is shipped with the Data Management Framework. In case of the discovery process, the Purge Object reverts to the first Purge Object that you discovered and saved.

Fields (Remove table pane)

Field	Description
Check boxes	Select tables to be removed from the relationship tree
Table name	Name of the table
Configuration key	The configuration key of the table
Level	The level of the table in the relationship tree
Row count	Number of rows in the table

Walkthrough: Create a Purge Object

This section provides a walkthrough on working with a Purge Object.

1. If you have not created the Purge Object already, see [Create a new Purge Object](#).
2. Use the **Zoom** slider to change the diagram size so you can read table names easily.
3. In the relationship tree, move the mouse pointer over table **PurchParmTable** to see the information that is displayed in the tooltip. In the relationship tree, move the mouse pointer over **1:N** relationship description above the **PurchParmLine** table to see the information that is displayed in the tooltip.
4. In the **Properties** pane, expand each node in the tree view to see the information that is provided.
5. If the **Advanced filter** control does not show the selection criteria, click the **Advanced filter** arrow. Type "1" in the level box and click **Search**. This will change the grid to show only those tables that are in level 1.
6. Select the **PurchParmLine** table by clicking the check box next to it.
7. In the **Properties** pane, click **Remove**. Note that table **PurchParmLine** and all its child tables in levels 1 and 2 are marked with a red border, indicating that these tables are removed from the relationship tree.
8. Click the **Show all** button in the **Properties** pane. The diagram shows all the tables, with the removed tables in a grey colored shape with solid red lines.

9. Click **Revert** to restore the original Purge Object. The diagram now reverts to the state it was in before you deleted the **PurchParmSubTable** in step 7.
10. Repeat steps 6, 7, and 8 to remove the tables again. You will see all the removed tables in a grey colored shape with solid red lines.
11. Click **Show selected** in the properties pane to see the revised Purge Object. Verify that removed tables are not in the Purge Object.
12. Save the Purge Object. On the toolbar, click **Save**.
13. In the **Warning** dialog, read the warning carefully. The Data Management Framework allows you to create a Purge Object that defines a hierarchical relationship tree among the Microsoft Dynamics AX application tables. You can then apply rules to select records based on specific criteria. Records matching the selection criteria are deleted from the entire relationship tree in the Purge Object. Improper use of the Data Management Framework can cause unexpected results, database corruption, and application downtime requiring full database and application recovery. Exercise extreme caution and thoroughly test your recycling strategy in a test environment before working in the production environment. Click **Yes** to continue saving the Purge Object or **No** to cancel the save operation. Click **Yes** to continue the walkthrough.
14. In the **Save as** dialog, verify that the **Name** is PurchParmTable. Click **Save**.
15. On the **Overwrite Purge Object** dialog, click **Yes** to overwrite the Purge Object with the new version. On the **Save status** dialog, click **OK** to continue.
16. On the toolbar, click **Purge template/Purge Objects** and then click **PurchParmTable**. Notice that the icon of the Purge Object has changed. This change in the icon differentiates a Purge Object from a purge template.
17. Verify that the Purge Object diagram does not contain the tables you removed in step 10. Toggle **Show all/Show selected** to confirm that the removed tables still show up in the red border.
18. In the **Properties** pane, click **Revert**. The Purge Object will not change. The revert action cannot undo the changes because the Purge Object has been saved.
19. In the **Properties** pane, click **Restore** to revert the Purge Object. The Purge Object reverts to the to the purge template that is included in Data Management Framework, if one exists. If there is no default purge template, the Purge Object reverts to the first version you created through the discovery process. The restored Purge Object may be same as or different from the Purge Object you started working with in step 1. The default purge template contains a modified or adjusted relationship tree based on the functional assessment of Microsoft Dynamics AX application.
20. On the toolbar, click **Save**. Click **Yes** to overwrite the object. On the **Save status** dialog, click **OK**. Keep the Purge Object open as we will continue to use this object in the next section.

Walkthrough: Disable a relation

This section provides a walkthrough on disabling relations in a table with multiple relationships.

1. Click the **Configure** menu. In the toolbar, click **Purge template/Purge Objects** and then select **InventJournalTable** from the list.
2. In the relationship tree, move the mouse pointer over the **1:N** relationship description above the **InventJournalTrans** table to see the information displayed in the tooltip. Notice that this table has only one relation.
3. In the relationship tree, move the mouse pointer over the **1:N** relationship description above the **WMSJournalTable** to see the information displayed in the tooltip. Notice that this table has many relations.
4. In the relationship tree, select **WMSJournalTable**.
5. In the **Properties** pane, expand **WMSJournalTable > Relations**. This table has many relations, and you can use the **Properties** pane to disable one or more relations. Right-click the **InventBOM** relation to see the **Disable relation** pop-up menu. The **Disable relation** pop-up menu appears only when you have multiple relations.
6. In the **Relations** node, right-click **InventCount** and select **Disable relation**. Note that the disabled relation appears in red. Right-click **InventCount** and select **Enable relation**. Note that the enabled relation appears in black. An excluded relation does not become part of the Purge Object. Multiple enabled relations form the “or” clause in the SQL statement that the Purge Object generates.
7. Do not save the purge template.

Warning: Exercise care and thoroughly test the excluded relations. An erroneous or improper exclusion will result in data corruption.

Create Archive Object

This command allows you to define and create your own Archive Objects.

An Archive Object is a hierarchical relationship tree that you create based on a driver table. The Archive Object archives transactional records from the production database based on the rules and selection criteria you created in the Archive Object. This document refers to *archiving* as copying records from the production database to the archive database and then deleting these records from the production database.

You can create an Archive Object based on the templates that are included in the Data Management Framework or the discovery process.

Warning: To use this white paper, you must have experience in maintenance and administration of the Microsoft Dynamics AX application and database. The Data Management Framework allows you to create an Archive Object that defines a hierarchical relationship tree among the Microsoft Dynamics AX application tables. You can then apply rules to select records based on specific criteria. Records matching the selection criteria from the entire relationship tree in the Archive Object are moved from the production database to the archive database. Improper use of this framework can cause unexpected results, database corruption, and application downtime requiring full database and application recovery. Exercise extreme caution and thoroughly test your recycling strategy in a test environment before working in the production environment.

Note: The Data Management Framework provides archive templates for only versions 4.0 and 2009 of Microsoft Dynamics AX. In Microsoft Dynamics AX 3.0, you must create your own Archive Objects by using the discovery process to use the archive functionality.

Overview of the data archival process

The archive function archives records through the use of an archive schedule. At runtime, an archive schedule will:

- Select qualifying records from the driver table based on rules that are defined for the driver table.
- Selects qualifying records from all child tables based on the parent-child relationship.
- Copy the qualifying records from the production database to the archive database.
- Delete the qualifying records from the production database.

Considerations for Archive Object, driver table, relation, and rules

Consider the following points when working with an Archive Object.

- Verify that the table you select as the driver table for an Archive Object is a header table such as the SalesTable. A good indication, although not always accurate, is that the TableGroup value for such a table should be WorkSheetHeader, Miscellaneous, or Transaction.
- Tables with the TableGroup value of Main, Parameter, Group, and WorkSheetLine cannot be driver tables.
- The driver table becomes the root parent in the hierarchical relationship tree. The Data Management Framework queries the XRefTableRelation table to determine the parent-child relationships as detailed in [Introduction to the discovery process](#).

- Archiving records based on an Archive Object may result in data inconsistency in your application if:
 - The driver table has known parent tables.
 - The driver table has known child tables that are not added to the Archive Object as a relation.
- Always minimize the Archive Object and keep only necessary relations and rules for your implementations. Consider removing a table from the Archive Object if:
 - The table does not contain any data and it is not likely to contain data in the future.
 - The table stores only transient data such as SalesParmTable or SpecTrans.
 - The table stores open transactions such as CustTransOpen or VendTransOpen.
 - The table creates nested relationships. If a table appears multiple times on different levels in the relationship tree with the same qualifying relationship condition or set of primary keys, then you should consider keeping the occurrence only in the lowest-numbered level in the relationship tree.
 - There may be multiple relationships between two tables. In that case, evaluate the relationships and functionality carefully. We recommend that use only a single relationship in the Archive Object. However, you can use multiple relationships if necessary. Be sure to select a valid relationship using a unique index or primary key if one exists.
 - A table is related to the driver table and will make its own Archive Object because it is a good candidate to be a driver table itself. For example, if you use the SalesTable as a driver table, it will discover PurchTable for intercompany. In this case, consider making PurchTable a separate archive object. When the driver table and the child table are from completely different functional areas, and if the child table is a good candidate to be archived separately, create a separate, independent Archive Object. Consider making this child table a Related Archive Object if:
 - a) This driver table and its child table must be archived together due to functional or relational dependencies.
 - b) If you need to add rules for proper archival of the child table. You can only add rules for the driver table. To add rules for the child table, it needs to be the driver table of a Related Archive Object. Be sure that the rules you add to the Related Archive Object do not adversely affect the archival from the parent Archive Object.
- Do not use the archive templates that are included with the Data Management Framework to create your Archive Objects. These templates are created in a specific Microsoft Dynamics AX application and may or may not match your implementation. Instead, use the discovery option to create your own objects and compare them with the templates. Tables in the templates may or may not be valid for your implementation. Some valid and known tables may not appear in your discovered object. Carefully analyze any discrepancy between the objects you created through the discovery process and the templates to determine the cause of the discrepancy. Manually add or remove tables from the purge and archive objects to fit your requirements.
- Be sure that relationships among tables are set properly. Be sure that the TableGroup property is set for all custom tables.

- Verify that the relationship tree that you have created in the Archive Object archives the intended data from your system. An archive Object that you create must look at all related records. Missing relationships will lead to orphaned data.
- A table in the relationship tree becomes part of the archive function and gets disabled from the master data tables list. For more information, see [Master data](#).

Warning: Thoroughly test the impact of your Archive Objects and your archive strategy on your database and application in a test environment that is similar to your production environment. Test all business processes and obtain user sign-off before attempting to implement the Archive Objects in production environment.

Overview of the Related Archive Objects

The relationship tree of an Archive Object can be quite complex. In order to archive data from all the related tables, complex relationship trees need to be managed. In some cases, one relationship tree needs to include another relationship tree to ensure data integrity and functionally valid archival of data. Due to this complexity, you can create a relationship between two independent Archive Objects by adding one Archive Object within the other Archive Object. The child Archive Object is called the Related Archive Object and forms a relationship with the parent Archive Object. An archive schedule archives selected records from all the tables in the Archive Object and the Related Archive Object.

The following steps provide a walkthrough of an Archive Object with a Related Archive Object:

1. Use the Microsoft Dynamics AX Windows client to work with the Application Object Tree (AOT). In the AOT, expand **Data Dictionary > Tables > CustInvoiceSalesLink > Relations**. Note that the CustInvoiceSalesLink table has two parents, SalesTable and CustInvoiceJour table. Expand the **Relations** node for the CustInvoiceJour table. Note that the CustInvoiceJour table does not have a relationship with SalesTable in the AOT.
2. In the Data Management Framework, click **Configure > Archive templates/Archive Object > SalesTable** to open the SalesTable Archive Object.
3. Review the hierarchical tree and understand the parent-child relationship. Notice that the CustInvoiceSalesLink table from step 1 is not in the relationship tree as a child of the SalesTable even though the AOT shows a relationship between the two tables.
4. Notice that the relationship tree contains two oval shapes titled CustInvoiceJour and CustPackingSlipJour. These oval shapes represent Related Archive Objects for SalesTable. A Related Archive Object is an Archive Object that is part of the relationship tree of a parent Archive Object.
5. Double-click **CustInvoiceJour**. Notice that the CustInvoiceJour Archive Object opens and shows the SalesTable in a green rectangle next to the root table, CustInvoiceJour. The green rectangle highlights SalesTable as the parent Archive Object of CustInvoiceJour Archive Object.
6. Note that the CustInvoiceSalesLink table is a child table of CustInvoiceJour table. Double-click **SalesTable** to return to the parent Archive Object.
7. Right-click **CustInvoice Jour** and select **View relation**. On the **Configure Related Archive Object** window, review the relationship between SalesTable and CustInvoiceJour table and then click **Close** to close the window.
8. In the SalesTable Archive Object, right-click **CustInvoiceJour** and select **Remove**. Confirm that the Related Archive Object CustInvoiceJour is removed.
9. Right-click **SalesTable** and select **Add Related Archive Object**.

10. In the **Configure Related Archive Object** window:

- a. In the **Select Related Archive Object** list, select **CustInvoiceJour**. Click **New relation**.
- b. In the **Configure relations for Related Archive Object** pane, enter SalesTable in the **Relation name** field.
- c. Enter a valid description in the **Relation description** field.
- d. In the **Configure relations** pane, use the following table to select fields, and then click **Add**.

From the list	Select value
Table name	SalesTable
Field name	dataAreaId
Condition	=
Related table name	CustInvoiceJour
Related field name	dataAreaID

- e. Click **New Relation**. Using the preceding step, use the SalesID field to create a relationship between SalesTable and CustInvoiceJour tables. Click **Add**.
 - f. Verify that SalesTable and CustInvoiceJour tables are related using DataAreaId and SalesID fields.
 - g. Click **Save**. Click **OK**. Click **Close** to return to the parent SalesTable Archive Object.
11. Verify that the SalesTable Archive Object contains the CustInvoiceJour Related Archive Object.
12. Click **Save**. In the **Save as** dialog, type SalesTable_a and click **Save**. Click **OK** to continue. You must save an Archive Object with a new name.
13. Click **Properties** (top right-hand side of the relationship tree workspace) and pin the **Properties** window to the workspace. In the relationship tree, right-click the Related Archive Object CustInvoiceJour and click **Remove**. In the **Properties** window, click **Revert**. Reverting action reverts the changes you have made until the last time you clicked **Save**. The revert action here bring back the Related Archive Object CustInvoiceJour because you did not click **Save** after you removed the Related Archive Object.
14. Remove the Related Archive Object CustInvoice Jour again. Save the Archive Object as SalesTable_deleted.
15. From the toolbar, Click **Archive templates/Archive Object**. Compare the icon of the SalesTable Archive Object with other objects in the list and notice the difference. The icon indicates that you have opened the default template and saved it as an Archive Object. You can only use an Archive Object, and not an archive template, in an archive schedule. Click **SalesTable**. Notice that the Data Management Framework opens the Archive Object you just saved as SalesTable_deleted. The Data Management Framework allows you to save multiple versions of the Archive Object and uses the most recently saved version of the Archive Object.
16. In the toolbar, click **Show versions**. In the **Version history** window, select SalesTable in the **Archive template** list. Notice that the **Archive Object** list displays the different versions of SalesTable, with the most recent version at the top. An archive schedule uses the most recent version of the Archive Object.

The preceding walkthrough explained the concept of including a Related Archive Object within an Archive Object. You must understand the Microsoft Dynamics AX metadata and functionality thoroughly before working with an Archive Object or including a Related Archive Object within a parent Archive Object.

Understanding the Archive Object exception list

The Data Management Framework allows you to create an exception parameters list as detailed in the later section of this document. The exception parameters list applies globally to all purge templates, Purge Objects, archive templates, and Archive Objects.

In addition to the exception parameters list, the Data Management Framework allows you to maintain an exception tables list specifically for Archive Objects. When you create a new Archive Object and select the driver table, the Data Management Framework displays **Archive Object exception tables** window. The exception list contains tables that you can use as driver tables to create a separate and independent Archive Object. Carefully evaluate these tables to determine if they should be included in the relationship tree as a related child table or as a Related Archive Object. The discovery process ignores all tables in the **Archive Object exception tables** list and does not include them in the relationship tree even if a relationship exists.

To include a table from the **Archive Object exception tables** list, unselect the table and click **Save** before you click **Continue** to begin the discovery process. In this case, the discovery process will include the unselected table if a relationship exists.

You can also add other tables to the list and select the newly added table to be excluded from the discovery process.

Navigation of the Archive Object exception tables window

The following tables provide descriptions for the controls in this window.

Panes

Pane	Description
Main window	Provides the driver table and a list of exception tables for this driver table. The list of exception tables will change depending on the driver table.
Add new exception table	Provides controls to add a table to the exception tables list.

Buttons (Main window)

Pane	Description
Save	Save the changes to the exception tables list, including change to the status field and any added tables.
Continue	Continue with the discovery process for the Archive Object.

Buttons (Add new exception tables pane)

Pane	Description
Add	Add the selected table to the exception tables list. You must select a table in the Table name list before you click Add .
Clear	Clear the Table name list.

Fields (Main window)

Pane	Description
Table name	Name of the driver table
Exception tables	Names of the exception tables. Each table in the list can become a driver table.
Status	A table with the selected Status field will not be added to the relationship tree during the discovery process, even if a relationship exists. If the Status field is not selected, the table will be added to the relationship tree if there is a relationship.
Modified by	Contains the Windows ID of the user who changed the status of an existing table or added a new table to the list.

Fields (Add exception list)

Pane	Description
Table name	Select the table to be added to the exception tables list.

Create a new Archive Object

You can create an Archive Object by selecting a driver table and creating a relationship tree based on the driver table instead of selecting a template that is included with the Data Management Framework. Use the following steps to create a new Archive Object:

1. On the toolbar click **Create Archive Object**.
2. In the **Create Archive Object** dialog, select a table from the **Driver table** list. (The list excludes tables from the exception parameters list, which is covered later in this document.) Navigate to and select **WMSOrder**. A driver table is the parent table that defines the relationship tree for a specific Archive Object. In an Archive Object, the driver table may have child entities but does not have a parent entity in the relationship tree.

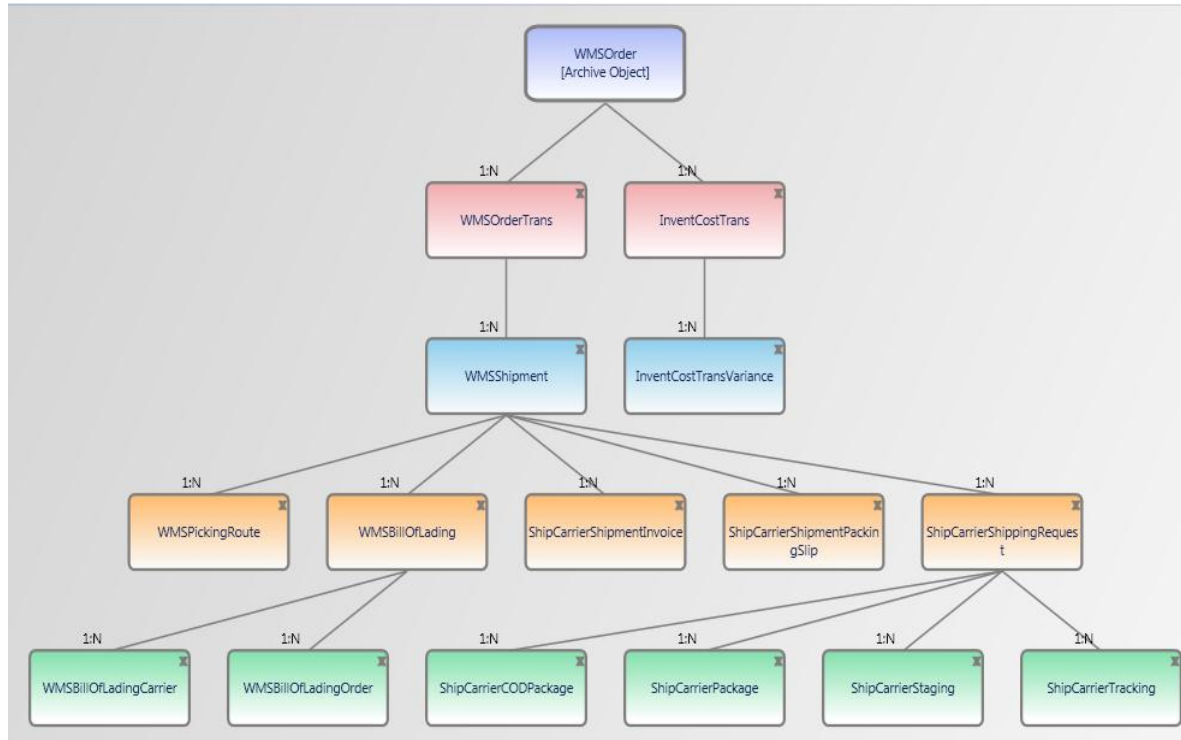
To continue, you must select the most unique index for the table. This index will be used to create the parent child relationship in the Archive Object. In the **Unique keys** box, select **orderid**. The WMSOrderIdx index is created using the orderid field.

3. Click **Discover**. In the **Archive Object exception tables** dialog, review the exception tables list carefully. If a table from the list should be included in the relationship tree, deselect the **Status** field. Use the **Add new exception table** pane to add a new table to the list. The **Status** field for the newly added table is selected by default. Click **Save** to save your changes in the **Archive Object exception tables** window. Click **Continue** to start the discovery process and create the Archive Object.

Note: The **Administer > Discovery** command allows you to maintain an exception parameters list. A table that belongs to the exception parameters list with the **Archive discovery** selected is excluded from the relationship tree when the Archive Object is created, even if there is a relationship between the driver table and excluded table.

4. Using the metadata that is imported from the Microsoft Dynamics AX database, the exception parameters, and the Archive Object exception tables list, the Data Management Framework generates a hierarchical relationship tree. The relationship tree starts with the driver table (Level 0) and creates a hierarchy of parent-child relationships. Your Archive Object should look similar to the following screen shot diagram. Do not save the Archive Object.

Figure 11. The Archive Object



The relationship tree in the preceding Archive Object is five levels deep. Level 0, the topmost level, contains the driver table, WMSOrder. Level 1 contains the child entities of the driver table. Level 2 contains child entities of tables in Level 1. The last level, level 4, contains the child entities of tables in Level 3.

This relationship tree is created based on:

1. The Microsoft Dynamics AX metadata. The discovery process uses the metadata to create a list of related tables that form the parent-child hierarchy based on the driver table.
2. The exception parameters. Tables that belong to the exception parameters list are filtered out from the relationship tree. Use the **Administer** menu to configure the exception parameters
3. The tables listed in the **Archive Object exception tables** dialog. The tables you selected in this dialog are also filtered out of the hierarchical relationship tree.

Warning: Tables that you configure as master tables (**Administer > Master data tables**) are replicated from the production database to the archive database by a master data synchronization schedule. The relationship tree created by the discovery process may include tables that are configured as master tables. Verify the tables in the relationship tree and remove any tables that must be treated as master tables. Right-click the table and select **Add to data replicate** to remove the table from the relationship and configure it as a master data table. If you do not remove these tables, tables in the relationship tree will be automatically removed from that master data synchronization list when you save the Archive Object.

Verify the relationship tree and assess the functional impact of the tree on your Microsoft Dynamics AX application. As a result of your assessment, you may require to manually add, edit, or remove some relations and rules. For example, click **Show versions** to display the Archive Object from the default template that is included with the Data Management Framework. Compare the relationship tree from the template with the relationship tree you created in the preceding step. The relationship tree in the WMSOrder archive template is different than the WMSOrder Archive Object you created. The difference is caused by the modification made to the default template based on the functionality assessment of the Microsoft Dynamics AX application. As demonstrated by this walkthrough, an Archive Object you create through discovery may or may not match with the default template that is included with the Data Management Framework.

Be careful with your selection of the driver table. If you select a table such as the ProdTable, the relationship tree can go many levels deep, spanning many tables on each level. Such an Archive Object creates a complex relationship tree and increases the potential for error and resulting data corruption and application downtime.

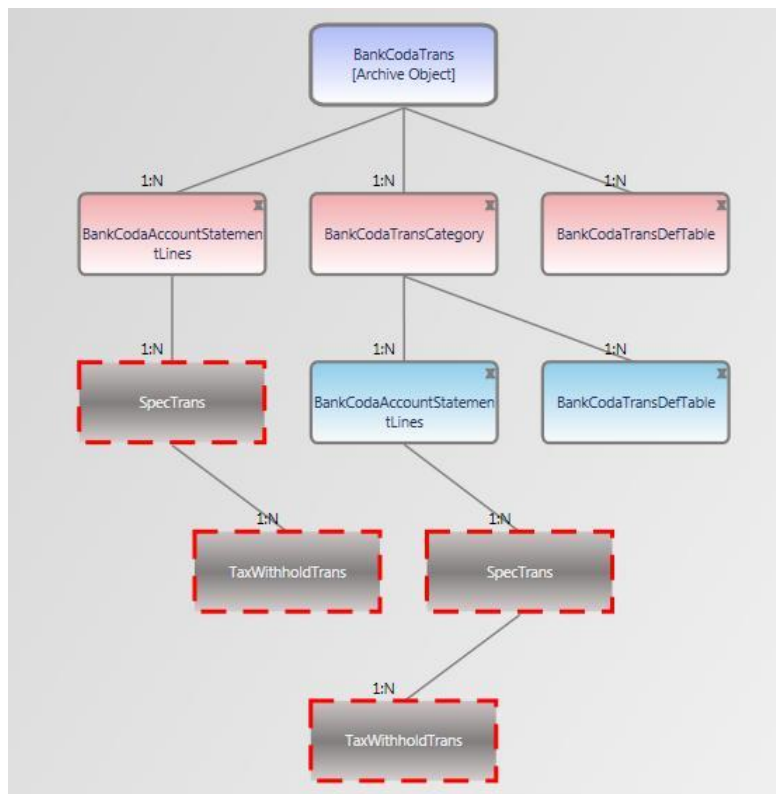


Figure 12 – Invalid tables in the Archive Object

The default template that is included with the Data Management Framework may not match the metadata of your Microsoft Dynamics AX implementation. A table in the default template that is not found in the metadata from the production database is marked invalid. An invalid table appears with a dotted red border, as shown in Figure 12. You must remove invalid tables. When you use fields from valid tables to create relations and rules, the Data Management Framework validates those fields. A field with a disabled configuration key is considered invalid. A valid table with invalid fields is shown with a yellow dotted border. You must either remove valid tables with invalid rules from the relationship tree or fix the relationships before you can save the Archive Object.

Navigation of the Create Archive Object workspace

The following tables provide descriptions for the controls in this workspace.

Panes

Pane	Description
------	-------------

Pane	Description
Relationship tree pane	<p>Provides a graphical view of all the tables in the Archive Object.</p> <p>The following list provides controls and command that you can use to navigate the Relationship tree pane:</p> <ul style="list-style-type: none"> • With the Level slider you can select a level and show only the tables at the selected and higher levels. • With the Zoom slider you can change the diagram size. • Right-click the driver table to work with the following commands: <ul style="list-style-type: none"> • Click Add Related Archive Object to add a Related Archive Object. For more information, see Overview of the Related Archive Object. • Click Remove all invalid tables to remove all invalid tables with all their related tables from the nested hierarchy. • Right-click a child table to work with the following commands: <ul style="list-style-type: none"> • Click Remove to remove the selected table from the level along with its nested child tables from the hierarchical tree. For example, if SpecTrans table appears in level 2 and level 3. If you right-click the table in level 2 and then select Remove, the table will be removed from level 2 along with its nested child relations. The occurrence of the table will remain intact in level 3. • Click Select all occurrences to display all occurrences of the selected table and their related tables in a dotted green border. In this case, the table SpecTrans will be selected in levels 2 and 3 along with all its nested child relationships. • Click Remove all occurrences to remove all occurrences of the selected table and their related tables. In this case, the table SpecTrans will be removed from levels 2 and 3 along with all its nested child relationships. The removed tables will appear in a grey colored shape with a solid red border. • Click Add to data replicate to remove the selected tables from the relationship tree and add it to the master tables list. This command will: <ol style="list-style-type: none"> 1. Remove the selected table, along with all its nested child tables, from the relationship tree of all existing Archive Objects and archive templates. 2. Add the selected table to the Recommended tables tab in the master data tables list (Administer > Master data tables). • Click Discover to regenerate the hierarchical parent-child relationship for the selected table. This command is available only when you select a table with the table group property of WorkSheetHeader. <p>Note: A table with the symbol X in the top right-hand corner table meets the row size or table size configuration in the Framework options window (Administer > Framework options). Review these tables carefully to determine if such tables should be removed from the relationship and added to the master data tables list (Administer > Master data tables). To add a table to the master data tables list, right-click the table and then select Add to data replicate.</p>

Pane	Description
Properties	Shows properties for the selected table and provides commands that you can use for the selected tables. When a table in the Archive Object contains multiple relations, you can use the Relations node in the tree view to disable one or more relations, as detailed later in this section.
Remove table	Provides a data grid that you can use to select tables to be removed from the relationship tree. This pane also provides an advanced filter control that you can use to filter the data grid.
Related Information	Provides additional information for some of the tables. If related information is provided, read it carefully to understand the archive relevance of the table.

Buttons

Button	Description
Remove	Removes tables that are selected in the data grid in the Remove table pane. When you remove a table, all of the child tables for the selected tables are removed. The removal spans the entire relationship tree, and all nested parent-child relations for the selected table's child tables will be removed.
Revert	Reverses the modifications made to the Archive Object back to the last save. This is similar to an "undo" command.
Show all\Show selected	Shows all tables or selected tables and toggles the label of the command button. Removed tables appear in grey colored boxes with a solid red line around them.

Fields (Remove table pane)

Field	Description
Check boxes	Select tables to be removed from the relationship tree
Table name	Name of the table
Configuration key	The configuration key of the table
Level	The level of the table in the relationship tree
Row count	Number of rows in the table

Walkthrough: Work with an Archive Object

This section provides a walkthrough on working with an Archive Object.

1. If you have not created the Archive Object already, see [Create a new Archive Object](#). On the toolbar, click **Save**. In the next dialog, read the warning message carefully and then click **Yes**. In the **Save as** dialog, enter WmsOrder1 and click **save**. Click **OK**.
2. Click **WMSOrder** table in the hierarchy diagram.
3. Use the **Zoom** slider to change the diagram size so you can read table names easily.
4. In the relationship tree, move the mouse pointer over table **WMSOrder** to see the information that is displayed in the tooltip. In the relationship tree, move the mouse pointer over **1: N** relationship description above the **WMSOrderTrans** table to see the information that is displayed in the tooltip.
5. In the **Properties** pane, expand each node in the tree view to see the information that is provided.
6. If the **Advanced filter** control does not show the selection criteria, click the **Advanced filter** arrow. Type "1" in the level box and click **Search**. This will change the grid to show only those tables that are in level 1.
7. Select the **WMSOrderTrans** table by clicking the check box next to it.
8. In the **Properties** pane, click **Remove**. Note that table **WMSOrderTrans** and all its child tables in levels 2, 3, and 4 are removed.
9. Click the **Show all** button in the **Properties** pane. The diagram shows all the tables, with the removed tables in red lines.
10. Click **Revert** to restore the original Archive Object. The diagram now reverts to the state it was in before you deleted the **WMSOrderTrans** in step 7.
11. Repeat steps 6, 7, and 8 to remove the tables again. You will see all the removed tables in grey shapes with solid red lines.
12. Click **Show selected** in the properties pane to see the revised Archive Object. Verify that removed tables are not in the Archive Object.
13. On the toolbar, click **Save**. Click **Save** in the dialog to save the Archive Object with a new version. You must provide a unique name when you save an Archive Object. The Data Management Framework treats the saved Archive Object as a new version of the original Archive Object or archive template.
14. On the toolbar, click **Archive template/Archive Objects** and then click **WMSOrder**.
15. Verify that the Archive Object diagram does not contain the tables you removed in step 10. Toggle **Show all/Show selected** to confirm that the removed tables still appear in grey shapes with solid red lines.
16. In the **Properties** pane, click **Revert**. The Archive Object will not change. The revert action cannot undo the changes because the Archive Object has been saved.
17. On the toolbar, click **Show versions**. In the **Version history** dialog, select **WMSOrder** from the list and select the version of the Archive Object in the **Archive Object** list. Each version of the Archive Object is saved with a unique name. Therefore, each Archive Object in the list provides a different version of WMSOrder.

18. Select the last item in the list (WMSOrder) and click **Show** to restore the Archive Object to the first version that is included as the archive template in the Data Management Framework. The restored Archive Object is different from the Archive Object you started working with in step 1. The relationship tree in the default archive template has been adjusted based on the functional assessment of the Microsoft Dynamics AX application.
19. On the toolbar, click **Save**. Click **Yes** to overwrite the object. On the **Save as** dialog, enter the name for the Archive Object and then click **Save**. Click **OK** to close the message window indicating the success of failure of the save operation. Keep the Archive Object open as we will continue to use this object in the next section.

Purge templates/Purge Object

This command provides a list of purge templates and Purge Objects in your system. A purge template is a template that is included with the Data Management Framework with a predefined relationship tree. You can use a purge template as a starting point to create a Purge Object. You must review a purge template to be sure that it meets your requirements and save it before you can use it in a purge schedule.

This command uses the same navigation and toolbar commands as the [Navigation of the Create Purge Object workspace](#).

Walkthrough: Working with a purge template

This section provides a walkthrough on working with a purge template.

1. On the toolbar, click **Purge templates/Purge Object** and then select **PurchReqTable** from the list.
2. Review the tables, table properties, and relationship tree to understand how the delete action on the **PurchReqTable** will cascade through the relationship tree in the purge template. Make sure this relationship tree covers your application and any customizations you may have made to the Microsoft Dynamics AX application.
3. On the toolbar, click **Save**. You can either keep the same name or change the name. Change the name to **PurchReqTable_object** to identify this saved object as a Purge Object. Click **Save**.
4. Click **OK** in the dialog to continue.
5. On the toolbar, click **Purge templates/Purge Objects** and compare the icons of **PurchReqTable** and **PurchReqTable_object**.
6. Repeat steps 1 through 3 but save the object as **PurchReqTable**, the original name.
7. In the **Overwrite Purge Object** dialog, click **Yes** to overwrite with the new version. Click **OK** to continue.
8. Repeat step 5. Note that the icons of a Purge Object and a saved purge template are same. This is because when you save a purge template, it becomes a Purge Object.

Archive templates/Archive Object

This command provides a list of archive templates and Archive Objects in your system. An Archive template is a template that is included with the Data Management Framework with a predefined relationship tree. You can use an Archive template as a starting point to create an Archive Object. You must review an archive template to make sure it meets your requirements and save it before you can use it in an archive schedule.

This command uses the same navigation and toolbar commands as [Navigation of the Create Archive Object workspace](#).

Walkthrough: Working with an archive template

This section provides a walkthrough on working with an archive template.

1. On the toolbar, click **Archive templates/Archive Object**. Observe the icon of **BankChequeTable** template.
2. Select **BankChequeTable** from the list.
3. Review the tables, table properties, and relationship tree to understand how the data archival for **BankChequeTable** will cascade through the relationship tree in the archive template. Make sure this relationship tree covers your application and any customizations you may have made to the Microsoft Dynamics AX application.
4. On the toolbar, click **Save**. You must provide a unique name when you save an Archive Object. Change the name to **BankChequeTable_object** to identify this saved object as an Archive Object. Click **Save**.
5. Click **OK** in the dialog to continue.
6. On the toolbar, click **Archive templates/Archive Objects**. You will not see **BankChequeTable_object** in the list. Unlike the Purge Objects, archive objects will be saved with the same name but with a new version. In this case **BankChequeTable** object is created with a version name **BankChequeTable_object**. Always, the latest version will be effective. See [Show versions](#) for more information about working with different versions of Archive Objects.
7. Observe that the icon of **BankChequeTable** is changed compared to earlier indicating it is saved as an object.

Add/Edit rules

A rule is the criterion that is used to filter records when a purge schedule or an archive schedule runs. This command allows you work with rules in a Purge Object or an Archive Object. The rules generally apply to the driver table but can be applied to a related child table. For example, the Purge Object **ProdJournalTable** includes a rule for the **ProdTable**, which is not the driver table. Regardless of how you create these rules, the where clause always applies on the driver table. You must understand the impact your rule will have on the Purge Object or Archive Object before applying any rule.

On the toolbar, click **Add/Edit rules** to open the **Add/Edit rules** window.

Navigation of the Add/Edit rules window

The following tables provide descriptions for the controls in the **Add/Edit rules** window.

Panes

Pane	Description
Rule collection	Contains a list of rules in the Purge Object or Archive Object and commands to create a new rule, add an expression to the existing rule, or delete a rule or an expression.
Configure rules for purge	Provides an area to enter or modify rule name, description, and conditions. Also provides commands to add or update the rule to the list in the Rule collection pane.
Conditional information	Allows you to specify conditions, or the selection criteria, for the Purge Object.

Buttons

Add/Edit rules window

Button	Description
Save	Save the changes that you made to the list in the Rule collection pane to the database.
Close	Close the Add/Edit rules window. You will be prompted to confirm the close or to save changes if you attempt to close the window without saving changes.

Rule collection pane

Button	Description
New rule	Create a new rule. When you have multiple rules in the Purge Object or Archive Object, they form the "and" condition in the "where" clause of the SQL statement.
Add expression	Create a new expression. An expression is a condition that is added to an existing rule. An expression creates an "or" condition in the "where" clause of the SQL statement.
Delete	Delete the selected rule or expression.

Configure rules pane

Button	Description
Add	Add the new rule or expression to the list in the Rule collection pane.
Update	Update the selected rule or expression to the list in the Rule collection pane.
Cancel	Cancel the add or update action.

Fields (across all panes)

Field	Description
Rule name	Name of the rule or expression.
Rule	The condition that you create based on the table name, field name, and a conditional operator. This condition forms the “where” clause in the SQL statement.
Rule description	Description of the rule. Enter a new description or modify an existing description here.
Table name	The table that you select from the list of all tables in the Purge Object or Archive Object.
Field name	The name of the field that you select from the list of all fields for the selected table.
Condition	The condition that you place on the table and the field. The condition list will change depending upon the field selections.
Value	The value field toggles between a text box and a list. You can select a value from a list but cannot enter a value in the text box here. You enter the value when you create a purge schedule.

Walkthrough: Add or edit a rule in a Purge Object

This section provides a walkthrough to add or edit a rule in a Purge Object.

Caution: This walkthrough will delete and recreate an existing rule from the Purge Object for ease of learning. You must have a detailed understanding of the database design, data flow, process flow, and application functionality of the Microsoft Dynamics AX application to work with rules. An error can cause data corruption or application downtime requiring full database and application recovery.

1. Click **Configure > Purge templates/Purge Object > PurchParmTable** to open the PurchParmTable Purge Object. This walkthrough assumes that you are working with the Purge Object that was created from the default purge template. If you have modified the Purge Object in any way, click **Restore** to restore the Purge Object to the original version.
2. On the toolbar, click **Add/Edit rules** to open the **Add/Edit rules** window.
3. In the **Rule collection** pane, click the first row in the data grid to select the rule **PurchParmTable.ParmJobStatus = Executed**. Click **Delete**. Click **OK** in the **Delete** dialog.
4. Click **New rule**. In the **Configure rules for purge** set the following values in selected fields:
 - a. In the Rule name field, type “Clean up”.
 - b. In the Rule description field, type “What should be cleaned up?”
 - c. In the Table name list, select PurchParmTable.

- d. In the Field name list, select ParmJobStatus.
 - e. In the Condition list, select “=”.
 - f. In the Value field, select Executed from the list.
5. Click **Add** and verify that the rule is added to the data grid in the **Rule collection** area.
6. Click **Save**. In the **Rules** dialog box, click **OK** to continue.
7. Click **Close** to close the window and return to the Purge Object.

Walkthrough: Add or edit a rule in an Archive Object

This section provides a walkthrough to add or edit a rule in an Archive Object.

1. Click **Configure > Archive templates/Archive Objects > SalesTable** to open the SalesTable Archive Object. This walkthrough assumes that you are working with the Archive Object that was created from the default archive template. If you have modified the Archive Object in any way, use **Show versions** command from the toolbar to revert to the original archive template that is included with the Data Management Framework. For instructions, see [Walkthrough: Working with an archive template](#).
2. On the toolbar, click **Add/Edit rules** to open the **Add/Edit rules** window.
3. In the **Rule collection** pane, click the second row in the data grid to select the rule **SalesTable.SalesStatus In Canceled, Invoiced**. Click **Delete**. Click **OK** in the **Delete** dialog.
4. In the **Rule collection** pane, click **Add expression** and set the following values in selected fields:
 - a. In the **Rule name** field, type “Status”.
 - b. In the **Rule description** field, type “Status values considered for archival”.
5. In the **Conditional information** pane, set the following values in selected fields:
 - a. In the **Table name** list, select **SalesTable**.
 - b. In the **Field name** list, select **SalesStatus**.
 - c. In the **condition** list, select **“In”**.
 - d. In the **Value** field, check **Cancelled and Invoiced** from the list.
6. Click **Add** and verify that the rule is added to the data grid in the **Rule collection** area.
7. Click **Save**. In the **Rules** dialog box, click **OK** to continue.
8. Click **Close** to close the window and return to the Archive Object.

Add relations

This command allows you to manually add a table to the Purge Object or Archive Object and establish a relationship. You may need to manually add a relationship if you have custom tables in the Microsoft Dynamics AX application without a metadata relationship in the Application Object Tree (AOT).

To add a relation, select a table by clicking it in the relationship tree diagram. The table you select becomes the parent table and the table you add becomes the child in the relationship. On the toolbar, click **Add relations** to open the **Add relations** window.

Navigation of the Add relations window

The following tables provide descriptions for the controls in the **Add relations** window.

Panes

Pane	Description
Relationships	Provides a list of relations that you created in the Purge Object or Archive Object. Also provides commands to create a new relation or delete an existing relation.
Configure relations	Provides an area to enter or modify the relation name, description, and conditions. Also provides commands to add or update the relation to the list in the Relationships pane or to cancel the add or update action.

Buttons

Add relations window

Button	Description
Save	Save the changes that you made to the list in the Relationships pane to the database.
Close	Close the Add relations window. If you attempt to close the window without saving changes, you will be prompted to confirm the close or save changes.

Relationships pane

Button	Description
New relation	Create a new relation or a new condition for an existing relation.
Delete	Delete the selected relation or condition. You cannot edit added relations. You must delete the relations and then add them again to make any changes.

Configure relations pane

Button	Description
Add	Add a new relation or condition to the list in the Relationships pane.
Update	Update the selected relation or condition to the list in the Relationships pane.
Cancel	Cancel an add or update action.

Fields

Add relations window

Field	Description
Include child relations	Select this field if you need to add all the child entities of the relation that you are adding to the relationship tree. If you clear field, the child entities of the table you are adding will not be added to the Purge Object or Archive Object.

Relationships pane

Field	Description
RelationsDefined	Provides a list of relations in the Purge Object or Archive Object and conditions for each relation. This field becomes available in the list when you add a relation.

Configure relations pane, Table relations area

Field	Description
Relation name	Name of the relation. You must enter a relation name to continue.
Table name	The table that you selected in the Purge Object or Archive Object. You cannot change this value in the window. This table is the parent entity in the relationship.
Field name	Select the field to be used in the relationship.
Condition	The condition defaults to "=". You cannot change the condition.
Related table name	Select the table that you are adding from the list. This table is the child entity in the relationship.
Related field name	Select the field to be used in the relationship.

Configure relations pane, Conditional information area

Field	Description
Table name	Select a table from the list. You can select either the parent or the child from the relationship.
Field name	Select the field to be used in the condition.
Condition	Select the condition to be used.
Value	Enter the value for the condition.

Walkthrough: Add a relation in a Purge Object

This section provides a walkthrough to add a relation in a Purge Object.

1. Click **Configure > Purge templates/Purge Objects** and then select **ProjJournalTable** from the drop-down list.
2. On the Purge Object, click **ProjJournalTable** table in level 0 and then click **Add relations** on the toolbar.
3. In the **Add relations** dialog box, click **New relation**.
4. In the **Relation name** field, enter a valid name for the relation.
5. In the **Table relations** area, use the following steps to add the relation.
 - a. From **Field name** list, select **JournalId**.
 - b. From the **Related table name** list, select **JournalError**.
 - c. From the **Related field name** list, select **JournalId**.
6. Click **Add**.
7. In the Add Relations dialog, click **New relation** to add a condition to the relation you just added.
8. In the **Conditional information** area, use the following steps to add a condition.
 - a. From the **Table name** list, select a table to be used in the condition.
 - b. From the **Field name** list, select the field to be used in the condition.
 - c. From the **Condition** list, select the condition.
 - d. From the **Value** list, select the value.
9. Click **Add** to add the condition to the data grid.
10. Click **Save**. Click **OK** to continue.
11. In the Purge Object, verify that the JournalError table you just added is shown as a child entity in level 1.
12. Save the Purge Object.

Walkthrough: Add a relation in an Archive Object

1. Click **Configure > Archive templates/Archive Objects** and then select **SalesTable** from the drop-down list.
2. On the Archive Object, click **DocuRef** table in level 1 and then click **Add relations** on the toolbar.
3. In the **Add relations** dialog box, click **New relation**.
4. In the **Relation name** field, enter a valid name for the relation.
5. In the **Table relations** area, use the following steps to add the relation.
6. From **Field name** list, select **ValueRecId**.
7. From the **Related table name** list, select **DocuValue**.
8. From the **Related field name** list, select **RecId**.
9. Click **Add**.
10. In the **Add Relations** dialog, click **New relation** to add a condition to the relation you just added.
11. In the **Conditional information** area, use the following steps to add a condition.
 - a. From the **Table name** list, select a table to be used in the condition.
 - b. From the **Field name** list, select the field to be used in the condition.
 - c. From the **Condition** list, select the condition.
 - d. From the **Value** list, select the value.
 - e. Click **Add** to add the condition to the data grid.
12. Click **Save**. Click **OK** to continue.
13. On the Archive Object, verify that the **DocuValue** table you just added is shown as a child entity in level 1.
14. Save the Archive Object

Export

This command allows you to export a Purge Object or an Archive Object to a file in the XML format.

Use the following steps to export a Purge Object.

1. Click **Configure > Purge template/Purge Objects** and then select **PurchParmTable** from the list.
2. On the toolbar, click **Export** to open the **Export Object** dialog.
3. The file name defaults to the driver table which is **PurchParmTable** in this case. Navigate to a location and click **Save**. Click **OK** to continue.
4. Locate the saved file and open it in a Web browser. Review the file to understand the schema and its relationship to the graphical representation of the Purge Object.

Note: You can take similar steps to export an Archive Object.

Note: In the toolbar, the **Import** command appears before the **Export** command.

Import

This command allows you to import a Purge Object or an Archive Object in XML format.

The following walkthrough provides hands-on instruction to reinstate a modified Purge Object to the state it was in at the time of export.

1. Click **Configure > Purge template/Purge Objects** and then select **PurchParmTable** from the drop-down list.
2. Use the **Advanced filter** to filter the data grid in the **Remove table** pane to show tables in only level 1.
3. Select all tables in the data grid and click **Remove**.
4. Confirm that relationship tree contains only the **PurchParmTable**.
5. On the toolbar, click **Save** to save the Purge Object. Respond to the prompt and overwrite the object.
6. On the toolbar, click **Import**. In the **Select a valid XML file** dialog box, navigate to the location of the file from the preceding section and click **Open**. Click **OK** to continue.
7. The Purge Object is now in the same state it was in when you exported it in the following section.

Note: You can take similar steps to import an Archive Object.

Save

This command allows you to save a Purge Object or save a newer version of an Archive Object. You must save a purge template as a Purge Object before you can use it in a purge schedule. You must save an archive template as an Archive Object before you can use it in an archive schedule. The archive function will not let you overwrite an existing Archive Object. When you save an Archive template or an Archive Object, you always create a new version of the Archive Object. The archive schedule always uses the most recent version of the Archive Object you save. Working with different versions of Archive Objects is covered in a later section of this document.

Show versions

This command allows you to work with different versions of an Archive Object. Use the following walkthrough to understand this functionality:

1. Click **Configure > Archive templates/Archive Object > BankDeposit** to open the BankDeposit Archive Object.
2. Review the Archive Object to ensure the relationship hierarchy, tables, and rules contained in this Archive Object are applicable to your Microsoft Dynamics AX implementation.
3. From the toolbar, click **Save**. In the **Save as** dialog, enter a name for the Archive Object. You must provide a new name for this version of the Archive Object. If the name you provide is already used by a different version of the Archive Object, you will get an error message. Enter **BankDeposit_1** and click **Save**.
4. Click **OK** to continue.
5. From the toolbar, click **Show versions**. In the **Version history** dialog, select **BankDeposit** from the **Archive template** list. The **Archive Object** list contains the different versions of the Archive Object you saved. The list contains the most recent version at the top, and you will see BankDeposit_1 at the top of the list.
6. Select any version and then click **Show** in the **Archive Object** list to work with that version of the Archive Object. If you make any changes to this version and save it, the saved version becomes the most recent version and will be used by the archive schedule at runtime.
7. Click **Close** to close the **Version history** dialog.
8. From the toolbar, click **Archive templates/Archive Object > BankDeposit** to open the BankDeposit Archive Object.

Note: The Data Management Framework opens the most recent version of the Archive Object that you saved. The archive schedule uses the most recent version of the Archive Object at run time.

Validate all

This command allows you to programmatically validate all templates and objects. The purge templates and archive templates that are included with the Data Management Framework are created based on functional validation of a standard Microsoft Dynamics AX application. Your installation may or may not have the same license, configuration, and security keys. As a result, you need to validate the default templates against your implementation before you can use them as Archive Objects or Purge Objects. This command provides you with a programmatic way to validate all purge templates or all archive templates with a single click. This functionality only work with templates that are not validated yet and will ignore any previously validated templates.

During the validation process, the Data Management Framework optionally removes tables from all templates and Objects that:

- Are not in the production database of your Microsoft Dynamics AX implementation.
- Have a relationship or a rule on a field that has a disabled configuration key in the production database.
- A record count of zero.

From the toolbar, click **Validate all** to open the **Validate all templates and Objects** window.

Navigation of the Validate all templates and Objects window

The following tables provide descriptions for the controls in the **Validate all templates and Objects** window.

Buttons

Button	Description
Save	Save the changes. When working with purge template and Purge Objects, all existing purge templates and Purge Objects are overwritten. When working with archive templates and Archive Objects, you must provide a suffix. The changes are saved as new versions, with suffix value being used for the new name.
Clear	Clear values for all the fields on the window.
Close	Close the window.

Fields

Field	Description
Select an object type	From the list, select Purge to validate all purge templates and Purge Objects. Select Archive to validate all archive templates and Archive Objects.
Suffix name	This field is only available when you select Archive from the Select an object type list. When validating objects, you will overwrite existing purge templates and Purge Objects. However, you must use a suffix value to save the new versions of archive templates and Archive Objects. The suffix value is used to create a new name for the archive templates and Archive Objects. For example, if you enter 1 as the suffix name, the BankDeposit archive template will be saved as BankDeposit_1.
Remove invalid tables	Select this field to remove all tables contained in templates and Objects that are not in the production database. Removing a table from the templates and objects also removes related child tables from the relationship hierarchy, if there are any.
Remove tables with invalid fields	<p>The templates and Objects may be using invalid fields in relationship and rules. An invalid field can be caused by your license, disabled security keys, disables configuration key, or incomplete post-installation tasks such as the database synchronization.</p> <p>Select this field to remove all tables with invalid fields from templates and Objects. Removing a table from the templates and objects also removes related child tables from the relationship hierarchy, if there are any.</p>
Remove tables with zero row	Select this field to remove all tables with zero rows from all templates and Objects. Removing a table from the templates and objects also removes related child tables from the relationship hierarchy, if there are any.

Schedule

Use this menu to schedule jobs. The Data Management Framework allows you to create different types of schedules from different areas of the application as shown in the following table.

Schedule type	Description	Available from	Frequency defaults to	Frequency Can be changed
Defragment index	<p>Defragment selected indexes. Use the following steps to defragment indexes:</p> <ol style="list-style-type: none">1. Click Analysis menu > Manage indexes.2. In the Fragmented indexes tab, select indexes to defragment and then click Schedule. <p>- or -</p> <ol style="list-style-type: none">1. Click Analysis menu > Analysis details > Index details pane > Fragmentation tab.2. Select indexes to defragment.3. Click Defragment index.	Analysis menu	One time only	No
Analysis snapshot	<p>Create a snapshot of the database including data size, index size, and index fragmentation.</p> <p>This schedule is available for the production and archive databases.</p>	Schedule menu		Yes
Ledger periods	<p>Create ledger periods that are used for the application health check queries.</p> <p>This schedule uses the production replica database.</p>			
Health check	<p>Create an application snapshot of selected measures from key modules.</p> <p>This schedule uses the production replica database.</p>			
Meta data synchronization	<p>Copy all database objects (without data) from production database to archive database based on the Microsoft Dynamics AX metadata.</p>			
Master data Synchronization	<p>Copy data from production database to archive database for selected tables.</p>			
Purge	<p>Delete records from the production database using a Purge Object and based on verification of rules within the Purge Object.</p>			

Schedule type	Description	Available from	Frequency defaults to	Frequency Can be changed
Archive	Move or archive records from production database to archive database using an Archive Object and rules within the Archive Object.			No
Restore archive	Restores data from the archive database to the production database.			
Restart	Restarts the following failed or aborted schedules: <ul style="list-style-type: none">• Archive• Restore archive• Revert schedule	Status > Job Status window		No
Revert	Reverts a failed or aborted archive schedule and brings the production database to the state it was in before the archive schedule started.	Status > Job Status window		No

Navigation

This section describes the toolbar commands available for the **Schedule** menu. These commands are explained in later sections of this document.

- Click [Production database](#) to create a new database analysis snapshot schedule for the production database,
- Click [Archive database](#) to create a new database analysis snapshot schedule for the archive database,
- Click [Ledger periods](#) to create a new ledger periods schedule.
- Click [System health check](#) to create an application health check snapshot schedule.
- Click [Metadata](#) to create a metadata synchronization schedule to synchronize the metadata of the archive database with the production database.
- Click [Master data](#) to create a master data synchronization schedule. The Master table synchronization schedule copies selected tables, called master tables, from the production database to archive database.
- Click [Purge](#) to create a new purge schedule. A purge schedule selects records from the production database based on rules within the Purge Object and then deletes these records from all the tables in the Purge Object.

- Click [Archive](#) to create a new archive schedule. An archive schedule selects records from the production database based on rules specified within the Archive Object. The schedule then moves or archives selected records from the production database to the archive database by:
 - Copying selected records from the production database to the archive database.
 - Deleting the copied records from the production database.
- Click [Restore archive](#) to create a new schedule to restore archived data. A restore schedule moves the data back into the production database from the archive database by:
 - a. Copy the selected records from the archive database to the production database.
 - b. Deleting the copied records from the archive database.
- Click [Export to Excel](#) to export selected scheduled tasks to Excel.

Navigation of the Scheduled tasks window

The following tables provide descriptions for the controls in the **Scheduled tasks** window. All schedules use this window.

Panes

Pane	Description
Scheduled tasks	Provides a list of schedules that you have created. To delete inactive current or future schedules, select the schedules. Right-click selected schedules and then select Delete schedule(s) . On the Delete schedule dialog, click Yes to continue.
Task details	Provides an area to enter the necessary information to create a schedule.
Related information	Provides a list of steps required to create a schedule.

Tabs

Tab	Description
General	Allows you to enter or modify the name and description of the schedule.
Configure rules for purge	Allows you to enter values for the rules that are created in the Purge Object. This tab is available only when you create a purge schedule.
Configure rules for archive	Allows you to enter values for the rules that are created in the Archive Object. This tab is available only when you create an archive schedule.
Configure archive restore	Allows you to select the archive schedule that is to be used to restore data. This tab is available only when you create a restore archive schedule.
Schedule	Allows you enter or select the frequency, start date, and start time of the schedule.

Buttons

Button	Description
Save/Update	Save the new or modified schedule. A saved schedule appears in the list in the Scheduled tasks pane. The label changes to Update when you select an existing schedule and Save when you are adding a new schedule.
Cancel	Cancel an add or edit action.

Fields (Scheduled tasks pane)

Field	Description
Schedule name	Name of the schedule
Type	Type of the schedule
Occurs	The frequency, start date, and start time of the schedule
Created by	The user ID that created the schedule. This value defaults to the login ID of the user that created the schedule, and you cannot change this value.
Description	Description of the schedule

Fields (Task details pane)

General tab

Field	Description
Name	Enter the name of the schedule in this field.
Description	Enter the description of the schedule in this field.
Author name	This value defaults to the currently logged on user, and you cannot change this value.
Index physical statistics	When this field is selected, the analysis snapshot schedule captures physical statistics for indexes such as the defragmentation percentage. This field is selected by default when you run the analysis snapshot for the first time. This capture is resource intensive and will impact the performance of your production system. Therefore, the field is not selected for subsequent snapshots after the initial snapshot. You must manually select the field to recalculate index statistics. Schedule the snapshot at times that will least impact your users.
Statistics by company	When this field is selected, the index and data statistics are grouped by each company in your Microsoft Dynamics AX system. This field is selected by default when you run the analysis snapshot for the first time. The grouping of data and index statistics by company is resource intensive and will impact the performance of your production system. Therefore, the field is not selected for subsequent snapshots after the initial snapshot. You must manually select the field to group statistics by company. Schedule the snapshot at times that will least impact your users.

Field	Description
Select archive type	<p>Select a value from the list to determine if the archival transactions will be row-based or set-based. Select Row-by-row to create a row-based transactional unit for the archival function. A row-based transactional unit will archive a single row from the Archive Object at a time. Similarly, a set-based transactional unit will archive all matching records from the Archive Object as a record set.</p> <p>The default value for this list is Set-based. We recommend that you use a set-based archival type for improved performance.</p> <p>Note: This field is available only for an archive schedule.</p>
Number of threads	<p>Select the number of threads to be used for master data synchronization ranging from 1 to 8. When you select a number that is greater than 1, each thread will process a certain number of tables.</p> <p>Note: This field is available only for a master data synchronization schedule.</p>

Schedule tab

Field	Description
Frequency	Select a frequency from the list.
Start date	Select or enter the start date in the date control.
Start time	Select or enter the time by using the hour, minute, and AM/PM lists.
Every day	Select a daily schedule. This field is available only for the daily frequency.
Every <n> days	Select the number of days between recurrences in a schedule. For example, select 10 to repeat the schedule every 10 days. This field is available only for the daily frequency.
Select days	Select the days of recurrence in a weekly schedule. This field is available only for the weekly frequency.
Select months	Select the months of recurrence in a monthly schedule. This field is available only for the monthly frequency.
Day	Select the date of recurrence in a monthly schedule. This field is available only for the monthly frequency.

Configure rules for purge tab (only appears when you create a purge schedule)

Field	Description
Purge Object	Select the Purge Object to be used in the purge schedule.
Table	Name of the table that is used in the rule.
Field	Name of the field that is used in the rule.
Operator	The condition that is used in the rule.
Value	Select or update values for the rules. Verify rules with existing values.

Configure rules for archive tab (only appears when you create an archive schedule)

Field	Description
Archive Object	Select the Archive Object to be used in the archive schedule.
Table	Name of the table that is used in the rule.
Field	Name of the field that is used in the rule.
Operator	The condition that is used in the rule.
Value	Select or update values for the rules. Verify rules with existing values.

Configure archive restore tab (only appears when you create a restore archive schedule)

Field	Description
Restore type	Select the filter to be applied to the schedule to be restored. On selection of Other schedules data grid will display all successfully archived schedule names.
Schedule	Name of the archive schedule. Select the schedule to restore. You can select multiple schedules. If you select a prior schedule, all more recent schedules are selected automatically. For example, 2005 for restore, 2006 and 2007 are automatically selected for restore. The Microsoft Dynamics AX application requires data for the more recent periods for transactional integrity.
Archive date	The date when the archive schedule was run.
Archive Object	The name of the Archive Object.

Production database

This command allows you to create a new database analysis snapshot schedule for the production database. The database snapshot schedule runs queries against the production database to capture database statistics such as the data size, index size, table size, and data and index usage.

You must have at least one database analysis snapshot to work with the Data Management Framework. You can schedule recurring database snapshots to create the trend analysis of your production database. The trend analysis requires a minimum of two database analysis snapshots and processes up to the latest 10 snapshots, ignoring earlier ones, if any.

On the toolbar, click **Production database** to create a schedule for the database analysis. Enter required information in the **Scheduled tasks** window and click **Save**. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) Window.

This information is used in the analysis dashboard, analysis details, and performance dashboard for the production database in the **Analysis** menu.

Archive database

This command allows you to create a new database analysis snapshot schedule for the archive database. The database snapshot schedule runs queries against the archive database to capture database statistics such as the data size, index size, table size, and data and index usage.

This is an optional step. This information is used in the analysis dashboard and analysis details, and performance dashboard for the archive database in the **Analysis** menu. You must have at least one database analysis snapshot to work with the analysis dashboard and analysis details for the archive database. The trend analysis requires a minimum of two database analysis snapshots and processes up to the 10 most recent snapshots. The Data Management Framework works with the 10 most recent snapshots and ignores earlier ones, if you have any.

On the toolbar, click **Archive database** to create a schedule for the database analysis. Enter required information in the **Scheduled tasks** window and click **Save**. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window.

Ledger periods

This command allows you to create a new schedule for the ledger periods that are required by the health check queries to capture and aggregate transactional information by company and year. You must successfully complete this schedule before running a health check schedule.

On the toolbar, click **Ledger periods** to create a schedule for the ledger periods. Enter required information in the **Scheduled tasks** window and click **Save**. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window.

This command works with the production replica database.

System health check

This command creates a snapshot of key measures from the Inventory management, Accounts receivable, Accounts payable, General ledger and Administration modules in the Microsoft Dynamics AX application. These queries capture an aggregate for each measure by company and by year. The number of years this information is captured for depends on the information in your ledger periods table.

A measure provides information about a specific business process. For example, the Inventory management module provides a measure called *Cancelled Inventory Settlement*.

You must have at least one application health check analysis snapshot to work with the application health check measures. You can schedule recurring application health check snapshots to create the trend analysis of your application health. The trend analysis requires a minimum of two analysis snapshots and processes up to the latest 10 snapshots, ignoring earlier ones, if any.

These application health check queries are extensive and will adversely affect the response time for online users. Therefore, these queries are run in a database called the production replica database by default. The accuracy of the application health check queries will depend on the relevance of the application data in the production replica database. Determine a synchronization strategy to keep your production replica database as close to the production database as possible.

On the toolbar, click **System health check** to create a schedule for the application health check snapshot. Enter required information in the **Scheduled tasks** window and click **Save**. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window.

This information is used to provide the application health check measures in the **Analysis > Show system health** command.

Warning: It is possible for you to configure the application to run the application health check queries against the production database. We caution you to not run the application health check queries against the production database as these queries adversely affect the response time of the Microsoft Dynamics AX application.

Metadata

This command synchronizes the metadata from the production database with the archive database. The metadata synchronization schedule copies database objects (and not data) from the production database to the archive database. You must complete this schedule successfully before running the master data synchronization process. Run the metadata synchronization schedule during your system maintenance period or at a time when the least amount of activity is taking place on the Microsoft Dynamics AX application.

The metadata synchronization schedule can fail while copying objects such as views, functions, and stored procedures. This failure may be due to invalid code in these objects that prevent creation, or due to cyclic dependencies. A successful schedule completes with a **Pass** status. If the schedule fails to copy database objects such as tables, indexes, and primary key, the job shows a **Fail** status. Click the **Status** menu and check the **Subtasks** list in the **Trace** pane to determine the exact step that resulted in error. The Data Management Framework creates trace files for these subtasks in the installation location. Each trace file maps to a specific subtask such as CreateColumn.txt, CreateIndex.txt, and DropTables.txt. Review the trace file corresponding to the failed subtask to determine the cause of the error. You must troubleshoot and manually resolve the issues that cause the jobs to fail. You must complete the metadata synchronization successfully to use the archival feature.

Upon successful completion of the initial metadata schedule, you must maintain the metadata synchronization between the production database and the archive database. After the first metadata synchronization, if there are changes in the production database schema, you must run the metadata synchronization to schedule to keep both the databases synchronized.

On the toolbar, click **Metadata** to create a schedule for the metadata synchronization. Enter required information in the **Scheduled tasks** window and click **Save**. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window.

Master data

This command synchronizes the master data tables from the production database with the archive database, as detailed in the later section of this document. The metadata synchronization schedule copies all the tables that you configure as master tables in the **Administer** menu. You must successfully complete the metadata synchronization schedule at least once before you can run the master data synchronization schedule. Whenever the metadata changes in the production database, you must successfully complete the metadata synchronization schedule before you run the master data synchronization schedule. The master data schedule takes the following steps to copy master tables:

1. Deletes all the existing records from each target table.
2. Copy the entire source table from the production database to the target database.

On the toolbar, click **Master data** to create a schedule for the master data synchronization. Enter required information in the **Scheduled tasks** window and click **Save**. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window.

Purge

This command allows you to create a new purge schedule. You must have a Purge Object before you create a purge schedule. At runtime, a purge schedule will use the rules in the Purge Object to filter records. The scheduled Purge Object will delete matching records from all the tables in the relationship tree of the Purge Object.

On the toolbar, click **Schedule > Purge** to open the **Scheduled tasks** window. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window. When you create a purge schedule, you work with an additional tab in the **Scheduled tasks** window. In the **Configure rules for Purge** tab, select the Purge Object and validate or edit values for the rules, in addition to working with the **General** and **Schedule** tabs.

Archive

This command allows you to create a new archive schedule. You must have an Archive Object before you create an archive schedule. At runtime, an archive schedule will use the rules in the Archive Object to filter records. The scheduled Archive Object will archive matching records from all the tables in the relationship tree of the Archive Object.

On the toolbar, click **Schedule > Archive** to open the **Scheduled tasks** window. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window. When you create an archive schedule, you work with an additional tab in the **Scheduled tasks** window. In the **Configure rules for archive** pane, select the Archive Object and validate or edit values for the rules in the Archive Object, in addition to working with the **General** and **Schedule** tabs.

Restart of the failed or aborted archive schedule

The Data Management Framework considers an unexpected error condition a failure and aborts the schedule if the scheduler service is not available. Both cases are logged in the error log files. For more information about restarting failed and aborted schedules, see [Navigation of Status workspace](#).

An archive schedule moves selected records from the production database to the archive database as detailed in the [Overview of the data archival process](#) section. The Data Management Framework archives each table as a single transaction. The Data Management Framework also tracks the completion of the archive process for each table. For example, you are processing a table called CustTable. If an archive schedule fails while archiving CustTable, it will roll back the transaction that was archiving CustTable. The Data Management Framework marks the tables being processed and is aware that the failure took place during the processing of CustTable. When you restart the failed archive schedule, the Data Management Framework starts archiving CustTable from the beginning and continues with the remaining tables in the Archive Object. For instructions, see the [Status](#) section.

For information about restarting a failed archive schedule, see [Navigation of Status workspace](#). If an archive schedule fails or aborts at the initial phase before it starts to copy data, you will not be able to restart the schedule. In that case, you must create and run a new schedule to archive data.

Revert the failed or aborted archive schedule

You can roll back a failed or an aborted archive schedule. Rolling back the schedule will restore the data that is partially archived. As detailed in the preceding section, in case of a failure, the Data Management Framework rolls back the current table that is being processed. However, the previously archived tables are not rolled back by default. For example, an archive schedule has successfully completed archival of the table WMSOrder. The schedule fails during the archival of WMSOrderTans table and rolls back the archival of the WMSOrderTans table. The database is now in a state where records from WMSOrder have been archived without archiving the corresponding records from WMSOrderTans and other related tables in the nested relationship tree. In order to maintain the integrity of the database, you must revert or restart the failed archive schedule. When you revert the failed schedule, it will attempt to restore the data it archived from the WMSOrder table. A successful revert will bring the database to a state it was in before the archive schedule started. For instructions, see the [Status](#) section.

If your archive schedule failed or aborted before it started to copy data, there is nothing to revert. If you attempt to revert such failed schedules, the Data Management Framework will display a message that there is nothing to revert for such schedules.

If a revert schedule fails or aborts, you can restart the failed revert schedule. Monitor the revert schedule to verify successful completion.

For information to revert a failed archive schedule, see [Navigation of Status workspace](#).

Restore archive

This command allows you to restore archived records from a successful archive schedule.

On the toolbar, click **Restore archive** to create a schedule for restoring the archived data. Enter required information in the **Scheduled tasks** window and click **Save**. For details on the **Scheduled tasks** window, see [Navigation of the Schedules tasks](#) window.

A successful archive schedule completes with the **Pass** status. Click **Status** menu to work with the status workspace. In the **Job** status pane, select an archive schedule with the Pass status. The **Trace** pane in the **Status** workspace lists the number of records archived in each table of the Archive Object.

Walkthrough: Create a restore archive schedule for the Archive Object

This section provides a walkthrough to create a restore archive schedule.

Use the following steps to restore the archived records to the production database:

1. Click **Schedule > Restore archive** to create a restore archive schedule. You will be working in the **Task details** pane.
2. In the **General tab**, enter the schedule name and description.
3. In the **Configure archive restore** tab, select **Other schedules** from the list to filter the data grid based on your selection. Selecting **Other schedules** will filter the data grid to show successfully completed schedules for all other Archive Objects such as SalesTable or BankDeposit.

4. Review the archive schedules shown in the **Schedule** column. You can select multiple schedules. When restoring, you can restore the most recent archive without having to restore a older archive for the same object.
5. In the **Schedule** tab, enter the values for start date and start time.
6. Click **Save**.
7. Carefully read the warning in the dialog window indicating that if you have performed a valuation model change or a standard cost revaluation. Restoring the Archive Object after a validation model change or standard cost evaluation may lead to data inconsistencies. In the production environment, do not run the restore archive schedule if you have performed any of these actions. For the purpose of this walkthrough, continue the walkthrough in the test environment. Click **OK** to close the dialog.
8. On the **Schedule** dialog, click **OK** to continue.
9. Using the **Status** menu, monitor the progress of the restore schedule and verify that it has completed successfully with a **Pass** status. Compare the information in the Trace pane from the restore schedule with the information from the archive schedule. Verify that the number of records restored match the number of records archived for all tables in the archive and restore schedules.
10. Test the restored tables in the test environment to verify data and application integrity.

Warning: Improper use of an Archive Object or restore of the archived records can cause unexpected results, database corruption, and application downtime requiring full database and application recovery. You must exercise extreme caution and thoroughly test your archival strategy in a test environment before working in the production environment.

Restart the failed or aborted restore archive schedule

The Data Management Framework considers an unexpected error condition a failure and aborts a schedule when the scheduler service is unavailable. Both cases are logged in the error log files. For more information about restarting failed and aborted schedules, see [Navigation of Status workspace](#).

A restore archive schedule restores records from the archive database to the production database as detailed in the preceding section. The Data Management Framework restores each table as a single transaction. The Data Management Framework also tracks the completion of the restore process for each table. For example, you are restoring a table called CustTable. If a restore archive schedule fails while restoring CustTable, it will roll back the transaction that was restoring CustTable. The Data Management Framework marks the tables being processed and is aware that the failure took place during the restore of CustTable. When you restart the failed restore archive schedule, the Data Management Framework starts restoring CustTable from the beginning and continues with the remaining tables in the Archive Object. For information about restarting a failed restore archive schedule, see [Navigation of Status workspace](#). If a restore archive schedule fails or aborts at the initial phase before it starts to restore data, you will not be able to restart the schedule. In that case, you must create and run a new schedule to archive data.

Export to Excel

This command exports the selected schedule to Excel.

Status

The **Status** menu allows you to refresh the status of all currently running jobs and provides details for all currently running and completed jobs.

On the toolbar, click **Status** to open the **Status workspace**.

Navigation of the Status workspace

The following tables provide descriptions for the controls in the **Status workspace**.

Panes

Pane	Description
Job status	Provides a list of all currently running, aborted, failed, or completed schedules. Each currently running or completed schedule is referred to as a job. This workspace uses the terms <i>job</i> and <i>schedule</i> interchangeably. Use the Show status details for previous option in the Administer > Framework options window to configure the number of jobs shown in this pane.
Status details	Provides status details of the selected job.
Trace	List the multiple steps and queries for the selected schedule, if there are any.

Fields

Job status pane

Field	Description
Schedule name	Name of the schedule
Type	Type of the schedule
Last run time	The date and the time when this schedule was last run

Field	Description
Last run result	<p>The result can be Pass, Fail, Running, or Aborted. The Data Management Framework and scheduler service log errors in their respective error log files as described in the Troubleshoot administration of the Data Management Framework.</p> <p>Use the following steps to restart a failed or aborted archive, revert, or restore archive schedule:</p> <ol style="list-style-type: none"> 1. Check the schedule details in Status workspace. 2. Check the error log files to determine conditions that caused the error and fix the error condition. For more information about the error log files, see Troubleshoot administration of the Data Management Framework. 3. To restart an archive or restore archive schedule, in the Status workspace > Job Status pane, right-click the schedule and then select Restart schedule. For all other schedules, create a new schedule and run the schedule. <p>Note: You will not be able to restart a schedule that failed or aborted before it started copying the data. If you cannot restart a failed or aborted schedule, you must create a new schedule.</p> <p>Monitor the schedule you started in the preceding step to verify successful completion.</p> <p>Use the following steps to revert a failed or aborted archive schedule:</p> <ol style="list-style-type: none"> 1. Check the schedule details in Status workspace. 2. Check the error log files to determine conditions that caused the error and fix the error condition. For more information about the error log files, see Troubleshoot administration of the Data Management Framework. 3. To revert an archive schedule, in the Status workspace > Job Status pane, right-click the schedule and then select Revert schedule. For all other schedules, create a new schedule and run the schedule. <p>Note: If you revert a schedule that failed or aborted before it started copying the data, the Data Management Framework will display a message indicating that there is nothing to revert.</p> <ol style="list-style-type: none"> 4. Monitor the schedule you started in the preceding step to verify successful completion.
Next schedule	The date and time for the next schedule for a recurring schedule

Field	Description
Description	<p>The Schedule description. By default, the Data Management Framework provides a brief description for the schedule such as <i>Purge data for the selected Purge Object</i>. The description does not provide any information about the rules in the Archive Object or Purge Object.</p> <p>You can modify the Data Management Framework configuration file to provide the rules and their values in the description. Use the following steps to modify the configuration file:</p> <ol style="list-style-type: none"> 1. Using an editing tool such as the Notepad, open AXDataManagementTool.exe.Config from the installation folder of the Data Management Framework. The default path for the installation folder is C:\Program Files\Intelligent Data Management Framework for Microsoft Dynamics AX. 2. Locate the configuration key IncludeRulesInDescription and change the value to "true" as shown in the following code: <pre><add key="IncludeRulesInDescription" value="true" /></pre> 3. Save the configuration file and restart the Data Management Framework. <p>With the modified configuration file, when you create or update an archive or purge schedule, the schedule description will include all the rules and their values for the selected Purge Object or Archive Object.</p>
SSIS trace	<p>This field is provides a hyperlink to view the SQL Server Integration Services (SSIS) trace output for this schedule.</p> <p>The hyperlink is only available when you set the trace level to Verbose. Use the Framework options menu to set the SSIS trace level.</p>

Status details pane

Field	Description
Type	Type of the schedule.
Status	The status can be Pass , Fail , Running , or Aborted . You must run a failed (aborted) job again. Determine the impact of the status on any currently running or future schedules
Start time	The start time of the schedule.
End time	The end time of the schedule.
Created datetime	The date and time when the schedule was created.

Trace pane (All schedules contain the following fields. The archive and restore archive schedules do not provide **Start time** and **End time** fields but provide the **Duration** field instead.)

Field	Description
Start time	The start time of the Subtask . This field is available for all schedules except for the ledger periods, archive, and restore schedules.
End time	The end time of the schedule. This field is available for all schedules except for the ledger periods, archive, and restore schedules.
Status	<p>The completion status for all schedules is either Pass, or Fail. The archive and purge schedules also show the status as Running when a schedule is active.</p> <p>For a failed schedule, check to see which subtask has failed before you run the schedule again. Determine the impact of the failed schedule or running the schedule again on any other schedules.</p>

Trace pane (Analysis snapshot schedule for production and archive databases)

Field	Description
Subtasks	A subtask is a step within a task such as Index analysis snapshot or Table analysis snapshot. The Trace pane lists all subtasks for a given schedule, where applicable.

Trace pane (Master data synchronization schedule)

Field	Description
Table name	The table name in the Purge Object.
Rows inserted	Number of rows inserted in this master table (from production database to archive database)
Rows updated	Number of rows updated in this master table (from production database to archive database)

Trace pane (Metadata synchronization schedule)

Field	Description
Subtasks	The name of the subtask that the metadata synchronization schedule performs. Each subtask corresponds to a database object such as functions, indexes, primary keys, stored procedures, tables, and views.
Production DB count	The number of objects found in the production database.
Archive DB count	The number of database objects copied to the archive database. Values from Production db count and Archive db count must match for subtasks Table, Index, and Primary key. In case of a mismatch in these counts, the metadata synchronization schedule will fail.

Trace pane (System health check schedule)

Field	Description
Period	The period is a four-digit calendar year based on ledger periods.
Measure	This field describes the measure.

Trace pane (Purge schedule)

Field	Description
Table name	Name of the table in the Purge Object.
Records recycled	The total number of records recycled from this table by the purge schedule.
Records inserted into the purge table	The total number of records inserted into the purge table. The purge table contains all records that are recycled from the production database until you remove them from the purge table. This value is same as the Number of records recycled .

Trace pane (Archive schedule)

Field	Description
Table name	Name of the table in the Archive Object or Related Archive Object.
Records inserted into the temporary tables	The total number of records that were inserted into the temporary table that was used as the source for this archival activity.
Records inserted into the archive database	The total number of records inserted into the tables in the archive database from the corresponding tables in the production database. This number must match the number of records in the source database field.
Records offlined	The total number of records archived, that is, the total number of records that were copied from the production database to the archive database and then deleted from the production database.
Summation count	The total number of adjustment records inserted into the production database you may need to make adjusting entries to balance the affected accounts.
Duration	The amount of time required to complete this archival step. Various stages of data archival process run at different time for each table in an Archive Object. As a result, the data grid provides duration instead of start time and end time
Status	The status of the subtask such as Pass or Fail.

Trace pane (Restore archive schedule)

Field	Description
Table name	Name of the table in the Archive Object or Related Archive Object that is being restored.
Archive date	The archive date and time.
Records qualified for restoration	The total number of records that were selected based on the company and period you selected for a calendar year.
Records restored	The total number of records restored from archival database to production database for this table.
Summation count	The total number of adjustment records deleted from production database.

Field	Description
Records deleted from archive database	The total number of records that are deleted from the archive database for this table.
Duration	The amount of time required to complete this archival step. Various stages of data archival process run at different time for each table in an Archive Object. As a result, the data grid provides duration instead of start time and end time
Status	The status of the subtask such as Pass or Fail.

Refresh

To refresh the **Status** workspace, click **Refresh** on the toolbar.

Export to Excel

This command allows you to export selected rows from the data grid from all workplaces where this command is available.

Administer

The **Administer** menu allows you to configure and administer the Data Management Framework. These administrative tasks include database configuration, e-mail configuration and alerts.

Navigation

This section describes the toolbar commands available for the **Administer** menu. These commands are explained in later sections of this document.

- The [Database](#) command allows you to configure the database connections for the production, archive, and production replica databases.
- The [E-mail](#) command allows you to configure the e-mail functionality.
- The [Thresholds](#) command allows you to set threshold values for tables, purge tables, and application measures.
- The [Alerts](#) command allows you to select events that generate an alert or an e-mail, if you have configured the e-mail parameters.
- The [Related information](#) command allows you to associate a description with application entities of the Data Management Framework such as tables, measures, schedules, Purge and Archive Objects. You can modify existing application entities or add new application entities.
- The [Discovery](#) command allows you to maintain the exception list. The exception list places restrictions on the use of a table in a Purge Object or an Archive Object
- The [Framework options](#) command allows you to configure the options for analysis details, SQL Server Analysis Services (SSIS) logging, display of Related Archive Objects, and master data tables.
- The [Application health check](#) command allows you to work with the application health check queries.
- The [Master data tables](#) command allows you to configure master tables that the master data synchronization schedule will copy from production database to archive database.
- The **Export to Excel** command allows you to export selected records to Excel.

Database

This command allows you to configure connection strings for the production replica, archive, and production databases.

The production database is the Microsoft Dynamics AX database in the production environment. The production replica database is used to run the queries for the health check snapshot. Although it is technically possible to run the health check queries against the production database, doing so will result in performance degradation. The archive database is used by the archive schedule to archive records based on Archive Object used in the schedule.

On the toolbar, click **Database** to work with the **Database connection settings** window.

Warning: Although it is technically possible to configure the application health check queries to run in the production database, doing so will result in performance downgrade and slower response time. We recommend that you run the health check queries in the production replica database and not select the **Same as production** field.

Navigation of the Database connection settings window

The following tables provide descriptions for the controls in the **Database connection settings** window.

Tabs

Tab	Description
Production	Configure the connection string for the production database. Note: The production database you select here must match the database in Microsoft Dynamics AX server configuration.
Production replica	Configure the connection string for the production replica database.
Archive	Configure the connection string for the archive database.

Buttons

Button	Description
Test connection	Test the database connection using the connection strings you configured.
Save	Save the database configuration.
Clear	Clear all the input fields.

Fields (across all tabs)

Field	Description
SQL Server name	Select the name of the database server from the list.
Use Windows integrated security	The scheduler service will connect to the database using the Windows integrated security.
Use specific username and password	The scheduler service will connect to the database using the User name and Password values that you enter in this window.
Username	The user name that the scheduler service will use to connect to the database.
Password	The password that the scheduler service will use to connect to the database.
Name of the database on the server	The database that you are connecting to.
Connection timeout (seconds)	The value of the connection timeout in seconds.
Maximum pool size	The maximum value of the connection pool.
Minimum pool size	The minimum value of the connection pool.
Same as production (Production replica tab)	When you select this field, the connection string from the production database is copied to the production replica database. Therefore, the health check queries will run against the production database instead of the production replica database.

E-mail

The Data Management Framework can send e-mail alerts to configured recipients. This command is used to configure the e-mail functionality.

On the toolbar, click **E-mail** to work with the **E-mail parameters** window.

Navigation of the E-mail parameters window

The following tables provide descriptions for the controls in the **E-mail parameters** window.

Buttons

Button	Description
Save	Save changes made to the e-mail configuration.
Cancel	Clear all the input fields.

Fields

Field	Description
Outbound SMTP server	The IP address of a valid SMTP Server.
SMTP port number	The port number that the Data Management Framework will use to send e-mail.
From address	The e-mail address that will appear as the sender of e-mail.
Recipient list	The e-mail address for the recipient. You can enter multiple e-mail addresses separated by a semicolon.
Username	The user name that is used to access the SMTP server. This is an optional field. Enter this value if you would like to receive e-mail alerts from the Data Management Framework.
Password	The password that is used to access the SMTP server. This is an optional field. Enter this value if you would like to receive e-mail alerts from the Data Management Framework.

Thresholds

This command allows you to determine the threshold values for index fragmentation growth, index to data ratio growth, and snapshot to snapshot growth ratio. When a table or a measure reaches the threshold value, an alert is generated. Threshold parameter values default to 0. Only a threshold parameter containing a non-zero value will generate an alert. You can see the alerts in the **Analysis details** workspace (**Analysis** menu > **Analysis details** command). You will also get an e-mail along with an alert if you have configured the e-mail functionality. On the toolbar, click **Threshold** to work with the **Configure threshold parameters** window.

Navigation of the Configure threshold parameters window

The following tables provide descriptions for the controls in the **Configure threshold parameters** window.

Tabs

Tab	Description
Tables	A list of tables and their threshold values.
Health check measures	A list of health check measures and their threshold values.
Purge	A list of tables and their threshold values.

Buttons

Button	Description
Configure	Add the entered or modified threshold values to the list.
Save	Save the changes to database.

Fields (across all tabs)

Field	Description
Filter criteria	<p>This list consists of types and categories of tables.</p> <p>You can select a value based of types of tables such as top 50, top 100, or all tables. You can also select a value based on the category of the table such as log tables or parameters tables.</p> <p>The data grid will filter based on the value you select in this list.</p> <p>The Filter criteria field defaults to Top 10 tables.</p>
Table name	Name of the table.
Index fragmentation growth (%)	<p>A threshold value for the index fragmentation growth percentage between snapshots for the selected table. For example, a threshold value of 10 will generate an alert if the index fragmentation percentages grow by 10 between now and the next snapshot.</p> <p>This value is only available for tables.</p>
Index to data ratio growth (%)	<p>A threshold value for the index to data ratio growth percentage between two snapshots for the selected table.</p> <p>This value is only available for tables.</p>
Snapshot-to-snapshot growth (%)	<p>A threshold value for the snapshot to snapshot growth percentage for the selected table.</p> <p>You can enter this value for tables or measures.</p>
Measure name	Name of the measure.
Data size (MB)	A threshold value for the purge tables. When a purge table reaches the specified size, an alert is generated.

Alerts

This command is used to configure the alerts functionality in the Data Management Framework. To receive an e-mail alert, you must configure e-mail parameters from the **Administer > E-mail** command.

On the toolbar, click **Alerts** to open the **Configure alerts** window. Select an **Alert event** to generate an alert. The following table describes the alert events you can select.

Alert event	Description
Job start	Send an alert when a schedule starts. This option is selected by default.
Job end	Send an alert when a schedule ends. Alert is sent regardless of the completion status of the schedule — pass, fail, or abort. This option is selected by default.
Thresholds	Send an alert when a specific threshold value is reached.

To view generated alerts, click **Analysis > Analysis details** and then click **Alerts** in the **Related information** pane. If you have configured e-mail functionality, the generated alerts will be sent via e-mail to configured recipients.

Related information

This command allows you to work with the related information. Related information provides contextual description for various application entities such as tables, measures, and schedules in the **Related information** pane throughout the application. These application entities are called related information parameters. These parameters are categorized in a grouping called the related information type as described in the following table.

Related information type	Description
Analysis	Provides a contextual description for a table in the Detailed analysis view.
Measure	Provides information for a measure including the measure description, archival relevance, and the formula used to calculate the measure. This information is used in the application health check analysis.
OfflineObject	Provides contextual description for a table in an Archive Object.
RecycleObject	Provides contextual description for a table in a Purge Object.
UnusedIndex	Evaluate this unused index to determine if the index is required or should be deleted.
FragmentedIndex	Provides instructions to create a schedule to defragment fragmented indexes.

Note: The Data Management Framework provides related information for a limited set of related information parameters. You can add your own related information parameters or modify existing related information parameters to suit your requirements.

On the toolbar, click **Related information** to launch the **Related information list** window.

Navigation of the Related information window

The following tables provide descriptions for the controls in the **Related information** window.

Panes

Pane	Description
Related information list	A list of related information parameters
Related information details	Fields that allow you to add or modify a related information parameter

Buttons

Button	Description
New	Add a new related information parameter.
Update/Save	The label toggles depending upon your selection. The label is Update for an existing description and Save for a new description.

Fields

Field	Description
Related information type	The related information type.
Language ID	The language ID of the related information parameter. The language ID defaults to en-US (English, United States) for the default related information parameters. You can select a language from the list when you add a new related information parameter.
Parameter	The name of the related information parameter. This is the name of the application entity such as a table name, a measure name, or a schedule type.
Description	Contextual description for the related information parameter. For example, the description provides the measure description, archive relevance, and the formula for a measure.

Discovery

This command allows you to work with the exception list. The exception list consists of exception parameters. The exception list restricts you from adding tables, relations, or rule to a Purge Object or an Archive Object. The type of restriction depends upon the type of exception you apply to an exception parameter.

The following table describes the exception parameters. For detailed explanation of configuration keys and table groups, refer to the Microsoft Dynamics AX documentation.

Exception parameter type	Description
Table name	Tables that are added explicitly by table name to the exception list. For example, the BOMCalcTrans table is configured to be an exception for the driver table. Therefore, you cannot use the BOMCalcTrans as a driver table in a Purge Object or an Archive Object.
Configuration key	A list of configuration keys in the exception list. All the tables that use these configuration keys are implicitly included in the exception list based on their relationships with the configuration keys.
Table group	A list of table groups in the exception list. All the tables that are part of the table groups in the exception list are implicitly included in the exception list based on their relationship with the table group.

Warning: Be careful when including or excluding tables from the exception list, as this can lead to improper discovery and hence an incorrect Purge Object or an Archive Object.

Note: Use the **Properties** pane in the **Purge Object** to see the table group and the configuration key for the selected table.

Click **Administer > Discovery** to work with the **Exception parameters** window.

Navigation of the Exception parameters window

The following tables provide descriptions for the controls in the **Exception parameters for discovery** window.

Panes

Pane	Description
Exception parameters	Select the exception parameter type and filter the data grid based on your selection. You cannot delete or modify the default exceptions in The Data Management Framework. You can add, modify, and delete your own exception parameters.
Save or update exception parameters	Create or modify your custom exception parameters.
Search	Search for a specific exception parameter.

Buttons

Button	Description
Add new parameter	Add a new exception parameter.
Add	Add the new parameter to the Exception parameters list.
Save/Update	Modify the data grid to reflect additions, deletions, and changes to the exception parameter and save the information to the database.
Search	Use the search value to filter the Exception parameters list.
Clear	Clear the search text in the Parameter name field.
Delete	Delete selected exception parameters.

Fields (Exception parameters pane)

Field	Description
Parameter name	Name of the exception parameter, a table name, configuration key name, or a table group name.
Parameter type	The exception parameter type: Table name , Configuration key , or Table group .
Rules	When this field is selected, you will not be able to use this parameter to create a rule in a Purge Object or an Archive Object.
Relations	When this field is selected, you will not be able to use this parameter to create a relation in a Purge Object or an Archive Object.
Purge discovery	When this field is selected, you will not be able to include any table from this parameter as a related (child) table in a new Purge Object. The discovery process will ignore all tables that are part of the exception parameter when generating the hierarchical relationship tree. Existing Purge Objects that use this parameter as a child table will continue to function normally.
Driver table	When this field is selected, you will not be able to use any tables from this parameter as a driver table in a Purge Object or Archive Object.
Archive discovery	When this field is selected, you will not be able to include any table in this parameter as a related (child) table in a new Archive Object. The discovery process will ignore all tables that are part of the exception parameter when generating the hierarchical relationship tree. Existing Archive Objects that use this parameter as a child table will continue to function normally.

Fields (Search pane)

Field	Description
Parameter type	This field is used to filter the data grid. The data grid will automatically filter based on the value you select. The default value is All .
Parameter name	Allows you to enter the search value.

Fields (Save or update exception parameters pane)

Field	Description
Parameter type	When adding a new parameter, select the parameter type from this list.
Parameter name	When adding a new parameter, select the parameter name from this list.
Rules	You cannot use tables that belong to this parameter to create rules in an Archive Object or a Purge Object. For more information, see the field description for the Exception parameters pane in the preceding section.
Relations	You cannot use tables that belong to this parameter to create relations in an Archive Object or a Purge Object. For more information, see the field description for the Exception parameters pane in the preceding section.
Purge discovery	You cannot use tables that belong to this parameter in the discovery process of a Purge Object. For more information, see the field description for the Exception parameters pane in the preceding section.
Driver table	You cannot use tables that belong to this parameter as a driver table to create Archive Object or a Purge Object. For more information, see the field description for the Exception parameters pane in the preceding section.
Archive discovery	You cannot use tables that belong to this parameter in the discovery process of an Archive Object. For more information, see the field description for the Exception parameters pane in the preceding section.

Framework options

This command allows you to configure options for the Data Management Framework such as the minimum and maximum database analysis snapshots used in the trend analysis.

On the toolbar, click **Administer > Framework options** to work with the configuration options for the Data Management Framework.

Navigation of the Framework options workspace

The following tables provide descriptions for the controls in the **Framework options** workspace.

Panes

Pane	Description
Framework options	Provides configuration options for the Data Management Framework.
Related information	Provides additional information for the selected configuration option.

Fields (Framework options pane, Database analysis options group)

Field	Description
Number of snapshots required for index analysis	The minimum number of snapshots required for meaningful analysis of index statistics. The default value is 5.
Minimum number of snapshots required for trend charts	The minimum number of snapshots required for trend analysis. The default is 2, and you cannot decrease this value.
Maximum number of snapshots allowed for trend charts	The maximum number of snapshots required for trend analysis. The default is 10. Enter a number between 2 to 10 to configure the maximum number of snapshots for trend analysis.

Fields (Framework options pane, Other configuration options group)

Field	Description
Number of Related Archive Objects per row	The number of Related Archive Object that will be displayed per row in the Configure workspace. The default value is 6.
Table row count	The threshold value for the row count. The default value is 500,000. An Archive Object highlights each table with a row count that is less than the configured threshold value with an "x" in the upper right-hand corner. You can configure the highlighted table to be a master data table and then remove it from the relationship tree. A master data table is copied and not archived.

Field	Description
Table size	<p>The threshold value for the table size. The default value is 512,000 (MB). An Archive Object highlights each table with the size that is less than the threshold value with an "x" in the upper right-hand corner. You can configure the highlighted table to be a master data table and then remove it from the relationship tree. A master data table is copied and not archived.</p> <p>Enter the size, a non-zero value, in this field and then select a value from the adjoining list to specify the size in megabytes, kilobytes or gigabytes. The default value in the list is megabytes (MB).</p>

Fields (Framework options pane, Status display)

Field	Description
Show status details for previous	This field allows to you select the period for which the status details are displayed in the Status workspace. The default value is 1 week. To change this value, enter a non-zero number in this field and then select days, weeks, or months from the adjoining list.

Fields (Framework options pane, Logging group)

Field	Description
Trace level	<p>Select a trace level from the list. Select Minimal to log only critical errors. Select Verbose to log all events, including informational events. Select None to stop logging events.</p> <p>If you select Verbose, the master data synchronization and analysis snapshot schedules will provide a hyperlink in the SSIS trace field in the Job status pane. Use the Status menu and locate the schedule in the Job status pane. Click the hyperlink in the SSIS trace field to view the trace output.</p> <p>Note: The SSIS trace field becomes available in the Status workspace only when you select the trace level of verbose.</p>
SQL Server Integration Services (SSIS) trace log refresh frequency	The frequency, in milliseconds, at which the SQL Server Integration Services (SSIS) trace output is refreshed. The default value is 1000.

Fields (Related information pane)

Field	Description
--- (Untitled)	This pane displays related information for the selected field on this workspace.

Application health check

This command allows you to work with the queries that are used to capture the application health information from selected modules. You can view and modify existing queries or create new queries.

On the toolbar, click **Administer > Application health check** to work with the **Application health check queries** workspace.

Navigation of the Application health check queries workspace

The following tables provide descriptions for the controls in the **Application health check queries** workspace.

Panes

Pane	Description
Application health check queries	Lists all the health check queries and their data sources in a tree view.
Additional health check information	Provides additional information for the selected query. For example, this pane lists the SQL statement when you generate a query.
Message	Provides an error message, for example, the error message from Generate SQL query or Save query command.
Table dictionary	Provides a tree view listing tables and their fields from the Microsoft Dynamics AX production database.
Properties	View or change the properties of the query. You can change some properties such as the query name. You can only view other properties such as the field name.

Buttons (Application health check queries pane)

Pane	Description
Validate queries	<p>The Validate queries button is available only when you launch the application for the first time. An invalid query is caused either by metadata mismatch or by incorrect security keys. When you expand the Queries node, the invalid queries appear in red. You must validate the queries before you create an application health check schedule as all invalid queries are ignored by the schedule.</p> <p>To validate queries, click Validate queries and wait for the validation to complete. Scroll through the Queries node and verify that all the queries are displayed in black.</p>

Pane	Description
Search	<p>Opens the Health check query search window.</p> <p>In the Health check query search window:</p> <ol style="list-style-type: none">1. From the Search list, select Query name or Table name.2. In the Containing text box, enter the search text. You can use * to form a wild card before and after the text. For example, you can search for *sales, sales*, or *sales*.3. Click Find now to begin the search. Depending upon your selection in step 1, the Data Management Framework will search for matching values in table names or query names. The Query name data grid displays matching results.4. Double-click the query of interest from the result set. The Health check query search window will close and take you to the query node that matches your selection.

Fields

Field	Description
Queries	<p>Expand the Queries node to work with the data source of the query. Right-click the query to work with the following commands:</p> <ul style="list-style-type: none">• Generate SQL query: Generates the query and shows it in the Additional health check information pane. You must generate a query before you can save it.• Save query: Saves the query to the database. Status of the save operation and errors are displayed in the Message pane.• Disable query: This command is available when you select an already saved query that is enabled. It disables the selected query and shows it in red. A disabled query is ignored by the application health check schedule and does not appear in the application health check analysis.• Delete query: This command is available for a new query that is not yet saved. Right-click the newly created query and then select Delete query to delete the query.• Enable query: This command is available when you select a disabled query. It enables the query and shows it in black. Enabled queries are run by the health check schedule and appear in the health check analysis.
Table Dictionary	Expand the Table dictionary node to see tables in the production database.
Fields in the properties pane	The fields in the Properties pane will change depending upon your selection in the Queries node.

Walkthrough: Create a new query

Use the following steps to create a new query.

1. Click **Administer > Application health check** to open the **Application health check queries** workspace.
2. In the **Application health check queries** pane, right-click the **Queries** node and then click **New query**.
3. The query name will be Query_nnn where “nnn” is a number. For example, the query name may be Query_70.
4. Click the newly created query. In the **Properties** pane, click the **Name** field and enter an appropriate name for the query. For the purpose of this walkthrough, type “SalesTable query” in the **Name** field. After you type the name, click outside of the **Name** field and verify that the name has changed to SalesTable query in the **Queries** node.
5. Expand SalesTable query, right-click **Data sources** and then click **New data source**.
6. The new data source defaults to the first table in the **Table Dictionary** pane. The data source name is Address(Address) where Address is the name of the data source and (Address) is the name of the database table that is being used as the data source.
7. In the **Queries > SalesTable Query > Data sources** node, click the data source **Address(Address)**.
8. In the **Properties** pane, click the **Name** field and enter an appropriate name for the data source. For this walkthrough, type “SalesTable DS” in the Name field.
9. Click the **Tables** list. Enter “SalesTable” or navigate to and select **SalesTable** from the list.
10. In the **Application health check queries** pane, click **Queries > SalesTable query > Data sources** and verify that the data source is now SalesTable DS(SalesTable).
11. Right-click **SalesTable query** and click **Generate SQL query**. Review the SQL statement in the **Additional health check information** pane.
12. Right-click **SalesTable query** and then click **Save query**.
13. In the **Select module** dialog, select **AR** from the list and click OK.
14. Click **OK** to continue.

Export to Excel

This command allows you to export selected information, such as master data tables, to Microsoft Office Excel. To export, select information in the active window that allows the export functionality and click **Export to Excel** on the toolbar.

Master data tables

This command allows you to classify selected tables as the master data tables. The master data synchronization schedule copies the master data tables from the production database to the archive database.

Warning: When a master data table becomes part of an Archive Object, the Data Management Framework disables the master data table and will not replicate it but will archive it.

On the toolbar, click **Administer > Master data tables** to work with the **Master data tables** window.

Navigation of the Master Data tables workspace

The following tables provide descriptions for the controls in the **Master data tables** workspace.

Tabs

Tab	Description
Master tables	A list of tables with the TableGroup value of Group, Main and Parameter.
Recommended tables	A list of tables with the TableGroup value of Miscellaneous, Transaction, and WorksheetLine. Top one hundred tables, based on data size or row count, are not included in this list regardless of their TableGroup value.
System and custom user tables	A list of tables that are not visible in AOT.

Buttons

Button	Description
Save	Save the changes within the list and in the database.
Clear	Clear the search criteria.
Search	Perform the search in the data grid.

Fields

Field	Description
Table name	Name of the table.
Index name	Name of the index.
Index fields	Fields that are used to create the index.
Table group	The TableGroup property of the table.
Rows	Number of rows in the table.
Modified by	The last user to select or deselect this table.
Search list (untitled)	Select a field from the list that will be used to search the data grid. The default value is All .
Search condition list (untitled)	Select a search condition from the list. The default value is = . The value you can select depends upon the field selected from the search list.
Search value (untitled)	Enter the value that you are searching for in this box. You can use the wild card character (*). For example, you can search for *sales, sales*, or *sales*.

Working with master tables

The **Master tables** tab lists all the tables with the TableGroup value of Group, Main and Parameter. By default, all tables in the master data tables list are selected. At run time, the master data synchronization schedule synchronizes all selected tables in this list using the production database as the source database and the archive database as the target database. Clear the checkbox for the table to deselect a table. A deselected table is not considered a master data table and will not be synchronized when the master data table synchronization table runs.

Tables that have indexes defined on RecId column are shown in red fonts. Evaluate these tables to determine if selecting another index will improve the synchronization performance. To select a different index:

1. Double-click the **Index name** column for the table in the grid.
2. Select another index from the list, if there are other indexes.
3. Click **Save**.

To search for a particular table, use the search utility at the top of the window. Select the field, select the search criteria, provide the search value, and click **Search**. The search utility will filter the data grid to display matching records. You can use wild card character (*) in your search value. Click **Clear** to clear the search criteria. Select **All** from the **Search list** and then click **Search** to display all tables.

Working with the recommended tables

Recommended tables typically belong to the TableGroup type Miscellaneous, Transaction WorkSheetHeader, and WorksheetLine. The recommended tables list excludes top one hundred tables in the database by size or by row count. By default, none of the tables in the **Recommended tables** list are selected. Review the recommended tables list and select the tables that need to be treated as master data tables. To select a table, click the checkbox next to the table. The synchronization schedule will treat all selected tables from the **Recommended tables** tab same as the master tables and will copy data from the production database to the archive database.

The data grid shows tables with index based on the RECID field in red fonts. If the table has multiple indexes, you can select another index, one with a unique key based on a field other than RECID. To change the index, double-click the **Index name** and select another index from the list, if there is one. The search and filter functionality works the same as the **Master tables** tab. For more information, see preceding section.

Working with system tables and custom user tables

The Microsoft Dynamics AX application uses certain tables for initialization of the Application Object Server (AOS) and to maintain metadata. These tables are part of the Microsoft Dynamics AX database but do not appear in the Application Object Tree (AOT). The Data Management Framework considers these tables as system tables. However, you may have created some custom tables in the Microsoft Dynamics AX database for integration purposes. These tables may not appear in the AOT. The Data Management Framework does not differentiate between the system tables and custom user tables.

The Data Management Framework lists all system tables and custom user tables in the **System and custom user tables** tab. Review the tables and be sure that all system tables are selected for data synchronization. You can unselect any custom user tables that are not required in the archive database.

This list also includes top 100 tables, based on number of row, with the table group property of WorkSheetHeader, WorkSheetLine, Transaction, and Miscellaneous. However, these tables are not selected by default. Determine if you need to synchronize these tables with the archive database.

Warning: You need to configure an Application Object Server (AOS) to connect with the archive database to view archived transactions. If the system tables are not replicated properly, the AOS will not be able to connect with the archive database. Be sure that users have only read access to the archive database. Modifying archived transactions will result in data inconsistency.

Database maintenance

You must maintain the management, archive, production, and production replica databases in accordance with your database maintenance, backup and restore practices.

Consider the following points for the configuration of the databases that are used by the Data Management Framework or when you schedule jobs:

1. Determine the location and initial size for data files and log files for the each database. Be sure that the initial size for the database and log files is sufficient for optimal performance and future growth. For more information about database configuration, see [Planning database configuration for Microsoft Dynamics AX](#) and the Microsoft Dynamics AX Performance team's [blog](#).
2. Schedule the set-based archive and purge operations during the regular maintenance window to maintain optimal user experience. You can schedule row-by-row operation during normal user activities but know that these operations can result in slower response time for users.
3. Run only a single archive schedule at any given time. Be sure to delay the start time of any future schedules if the current schedule seems to be taking longer than expected time. Use the **Schedule** menu to work with your archive schedules.
4. The purge and archive schedules generate disk, processor, and memory intensive operations.

Restore recycled records

This section describes how to restore recycled records after a successful purge schedule.

A successful purge schedule completes with the **Pass** status. The **Trace** pane in the **Status** workspace lists the number of records recycled in each table of the Purge Object.

To restore the recycled records to the production database, use the following steps:

1. Backup the product and management databases.
2. The stored procedure in step four runs a distributed query in the production database. You must create a linked server from the production database to the management database. For instructions, see [SQL Server 2008 Books Online](#).
3. In Microsoft SQL Server Management Studio, click **New Query**. Connect to the production database by selecting the production database from the **SQL Editor** Toolbar. Open and run the query **Cleanup_restore_purge.sql** from the installation path. The default installation path is C:\Program Files\Intelligent Data Management Framework for Microsoft Dynamics AX. This query creates the two stored procedures that are required for restoring and deleting the recycled records.
4. In Microsoft SQL Server Management Studio, click **New Query**. Connect to the management database by selecting the management database from the **SQL Editor** toolbar. Run the following query to find the **traceid** from AXScheduleTrace table. Replace 'Purge_schedule_name' in the query with the name of your purge schedule from the **Status** window. Copy the **traceid** from the **Results** pane to clipboard or Notepad.

```
SELECT B.TRACEID AS JOBIDENTIFIER , B.CREATEDDATETIME AS RUNTIME
FROM AXSCHEDULES A, AXSCHEDULETRACE B
WHERE A.SCHID=B.SCHID AND A.SCHEDULENAME = 'Purge_schedule_name'
ORDER BY B.CREATEDDATETIME
```

5. In Microsoft SQL Server Management Studio, click **New Query**. Connect to the production database. Replace the parameter values in the query with appropriate values from your environment and **traceid** from the previous step. Run the stored procedure.

```
EXEC IMPORT_PURGERECORDS
@PRODDB = 'Production database name',
@MANAGEMENTSERVER = 'Management db Server Name',
@MANAGEMENTDB = 'Management database name',
@JOBIDENTIFIER = 'Selected TRACEID from Step 3',
@AXVERSION = 'Vesion'
```

In the preceding query, replace Version with 3.0 for version 3.0, 4.0 for version 4.0 , and 5.0 for version 2009 of Microsoft Dynamics AX.

6. Test the tables in the Purge Object to verify that the restore is successful.

Warning: Improper use of a Purge Object or the restore of the purged records can cause unexpected results, database corruption, and application downtime requiring full database and application recovery. You must exercise extreme caution and thoroughly test your recycling strategy in a test environment before working in the production environment.

Delete recycled records

This section describes how to permanently delete recycled records after a successful purge schedule.

A successful purge schedule completes with the **Pass** status. The **Trace** pane in the **Status** workspace lists the number of records recycled in each table of the Purge Object.

To permanently delete records from the recycled tables on the production database, use the following steps:

1. Back up the production and management databases.
2. Run steps two and three from the preceding section to create the stored procedures and obtain the **traceid**.
3. In Microsoft SQL Server Management Studio, click **New Query**. Connect to the production database by selecting the production database from the **SQL Editor** toolbar. Replace the parameter values in the query with appropriate values from your environment and **traceid** from the previous step. Run the stored procedure.

```
EXEC DELETE_PURGERECORDS
@MANAGEMENTSERVER = 'Server Name',
@MANAGEMENTDB = 'Management DB',
@JOBIDENTIFIER = 'Selected TRACEID from Step 1'
```

4. A purge schedule deletes records from the tables in the relationship tree of the Purge Object and inserts them into purge tables in the production database. The naming convention of purge tables is purge_nnnn where nnnn is the ID of the table in the production database. For example, table purge_343 contains records that are deleted from PurchParm table. Look at the properties pane in the **Configure > Purge templates/Purge Object** workspace to get the table ID for a table in the Purge Object. Launch SQL Server Management Studio and navigate to the **Tables** node for the production database in the left pane. Verify that records for the **traceid** you used in step 1 do not exist in the purge tables that correspond with the tables in the **Trace** pane of the **Status** workspace.

Schedule this activity during a scheduled maintenance window.

Warning: These records are permanently deleted from the database. You must exercise extreme caution and thoroughly test your recycling strategy in a test environment before working in the production environment.

Troubleshoot administration of the Data Management Framework

This section contains information about troubleshooting the Data Management Framework.

View log files

The Data Management Framework logs error events in a log file in a folder named Log under the installation folder. The default installation path is C:\Program Files\Intelligent Data Management Framework for Microsoft Dynamics AX. The Data Management Framework creates the log file when the first error message is generated. The file is named trace_mm-dd-yyyy.log with mm-dd-yyyy providing the current month, day, and year. The Data Management Framework scheduler creates an error log file called servicetrace_mm-dd-yyyy.log. The error log files are created for each day. When the first error occurs during the day, the error log file is created and the error message is appended to the newly created error log file. All subsequent error messages are appended to the existing error log file for the day.

The Data Management Framework fails to start

When starting the Data management Framework, you may encounter the following error: "An unhandled exception occurred and has been logged. Please contact support."

The preceding error message is a generic message that the Data Management Framework displays when the error condition is caused by an environmental issue such as permissions. This particular error condition typically occurs when the user does not have read and write permission on the installation folder of the Data Management Framework. To fix the error condition, provide read and write permission to the installation folder of the Data Management Framework. The default path is C:\Program Files\Intelligent Data Management Framework for Microsoft Dynamics AX.

Index defragmentation does not reduce the fragmentation percentage to zero

The index defragmentation schedules do not always result in a complete defragmentation (zero percent fragmentation). Some of the indexes may still have a smaller percentage of fragmentation. The defragmentation process attempts to defragment selected indexes to the maximum possible level. However, Microsoft Dynamics AX uses a fill factor of zero. The fill factor of zero breaks records out to extra pages without completely filling an index page. As a result, fragmentation level for some of the indexes will remain at a non-zero level even after the defragmentation schedule completes successfully.

Distributed transaction error

You may encounter the following error from SQL Server Integration Services (SSIS) runtime, "The SSIS Runtime has failed to enlist the OLE DB connection in a distributed transaction with error 0x8004D025. The partner transaction manager has disabled its support for remote/network transactions." To fix this error condition, you must configure your servers so that Microsoft Distributed Transaction Coordinator (MSDT) communication flows between the firewall. For instructions, see the [Knowledge Base article 306843](#).

Export to Excel fails

The export to Excel functionality may not work in some environments. To fix this problem, you must have Excel and the Excel PIAs installed on the computer. Verify that Excel is properly installed and working on the computer. If required, download and install the Microsoft Office Primary Interop Assemblies (PIAs). For Excel 2003, see [Office 2003 Update: Redistributable Primary Interop Assemblies](#). For Excel 2007, see [2007 Microsoft Office System Update: Redistributable Primary Interop Assemblies](#). Restart the Data Management Framework after successful installation of the PIAs.

The discovery of a driver table fails

When starting a discovery process for an Archive Object or a Purge Object, you may encounter the following error: "Unable to discover the driver table for the Archive Object or Purge Object."

This error message is usually caused due to metadata synchronization issues. To fix this error condition, you need to run the post-installation tasks. Run the post-installation application (**Start > All programs > Intelligent Data Management Framework > Post-installation tasks**) or use the following steps to resolve this error manually. For details, see the [Data Management Framework Installation Guide](#).

1. Verify that the AXDataManagementToolProject nn .xpo exists in the in the Application Object Tree (AOT) where nn is 30, 40, or 50 depending on the version of your Microsoft Dynamics AX application.
- 2.
3. Verify that you have successfully completed the post-installation checklist.
4. Verify that all member of the implementation team are working on the same layer, for example CUS, VAR, or USR.
5. Verify that the XPOs mentioned in steps 1 and 2 are imported in the same layer such as the VAR layer or the USR layer. Verify that all methods from these XPOs are imported in the same layer such as USR or VAR.
6. Use the AOT to correct any discrepancy you find between your AOT and the XPOs that are included with the Data Management Framework. You may delete these XPOs from the layer they are currently in and manually import them to another layer as described in the [Data Management Framework Installation Guide](#).

Database snapshot schedule fails with a permission error

In a distributed SQL Server 2005 environment, you must install service pack 2 or above on all SQL servers when you store your databases on multiple servers. For example, when you store your production database and archive database on two separate servers, you have a distributed SQL Server environment. Both the servers must have SQL Server 2005 SP2 or above before you install and launch the Data Management Framework. Failure to do so will result in the following error message from a database snapshot schedule, "Syntax error, permission violation, or other nonspecific error."

Discovery process fails with an error message

The discovery process displays an error message, "Cannot create archive object or purge object."

This message may appear if the Application Object Server (AOS) is not available or the connection with the AOS has failed. Check the availability of AOS. Restart the AOS if necessary and then restart the Data Management Framework application.

Application Object Server (AOS) error when connecting to archive database

In some environments, you may get an error when the Application Object Server (AOS) connects to the archive database, "An internal error occurred while creating session for the user."

Verify that the account used for the AOS service has sufficient permissions to the archive database. For more information, see the **Rights required for installation** section in the [Data Management Framework Installation Guide](#). Verify that the account used for the AOS service has the execute permission for stored procedures CREATESERVERSESSIONS and CREATEUSERSESSIONS in the archive database.

Analysis snapshot schedule fails with a permission error

The analysis snapshot schedule fails with an error message, "An OLE DB error has occurred. Error code: 0x80040E14. An OLE DB record is available. Source: "Microsoft SQL Native Client" Hresult: 0x80040E14 Description: "The user does not have permission to perform this action".

This error message appears due to insufficient database permission. Be sure that the database permissions are assigned by following the instructions in the see the **Rights required for installation** section in the [Data Management Framework Installation Guide](#).

Analysis snapshot schedule fails due to SQL Server Integration Services error

The analysis snapshot schedule fails with an error, "Message: Package Error: The file exists. The buffer manager could not get a temporary file name. The call to GetTempFileName failed. The buffer manager could not create a temporary file on the path C:\DOCUME~1\ADMINI~1\SAM\LOCALS~1\Temp. The path will not be considered for temporary storage again".

To resolve this error, you must install SQL 2005 SP3 Cumulative Update 5 or higher. For more information, see the Microsoft Knowledge Base article [972365](#).

Unable to create an archive schedule or restore schedule

If an archive schedule is failed or aborted, you must attempt to restart or revert it before you can create another schedule with the same object. This restriction applies to all Archive Objects.

Unable to select an archive schedule for restoration

The archive schedule that you want to restore may be disabled and you cannot select it. This occurs when your previous attempt to restore the schedule has either aborted or failed. In the **Status** window, navigate to the aborted or failed restore schedule. Right-click the schedule and then select **Restart schedule** or **Revert schedule** depending on the action you need to perform.

If the schedule failed during initialization state, no data was archived. In that case, there is no need to restore or revert the schedule, and the Data Management Framework will display a warning that asks you to create a new schedule. Use the **Schedule** menu to create a new archive schedule.

NOTE: Always restart or revert a failed archive schedule before you create a new archive schedule for the same Archive Object.

A Schedule fails due to insufficient space

Archive and Purge schedules will fail if the database or log files become full. These schedules require significant amount of free database and log space depending on the volume of data you archive or purge. By default, the purge and archive schedules process 100,000 records in a batch. Reduce the batch size to decrease the database and log space required when a schedule processes a batch. Use the following steps to configure the batch size:

1. Using Notepad, open the AXDataManagementSchedulerService.exe.config from the installation folder of the Data Management Framework. The default path for the installation folder is C:\Program Files\Intelligent Data Management Framework for Microsoft Dynamics AX.
2. To decrease the batch size for an archive schedule, locate the configuration key BatchSizeForPurge <add key="BatchSizeForArchive" value="100000" /> and change the key value to a lesser number.
3. To decrease the batch size for a purge schedule, locate the configuration key BatchSizeForPurge <add key="BatchSizeForPurge" value="100000" /> and change the key value to a lesser number.

You must carefully determine the location and initial size for data and log files for the archive, production, management, and production replica databases. Be sure that the initial size is sufficient for optimal performance and future growth. For more information about database configuration, see [Planning database configuration for Microsoft Dynamics AX](#) and the Microsoft Dynamics AX Performance team's [blog](#).

Troubleshoot failed schedules

Use the following steps to troubleshoot failed schedules:

1. View the log file and a list of common causes and associated solutions that are described in the preceding sections.
2. Be sure to synchronize any metadata changes in the production database with the archive database via the metadata synchronization schedule. The master data synchronization schedule may fail if the metadata is not synchronized between the two databases.
3. The purge schedule may fail if the Microsoft Distributed Transaction Coordinator (MSDTC) is not configured properly.
4. A schedule may fail if the data file or log file in the database becomes full.
5. You must fix the error conditions before you rerun the schedule.

Appendix A: Terminology

This section explains the terminology and concepts used in the Data Management Framework.

Term	Description
analysis snapshot	A schedule that captures database analysis information such as the data size, index size, and index usage.
archive	Archive the records matching selection criteria from all tables in the relationship tree of an Archive Object. The Archiving process moves records from the production database to the archive database. The terms archive and offline are used interchangeably in this document.
Archive Object	A relationship tree that you create by selecting a driver table and related tables.
archive template	An Archive Object that is included with the Data Management Framework. You must review and save the archive template before you can use it to archive data. A saved archived template becomes an Archive Object.
baseline	A snapshot of the database or the application that is captured through specific queries.
discovery	The discovery process begins with the driver table and searches through the Microsoft Dynamics AX application metadata to capture all tables in the relationship tree.
driver table	A parent table in the parent-child relationship of related tables. The driver table is the root parent.
entity	A database table.
exception	A list of tables, table groups, and configuration keys with defined restriction from use in a Purge Object or an Archive Object. An exception parameter can be a table, a table group, or a configuration key.
expression	An expression is a condition you place within a rule in a Purge Object or an Archive Object. When you have multiple expressions, they form the "or" clause in a query.
level	A grouping of tables in a Purge Object or an Archive Object. The driver table starts at level 0, which is the root level. All the tables related to the driver table are placed in level 1. All the tables related to a table in level 1 are placed in level 2, and so on. The levels form the parent-child hierarchy in a relationship tree.
management database	The database that is used by the Data Management Framework.
master data synchronization	The process of replicating master data tables from the production database to the archive database. This process copies the master data tables from the production database to the archive database to keep them synchronized.
master data tables	The master data synchronization schedule copies the master data tables from the production database to the archive database. These tables are not archived, meaning they retain their data in production database after the records are copied to the archive database.

Term	Description
measure	A measure captures aggregated statistics for key business processes by company and by year in the Microsoft Dynamics AX application.
meta data synchronization	The process of synchronizing database schema and database objects from the production database to the archive database.
offline or offlining	Same as archive.
performance dashboard	The performance dashboard displays the database configuration details for all databases used by the Data Management Framework and table statistics and query statistics for the production and archive databases.
purge	Delete the records matching selection criteria from all tables in the relationship tree of a Purge Object. Purge and recycle are used interchangeably in this document.
Purge Object	A relationship tree that you create by selecting a driver table and related tables.
purge schedule	A purge schedule recycles all the records from the database using a Purge Object, based on values set for the rules of the Purge Object.
purge template	A Purge Object that is included with the Data Management Framework. You must review and save the purge template before you can use it for recycling. A saved purge template becomes a Purge Object.
recycle	Same as purge.
relation	A relation is a child table that you add to a Purge Object or an Archive Object. All tables in a relationship tree have a parent table except for the driver table.
relationship tree	A parent-child relationship that starts with the driver table. The driver table is placed in level 0, or root level, of the relationship tree. All the tables that refer to the driver table as a parent are placed in level 1. All the tables that refer to a table in level 1 as a parent are placed in level 2. The child-parent relationship keeps going deeper until all the related tables are grouped in the relationship tree.
restore	The process of restoring archived data from the archive database to the production database. The restore process moves qualifying data from the archive database to the production database.
rule	A rule is a condition that you place on a field in one of the entities in a Purge Object or an Archive Object. When you have multiple rules, they form the "and" clause in a query.
system health (also referred to as the application health)	System health provides analysis for selected measures from the Inventory management, Accounts receivable, Accounts payable, General ledger and Administration modules of Microsoft Dynamics AX.

Microsoft Dynamics is a line of integrated, adaptable business management solutions that enables you and your people to make business decisions with greater confidence. Microsoft Dynamics works like and with familiar Microsoft software, automating and streamlining financial, customer relationship and supply chain processes in a way that helps you drive business success.

U.S. and Canada Toll Free 1-888-477-7989

Worldwide +1-701-281-6500

www.microsoft.com/dynamics

This document is provided “as-is.” Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes. You may modify this document for your internal, reference purposes.

Microsoft, Microsoft Dynamics, Microsoft Dynamics logo, Microsoft SQL Server, Windows, and Windows Server are trademarks of the Microsoft group of companies. All other trademarks are property of their respective owners.

© 2010 Microsoft Corporation. All rights reserved.

Microsoft