CA XOsoft® r12

User Guide

Edition no. 3

Contents

Chapter 1: Introduction

This Chapter contains general information about the CA XOsoft products, and their various modules. It briefly lists what is new in CA XOsoft r12, describes how CA XOsoft Replication and CA XOsoft High Availability work, and how the various modules function in the replication process.

About this Guide

This Guide contains all of the necessary information for configuring and running the CA XOsoft application. It describes and provides instructions on how to perform the following procedures: synchronizing, replicating and recovering data, monitoring the procedures and generating reports, switching over from the production server to the Replica standby server, and switching back.

This Guide focuses on a generic File Server Disaster Recovery solution, but it also provides information about other application and database servers and the High Availability solution. For more detailed instructions involving High Availability scenarios or scenarios tailored to specific applications such as Exchange or SQL, see the appropriate Operations Guide. You can find the most up to date Operations Guides for each application on http://supportconnectw.ca.com/public/xosoft/manuals/xosoft_docman.asp

Important! This Guide applies to CA XOsoft File Server, CA XOsoft Exchange, CA XOsoft Oracle, CA XOsoft SQL, CA XOsoft IIS and CA XOsoft High Availability products (which include all the capabilities of CA XOsoft). Throughout this document, the term CA XOsoft refers to all products, unless otherwise specified.

Terms and Abbreviations

The following terms and abbreviations are used throughout the document.

Bookmark

Checkpoint that is manually set in the Rewind log to mark a state that you may want to rewind back to. Usually set before a major set of operations is performed on the data.

Data Rewind

Data recovery solution that uses the rewind points or bookmarks in the Rewind log to reset the current data back to a previous state.

Master

Main or production server where database and file activities are taking place. The server in the replication scenario that you want to replicate.

Node

A Master or Replica host in a replication scenario tree.

Recovery

Process of retrieving lost or corrupted data on the Master from any Replica by activating a synchronization process in the reverse direction.

Replica

Server set up to receive replication data from a Master server.

Replication

Replication consists of several actions that operate simultaneously and continuously: real-time capture of changes to files and databases on a Master server; transfer of the capture changes to the Replica server(s) of the Master; applying the changes to the Replica(s) data. Replication is a continuous process. It stops when a scenario is stopped.

Replication Tree

The connection scheme of replication from Master host to its Replica host(s) represented as a tree.

Resume replication

Begin transferring accumulated changes to a Replica that was previously suspended.

Rewind point

Checkpoint in the Rewind log marking an event or operation.

Scenario

Full Replication process definition including Master and Replica servers and their connectivity (replication tree), report and event handling rules, properties of nodes, directories, sub-directories, databases and files that will be participating in the replication process. Each scenario is saved as an XML file.

Suspend replication

Temporarily cease delivering changes to the suspended Replica. Changes are accumulated in a spool until replication is resumed so that re-synchronization is not required.

Synchronization

The process of producing an exact copy of the Master server's contents on remote Replica servers.

Related Documentation

Use the following Guides along with this User Guide:

- *CA XOsoft r12 Installation Guide*
- *CA XOsoft SQL Server Operations Guide*
- *CA XOsoft Exchange 200x Operations Guide*
- *CA XOsoft Red Hat Server Operations Guide*
- *CA XOsoft Solaris Server Operations Guide*
- *CA XOsoftHA Exchange 200x Operations Guide*
- *CA XOsoftHA SQL Server Operations Guide*
- *CA XOsoftHA Oracle Server Operations Guide*
- *CA XOsoftHA File Server Operations Guide*
- *CA XOsoftHA IIS Server Operations Guide*
- *CA XOsoftHA for Blackberry Enterprise Server Operations Guide*

White papers offering a broader introduction to disaster recovery, business continuity and content distribution solutions using CA XOsoft and CA XOsoftHA are also available on the CA XOsoft website: www.caxosoft.com

CA XOsoft Replication and High Availability Overview

CA XOsoft is a powerful, easy-to-use product for creating and maintaining full online backups of your business or other important data. CA XOsoft monitors your file system online, intercepts and copies all additions or changes to it, and transmits this information online to your backup servers.

CA XOsoft Replication

CA XOsoft Replication is a data protection solution that uses asynchronous real-time replication to provide cost-effective Disaster Recovery capabilities. The asynchronous data replication is performed over LAN, the Internet or WAN, if participating servers have the required software installed and TCP connections between them. CA XOsoft allows the replication of databases and other files from a production or Master server to one or more local or geographically distant Replica servers. The Disaster Recovery capabilities offer rapid data recovery, integrated continuous data protection as a guard against data corruption, and completely non-disruptive automatic testing of your Disaster Recovery system.

The replication process requires first setting up a valid replication scenario. Then, synchronization between the Master and Replicas data is performed, and once the Master and Replica servers have an identical set of data, the replication process starts. Once the replication process has started, the data may be retrieved from the Replica (standby) server at any time. The CA XOsoft Manager manages the process, and also monitors both the replication and the retrieval of data.

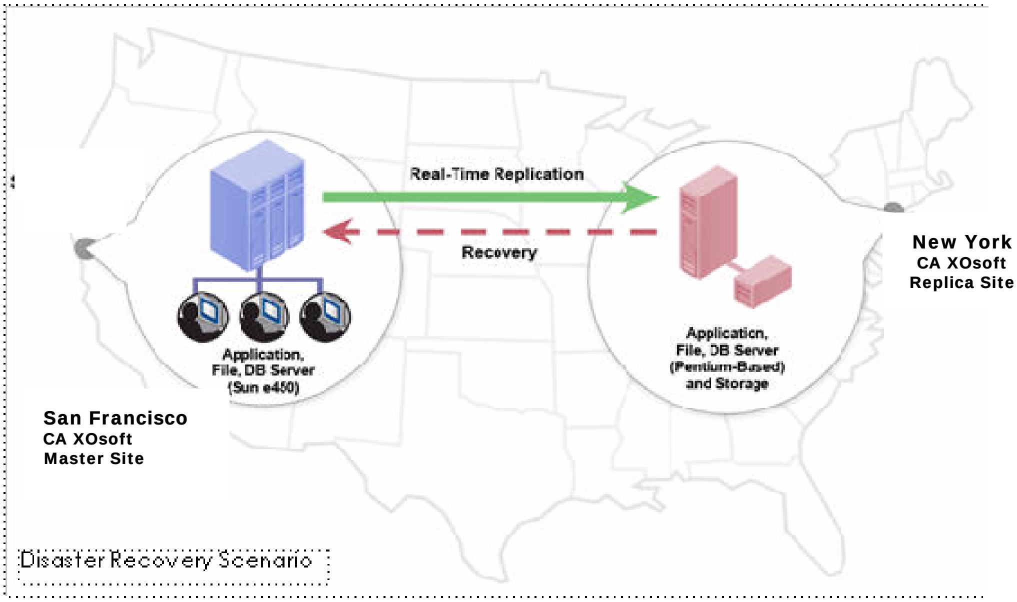

CA XOsoft High Availability

CA XOsoft High Availability (CA XOsoftHA) is based on asynchronous real-time replication and automated switchover and switchback. It provides cost-effective business continuity for Microsoft Exchange, Microsoft SQL Server, Oracle, Microsoft IIS Server, File Server, and other applications on both 32-bit and 64-bit Windows servers.

CA XOsoftHA offers push-button or fully automatic switchover of mission-critical application servers over a LAN or WAN. It features application-aware status monitoring, integrated continuous data protection as a guard against data corruption, and completely non-disruptive automatic testing of your Disaster Recovery system.

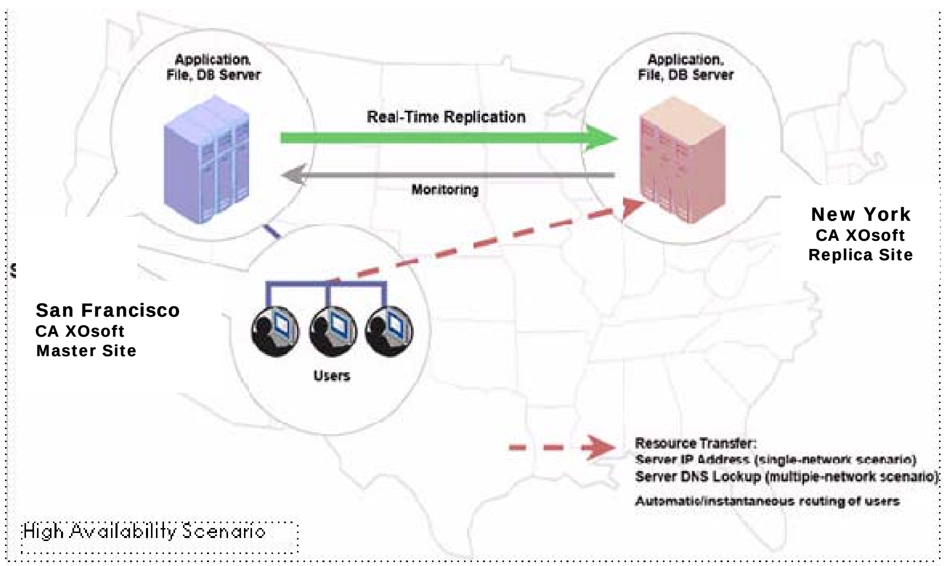

Limitations

Only one-way, asynchronous, replication is permitted, and the Replica database must be offline. Bidirectional replication is not supported. However, cross replication with different data sets is supported. A server running CA XOsoft can act as a Master and Replica for an unlimited number of scenarios so long as each data set only has a single Master server, i.e., one way replication CA XOsoft.

Supported Application and Database Servers

These replication and Disaster Recovery capabilities are custom-tailored for the following application and database servers, for both 32-bit and 64-bit Windows:

- Microsoft Exchange
- Microsoft SQL Server
- Microsoft IIS
- Oracle
- File Server

For an up-to-date list of supported platforms and applications, see http://supportconnect.ca.com/sc/kb/techdetail.jsp?searchID=TEC408153&docid=408153&bypass=yes&fromscreen=kbresults

CA XOsoft Solutions

How Synchronization Works

Synchronization of files is the process of making the set of files to be protected identical on the Master and Replica servers. It is usually necessary to synchronize the Master and Replica as the initial step of a replication scenario.

Important! We strongly recommend running the initial synchronization during off-peak hours.

Synchronization Methods

In order to properly synchronize the Master and the Replica, first it is necessary to perform a comparison of their two file structures. This comparison determines what content (files and folders) on the Master is missing or different from that on the Replica. CA XOsoft supports two different synchronization modes, each with a different comparison algorithm and method of operation: File and Block synchronization.

In File synchronization, the Replica server sends its snapshot to the Master server, which uses it to drive the information and content for the data comparison. After the comparison is performed, the CA XOsoft Engine on the Master server sends a sequence of commands to the Replica. These commands consist of (a) deletes of files that exist only on the target, and (b) the entire file contents of files that either exist only on the Master, or that exist on the Replica but differ from the version on the Master. This method is best suited for File Servers or application servers with large number of relatively small files.

In Block synchronization CA XOsoft performs a block-by-block comparison of the Master and Replica files, and copies over only those blocks that are different. When differences exist between files, instead of requiring the transfer of the entire file, the block synchronization transfers changes only. This method is best suited for database applications, such as MS Exchange, Oracle, and SQL Server, or application servers with very large files.

The comparison may be configured to consider only file size and modification time to determine whether two files differ, or it can perform a check of the actual contents of the data. The former approach, which is not valid in the case of database applications, can be a legitimate way to significantly speed up the comparison process on a File server scenario.

Synchronization Filter

Before the synchronization starts, CA XOsoft offers you the possibility to filter the synchronization process. This filter, called **Ignore files of same size/time,** appears on the **Run** dialog, and it can be either enabled or disabled:

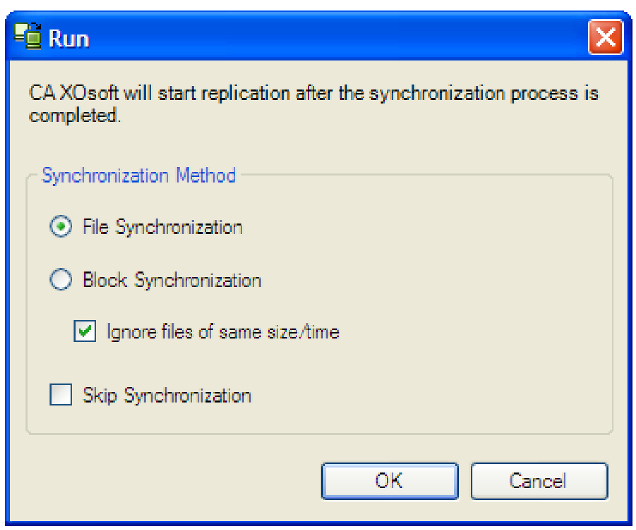

The **Ignore files of same size/time** filter allows the data comparison between the Master and the Replica to consider only file size and modification time, when determining whether two files differ. It skips comparison of files with the same path, name, size and modification time, under the assumption that these files are identical. This approach is not valid in the case of database applications, but it can be an efficient way to significantly speed up the comparison process for a File server solution, and reduce the overall synchronization time dramatically.

Important! Do NOT skip synchronization unless you are absolutely sure that files on the Master and Replica are identical.

Automatic Synchronization

CA XOsoft enables you to configure your system to perform automatic synchronization once certain events occur. The Automatic Synchronization property causes the Master and Replica servers to be automatically resynchronized in the event one of the following occurs:

- If the Replica is rebooted, CA XOsoft resynchronizes the Master and the Replica automatically after reboot.
- If the Master is rebooted, CA XOsoft resynchronizes the Master and the Replica automatically after reboot.
- If the Master spool overflows because of network failure, CA XOsoft resynchronizes the servers automatically upon resumption of the connection.

Simultaneous Synchronization and Replication

CA XOsoft allows simultaneous synchronization and replication. This means that servers can be synchronized while files are in active use and being updated; all changes that occur while the initial synchronization is performed are replicated without any need for administrative intervention.

Reporting Synchronization Differences

The data sets on the Master and Replica servers may be checked for differences without actually performing a full synchronization through the **Difference Report** option.

How Replication Works

The replication mechanism maintains identical copies of files and databases on the Master and Replica. This replication is done by real-time capture of byte-level changes in files on the Master server, using a file-system filter-driver. The captured changes are asynchronously transmitted to the Replica servers using the CA XOsoft Engine. The replication process does not interfere with write operations.

To accommodate real-time replication of all types of files, CA XOsoft supports the following replication modes:

- Online mode - Replicates captured changes of files, even if files are always open (as is the case for most database and mail servers). This mode maintains the order of file system operations. In this mode, the Engine records all I/O operations related to the root directories in journal files. The journal files are then sent to the Replicas where the operations that were recorded in the journal are replayed on the replicated files.
- Scheduled mode - Synchronization of servers takes place at fixed times. There is no online replication in this mode, however, online changes made during synchronization are replicated.

CA XOsoft also enables you to assess the accurate bandwidth usage and compression ratio benchmarking that is needed for replication, without actually replicating data. When you select the Assessment mode, no replication occurs but statistics are gathered. A report is provided once the assessment process is completed.

How Recovery Works

When Master data is lost or corrupted for any reason, CA XOsoft can recover the data from any of the Replica servers participating in the scenario. The Restore Data option activates a synchronization process in the reverse direction: from Replica to Master.

When recovery is initiated, CA XOsoft Manager builds a temporary tree, containing a single branch. In this tree, the Replica becomes the source of data and the original Master becomes the target (i.e., the terminating Replica). Once the synchronization process completes, the CA XOsoft Manager reverts to the original replication scenario, and continues working.

Important! All file system activity must be halted on the Master host until the recovery process finishes.

How Data Rewind Works

Because replication continuously updates source data to another machine, a Replica always holds the same data as in the Master. In the case of data corruption, recovering files from the Replica will not help, since chances are high that data on the Replica is also corrupted. Data Rewind is a technology that allows repairing a corrupted file by rewinding it back in time as if it were a tape.

Data Rewind can be compared to the Undo feature of Microsoft Office where user actions can be cancelled, thus bringing the file to a previous state in time. Data Rewind is based on rewind journals that store I/O operation information that result in modified files. Using the rewind journal, it is possible to *undo* I/O operations, thus rewinding the file to a previous point in time, supposedly to a valid, non-corrupted state.

How Replication Suspension Works

At times it may be necessary to suspend updates on a Replica machine in order to perform system maintenance or some other form of processing that does not modify the replicated data there. It is not desirable to stop replication since this requires a full resynchronization afterward.

The replication suspension feature of CA XOsoft makes this possible. Replication may be suspended either manually or on a scheduled basis. During the suspension period, all changes are spooled on the Master or on the Replica located upstream of the suspended Replica. In other words, changes continue to be recorded for update on the suspended Replica, but are not actually transferred until replication is resumed. Once replication is resumed, the accumulated changes are transferred and applied without any need to perform a full resynchronization of the data.

How High Availability Works

CA XOsoft HA monitors all critical events, including global server failure and all database service failures, and either automatically or with a push of a button initiates a switchover.

If the Master server becomes unavailable, its activities can be switched over automatically to a remote site (Replica). The switchover, which is transparent to the user, includes immediate startup of a synchronized standby database, and redirecting all users to it in minimum time. All this is done without any need to reconfigure either clients or the network.

Depending on the requirements of the application being protected, redirection can be based on the following methods:

- Move IP (if the standby site is implemented within the same network segment)
- Redirect DNS, can be used on a local network or when the remote standby site is located on a different IP network (cross-network switchover)
- Switching the server hostname/NetBIOS name

Note: You can also apply user-defined scripts that add or replace the built-in redirection methods. Identify Network Traffic Direction scripts are required to fully support custom, or customized, redirection methods. Custom scripts or batch files are used to identify the active server. This script determines if the forward or backward scenario will run when the scenario is started. The script runs on both the Master and Replica: the one that returns zero is active. If both return zero, a conflict is reported.

CA XOsoft Modules

The main CA XOsoft system components are as follows:

- **CA XOsoft Control Service**
- **CA XOsoft Engine**
- **CA XOsoft Management Center** – consists of three components: **Overview Page**, **Manager**, and **Report Center.**
- **CA XOsoft CDP Repository** – consists of five components: **CDP Storage, CDP Web Server**, **CDP Support**, **CDP Admin** and **E-mail Retrieval**
- **CA XOsoft PowerShell**

Each of the CA XOsoft components is described in the following sections.

CA XOsoft Control Service

The CA XOsoft Control Service functions as the single-point-of-control of the CA XOsoft operation, and it contains the entire data of the existing scenarios. The Control Service communicates with both the Engines and the Managers. It is responsible for the management of all scenario-related-tasks, such as, creation, configuration, monitoring, and running of the scenarios.

The Control Service receives requests from the Managers, processes them, converts them to particular commands, and passes them on to the Engines. Then, the Control Service receives up-to-date data and events from the Engines, and sends back information and statistics about the scenario's state to the Manager.

The Control Service is also responsible for the authentication and authorization of users. It can also serve as a central point for CA XOsoft report handling and storage. The information and statistics that are accumulated by the Control Service can be presented to the user through the Overview Page, Manager, Report Center and PowerShell.

All the scenario files are kept on the server that runs the Control Service. If the Control Service is down, the scenario functioning will not be affected. However, for receiving information about the scenario's state, the Control Service must be active.

CA XOsoft Engine

The CA XOsoft Engine is a Windows service that must be running before any scenario can start. It is installed on every server participating in any given scenario, meaning the Master (source) and Replica (target) hosts. Each Engine supports both a Master and Replica functionality, for both Disaster Recovery and High Availability scenarios. It may participate in multiple scenarios and serve in a different role in each scenario. Engines can be installed either locally on each host at a time, or through a remote installer on numerous hosts at once.

Management Center

The CA XOsoft Management Center consists of three components, none of which requires any manual installation:

- **Overview Page** - a statistical overview of the Disaster Recovery and High Availability scenarios' state:

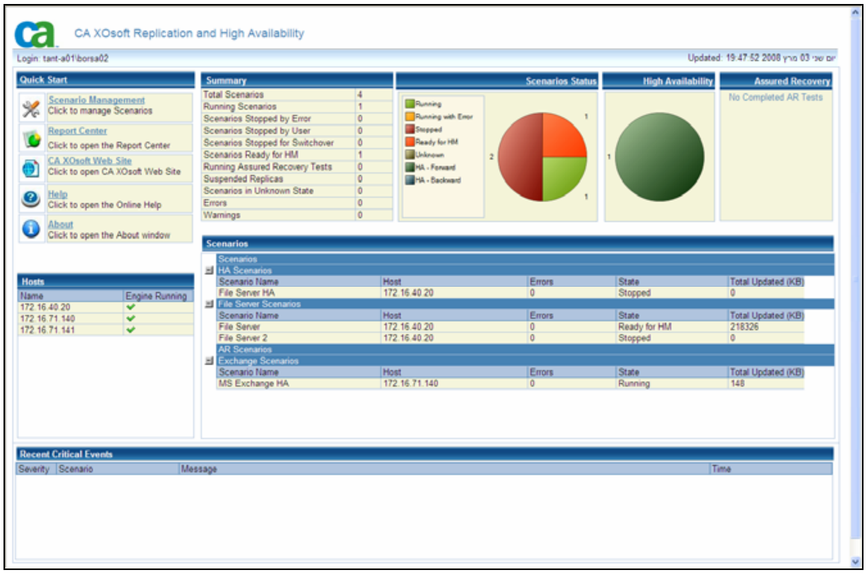

- **Manager** – a User Interface that enables the user to create, configure, manage and monitor scenarios. This is GUI application that is activated from the Overview Page by clicking the Scenario Management link:

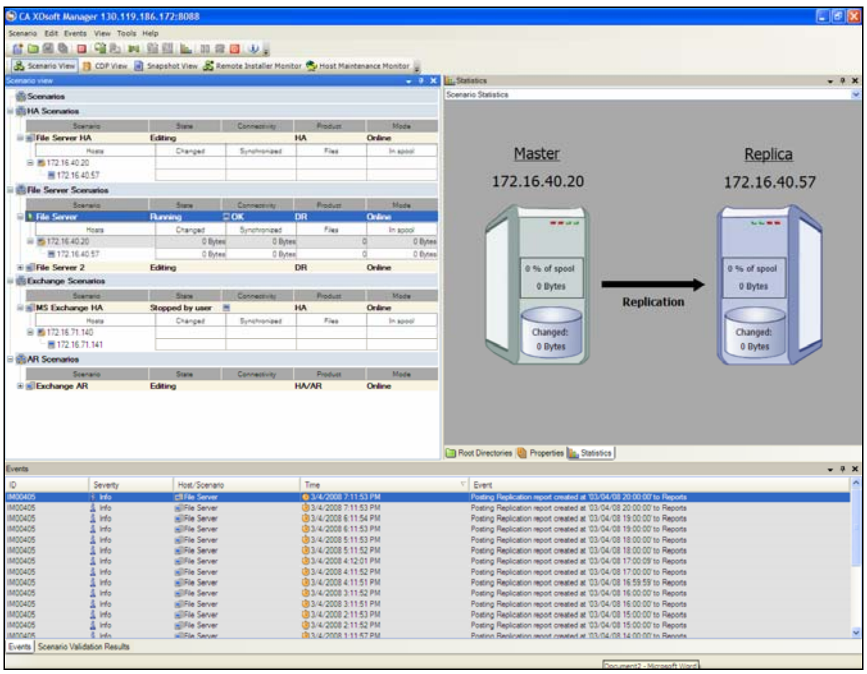

- **Report Center** – a User Interface that gathers all existing reports, along with information about the available reports per scenario. The user can decide where these reports will be stored, and for how long they will be displayed and saved in the Report Center:

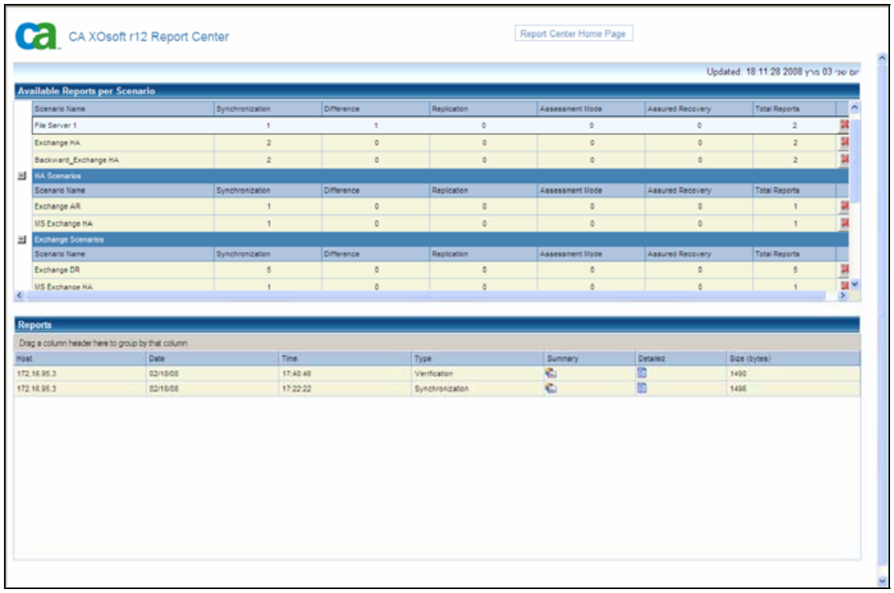

CDP Repository

The Continuous-Data-Protection (CDP) Repository module provides the ability to store deleted Outlook items, to search for certain items according to different criteria, and to retrieve them upon end-users requests. The types of Outlook items that can be retrieved are defined by the administrator, and they can include: e-mail messages, appointments, contacts, tasks, journal entries, notes and attachments. For more information about the CDP Repository module, see *CDP Repository* Components, page 226.

CA XOsoft PowerShell

The CA XOsoft PowerShell is offered as an alternative to users that do not want to manage the replication process using the CA XOsoft Manager graphic user interface. It enlarges and facilities the capabilities of the WANSync CLI that was provided in previous versions, and it supports both DR and HA operations.

The CA XOsoft PowerShell is a command-line shell and scripting environment that allows users to configure a replication scenario and control and monitor the replication process. All the scenarios that are managed by the CA XOsoft PowerShell look and operate exactly as the ones that are managed by the Manager, and they are automatically saved in the same default location: INSTALL_DIR/ws_scenarios.

CA XOsoft PowerShell is based on the standard Windows PowerShell™, which comes with a large set of built-in commands with a consistent interface. The CA XOsoft PowerShell component adds to this shell a number of scenario-related-commands, called snap-ins, which facilitates the scenario management.

CA XOsoft Deployment

The deployment of CA XOsoft components depends on the size of your IT enterprise network and your DR and HA needs. However, there are certain guidelines that you should follow when designing your Replication and High Availability environment and deploying CA XOsoft different components on a Windows platform. For information regarding an efficient deployment of CA XOsoft components, see the *CA XOsoft r12 Installation Guide*.

The following illustration shows a typical deployment of CA XOsoft components:

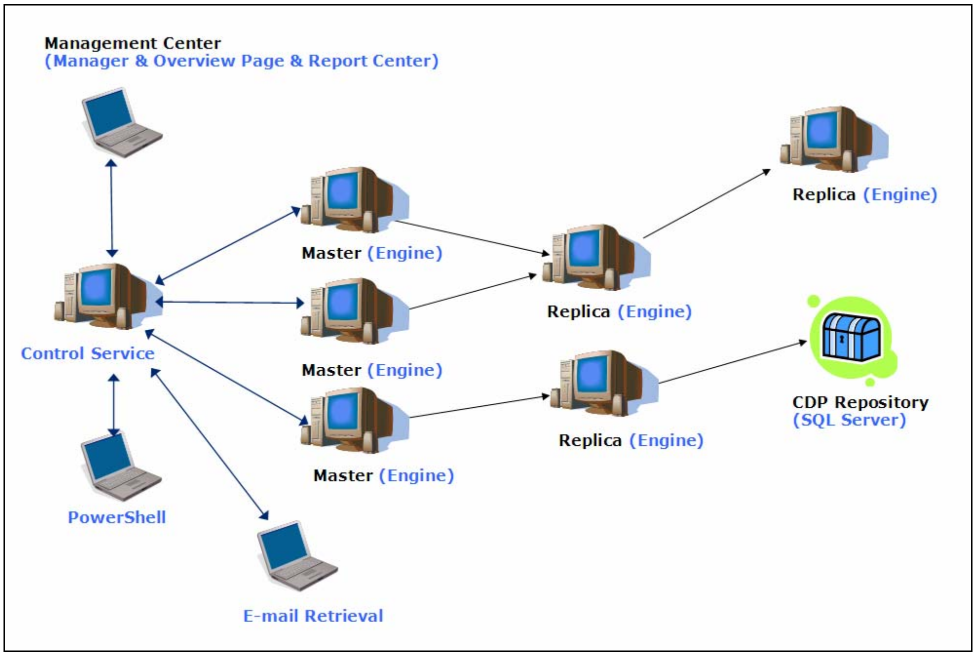

Chapter 2: Exploring CA XOsoft Manager

This chapter introduces CA XOsoft Manager and familiarizes you with its components and functionality. It instructs you how to log in to CA XOsoft Management Center and Manager, and describes the structure, menus, buttons and functions available in CA XOsoft Manager main window.

Logging In to the Management Center

CA XOsoft Management Center and Manager do not require any component or application installed in advance. It is based on a one-click-installation procedure that can be performed from any workstation that has a network connection and a Web browser. To log in, you will need your:

- Hostname/IP Address and Port Number of the server where the Control Service is installed.
- User Name, Password and Domain

To open CA XOsoft Manager:

1. Open Internet Explorer. On the **Address** box, enter the Control Service Host Name/IP Address and Port Number as follows: http://host_name:port_no/start_page.aspx

   Notes:

   - If you are opening the Management Center from the machine where the Control Service is installed, you can use the default parameters: http://localhost:8088/start_page.aspx
   - If you selected the **SSL Configuration** option during the installation of the Control Service**,** when you open the Overview page, you need to use the hostname of the Control Service machine (instead of its IP Address). Enter the Control Service Host Name and Port No. as follows:
     https://host_name:port_no/start_page.aspx

The **Login** dialog appears:

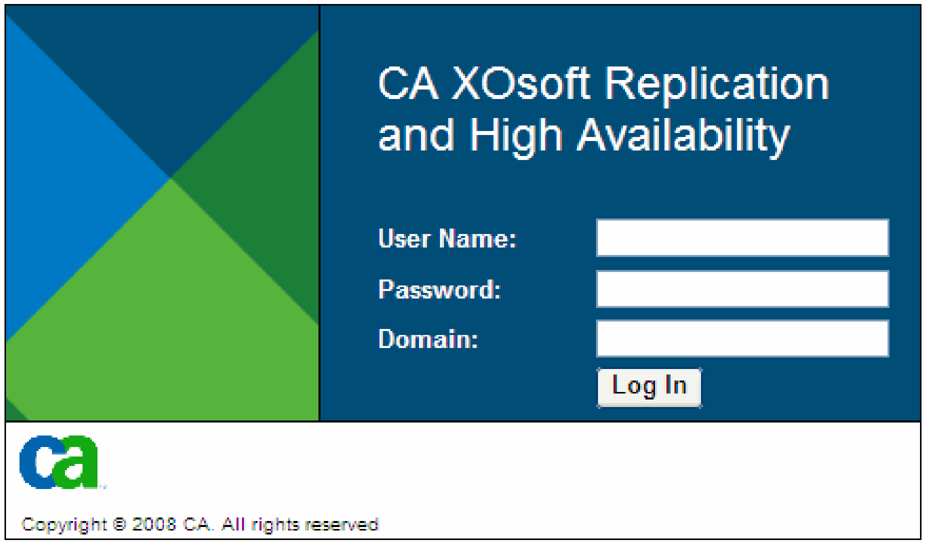

2. Enter your User Name, Password and Domain and click **Log In**.

   The **Overview page** appears:

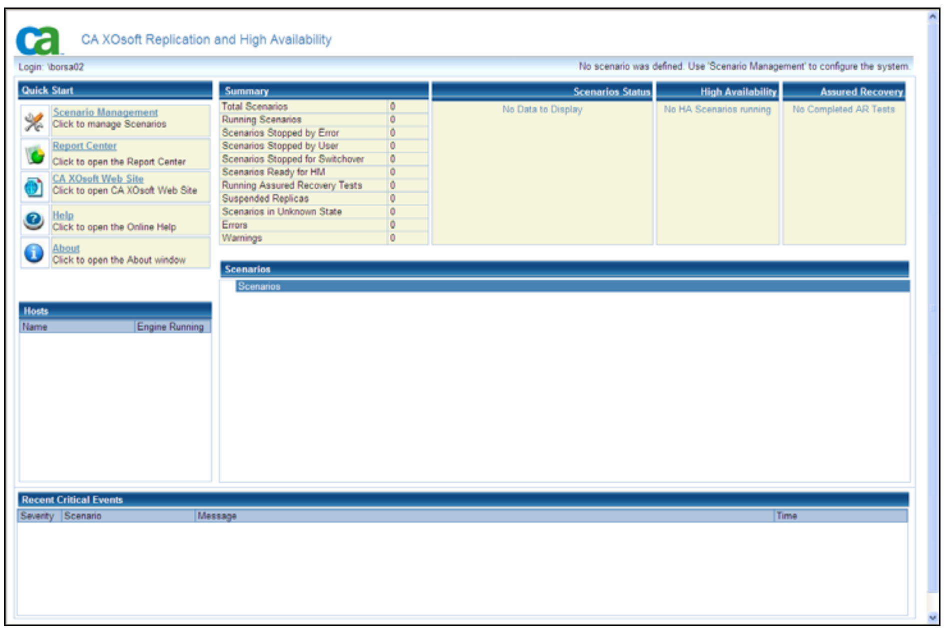

3. On the **Quick Start** toolbar on left, click the **Scenario Management** option.

   A progress bar appears, indicating that the Manager component is currently installed on the local machine.

4. Once the Manager installation is completed, the Manager appears:

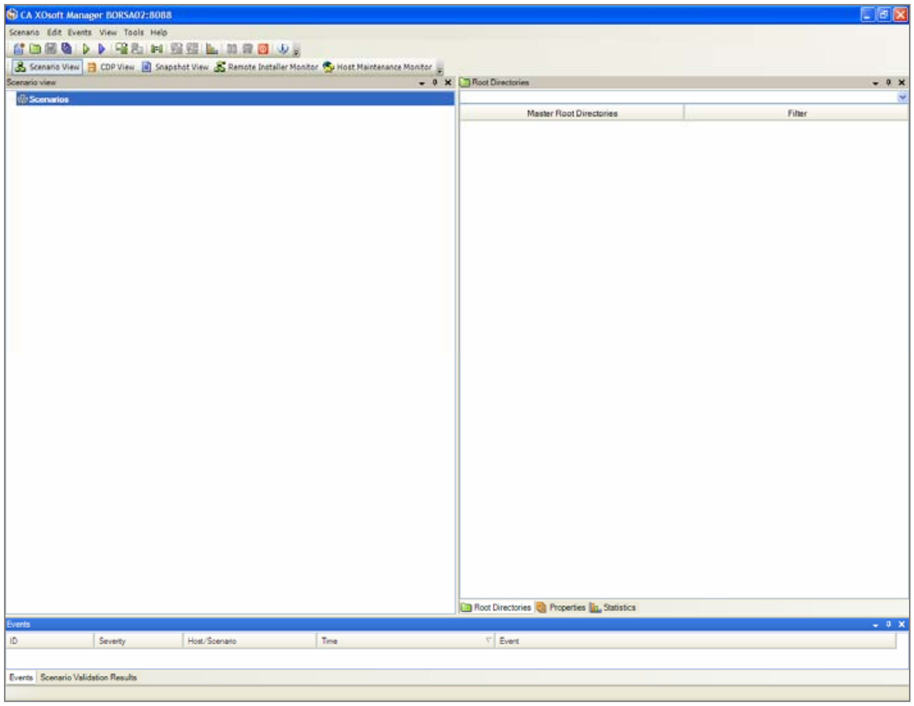

Exploring CA XOsoft Manager Screen

After logging in to the application, CA XOsoft Manager is displayed, enabling you to access all the Manager menus, toolbar functions and panes.

Unless a scenario exists, most of the user areas are blank. The following screen displays a Manager with active scenarios in it:

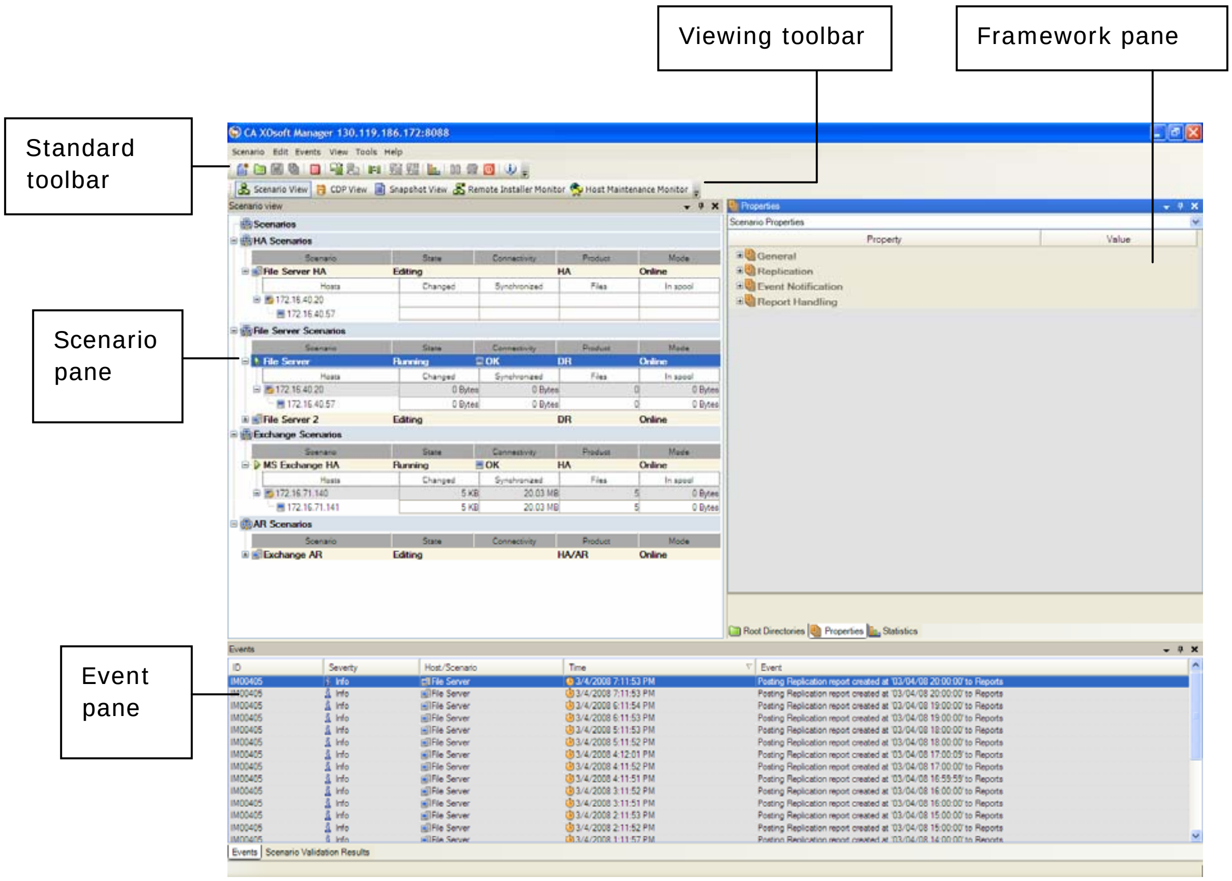

Note: Some of the panes and options are visible and enabled only with the appropriate product license.

The screen is divided into several areas:

- The application's name (CA XOsoft Manager) and the Control Service's connection details appear in the top left corner of the title bar; beneath it appear the menu line, the Standard toolbar and the Viewing toolbar:

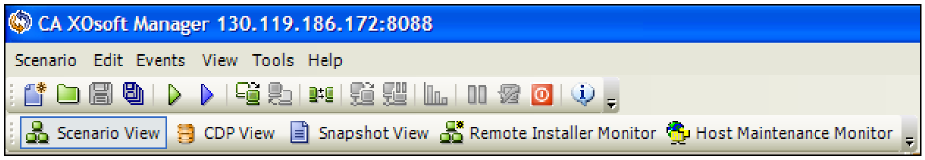

- Scenarios, including their replication tree, appear on the left pane:

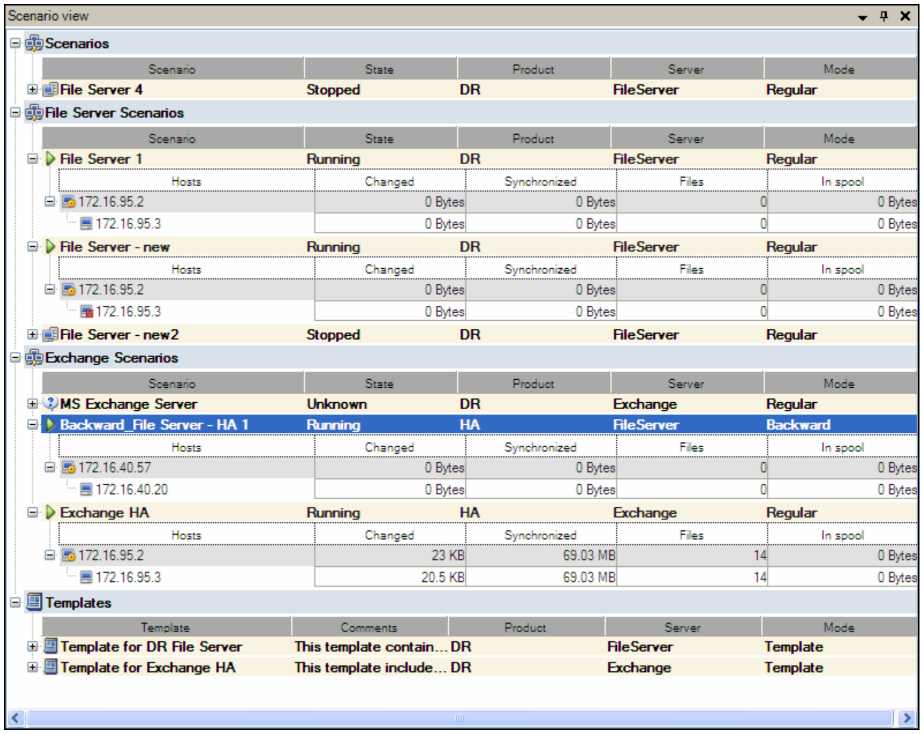

- The Framework pane appears on the right. In this pane, configurable property lists are displayed – the Scenario, Master, Replica, HA and Template properties. The displayed list depends on the item selected in the Scenario pane, or on the item chosen from the pane's drop-down list. The Framework pane also shows two, three or four tabs, depending on the selected CA XOsoft solution and on the scenario state. These tabs include:
  - Root Directories
  - Properties
  - High Availability Properties
  - Statistics

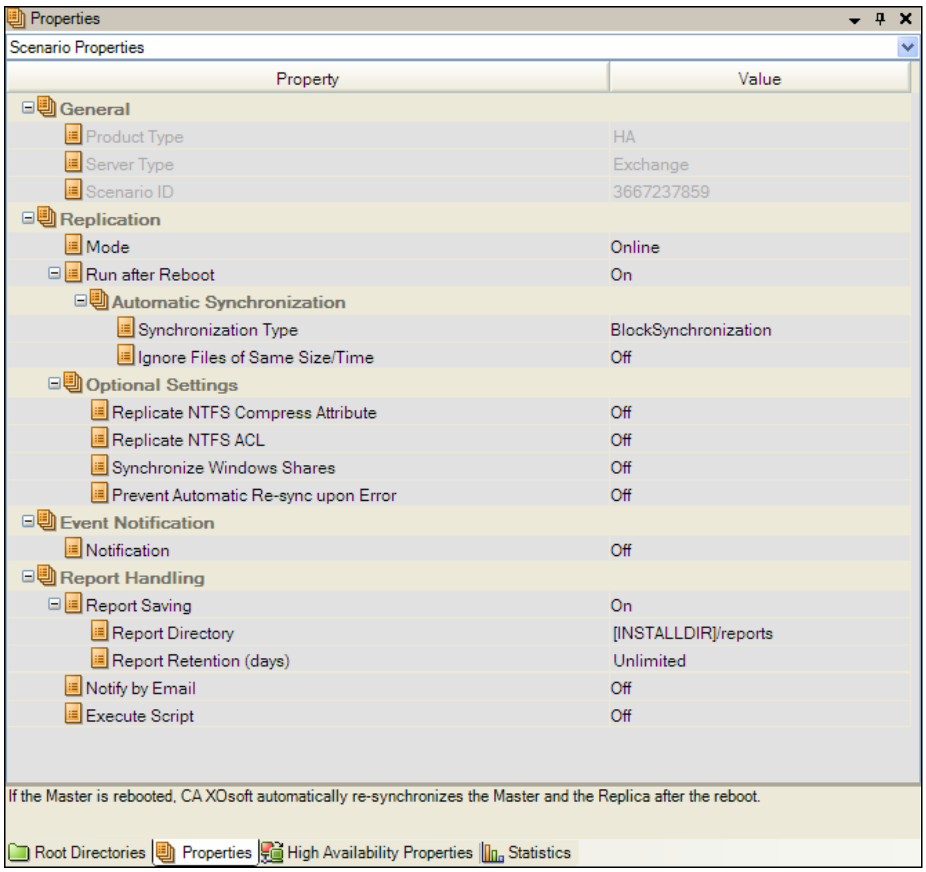

- The Events pane is below the screen's dividing line:

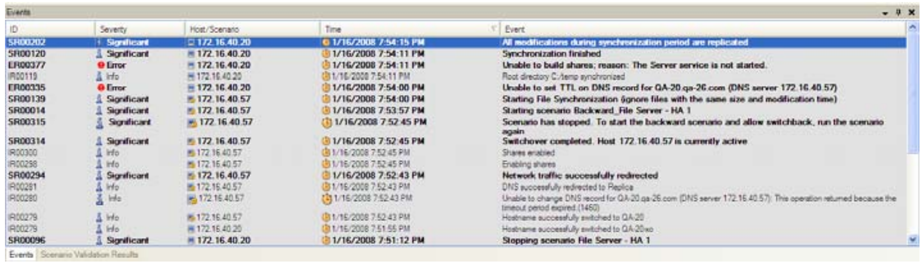

Note: The actual placement of the panes can vary, since they can be moved and resized. Also, the panes, including the toolbar and status bar, can be made invisible according to the selections in the **View** menu.

Viewing and Arranging the Manager Screen

CA XOsoft offers you different ways to get a good view of your work, depending on the task at hand. You can arrange your workspace according to your current needs.

Note: The last view setting you use is saved for the next session.

Viewing Options

The **View** menu contains all existing views. Selecting a view option from the menu will either open or close the indicated pane.

To reset your Manager layout

- From the **View** menu, select the **Reset** option.

  The original view settings are restored.

Customizing the Scenario Pane

The Scenario pane shows the current status for all scenarios in one pane, thus enabling you to monitor multiple scenarios at once. You can customize the way the scenario information columns are displayed.

To customize your scenario view:

1. From the **View** menu, select the **Customize Scenario View** option.

   The **Customize scenario view** dialog appears:

   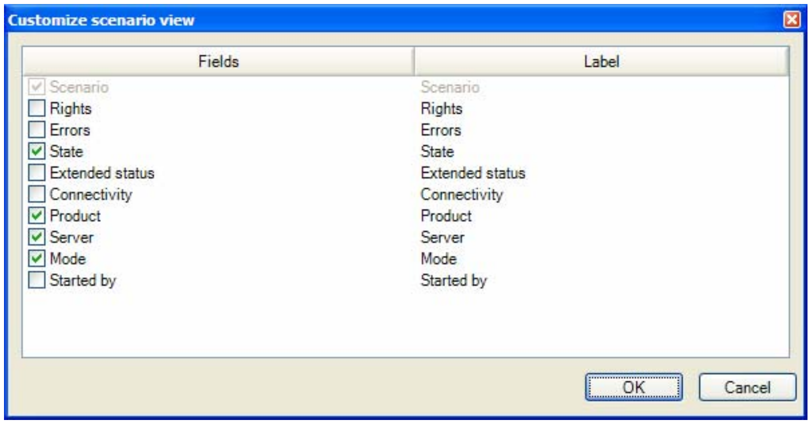

2. Select the fields you want to display in the Scenario pane and click **OK**.

   Note: The **Started by** field indicate the user that initiated the Run for the specific scenario.

   The fields you selected appear as columns in the Scenario pane.

Rearranging Panes

You can dock, stack, hide, show and float the Manager panes according to your needs. Once you are moving a pane, a built-in Pane Docking tool appears automatically. This tool allows you to quickly dock a pane on a specific location in the Manager window, and to stack panes on top of each other and present them as tabs. When a pane is undocked, it floats on top of the other panes.

Note: This MS tool is also called a "guide diamond".

To dock a pane

1. Click the pane's title bar and start dragging it. The Pane Docking tool appears:

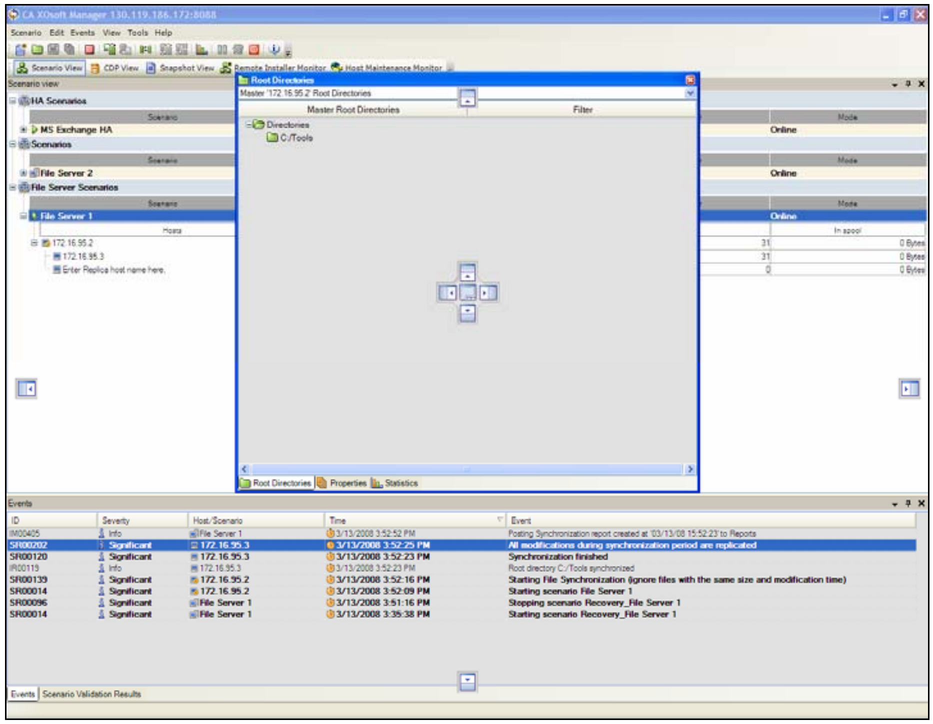

2. When the pane you are dragging reaches the position in which you want to dock it, move the pointer over the corresponding portion of the Pane Docking. You can use either the central arrows or the ones at the edge of the window.

   Note: To move a pane without snapping it into place, press **Ctrl** while dragging it.

3. To dock the pane in the indicated position, release the mouse button.

To stack a pane

1. Click the pane's title bar and start dragging it. The Pane Docking tool appears.
2. When the pane you are dragging reaches the docked pane on which you want to stack it, move the pointer to the center box of the Pane Docking tool.
3. To stack the pane, release the mouse button. The pane is now accessible by clicking its tab.

To hide a pane

1. Right-click the pane's title bar. A pop-up menu appears.
2. Select one of the following:
   - **Hide** – to completely hide the pane.
   - **Auto Hide** – to hide the pane while you are working on other panes, and to automatically return it when the pointer is placed on its tab.

To bring back hidden panes

- From the **View** menu, select the **Reset** option.

  All hidden panes are restored.

Toolbars

CA XOsoft provides you two toolbars to facilities your work: the Standard toolbar and the Viewing toolbar.

Standard Toolbar

The Standard toolbar buttons provide quick access to the most commonly used functions in the CA XOsoft Manager. The following list contains a brief description of each toolbar option:

New

Create a new scenario using the Scenario Creation Wizard. See *Creating a Scenario*, page 42.

Group

Create a new scenario group. See *Creating a Scenario Group*, page 41.

Save

Save a selected scenario. See *Save Scenarios*, page 113.

Save All

Save at once all existing scenarios. See *Save Scenarios*, page 113.

Run

Run the selected scenario to start the replication process. See *Initiate Replication*, page 61.

Run (Assessment mode)

Run the selected scenario in Assessment Mode. See *How Replication Works*, page 18.

Synchronize

Activate the synchronization process (whether replication is running or not). See *Synchronization*, page 68.

Restore Data

Recover lost or corrupted Master data from any Replica by activating a synchronization process in the reverse direction. See *Data Recovery*, page 145.

Difference Report

Generate a Difference Report, which displays the difference between a Master and its Replica(s) at a certain point in time. The comparison is performed using the same algorithms that are used in the synchronization process, but no data is transferred. See *Difference Report*, page 92.

Perform Switchover

[For HA only] Switch the active and passive roles between the Master and Replica servers following their current status. See *Switchover*, page 168.

Suspend Is Alive Check

[For HA only] Suspend the Is Alive check that verifies that the active sever is operational. See *Is Alive*, page 184.

Refresh Statistics

Update the scenario state information and live statistics display. See *Refresh Statistics Display Manually*, page 81.

Suspend Replication

Suspend replication updates on the Replica host in order to perform system maintenance or some other form of processing that does not modify the replicated data there. Changes continue to be recorded for update on the suspended Replica, but are not actually transferred until replication is resumed. You cannot suspend replication during synchronization. See *Suspend Replication*, page 70.

Replica Integrity Testing

Perform Assured Recovery test in a non-scheduled mode. See *Performing Assured Recovery Test in a Non-Scheduled Mode*, page 206.

Launch Host Maintenance

Prepare a node in your replicated system to planned maintenance procedures, while avoiding resynchronization once these procedures are completed. See *Host Maintenance*, page 116.

License Registration

CA XOsoft licensing policy is based on a combination of several parameters, which include: the operating systems involved, the required solution, the supported application and database servers, the number of participating hosts, and the additional modules - Assured Recovery and CDP Repository. The license key that is generated for you is therefore tailored to your exact needs.

After logging in for the first time, or if your old license has expired, you need to register CA XOsoft product using your license key. To register the product, you need to open CA XOsoft Manager, which does not depend on the existence of a valid registration key. Once the Manager opens, a License Warning message appears, prompting you to register the product. A License Warning message also appears when your license is about to expire during the next 14 days.

When you are creating a scenario, some of the options might be disabled following the terms of your license. However, you can create as many scenarios as you wish, since the validity of your license key is first checked when you try to run a specific scenario. Only when you click the **Run** button, the system checks whether you are allowed to run the selected scenario according to your license key. If the system determines that you do not have the required license for running this scenario, the scenario will not run and a message will appear on the Event pane informing you of the type of license you need.

To register CA XOsoft using the license key

1. Open the Manager. The **Welcome** message appears. Then, a **License Warning** message appears informing you that your product is not registered and prompting you to register it:

2. Click **OK** to close the message. Then, open the **Help** menu and select the **Register** option.

   The **Register CA XOsoft Replication and High Availability** dialog appears:

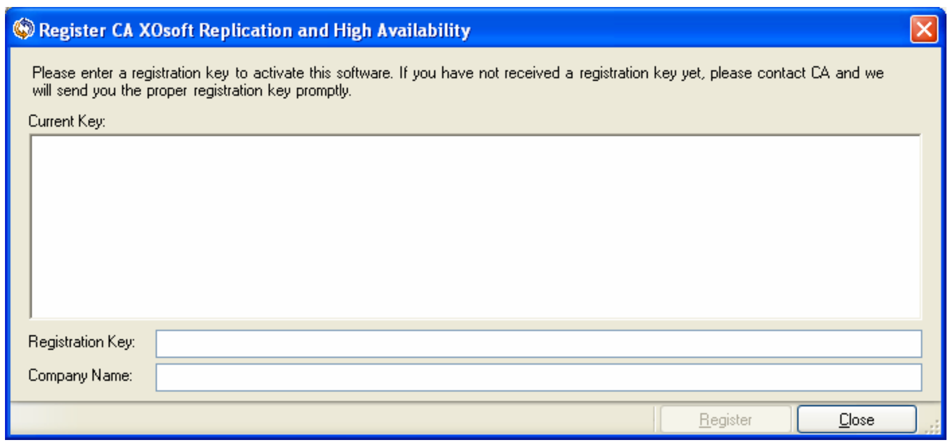

3. Enter the following information:
   - In the **Registration Key** box – enter your registration key.
   - [Optional] **Company Name** box – enter your company name
4. Click the **Register** button to register the product and close the dialog.

   Now, You can now start working with the CA XOsoft Manager according to your license permissions.

Chapter 3: Creating Replication Scenarios

This chapter describes how to create the following: a scenario group, a scenario using the Scenario Creation Wizard and a template, and a scenario template, (and how to run and stop a scenario).

Creating a Scenario Group

Each scenario in the system is assigned to a certain scenario group. By default, there is one scenario group called ***Scenarios***. You can use this group for all the scenarios you create, or you can add new groups to organize your scenarios according to a common denominator. These scenario groups will be displayed on both the Manager and the Overview Page. They can help you to better manage and monitor your scenario operation.

There are two ways to create a scenario group:

- During the creation of a new scenario, through the **Scenario Creation Wizard** (see page 43).
- Before the scenario creation, through the **New Group** option, as described below.

Note: We recommend planning and creating the scenario groups you want to use in advance. After you assign a scenario to a certain group, you can not move it to another group.

To create a new scenario group:

1. On the Manager, select from the **Scenario** menu the **New group** option, or click the **New group**  button on the Standard toolbar.

   A **New group** folder is added to the Scenario pane:

   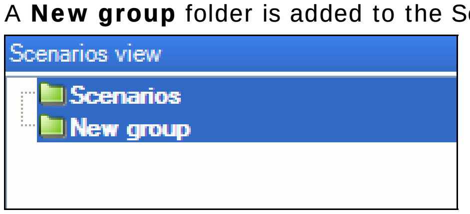

2. You can change the group's name by right-clicking it and selecting **Rename** from the pop-up menu, or by double-clicking the current name and entering a new name.

   The new group name appears on the following places: the Scenario pane, the **Group** drop-down list in the Scenario Creation Wizard, and the Overview Page.

   Note: When no scenario is defined, empty scenario groups will not appear on the Overview Page.

Creating a Scenario

CA XOsoft creates and maintains replications in the context of user-defined scenarios. A scenario is the basic unit of the CA XOsoft operation and it consists of a definition set that include:

- The type of application or database server to be protected.
- The type of data protection solution.
- Special tasks, such as Integrity Testing for Assured Recovery.
- The connection details of the Master and Replica hosts.
- The directories, sub-directories, databases and files that will be replicated and their location on the Master and the Replica.
- Configurable properties of the scenario and the Master and Replica hosts, which affect different settings and operations, such as, synchronization method, replication mode, spool size, report and event handling rules, and more.
- Recovery and Switchover parameters.

Each scenario defines a replication tree that set the flow of information from the Master server to any number of designated Replicas. It establishes the data recovery procedure, and, if applicable, the switchover parameters. You can configure, add or remove servers from a scenario and select or modify directories. This enables easy, complete control of the replication process over any network, large or small. Each scenario is saved as an XML file.

Note: This section demonstrates the configuration of a generic File Server Disaster Recovery scenario. For more detailed instructions involving High Availability scenarios or scenarios tailored to specific applications such as Exchange or SQL, see the appropriate CA XOsoft Operations Guide.

To create a new scenario

1. Open the CA XOsoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New** button on the Standard toolbar.

   The **Scenario Creation Wizard** appears:

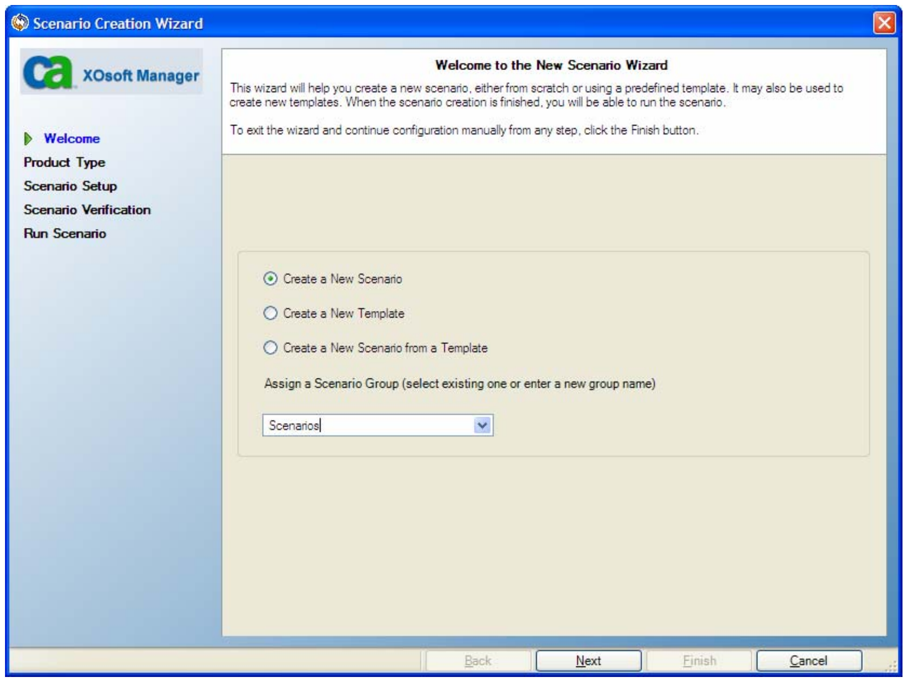

2. Select the required scenario options, as follows:

   - Select the **Create a New Scenario** option button.
   - From the **Group** drop-down list, select the group to which you want to assign the new scenario, or enter a name for a new scenario group.

3. Click **Next**. The **Select Server and Product Type** page is displayed:

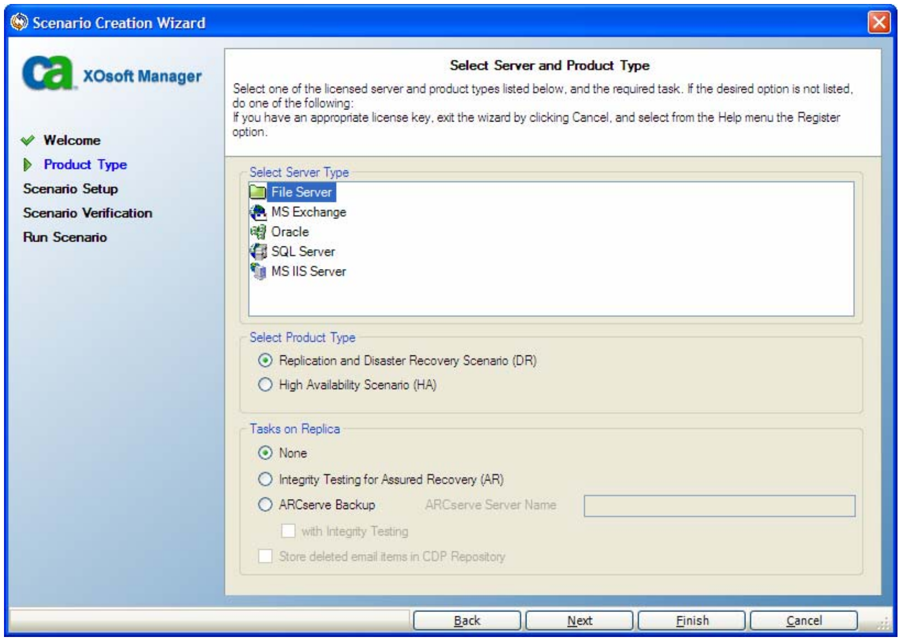

A list of available applications and scenario types is presented.

Note: The list of available applications depends on the licenses applied**.**

Select the required scenario options, as follows:

- From the **Select Server Type** list, select the type of server for which you want to create the scenario.
- From the **Select Product Type** options, select **Replication and Disaster Recovery Scenario**.
- [Optional – as license is needed] From the **Tasks on Replica** options, select the tasks you want to implement in this scenario.

Notes:

- To learn more about Assured Recovery and how it can help ensure success through automatic testing, refer to the About Assured Recovery, page 193.
- To learn more about ARCserve Backup and CA XOsoft r12 integration see the *CA ARCserve Backup for Windows CA XOsoft Integration Guide.*
- To learn more about the CDP Repository, see *Using the* CDP Repository, page 225.

4. Click **Next**. The **Master and Replica Hosts** page is displayed:

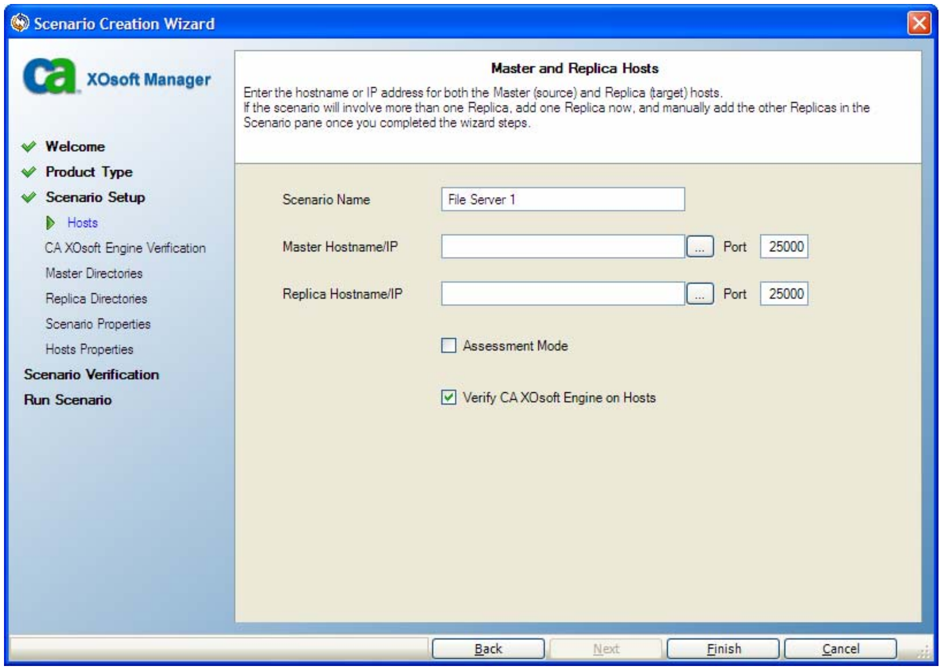

5. Enter the following information:

   - In the **Scenario Name** box – accept the default name or enter a new name for the scenario. When entering a name, choose a unique name, since you cannot use the same name for more than one scenario.
   - In the **Master** and **Replica Hostname/IP** boxes - enter the hostname or IP address of the Master (source) and Replica (target) servers, or use the **Browse** buttons to find them.
   - In the **Port** boxes: accept the default port no. (25000) or enter a new port numbers for the Master and Replica.

   Note: If you want to include more than one Replica in the scenario, enter here the details of the first or most upstream Replica. After you finished the scenario creation, manually enter the other Replicas, as described on *Add Additional Replica Servers*, page 97.

6. [Optional] Select the **Assessment Mode** check box, if you want to gather statistics about the accurate bandwidth usage and compression ratio benchmarking without actually replicating data. If you select this option, no replication occurs, but a report is provided once the assessment process is completed.

7. [Optional] Select the **Verify CA XOsoft Engine on Hosts** check box, if you want the system to verify whether Engines are installed and running on the Master and Replica hosts you specified in this page. If Engines are not installed on the selected hosts, you can use this option to remotely install the Engines on one or both hosts.

8. After you selected the desired options, click **Next**.

   Note: If you did not select the **Verify CA XOsoft Engine on Hosts** option, move to step 11, page 48.

   The **Host Verification** page appears:

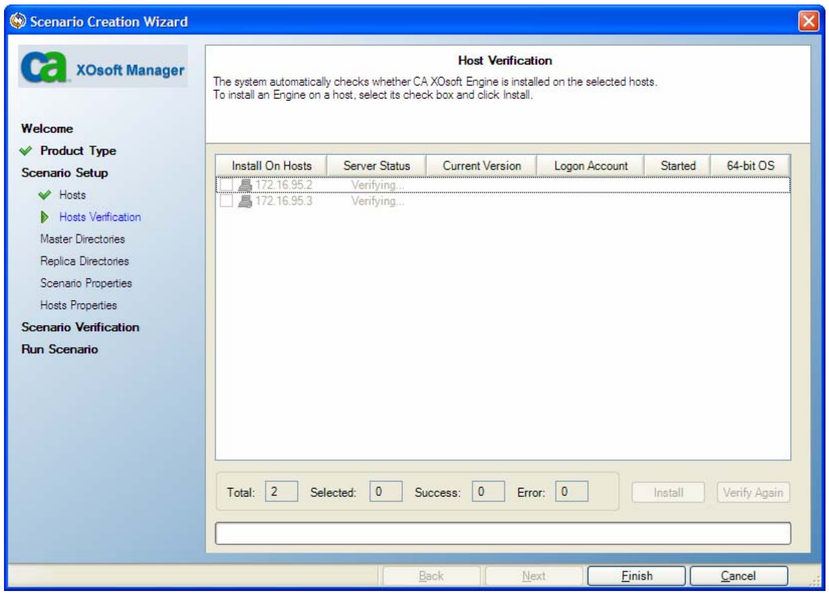

CA XOsoft verifies the existence and connectivity of the Master and Replica hosts you selected on the previous page. Once the connections are verified, the system checks whether an Engine is installed on each host.

Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the remote hosts, the **Server Status** is reported as **Not connected**. Then, another **User credentials** dialog appears, asking you to enter log on account details for each selected host:

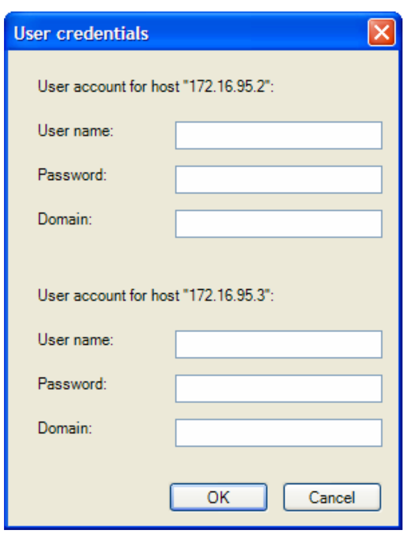

After you enter the details, the system repeats the verification process.

9. Check whether an Engine is installed on the selected hosts using the **Current Version** column:

   - If an **Installed** indication appears under the **Server Status** column in both rows, you can move to the next page.
   - If an **Installed** indication appears, but the version number is different than the version number of the Control Service you are using, you can reinstall the current version by clicking the **Install** button.
   - If a **Not Installed** indication appears, you need to install the Engine. Click the **Install** button to remotely install the Engine on the selected host.

     Note: You can install the Engine on both hosts at once. To perform this, select the check boxes of both hosts, and click the **Install** button.

10. Once you click the **Install** button, the **CA XOsoft Engine service log on account** dialog appears:

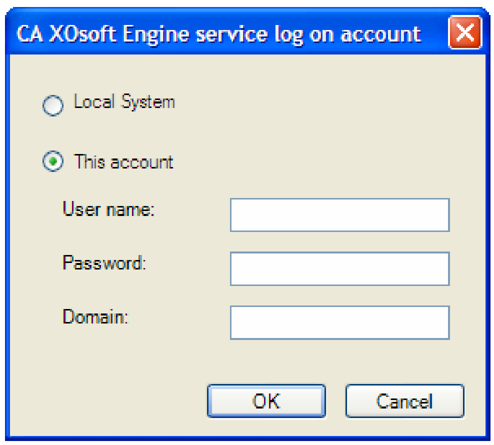

Enter the required information according to following guidelines:

- For Disaster Recovery purposes – it is sufficient to be a Local Administrator (Local System).
- For clusters (including DR scenarios) – you need to run under the same account as the Cluster Service account.
- For HA – you need to run under account with a Domain Administrative privileges.
- For [Exchange] CDP DR- you need to run under account with a Domain Administrative privileges.

Wait until the installation is complete, and the Engine's version number appears in the **Current Version** column:

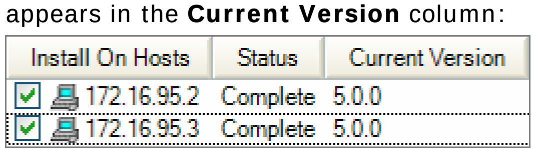

Install On Hosts	Status	Current Version
172.16.95.2	Complete	5.0.0
172.16.95.3	Complete	5.0.0

11. Click **Next**. The **Master Root Directories** page appears:

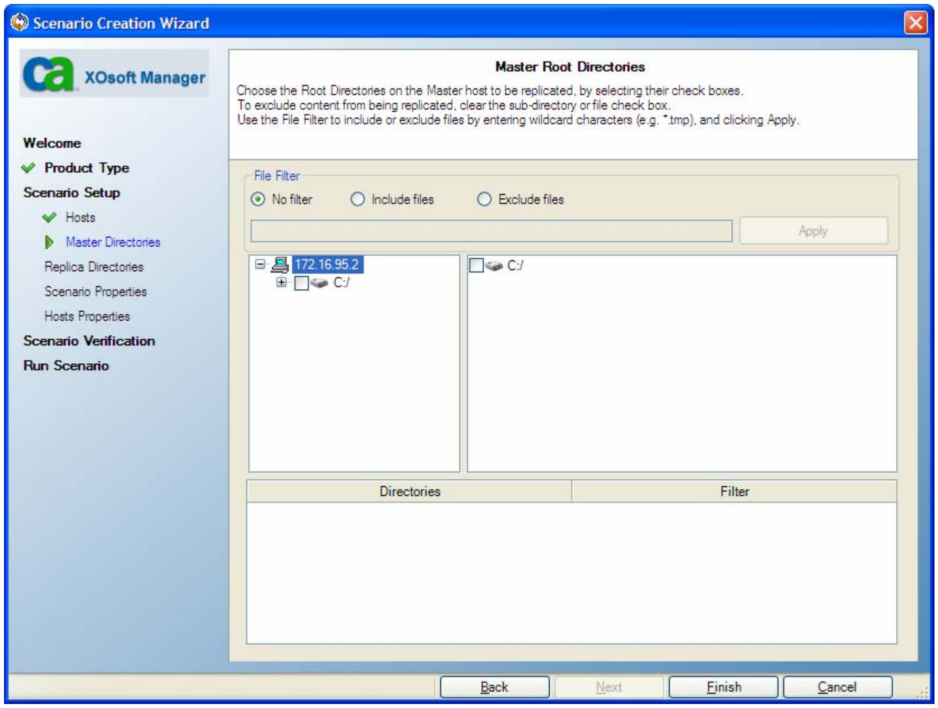

CA XOsoft displays the directories and files that are on the Master server. These directories and files are the data that can be replicated and protected. CA XOsoft automatically aggregates data that has a common path into one directory.

12. Choose the directories and files you want to replicate from the Master to the Replica by selecting their check boxes. You can exclude folders and files from replication by clearing their check boxes.

    Selecting a drive or a directory on the left side, displays its content on the right side:

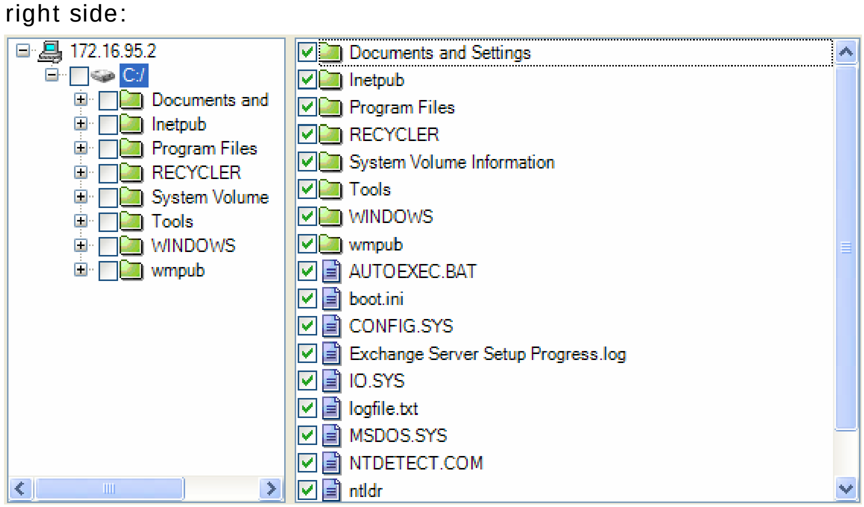

Notes:

- You can use wildcards to filter the data that will be replicated by selecting the appropriate option button on the **File Filter** section, and entering the expression in the box below (for more information, see *Filter Master Directory Files*, page 102):

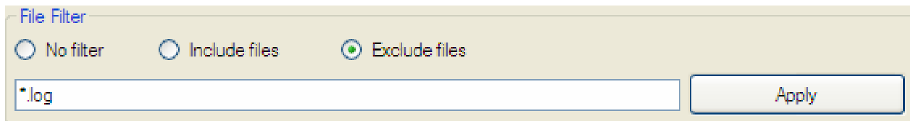

- Replication of mount points will succeed, only if those were added to the Master before the Engine began to run. If you included the mount points in the Master root directories when the Engine was already running, no error is reported but the replication does not start. In this case, you need to restart the Engine to initiate replication.

13. After defining the data to be replicated, click **Next**.

    The **Replica Root Directories** page is displayed:

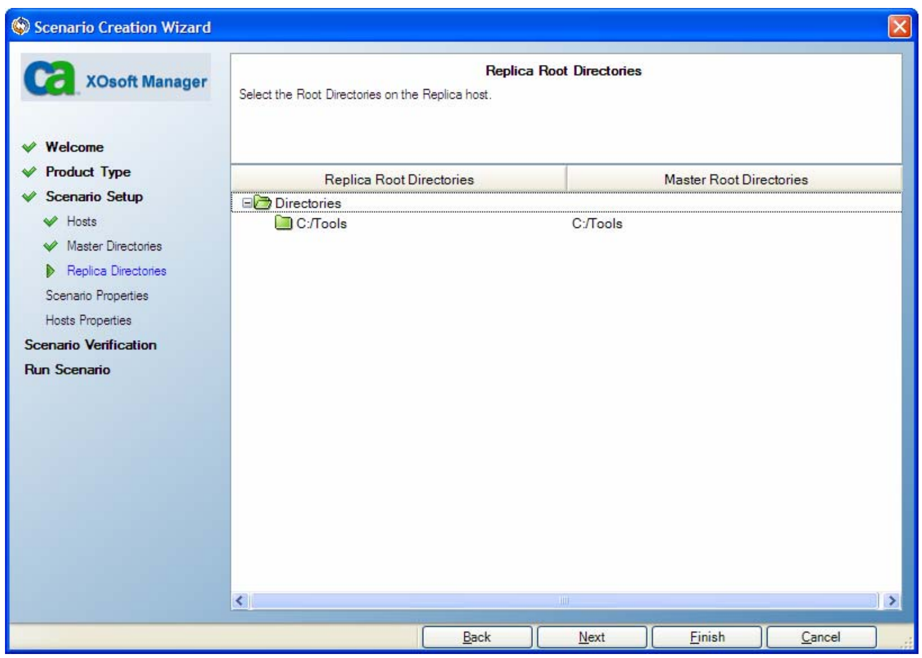

In this page you select the directories on the Replica where the replicated data will be stored.

Important! The Scenario Creation Wizard automatically configures the Replica root directories to be the same as the Master root directories. If you want to keep this configuration, ensure that your Replica server has the same drive letters as the Master server, and that the selected directories on the Replica do not contain data you want to save. You can change the default configuration at a later stage, as described on *Select Replica Root Directories*, page 111.

14. To change the Replica root directories, double-click the specified directories path. The **Browse and Select Replica Directory** appears:

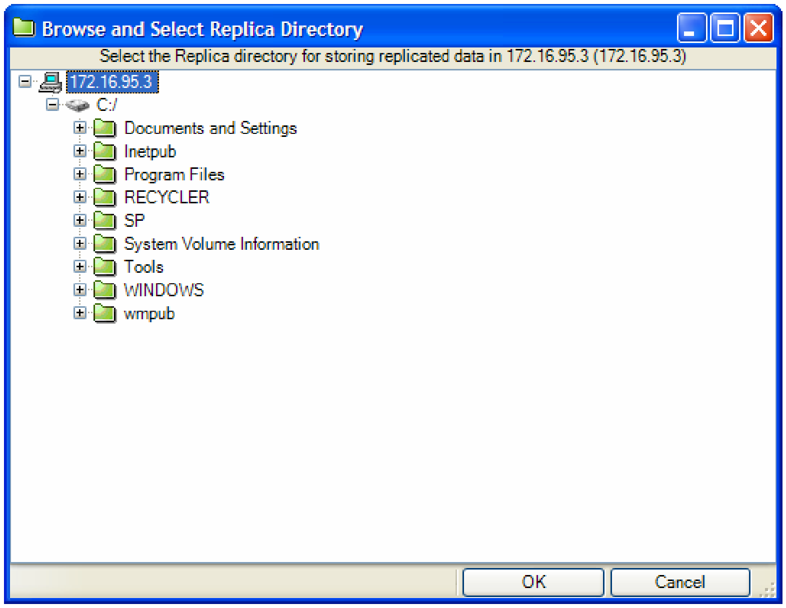

15. Select the directory on the Replica in which the replicated data will be stored, and click **OK**.

    You return to the **Replica Root Directories** page.

    Note: You can manually change the directory you selected for storing the replicated data, by clicking the selected directory name and entering a new directory. If you are entering a directory name that does not exist on the Replica, CA XOsoft creates it automatically:

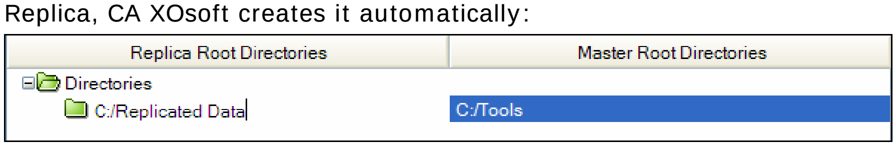

16. After defining the storage location of the replicated data, click **Next**.

The **Scenario Properties** page appears:

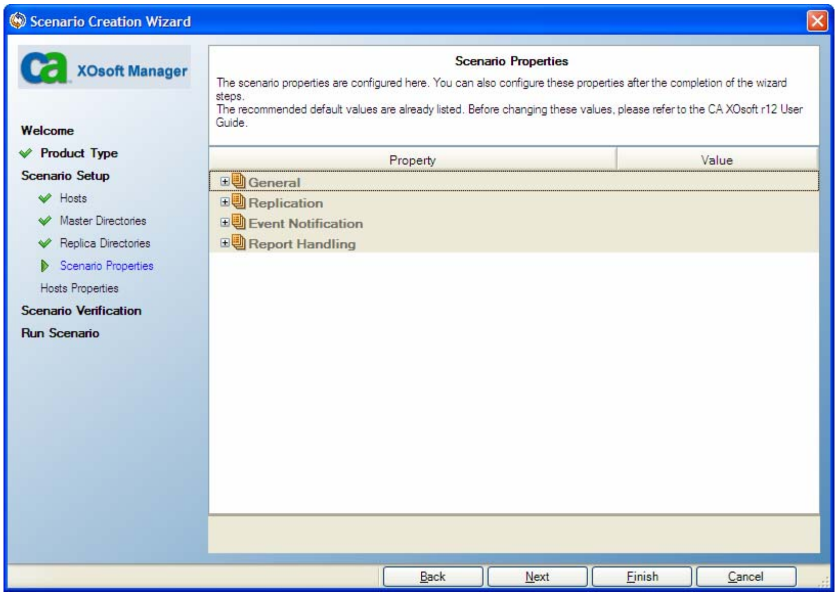

The **Scenario Properties** page enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

If you want to configure the scenario properties at this stage, refer to *Understanding Scenario* Properties, page 123. To configure the scenario properties at a later stage, refer to *Configuring Scenario Properties*, page 121.

17. Click **Next**. The **Master and Replica** Properties page appears:

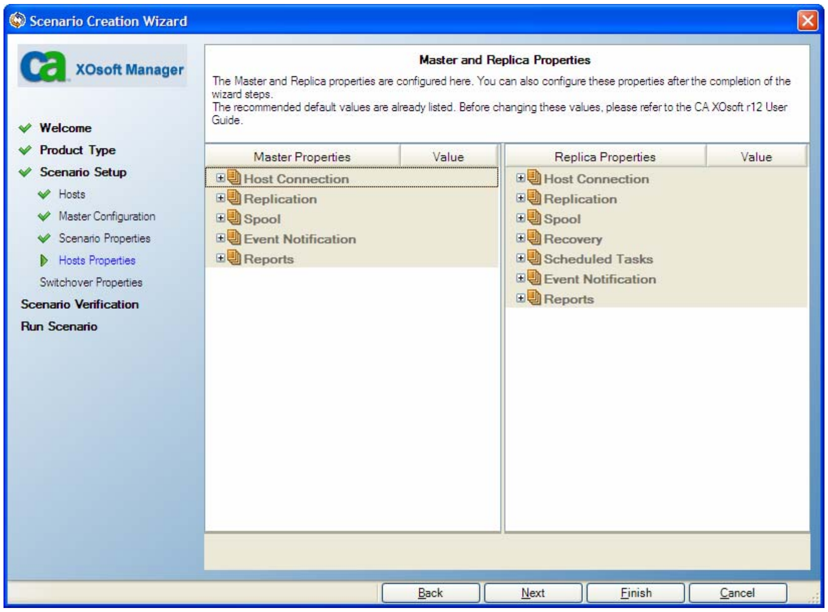

The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

If you want to configure the Master and Replica properties at this stage, refer to *Setting Master and Replica Properties*, page 133. To configure the Master and Replica properties at a later stage, refer to *Configuring Master or Replica Server Properties*, page 133.

Note: You can modify all the settings in this pane after the scenario is created. However, before changing any Spool properties (which can be configured here), review the *Spool* info on page 136 for configuration details.

18. Once you are satisfied with the Master and Replica properties, click **Next**.

    CA XOsoft verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful Disaster Recovery. Once the verification is completed, the **Scenario Verification** page appears:

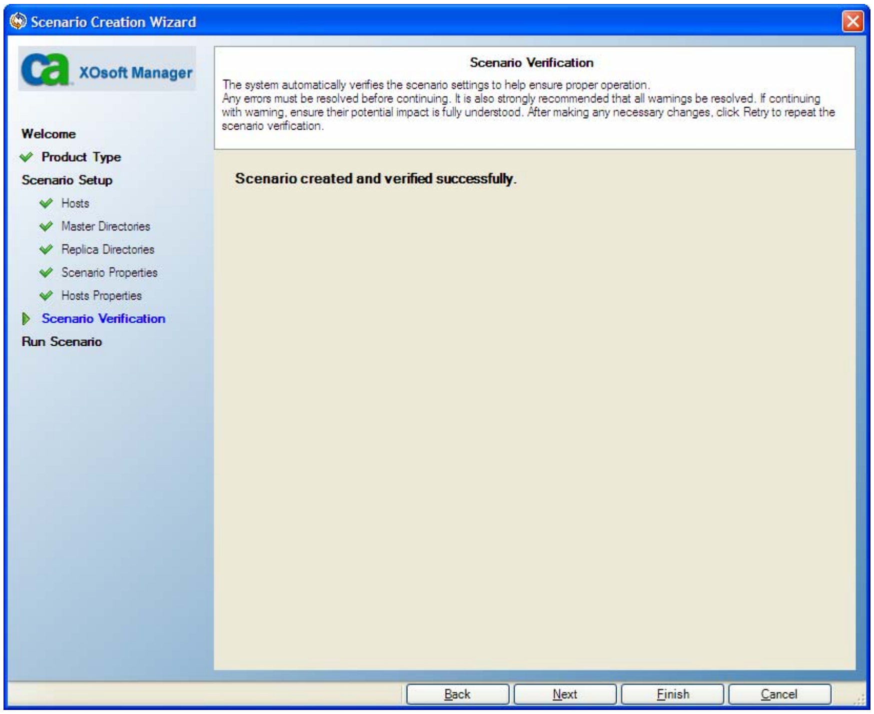

Note: Although the software allows you to continue with warnings, it is not recommended to do so. Resolve any warning situations before continuing to ensure proper operation of the application.

19. If the scenario is verified successfully, click **Next**.

The **Scenario Run** page appears:

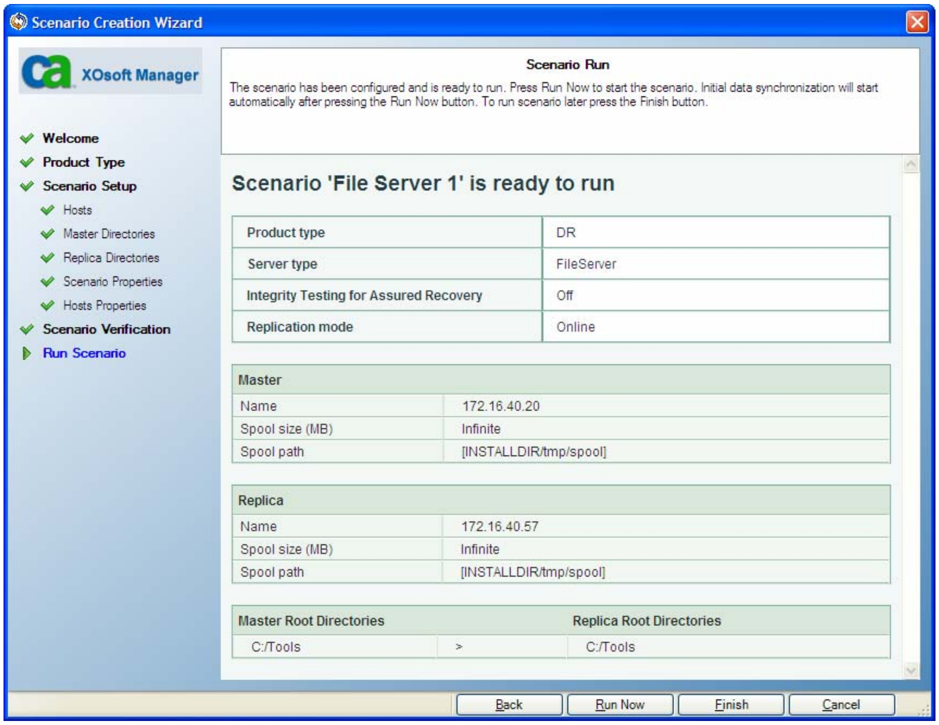

20. After the scenario is verified, you are prompted to run it. Running the scenario starts the data synchronization process.

    - To finish the scenario creation and run it later, select **Finish**.
    - To run the scenario now, click **Run Now**.

      The synchronization process starts.

21. Synchronization may take a while depending on database size and network bandwidth between the Master and Replica. You will receive the following message in the Event pane when synchronization is complete: *All modifications during synchronization period are replicated.* At this point, real-time replication is operational and the Disaster Recovery solution is installed and active.

22. By default, a synchronization report is generated:

    Note: For more information about opening a report, see page 63.

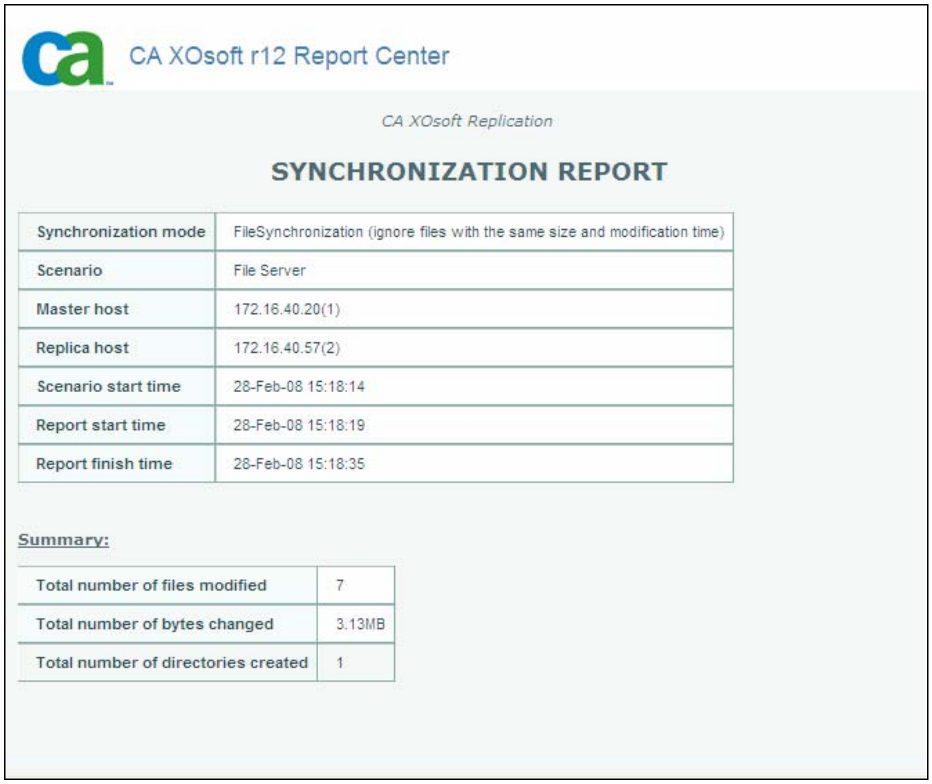

CA XOsoft r12 Report Center

CA XOsoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server
Master host	172.16.40.20(1)
Replica host	172.16.40.57(2)
Scenario start time	28-Feb-08 15:18:14
Report start time	28-Feb-08 15:18:19
Report finish time	28-Feb-08 15:18:35

Summary:

Total number of files modified	7
Total number of bytes changed	3.13MB
Total number of directories created	1

Using Templates

Templates are a powerful facility for customizing CA XOsoft for your own environment. Many applications allow the default values of individual parameters to be modified. Examples might be the default font to be used in new documents in Microsoft Word, or the default signature for new messages in a mail client application. Templates take this idea one step further.

Rather than provide a method to modify a single, global default value, templates offer the ability to create an entire scenario that can be used as the starting point for new scenarios in the future. These special template scenarios contain all the parameters of an actual scenario and all of them may be modified except those that clearly apply to a specific, individual scenario (such as the host names of the Master and Replica servers).

The second important advantage of templates over a set of global defaults is that they allow different sets of defaults to be created for different types of scenarios. For example, the default values appropriate for your High Availability Exchange scenarios are not identical to those for your DR File Server scenarios. With templates, you can create default settings and maintain them individually for each type of scenario your IT environment needs.

Create a Template

Templates are simple to create and use, and their creation process is basically similar to the creation of a new scenario. However, because a template is not associated with any actual servers, some values cannot be entered, such as the host names or IP addresses of the Master and Replica servers. Also, while default folder paths can be entered on the Directories tab, they must be typed explicitly rather than entered through a file browser.

All the templates are automatically stored in the **Templates** folder on the Scenario pane. This folder does not appear in the Scenario pane until at least one template is created.

To create a new template

1. Open the Scenario Creation Wizard by clicking the **New** button on the Standard toolbar, or selecting **New** from the **Scenario** menu.

   The **Scenario Creation Wizard** appears:

   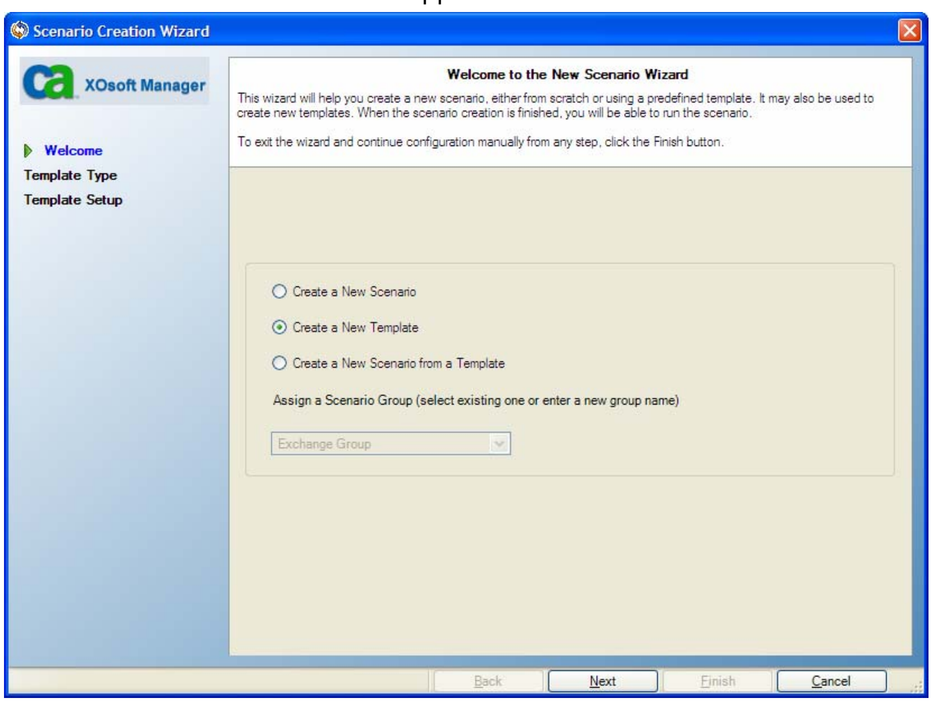

2. Select the **Create a New Template** option button, and click **Next**.

   The **Select Server and Product Type** page appears.

3. Select the required scenario options (for more information see page 44), and click **Next**.

The **Template Name and Comments** page appears:

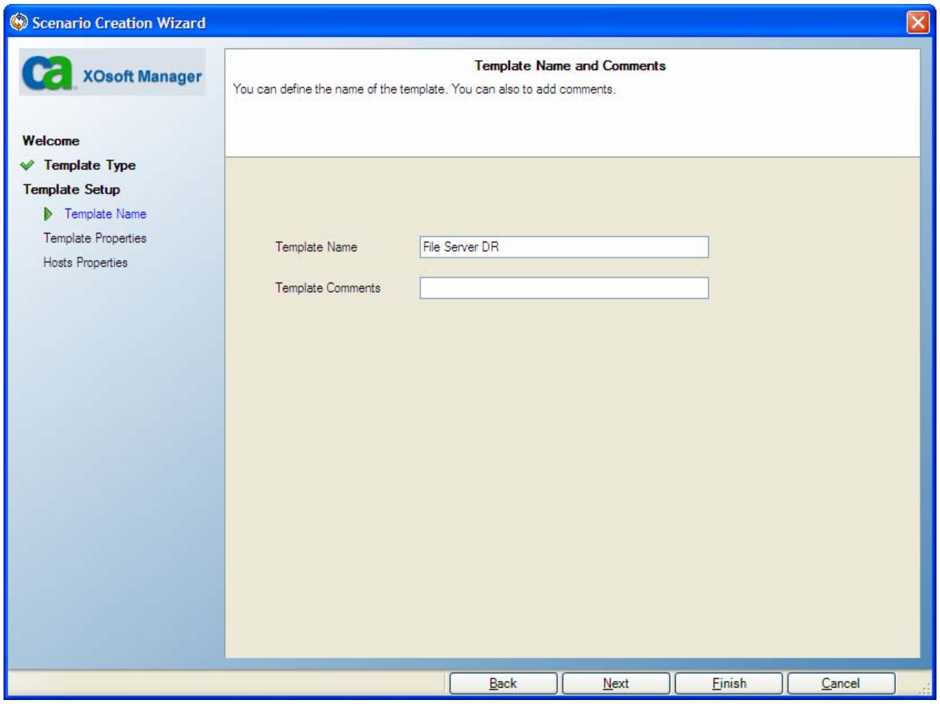

4. As with any scenario, you have the opportunity to rename the template. This name is the base for the default names of the scenarios that will be created using this template. You can also add description of the template in the **Template Comments** box. This description will appear on the Scenario pane as part of the template's details and can help you identify the template configuration.
5. Click **Next.** The **Template Properties** page appears. From this step, the wizard pages are similar to the ones you use in creating a new scenario (see page 51). Continue defining the template by following the wizard instructions.

Create a New Scenario using an Existing Template

Creating a new scenario using a template saves you the need to separately configure each required scenario. When you are using one of the existing templates, a new scenario is created with all parameter values taken from that template.

Note: Scenarios cannot be created from a template until the template has been saved. Changing parameter values in a template will not change the values of those parameters in a scenario that was previously created from it.

To create a new scenario using a template

1. Open the Scenario Creation Wizard by clicking the **New** button on the Standard toolbar, or selecting **New** from the **Scenario** menu.

   The **Scenario Creation Wizard** appears:

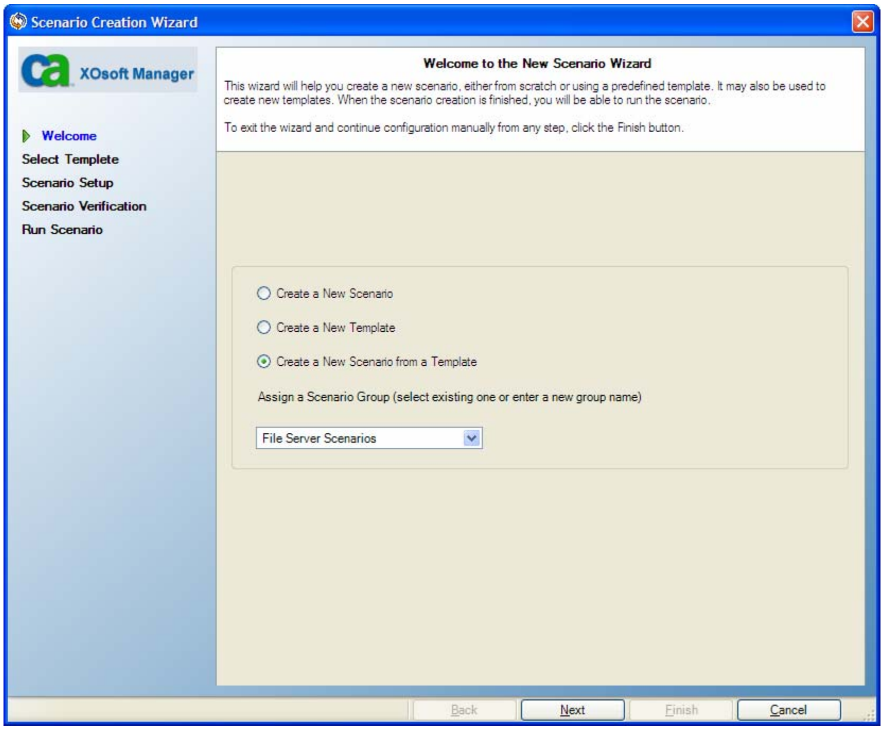

2. Select the **Create a New Scenario from a Template** option button, and assign the new scenario to a group. Then, click **Next**.

The **Select Template** page appears:

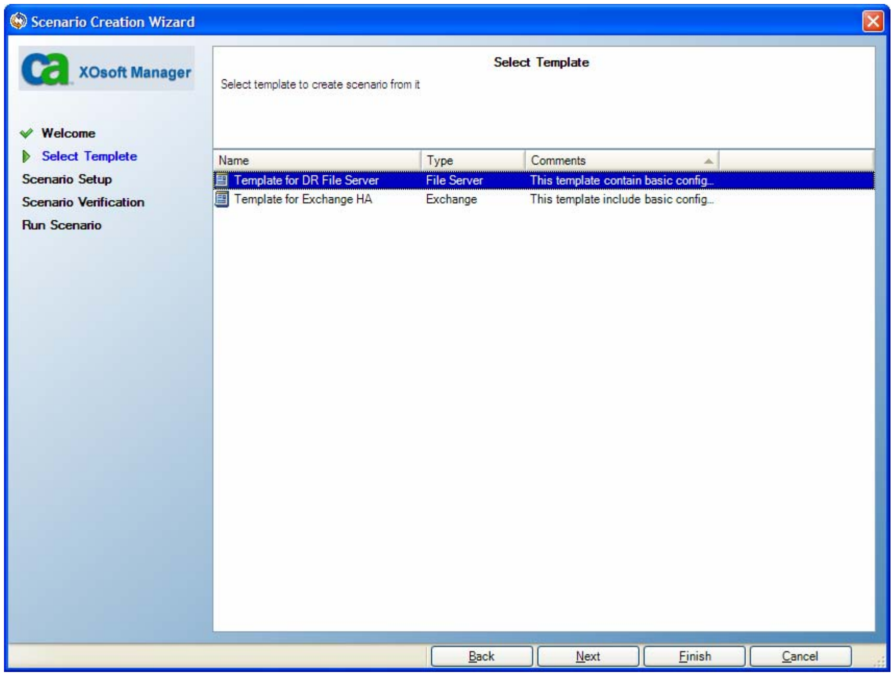

The available templates appear in this page.

3. Select the template that is best suited for your needs and click **Next**.

   The **Master and Replica Hosts** page is displayed:

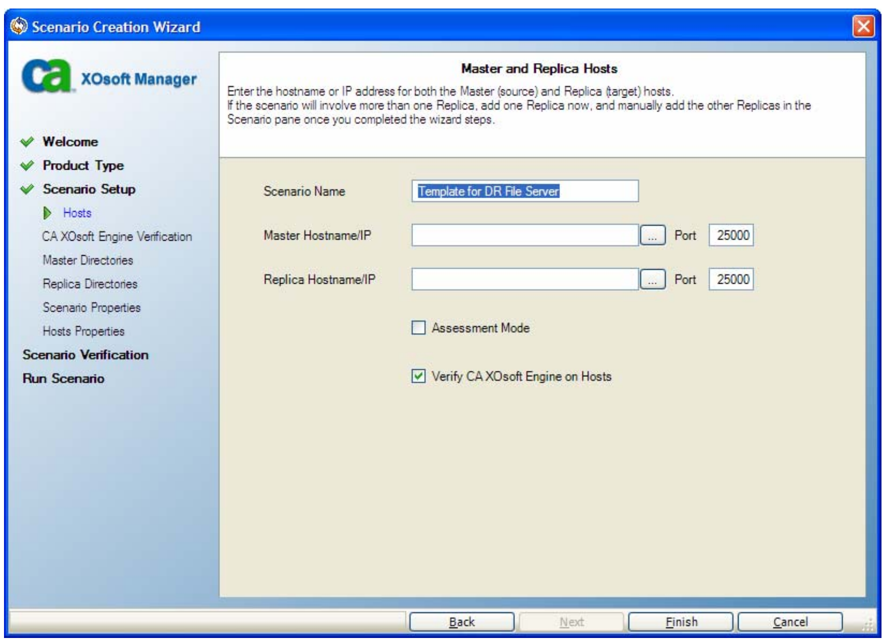

4. The default scenario name is the template name. You can either keep it or change it.

   From this step, the wizard pages are similar to the ones you use in creating a new scenario (see page 51). Continue defining the new scenario by following the wizard instructions.

Chapter 4: Running the Replication Process

This chapter describes the following procedures: initiating and stopping replication, changing configuration when replication is running, opening and closing the CA XOsoft Manager while replication is running, initiating synchronization, and suspending replication.

Initiate Replication

After you create a scenario, you need to run it to start the replication process. Normally, before data changes on the Master will begin to be replicated on the Replica, the Master and the Replica need to be synchronized. Therefore, the first step in initiating a replication is synchronizing the Master and Replica servers. After the servers have been synchronized, an online replication starts automatically, continuously updating the Replica with all of the changes that occur on the Master. CA XOsoft allows simultaneous synchronization and replication. In this case, the servers are synchronized while files are in use and being updated. All of these changes that occur on the Master are captured and being held in a spool, and once the synchronization is completed, the replication starts and the changes are updated on the Replica.

Note: In order for the replication process to succeed, verify that the user under which the CA XOsoft Engine is running has Read permission on the Master, and Read and Write permissions on each replication root directory and included files, and on all participating Replica hosts.

To initiate replication

1. From the Scenario pane, select the scenario you want to run.
2. To run the scenario, click the **Run**  button on the Standard toolbar:

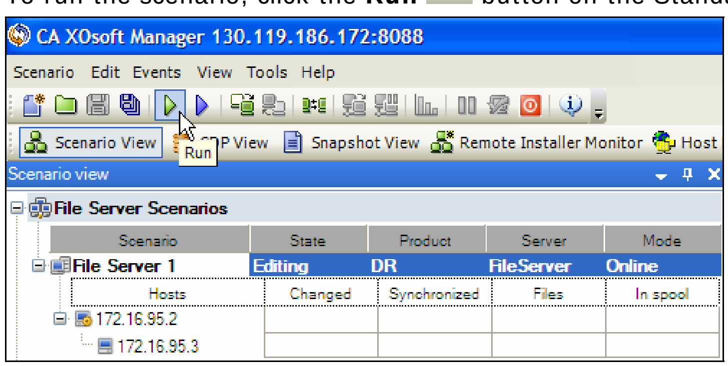

CA XOsoft verifies the scenario before running it.

3. If the scenario was not set up correctly or problems occurred in the participating hosts, errors are reported on the Event pane.

   Notes:

   - If any errors are displayed, you cannot run the scenario. These errors must be corrected before you can start the replication process.
   - Replication of mount points will succeed only if those were added to the Master before the Engine was started. If you included the mount points in the Master root directories when the Engine was already running, no error is reported but the replication does not start. In this case, you need to restart the Engine on the Master before initiating replication.

4. When no error is reported, the **Run** dialog appears:

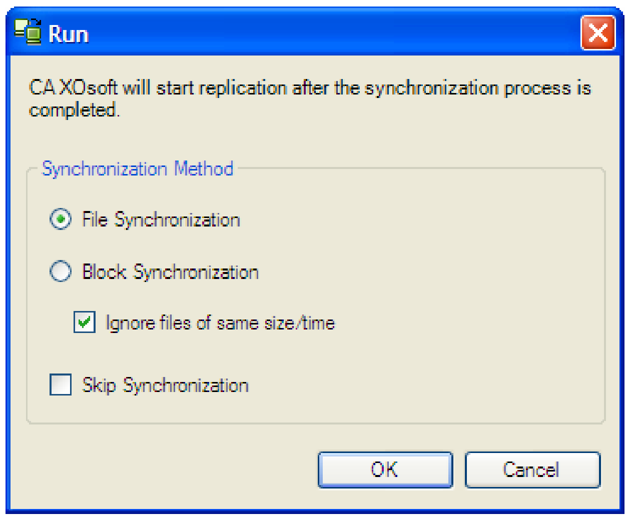

   The **Run** dialog contains the synchronization options.

   Note: In general, the default values are the most appropriate choice.

5. For File Server scenario, verify that the **File Synchronization** is selected as the synchronization method, and click **OK**.

   Important! Do not skip synchronization unless you are absolutely certain that the data in the Master and Replicas root directories is identical.

   The Manager now indicates that the scenario is running via the green play symbol to the left of the scenario, and via the scenario's state which turns into **Running**:

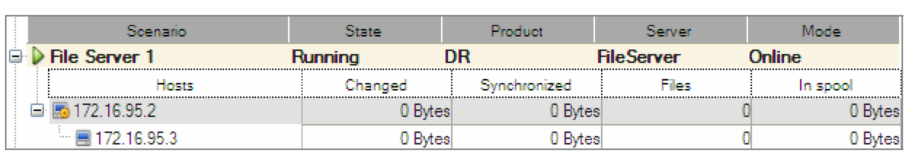

6. Once a scenario is running, a Statistics tab appears at the bottom of the Framework pane, displaying a graphical view of the replication:

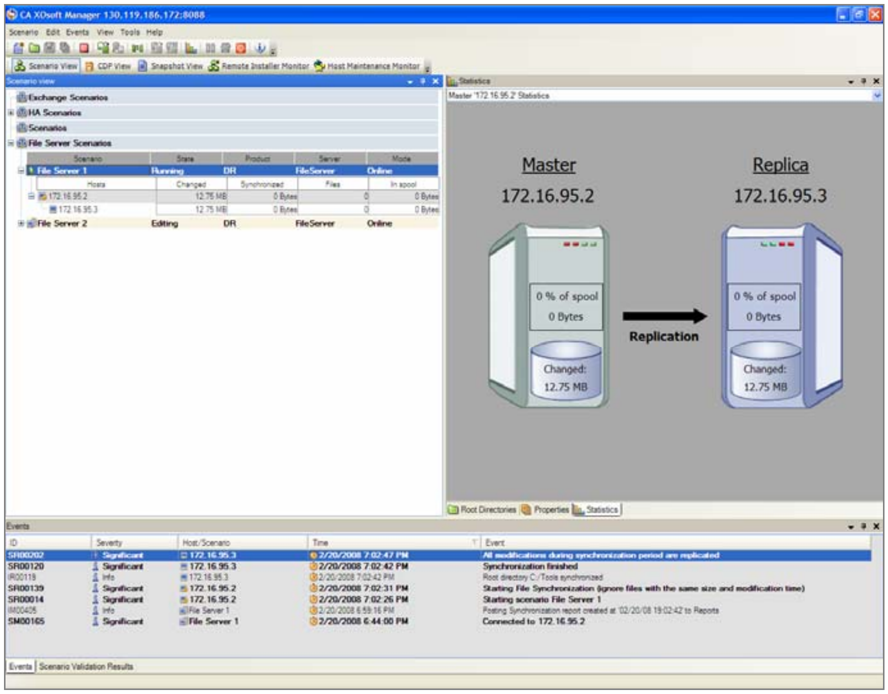

7. By default, once synchronization occurs, a Synchronization Report is generated. To view the report, refer to *Viewing a* Report, page 84.

   Note: You can also generate a Replication Report on a regular basis to monitor the replication process on each participating server. See *Replication Report*, page 90.

Run Mode

While replication is running and the running scenario is selected, the CA XOsoft Manager screen becomes gray. No configuration change or editing is possible at this time. Statistics are shown in the Framework pane on the right. Events are recorded in the Events pane on the bottom.

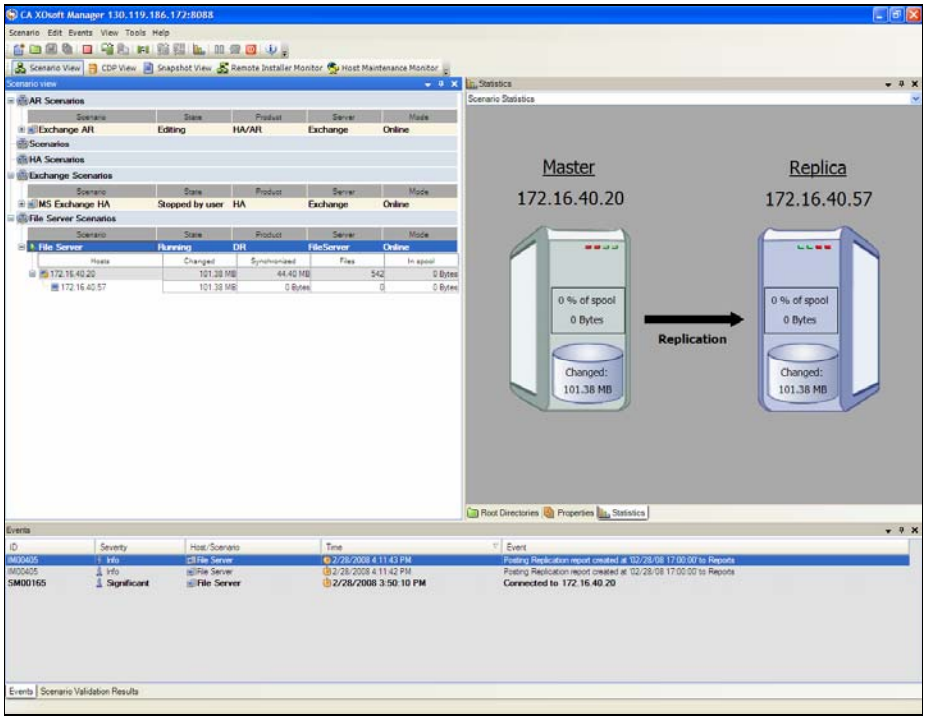

Stop Replication

To stop replication

1. From the Scenario pane, select the scenario you want to stop.

2. To stop the scenario, click the **Stop** button on the Standard toolbar.

   A confirmation message appears prompting you to approve the scenario stopping:

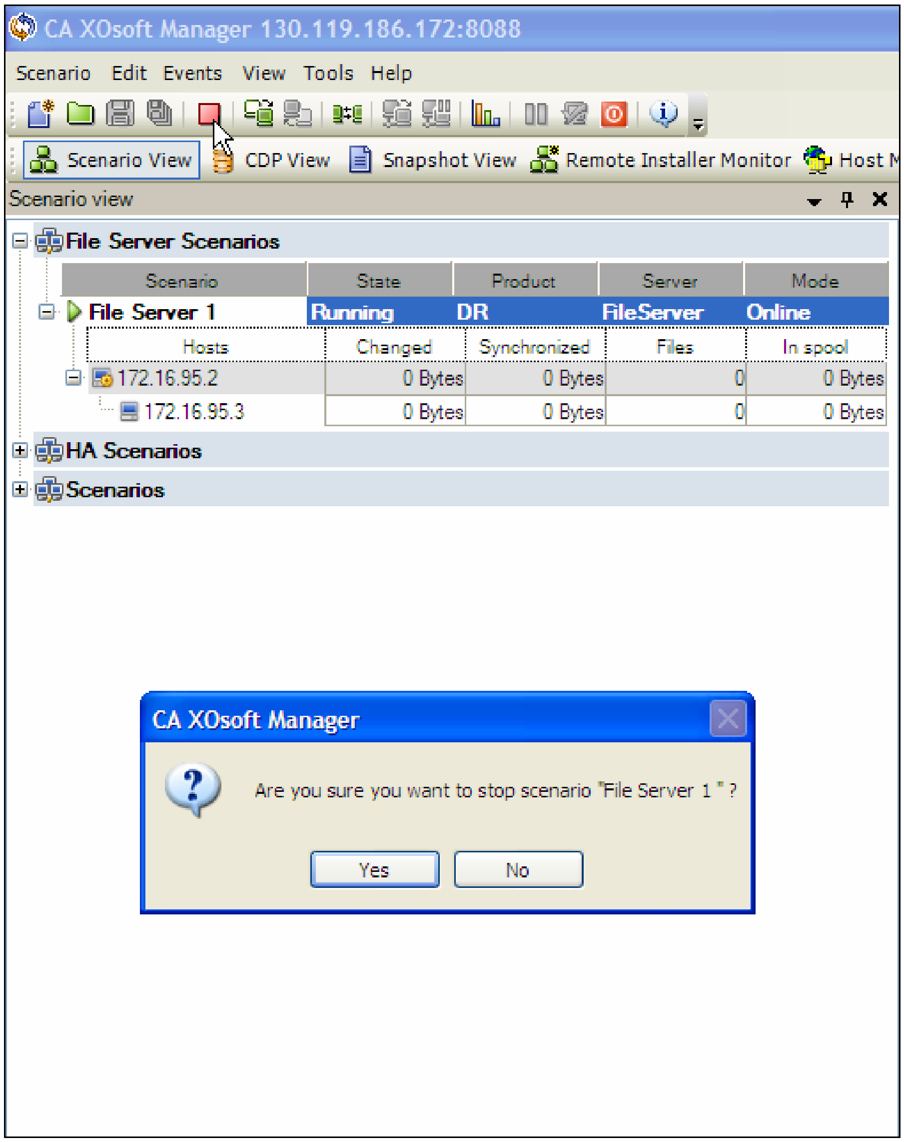

3. Click **Yes** in the confirmation message. The scenario stops.

After stopping the scenario, the Manager no longer shows the green play symbol to the left of the scenario, the scenario's state turns into **Stopped by user**, and the Statistics tab is no longer available on the Framework pane:

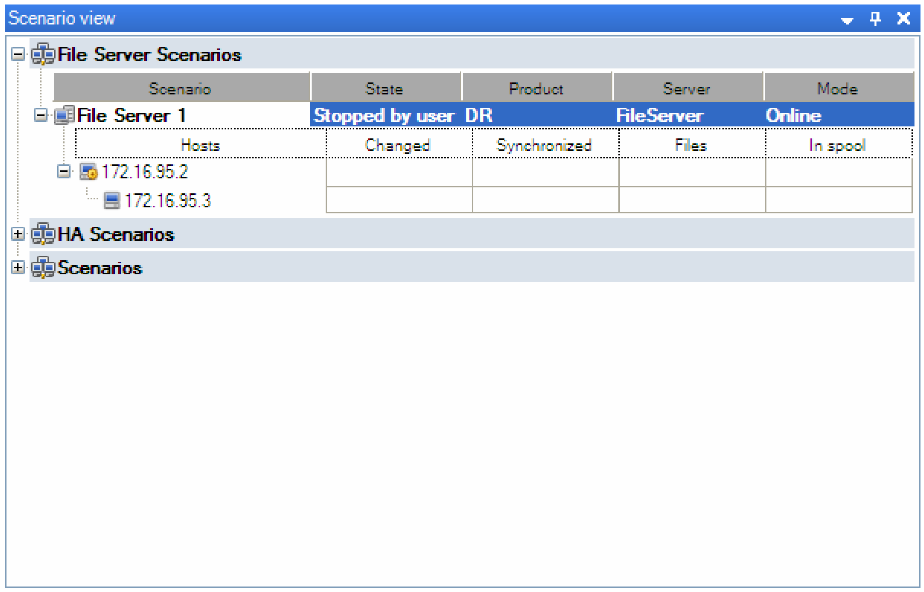

Change Configuration when a Scenario is Running

You cannot make changes to the scenario settings, nodes, or directories while the scenario is running.

To make changes to the scenario nodes or directories

1. Stop replication by clicking the **Stop** button.

   The CA XOsoft Manager screen background returns to its normal color, the Statistics tab is closed, and the functions of the Properties and Directories tabs in the Framework pane become active.

2. Make all necessary changes, and save them by clicking the **Save** button on the standard toolbar.

3. To restart the replication with the new configuration, click the **Run** button again.

Close and Open the CA XOsoft Manager during Replication

After the scenario has been defined, and replication has started, the CA XOsoft Manager can be closed. It may remain open only for real-time monitoring of the replication process. Closing the Manager does NOT stop the running scenarios. When it is opened again, it automatically uploads all of the saved scenarios and displays their status.

Note! Even when the CA XOsoft Manager is closed, it is possible to monitor the replicated system via the Overview page. You can also get notifications by email or by automatically running user-defined scripts when important events or errors occur. (For more information, see the *Event Notification* sections, on the *Scenario, Master and Replica Properties* chapters.*)*

Synchronization

The synchronization process can be manually activated at any time, whether replication is running or not.

Manual synchronization is recommended in the following situations:

- Before starting replication on servers with large amounts of data and a heavy update rate.
- After a lengthy network failure, if automatic synchronization is not activated.
- After rebooting one of the participating servers, if automatic synchronization is not activated.

To synchronize the Master and the Replica

1. On the Scenario pane, select the scenario you want to synchronize.
2. Click the **Synchronize**  button on the Standard toolbar, or select the **Synchronize** option from the **Tools** menu. The **Synchronization** dialog appears, displaying the synchronization options:

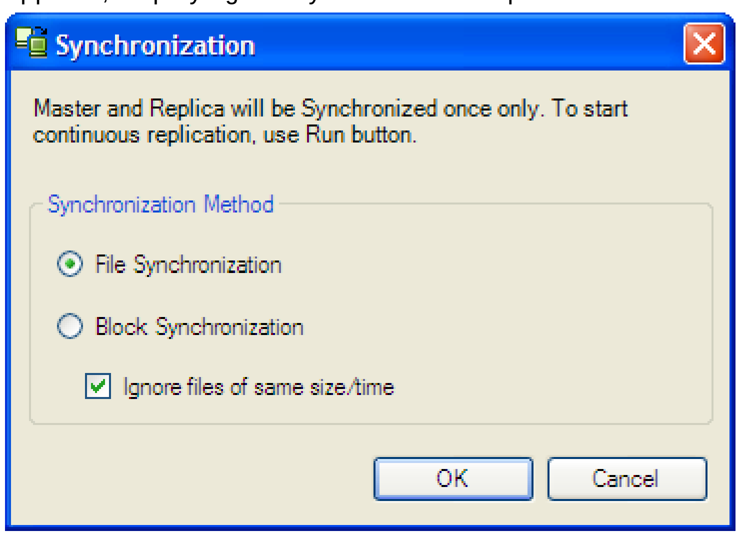

3. Choose the desired options. (See *How Synchronization Works*, page 15 and *Synchronization Type*, page 124).

   Note: You can also set synchronization to run automatically at pre-scheduled hours on specific days, and exclude it from specific dates. See *Schedule Synchronization*, page 128.

4. Click **OK** to start the synchronization according to the method you selected.

   Note: You can issue a synchronization command while a Replica is suspended; however it is performed only after the replication is resumed.

Once the synchronization process starts, the Statistics tab opens, informing you that **Synchronization in progress**:

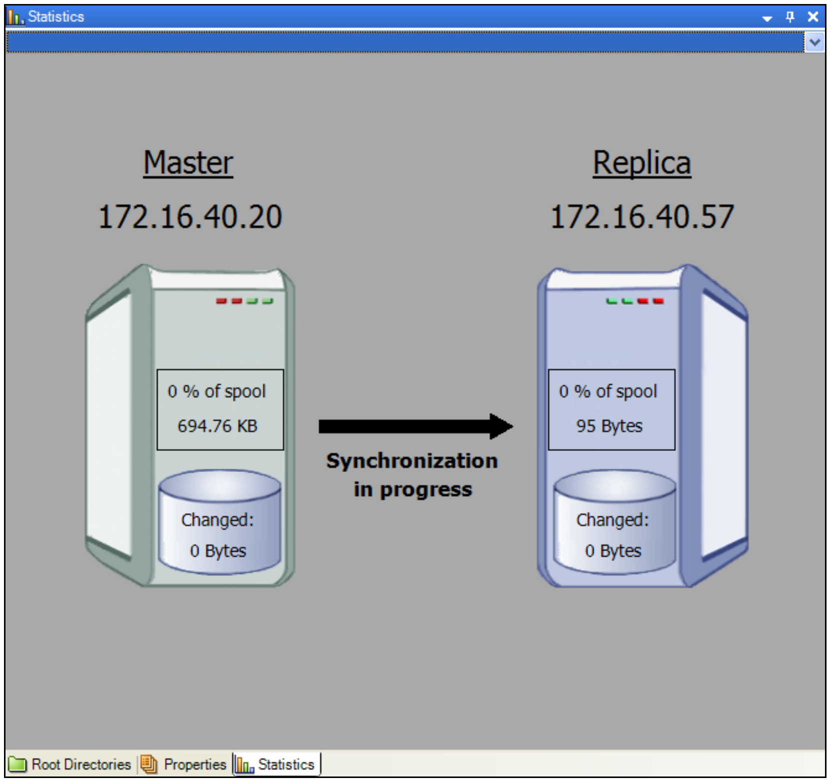

Suspend Replication

At times it may be necessary to suspend updates on a Replica machine in order to perform system maintenance or some other form of processing that does not modify the replicated data there. Usually, it is not desirable to stop replication since this requires a resynchronization afterward. The replication suspension feature of CA XOsoft solves this problem.

During the suspension period, all changes are spooled on the Master or on the Replica located upstream of the suspended Replica. In other words, changes continue to be recorded for update on the suspended Replica, but are not actually transferred until replication is resumed. Once replication is resumed, the accumulated changes are transferred and applied without any need to perform a full resynchronization of the data.

Replication may be suspended either manually or on a scheduled basis.

Important! It is imperative that during suspension, you do nothing on the Replica that causes the data to change in any way, including starting an application such as Exchange, SQL Server, or Oracle. If you need to start programs that will change data on the Replica, you may use the Assured Recovery option. (For more information about Assured Recovery, see page 193.)

Note: You cannot suspend replication during synchronization. You can suspend Replication only temporarily, since changes are accumulated in the spool directory of the Master or upstream Replica. Make sure that sufficient disk space is available for the spool to hold the changes during the time the Replica is suspended.

Suspend Replication Manually

To manually suspend replication to a particular Replica

1. On the Scenario pane, select the Replica you want to suspend. Then, click the **Suspend** button, or select the **Suspend Replication** option from the **Tools** menu.

   A confirmation message appears, informing you that any change of the Replica root directories content during suspension will require manual resynchronization:

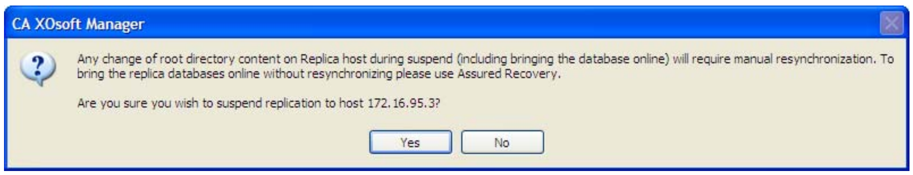

2. Click **Yes** to suspend the replication.

   Once the Replica is suspended, a red icon appears next to the Replica on the Scenario pane:

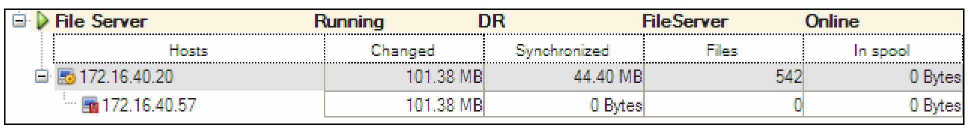

File Server	Running	DR	FileServer	Online
Hosts	Changed	Synchronized	Files	In spool
172.16.40.20	101.38 MB	44.40 MB	542	0 Bytes
172.16.40.57	101.38 MB	0 Bytes	0	0 Bytes

   Note: During suspension, the scenario's state does not change but stays **Running**, since it is only the replication to the Replica that is suspended.

   On the Scenario Statistics pane appear a suspension icon and a caption, informing you that the replication is suspended:

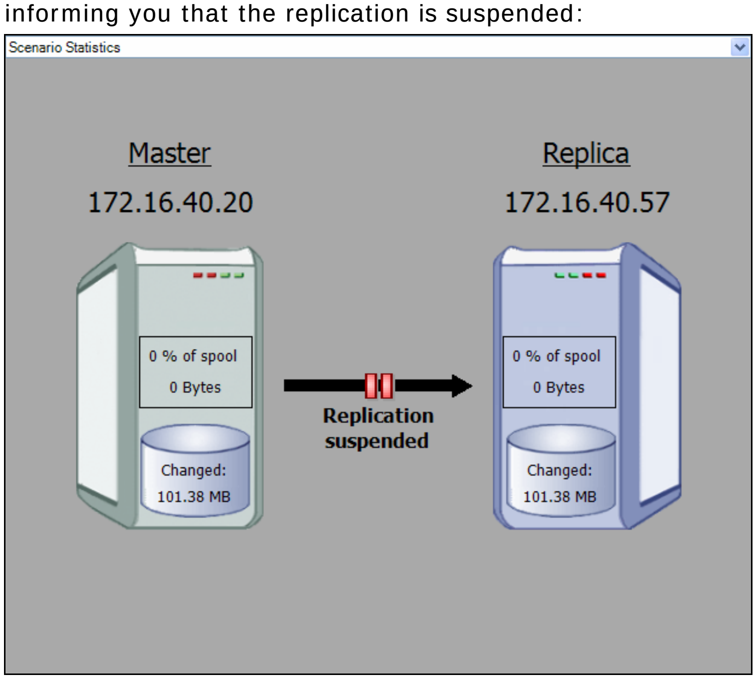

3. While replication is suspended, you may perform maintenance on the Replica server, including reboot. However, it is very important not to modify the replicated data in any way or a full resynchronization with the Master will be required.

   Note: You can issue a synchronization command while a Replica is suspended; however, it is performed only after the replication is resumed.

Resume Replication after Manual Suspension

To resume replication after manual suspension

1. After you suspend a Replica, the **Suspend**  button toggles to **Resume Replication**. When you are ready to resume the replication, click the **Resume Replication** button, or select the **Resume Replication** option from the **Tools** menu:

   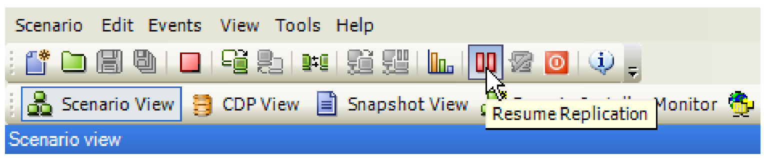

   A confirmation message appears:

   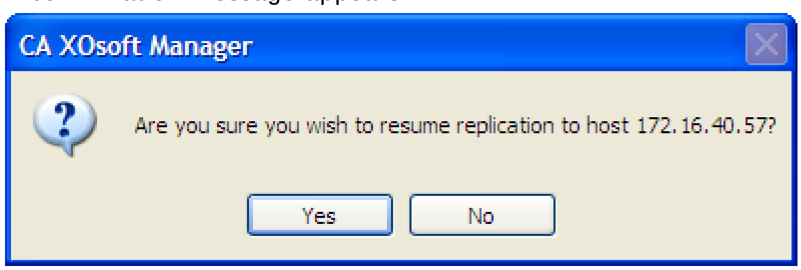

2. Click **Yes** to resume replication.

   Once the replication is resumed, the red icon disappears from the Replica on the Scenario pane, and the suspension symbol disappears from the Scenario Statistics pane.

Schedule Replication Suspension

To schedule automatic replication suspension

1. On the Scenario pane, select the Replica you want to suspend and stop the scenario in which it participates.
2. On the Framework pane, select the Properties tab to open the Replica Properties list.
3. On the Replica Properties list, open the **Scheduled Tasks** group. On the **Suspension** property, set the value to **On**:

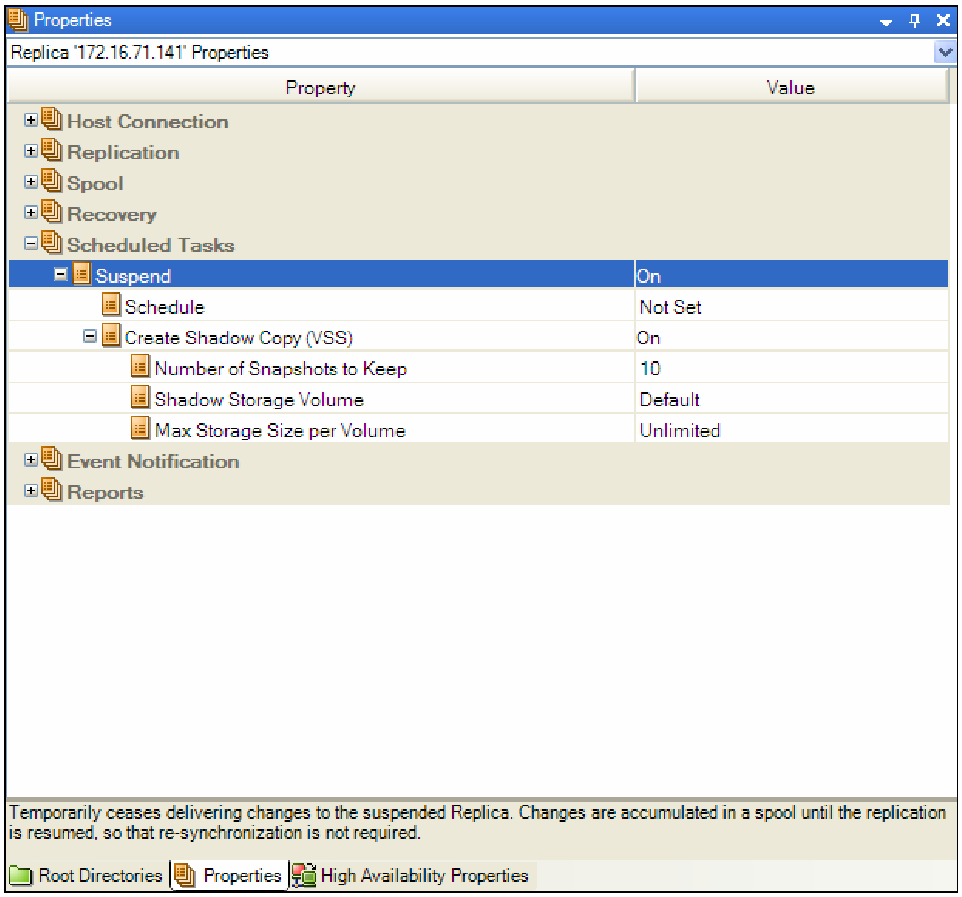

4. On the **Schedule** property, click the value box.

The **Suspend hours** dialog appears:

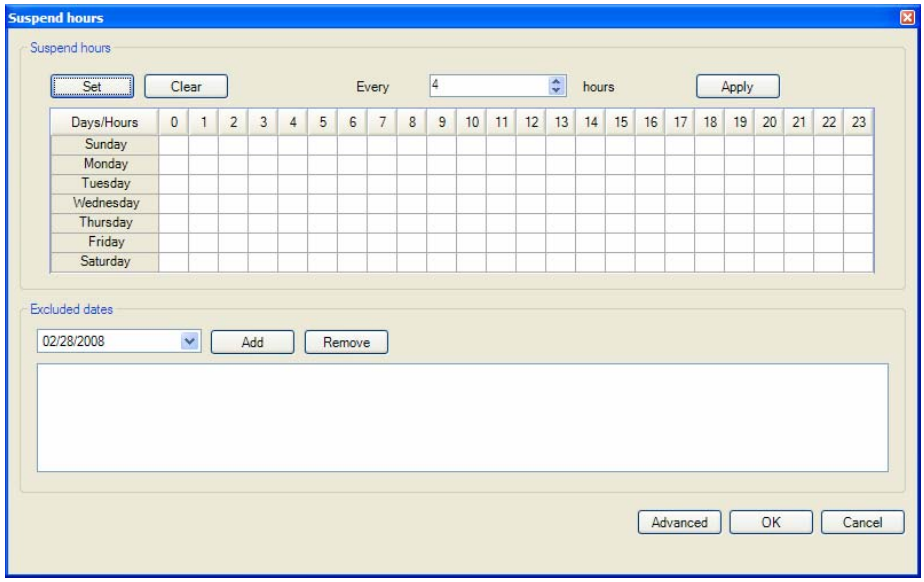

The **Suspend hours** dialog is similar to the **Schedule Setting** dialog, which is used for scheduling automatic synchronization. For information about setting a schedule, see *Schedule Synchronization*, page 128.

5. Set the schedule for automatic suspension in the **Suspend hours** dialog, and click **OK** to save your schedule and close the dialog.
6. To activate the schedule, click the **Save** button on the Standard tool bar and start the scenario.

   The Replica you selected for suspension will be suspended and resumed according to the set schedule.

Chapter 5: Monitoring Replication

This chapter describes the various monitoring tools of CA XOsoft that enable you to easily control and monitor your replication environment.

The CA XOsoft Overview Page

The CA XOsoft Overview Page allows simultaneous monitoring by any number of administrators or professionals who wish to be informed about the state of the replicated system:

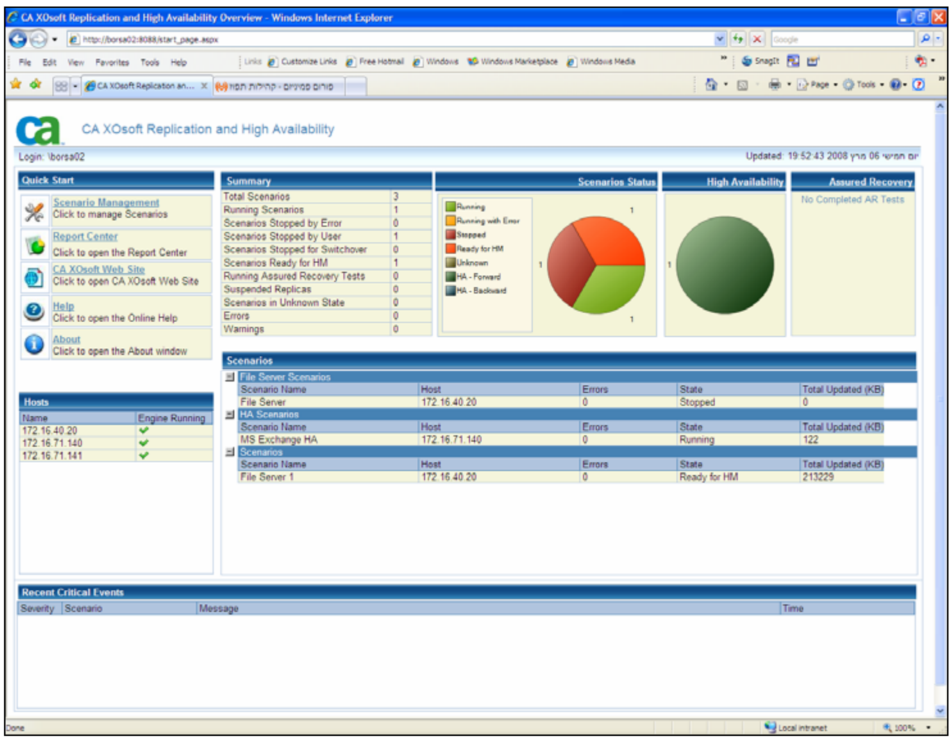

Note: There is no need to refresh the browser; the Overview page refreshes automatically.

The CA XOsoft Manager

The CA XOsoft Manager enables you to easily control and monitor your replicated system.

Monitor Multiple Scenarios

The Scenario pane shows the current status for all scenarios in one pane:

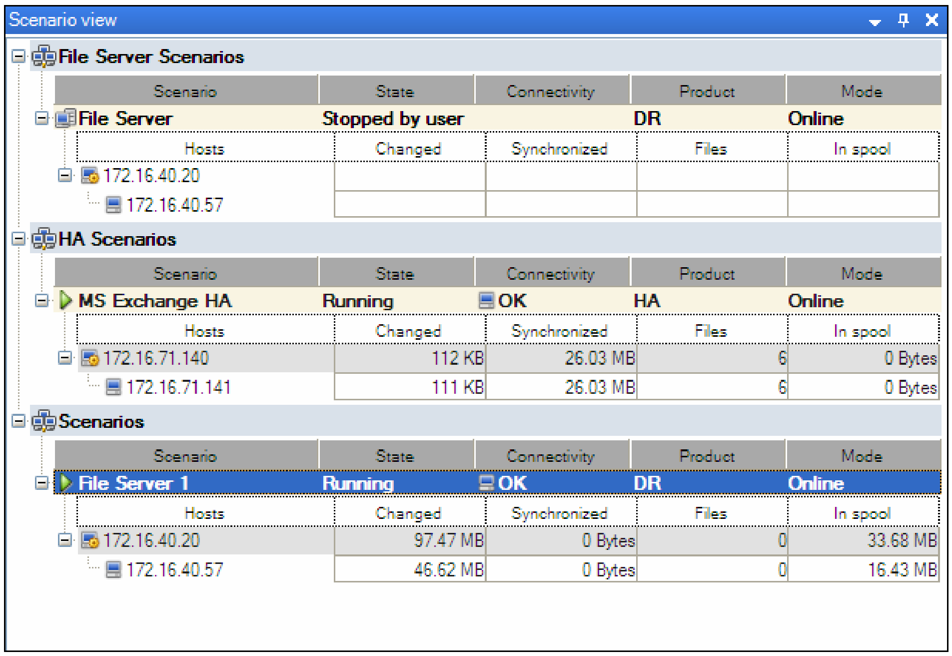

You can customize the way the columns are displayed. (For more information, see *Customizing the Scenario Pane*, page 34.)

State Information

State information is displayed beside each scenario name and beside each server in the replication tree, whenever a synchronization process has been started or completed, and whenever a replication process is underway.

The state information includes:

- A graphic indication next to the scenario name indicating the scenario is running , or is idle .
- A graphic indication next to the server name indicating that the server is a Master (active) server , or a Replica (standby) server .
- A graphic indication of whether the servers are connected: if the connection to any of the participating servers is lost, the server icon appears with a large red **X** marked over it:

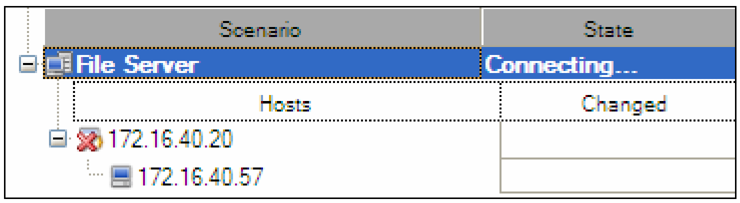

Live Statistics

Live statistics are displayed in two areas:

- Scenario pane
- Statistics pane

Scenario Pane

The live statistics displayed in the Scenario pane include the following information:

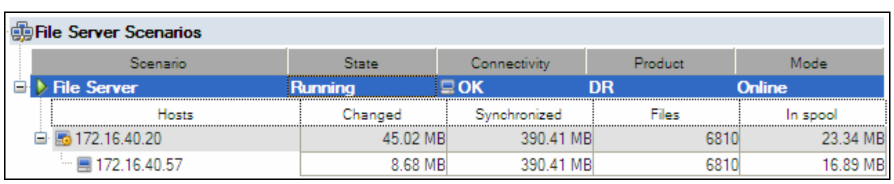

- **Changed** - total data replicated from this host since the last synchronization.
- **Synchronized** - total data synchronized from this host.
- **Files** - total number of files replicated from this host.
- **In Spool** - total (current) amount of data contained in the spool.

Statistics Tab

The Statistics tab in the Framework pane also displays live statistics. Different statistics information is displayed for a scenario, a Master and each Replica server.

Note: The Statistics tab on the Framework pane appears only when a scenario is running.

The available live information in the Statistics tab is as follows:

- **Statistics per scenario** – a graphical overview of the scenario state:

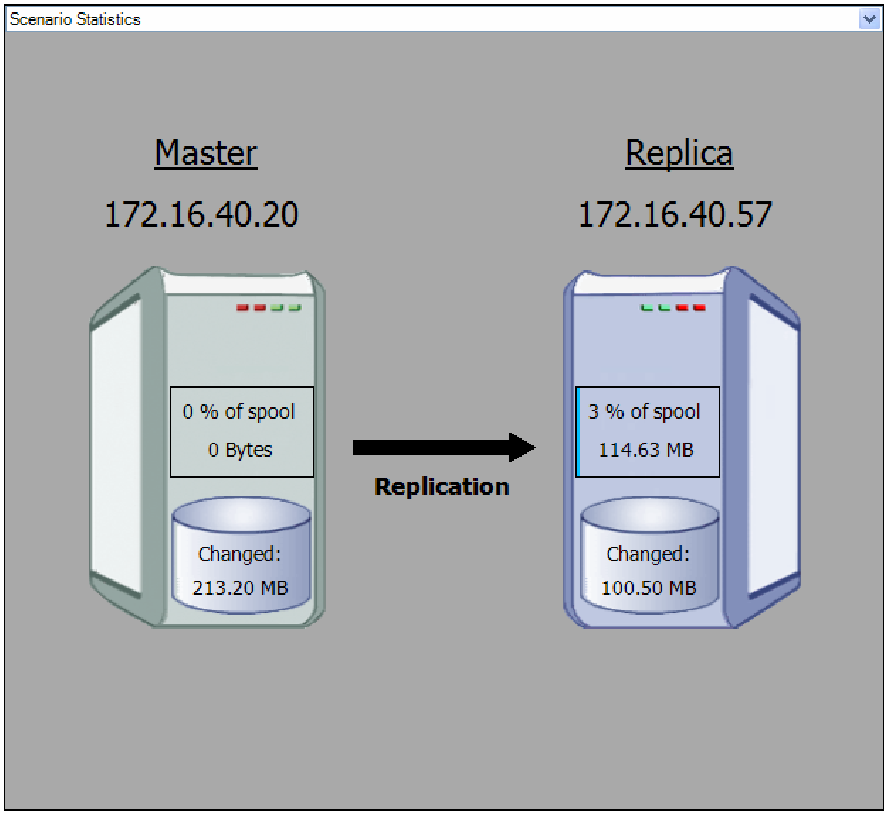

- **Statistics per Master** – a table that containing the following information: state; replication starting date and time; CA XOsoft Engine version no.; total amount of data contained in the spool; total amount of data changed in the root directories since the beginning of the replication process, including the number of created folders and changed, removed and renamed files; total size of files replicated from this host; synchronization method; and synchronization progress:

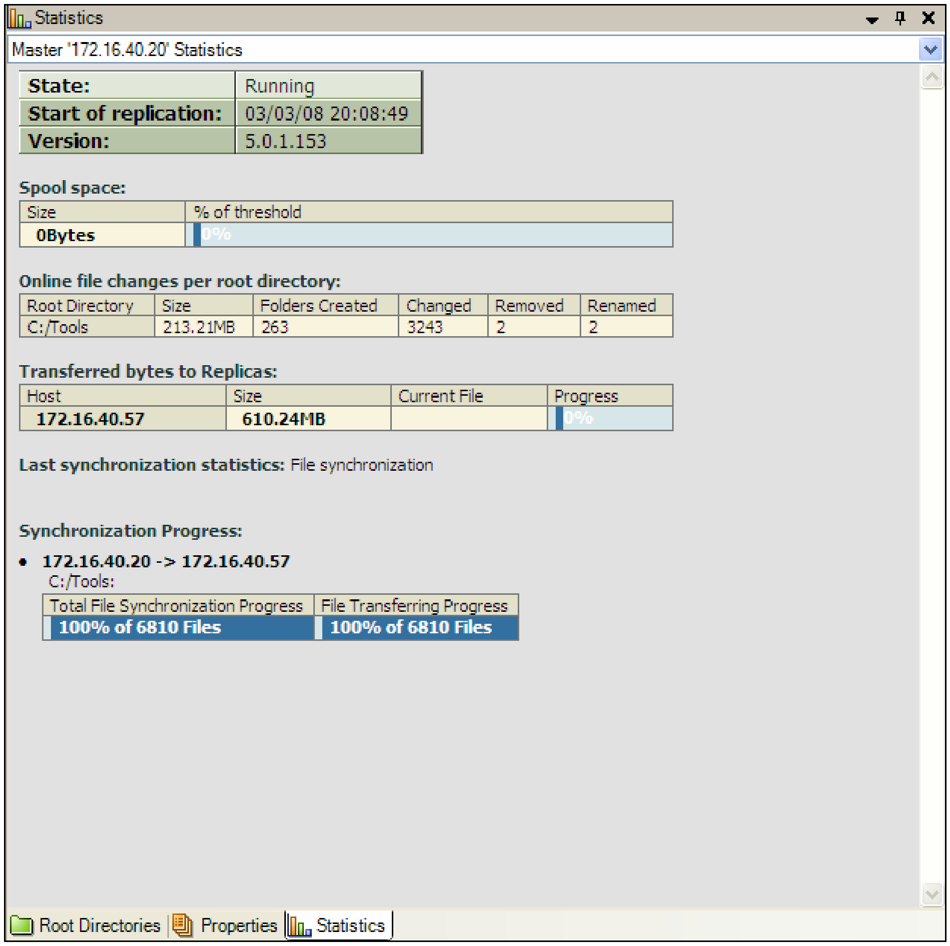

- **Statistics per Replica** - a table that containing the following information: state; replication starting date; CA XOsoft version no.; total amount of data contained in the spool; and total amount of data (in KB) changed in the root directories since the beginning of the replication process, including the number of files changed, removed and renamed:

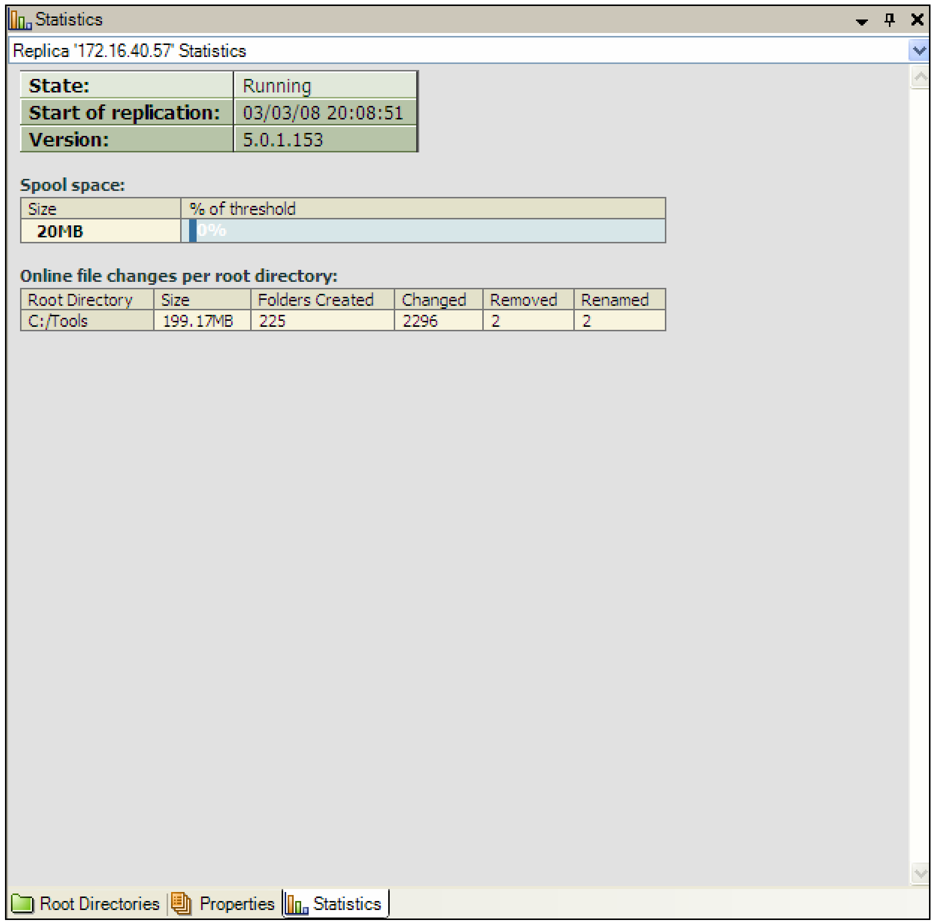

Refresh Statistics Display Automatically

The CA XOsoft Manager receives state information from all servers participating in the current scenario. You can set a default frequency for automatically updating the state information and live statistics display.

To define the refresh rate

1. From the **Tools** menu, select **Statistics**, **Refresh Rate**.

   The **Refresh Rate** dialog appears:

   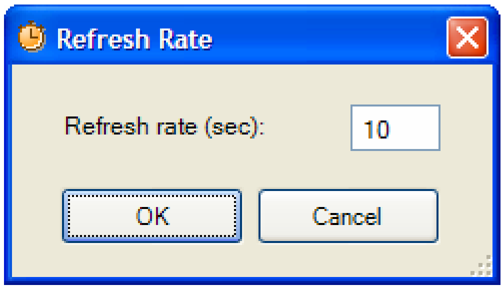

2. Enter the desired refresh rate in seconds and click **OK**. The Scenario pane updates accordingly.

   Note: The refresh rate interval can be between 10 to 99 seconds.

Refresh Statistics Display Manually

To manually refresh the displayed information

- Click the **Refresh Statistics**  icon on the Standard toolbar
- Press the **F5** key
- From the **Tools** menu, select **Statistics**, **Refresh**.

Events

The Events pane displays messages and general information about the selected scenario. This information is received from the servers participating in the running scenario. The information displayed can inform you, for example, that a directory is synchronized, a server is connected, synchronization started/finished, etc. The information includes the server name and time, and a brief explanation. Important events or error messages are shown in bold letters:

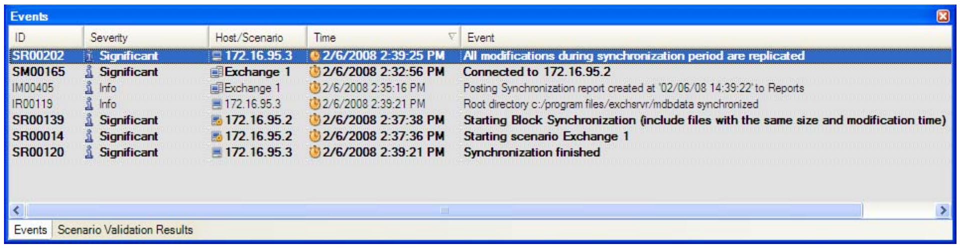

View Events in a Separate Window

Sometimes, event messages are very long, exceed the Event area, and are cut off (visually). In these cases, you may want to display the event messages in separate windows.

To display an event in a separate window

1. On the Event pane, select the event you want to view.
2. Double-click the event, or right-click it and select **View Event in other Window** from the pop-up menu. Alternatively, select from the **Events** menu the **View Event in other Window** option.

   A pop-up message appears, displaying the full message text of the selected event:

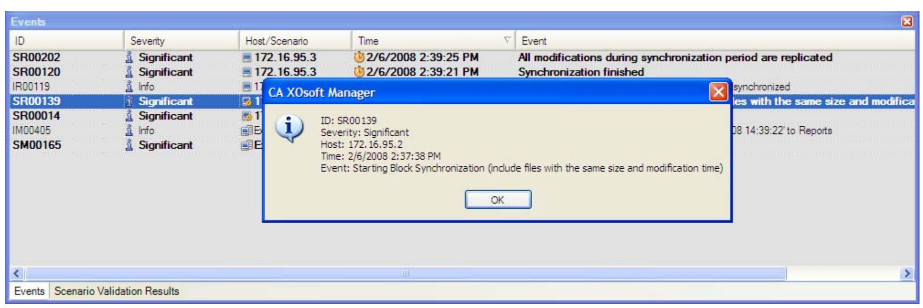

View Incoming Events

The CA XOsoft Manager can visually notify you when an incoming event occurs.

To view incoming events as they occur

- From the **Events** menu, select the **Pop-up on Incoming Event** option.

  When an incoming event occurs, the CA XOsoft Manager icon in the task bar flashes and the Manager is moved to the fore.

  Note: If you re-select this option and turn it off, the minimized application in the task bar does not flash during an incoming event.

Copy Events

When important events occur, you may want to copy their messages to other programs.

To copy the text of the events for use in other programs

1. On the Event pane, select any number of events using the **Ctrl** key.
2. Right-click in the Event pane and select **Copy**, or select the **Copy** option from the **Events** menu. Alternatively, press **Ctrl+C**.

   You can paste the copied event texts into any program you wish.

   Note: CA XOsoft also enables you to copy the event message directly to a file with CSV extension, such as Excel files. Once you select the **Copy To CSV** option, the application that is defined in your computer as CSV opens, displaying the copied message. (To set your default CSV application, see My Computer, Tools, Folder Options, File Types, CSV.)

   The following example shows event messages that were copied directly to Excel using the **Copy To CSV** option:

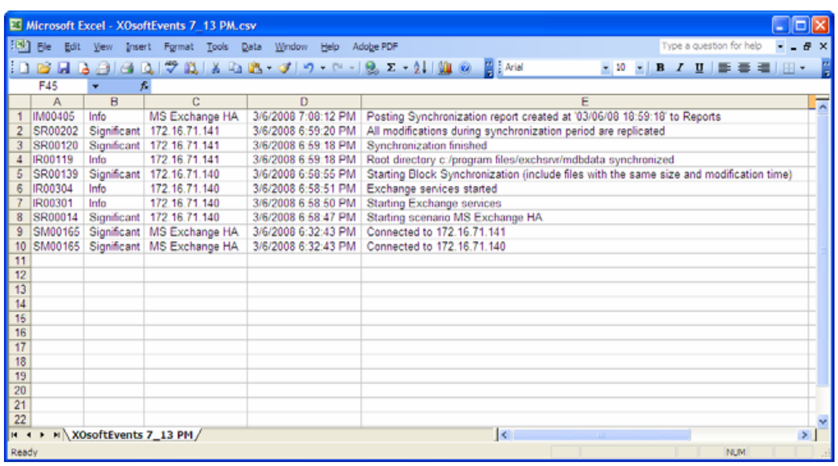

	A	B	C	D	E
1	IM00405	Info	MS Exchange HA	3/6/2008 7:08:12 PM	Posting Synchronization report created at '03/06/08 18:59:18' to Reports
2	SR00202	Significant	172.16.71.141	3/6/2008 6:59:20 PM	All modifications during synchronization period are replicated
3	SR00120	Significant	172.16.71.141	3/6/2008 6:59:18 PM	Synchronization finished
4	IR00119	Info	172.16.71.141	3/6/2008 6:59:18 PM	Root directory c:/program files/exchsrvr/mdbdata synchronized
5	SR00139	Significant	172.16.71.140	3/6/2008 6:58:55 PM	Starting Block Synchronization (include files with the same size and modification time)
6	IR00304	Info	172.16.71.140	3/6/2008 6:58:51 PM	Exchange services started
7	IR00301	Info	172.16.71.140	3/6/2008 6:58:50 PM	Starting Exchange services
8	SR00014	Significant	172.16.71.140	3/6/2008 6:58:47 PM	Starting scenario MS Exchange HA
9	SM00165	Significant	MS Exchange HA	3/6/2008 6:32:43 PM	Connected to 172.16.71.141
10	SM00165	Significant	MS Exchange HA	3/6/2008 6:32:43 PM	Connected to 172.16.71.140

Reports

CA XOsoft can generate reports on the replication and synchronization processes. These reports can be stored on your desired location, opened for view from the Report Center, sent by email to a specified address, or they can trigger script execution. To set these options, see the following*:*

- For defining the storage directory and the retention period of the reports, see the
-
-
- Report Handling section in *Understanding Scenario* Properties*,* page 127.
- For defining the automatic generation of synchronization and replication reports for the Master, see the *Reports* section in *Understanding Master Properties*, page 137.
- For defining the automatic generation of replication reports for the Replica, see the *Reports* section in *Understanding Replica* Properties, page 143.

The default storage directory of the generated reports is: *[ProgramFilesFolder]*\CA*XOsoft**Manager**reports*

Viewing a Report

To view a report

1. To view a report, first you need to open the Report Center. There are two ways to open it:
   - On the Overview Page, click the **Report Center** link on the **Quick Start** pane on the left:

- From the **Tools** menu, select the **Reports** option and then **Show Scenario Reports**.

The Report Center opens in a new window:

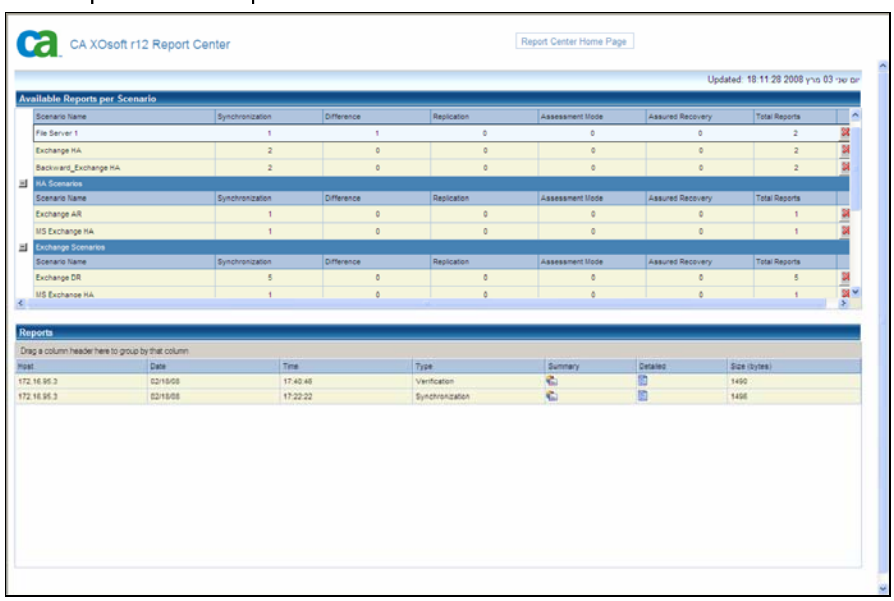

The Report Center consists of two tables:

- The upper table - **Available Reports per Scenario** – contains a list of all scenarios that have reports, along with the type and number of available reports for each scenario.
- The lower table – **Reports** – contains a list of all the reports that are available for the scenario selected in the upper table.

2. To view a specific report, select from the **Available Reports per Scenario** table the scenario that this report represents. Then, from the **Reports** table below, click the report you want to open:

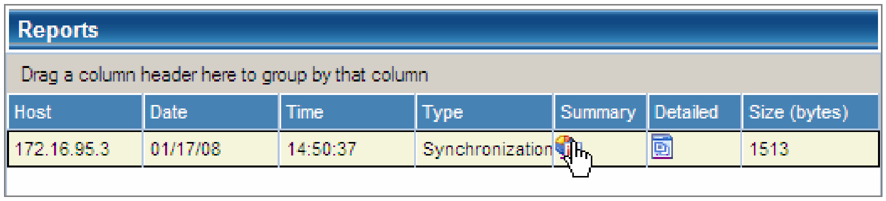

Note: Depending on your settings, for Synchronization and Replication reports a **Detailed** report can be generated in addition to the **Summary** report. Both reports represent the same process, but the **Detailed** report also provides a list of the files that participated in the process.

The report you selected appears:

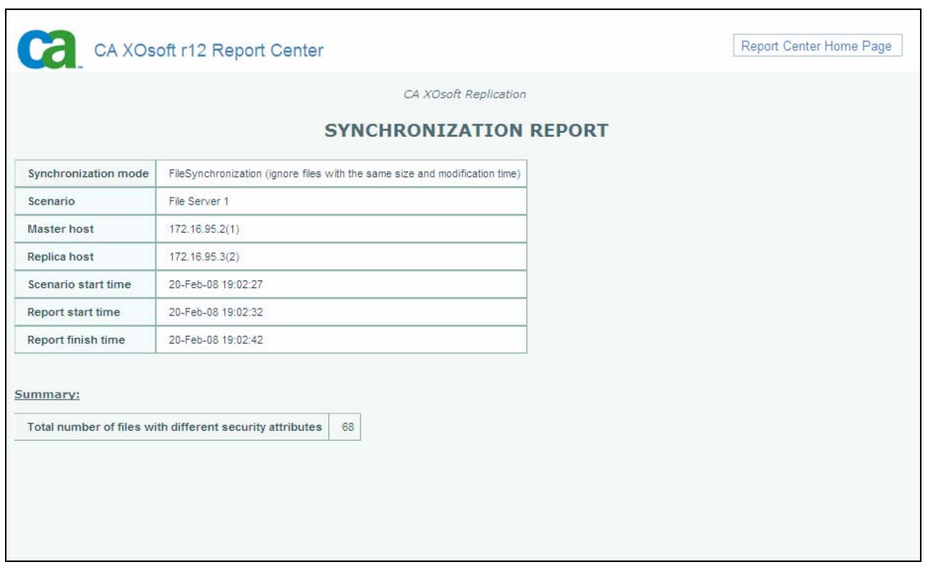

Deleting Reports

The reports that are displayed in the Report Center are stored for the period that is defined in their scenario properties (see Report Handling, page 127). The reports are displayed even if their scenarios were removed from the Manager. However, the Report Center enables you to delete reports that are no longer of use to you.

To delete a report

1. On the Report Center, select from the **Available Reports per Scenario** table the scenario whose reports you want to delete.
2. Click the **Delete**  button at the right end of the line.

   The selected scenario is removed from the Report Center, and all of the reports that belong to it are deleted.

Synchronization Report

Following synchronization, CA XOsoft creates a report that lists files that have been transferred or modified. The first few lines (at the top) include: the synchronization method, the scenario name, the names of the Master and the Replica(s), and the synchronization date.

The summarized Synchronization Report shows the sum total of the removed and modified files as well as bytes transferred. The report also provides information about the number of new directories that were created on the Replica following synchronization, as the following example illustrates:

CA XOsoft r12 Report Center

CA XOsoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server
Master host	172.16.40.20(1)
Replica host	172.16.40.57(2)
Scenario start time	28-Feb-08 15:18:14
Report start time	28-Feb-08 15:18:19
Report finish time	28-Feb-08 15:18:35

Summary:

Total number of files modified	7
Total number of bytes changed	3.13MB
Total number of directories created	1

Detailed Synchronization Report

The Detailed Synchronization Report presents the complete listing of the files that were transferred or modified during the synchronization process. For each file, the following information is provided:

- **Event** – the action that was performed on the Replica.
- **Bytes** – the size of the file.
- **Time Stamp** – modification time.
- **File Name** – the name and full path of the file.

The following example shows a typical Detailed Synchronization Report:

CA XOsoft r12 Report Center

CA XOsoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server
Master host	172.16.40.20(1)
Replica host	172.16.40.57(2)
Scenario start time	28-Feb-08 15:18:14
Report start time	28-Feb-08 15:18:19
Report finish time	28-Feb-08 15:18:35

EVENT	BYTES	TIME STAMP	FILE NAME
Change	1.33MB	28-Feb-08 15:14:09	C:/Tools/Backup5/REDIST.TXT
Create Directory	0Bytes	28-Feb-08 15:17:48	C:/Tools/Backup7
Create	1.71KB	24-Apr-98 00:00:00	C:/Tools/Backup7/2DARROW1.WMF
Create	2.71KB	24-Apr-98 00:00:00	C:/Tools/Backup7/2DARROW2.WMF
Create	376KB	30-Apr-98 00:00:00	C:/Tools/Backup7/ISWFRAME.DLL
Create	1.33MB	28-Feb-08 15:18:04	C:/Tools/Backup7/REDIST.TXT
Create	41KB	30-Nov-05 10:49:59	C:/Tools/Backup7/try1.ncb
Create	52.5KB	30-Nov-05 10:49:58	C:/Tools/Backup7/try1.opt

Summary:

Total number of files modified	7
Total number of bytes changed	3.13MB
Total number of directories created	1

Replication Report

The Replication Reports are generated periodically, according to a pre-defined frequency, and they are generated individually for the Master and the Replica servers.

The Replication Reports include statistics on data replicated since the beginning of the replication process, as well as statistics on data replicated since the last report. The data includes the number of replicated bytes, and the number of files created/updated/removed/and renamed. You can view either a summarized or a detailed report. A summarized report looks like this:


CA XOsoft r12 Report Center

CA XOsoft Replication

REPLICATION REPORT

Scenario	File Server
Host	172.16.40.20(1)
Scenario start time	28-Feb-08 15:18:12
Report start time	28-Feb-08 15:18:12
Report finish time	28-Feb-08 16:00:01

Summary:

Directory C:/Tools:

replicated bytes	3.64MB
changed files	8
removed files	1
renamed files	1
last change	28-Feb-08 15:34:59

Total number of bytes changed	3.64MB

Total statistics:

Directory C:/Tools:

replicated bytes	3.64MB
changed files	8
removed files	1
renamed files	1
last change	28-Feb-08 15:34:59

By default, Replication Reports are NOT automatically generated. To schedule CA XOsoft to generate Replication Reports, refer to the following:

- Generating Replication Reports for the Master server, see page 137.
- Generating Replication Reports for the Replica servers, see page 143.

Backward Scenario Report

A Backward Scenario Report is actually a Synchronization Report, which displays synchronization information about a Backward scenario.

To open a Backward Scenario Report

1. On the Report Center, select the Backward scenario from the **Available Reports per Scenario** table:

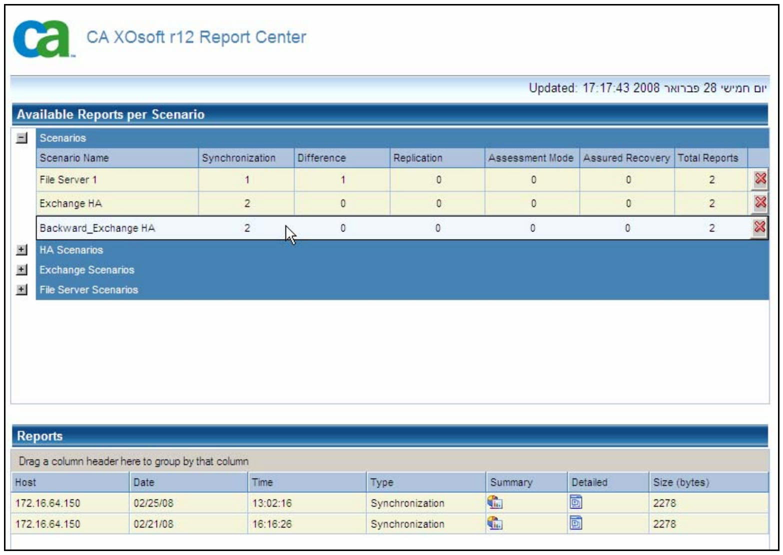

2. On the **Reports** table, select the Synchronization Report you want to view.

   The Synchronization Report for the Backward scenario appears:

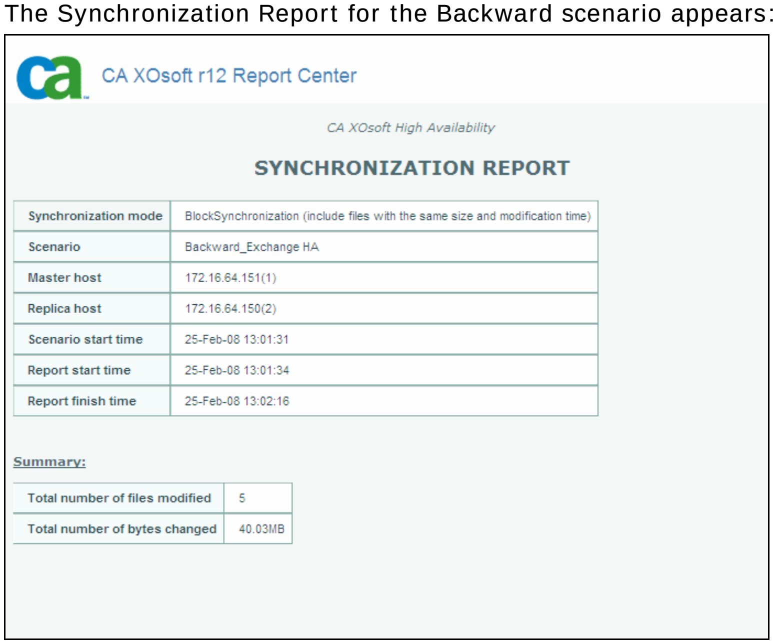

Difference Report

Sometimes, it is convenient to know the difference between a Master and its Replica(s) at a certain point in time. The comparison is performed using the same algorithms that are used in the synchronization process, but no data is transferred. A Difference Report is generated for each Replica and sent to the Manager at the end of the process. This can be produced at any time.

Important! We do not recommend initiating the Difference Report when data is being updated on the Master, since all updates that are not yet applied to the Replica will be shown as difference.

To create the difference report

1. Click the **Difference Report**  icon on the Standard toolbar, or select from the **Tools** menu the **Report**, **Difference report** option. The **Difference report** dialog opens with the same comparison methods as in the Synchronization method dialog:

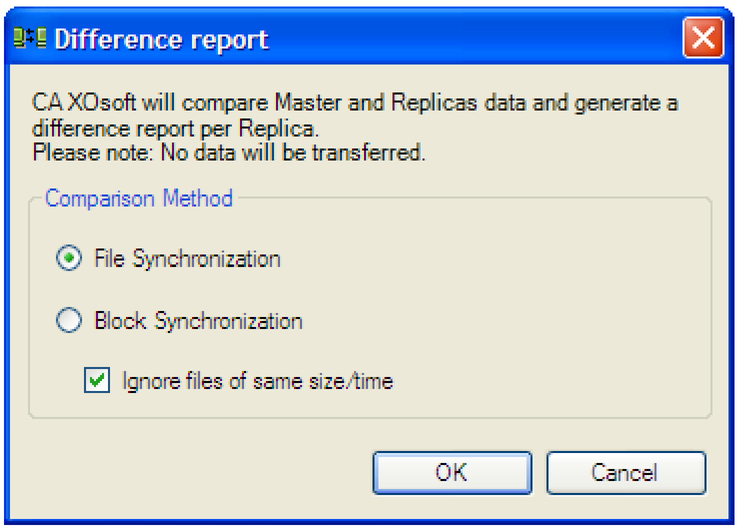

2. Choose the desired options (see *How Synchronization Works*, page 15). Then, click **OK**.

   At the end of the process, a Difference Report is generated for each eplica, and sent to the Report Center.

A Difference Report compares the differences between the Master and the Replica:

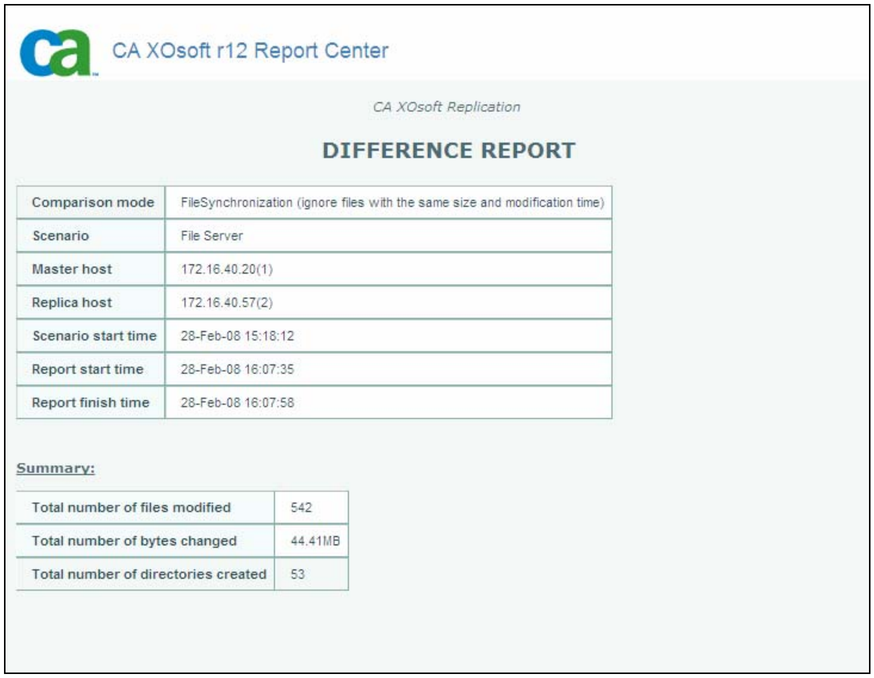

CA XOsoft r12 Report Center

CA XOsoft Replication

DIFFERENCE REPORT

Comparison mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server
Master host	172.16.40.20(1)
Replica host	172.16.40.57(2)
Scenario start time	28-Feb-08 15:18:12
Report start time	28-Feb-08 16:07:35
Report finish time	28-Feb-08 16:07:58

Summary:

Total number of files modified	542
Total number of bytes changed	44.41MB
Total number of directories created	53

Chapter 6: Editing and Managing Scenarios and Hosts

This chapter demonstrates the manual configuration of a generic File Server DR scenario, and explains the auto-discovery process for database applications. For more detailed instructions involving High Availability scenarios or scenarios tailored to specific applications such as Exchange or SQL, see the appropriate Operations Guide.

After you created a scenario using the Scenario Creation Wizard, you can manually edit most of its definitions. Although you cannot manually create a scenario from scratch, once you completed the wizard's first page, you can click the **Finish** button at any point, close the wizard and continue the configuration manually:

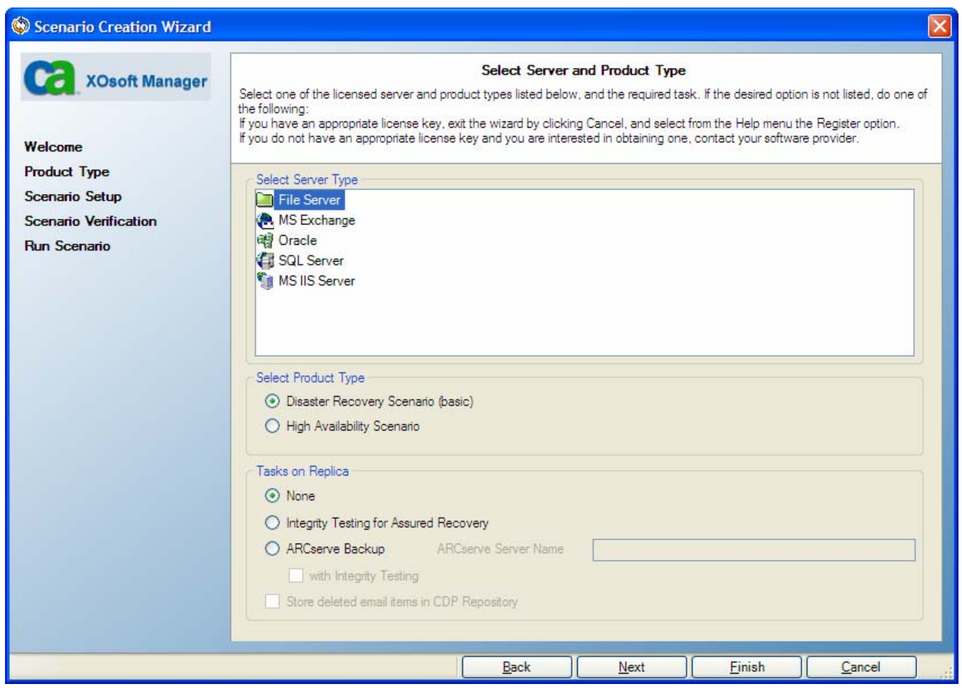

Define the Master and Replica Servers

Every scenario is automatically created with one Master and one Replica.

To define the Master or Replica host

1. On the Scenario pane, right-click the **Enter Master/Replica host name here** text and select **Rename** from the pop-up menu. Alternatively, double-click the required text:

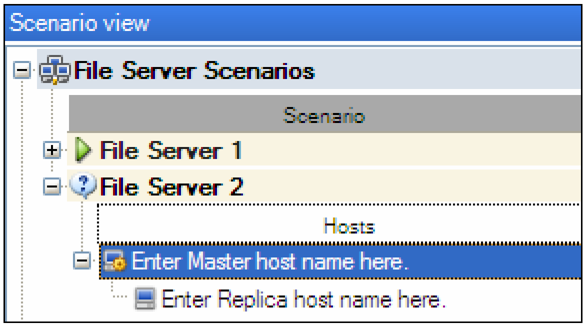

2. Enter the hostname or IP address of the host.
3. Press the **Enter** key, or click anywhere outside of the text field.
4. Save your changes by clicking the **Save** button.

   After defining a new host, you need to define its root directories for the data replication.

   - To define the Master root directories, see Select *Master Directories and Their Contents for Replication,* page 98.
   - To define the Replica root directories, see *Select Replica Root Directories*, page 111.

Add Additional Replica Servers

When creating a scenario using the Scenario Creation Wizard, you can define only one Replica server for the scenario. To add more Replica servers to the scenario, follow the instructions below.

To add additional Replica servers

1. On the Scenario pane, select the host (Master or Replica) under which you want to add a server. Right-click it and select **Insert Host** from the pop-up menu, or select **Insert Host** from the **Edit** menu.

   A new Replica server entry appears:

   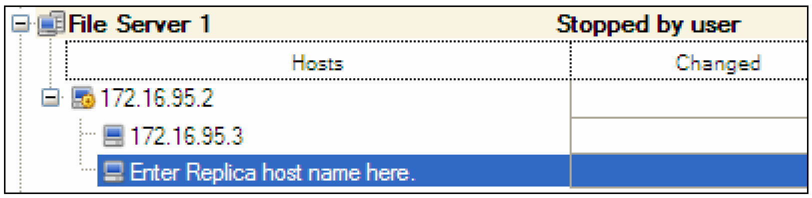

2. Define the new Replica as you defined the other Replica host, and set its properties and root directories.

3. Save your changes by clicking the **Save** button.

Select Master Directories and Their Contents for Replication

This section explains how to select directories and files on the Master for replication.

Note: Working with **Master Root Directories** dialog is possible only if the CA XOsoft Engine is installed and running on the host.

Important! Special limitations apply to remote root directories, such as mapped drives or UNC paths (\\server\share). These path types are not supported as a source for real-time replication. However, they can be the target for data replicated in real-time, meaning they can be used to store data on the Replica. In this case, these root directories can even support ACL replication. If these path types are used as the source, you are limited to scheduled synchronization only. Open files without shared read-permissions are not synchronized.

To select Master directories and their contents

1. In the Scenario pane, select the name of the Master server whose data you want to replicate.
2. In the Properties pane, click the **Root Directories** tab at the bottom. The Master Root Directories information appears in the pane:

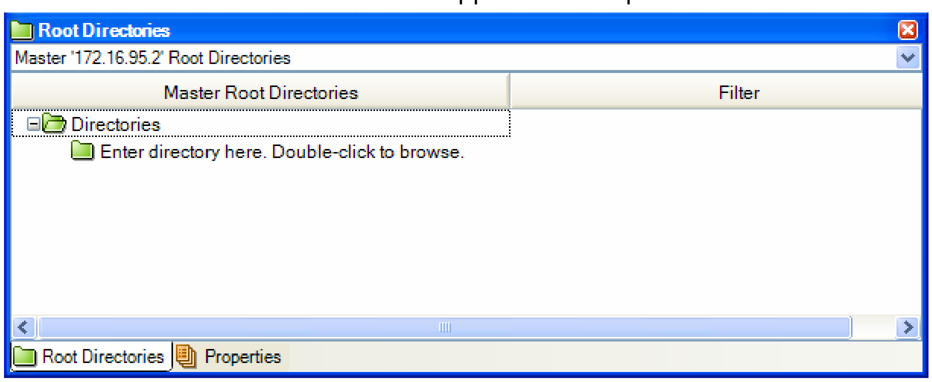

3. Right-click anywhere in the pane, and select **Browse and Select Directories**. Alternatively, double-click the Master root directory named **Directories**.

The **Master Root Directories** dialog appears:

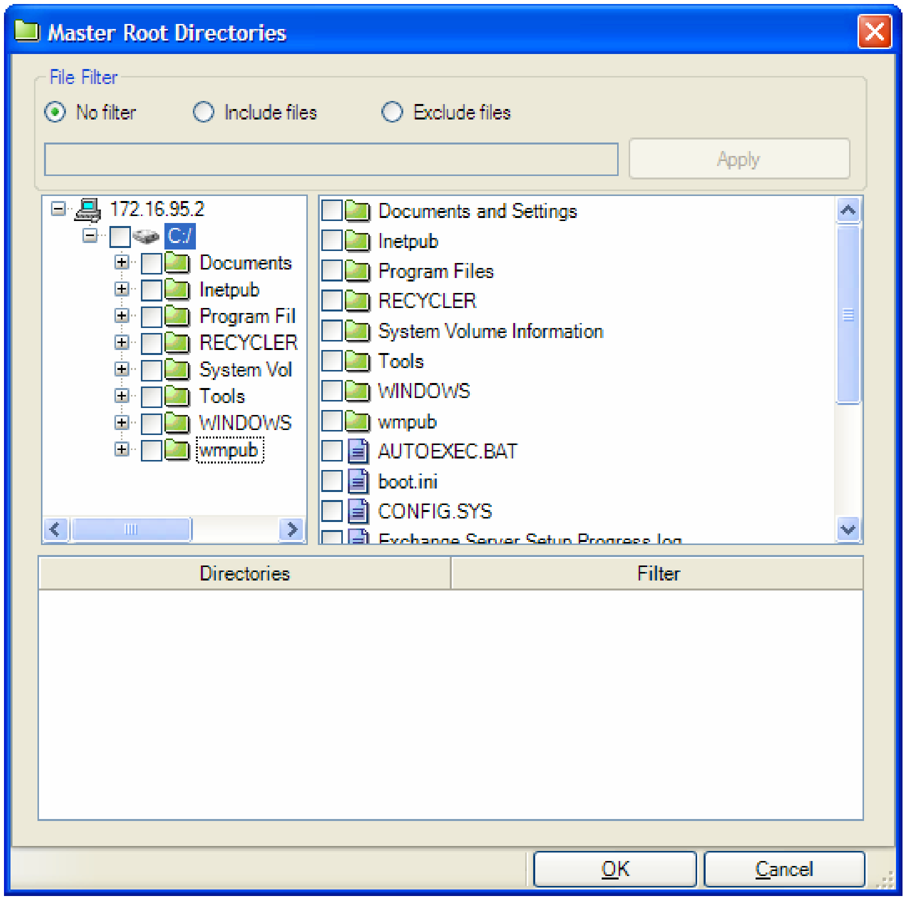

The **Master Root Directories** dialog has two areas. The left area shows only directories and sub-directories. The right area shows both directories and sub-directories, and files in those directories. The checkboxes are for you to select or clear. When selected, those directories or files will be replicated. Those not selected are ignored.

4. In the dialog's left area, select the directories that are to participate in the Master's replication scenario, by clicking on the relevant checkboxes. These are the Master root directories. The checkbox is selected, and the directory name is bold:

   Note: If a root directory is a sub-directory, it remains bold and marked, and its parent directory is marked with a grayed checkmark.

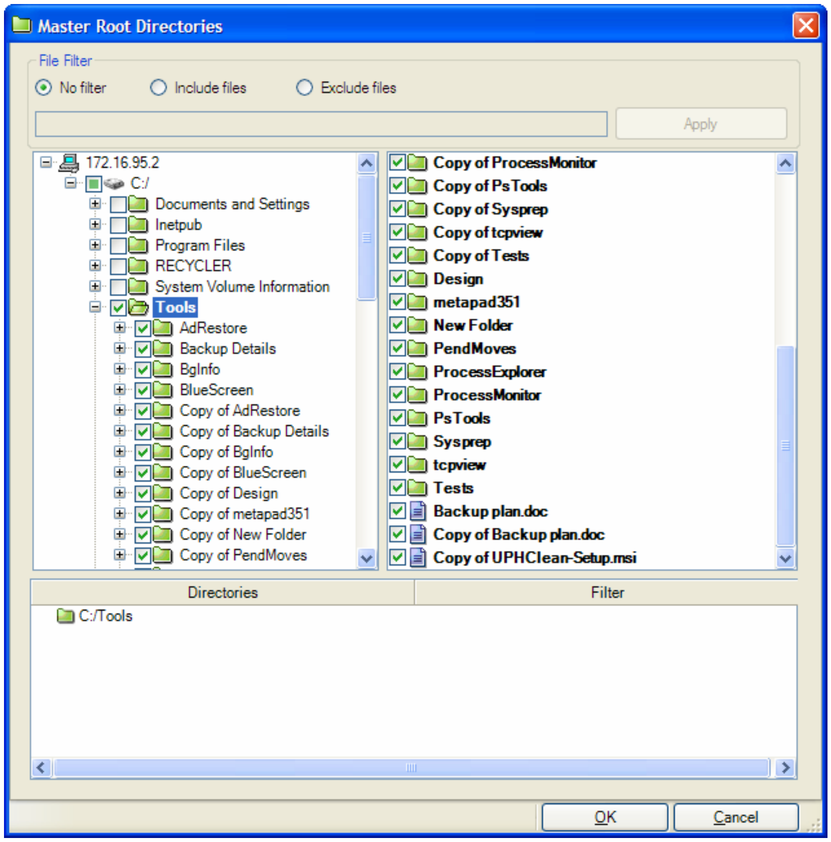

   All files and sub-directories belonging to the directory that is highlighted in the left area are displayed in the right area.

5. You can clear the check boxes of the sub-directories and specific files that appear on the right area. They are then ignored from replication.

   Note: If you clear any of the right area sub-directories and files, they are ignored, but the root directory is still selected. However, it is marked with a grayed checkmark.

6. When you have finished selecting all the directories and files you want to replicate, click **OK**.

   The selected directories now appear in the Root Directories pane under the Master root directories column.

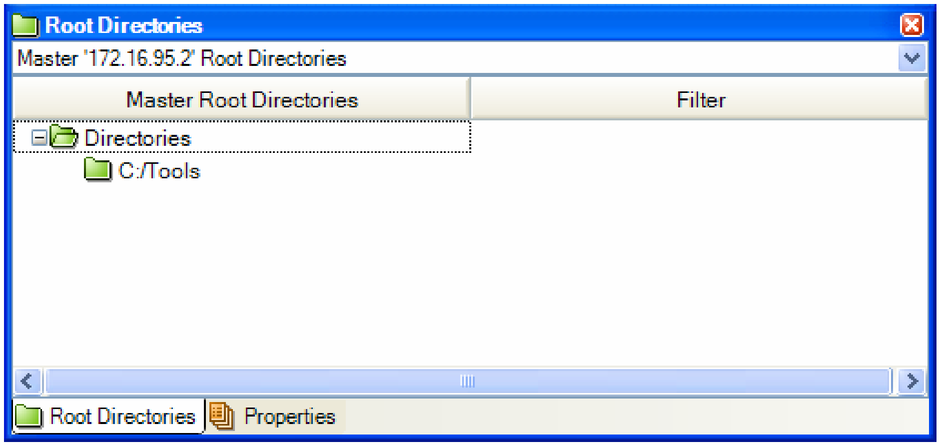

Note: When working with SQL Server DR, database that will be added to the already selected root directories after the replication starts, will not be replicated.

Edit Directory Names

You can edit the names of the Master root directories. However, when changing a root directory name, you need to verify that such a directory actually exists on the Master before you run the scenario. If you try to run a scenario with a non-existent Master root directory, the scenario will not run and a critical error will be reported.

To edit a directory name

- On the Root Directories tab, select the directory and enter a new name using Windows conventions;

  – or -

- Right-click the directory name, and select **Rename** from the pop-up menu.

Remove Master Root Directories

To remove a Master root directory

- Right-click a directory entry and select **Remove Directory** from the pop-up menu.

Filter Master Directory Files

The filtering options enable you to include or exclude files from the Master root directories. These options do not select (or clear) items in the **Master Root Directories** dialog. That has to be done manually. However, the filtering options enable you to fine-tune your directory selection and display only the files that will be replicated.

For example, if you choose to include only text files, you need to select the required directories and enter the filter parameter. Then, only text files that are stored in these directories will be displayed on the left area of the **Master Root Directories** dialog:

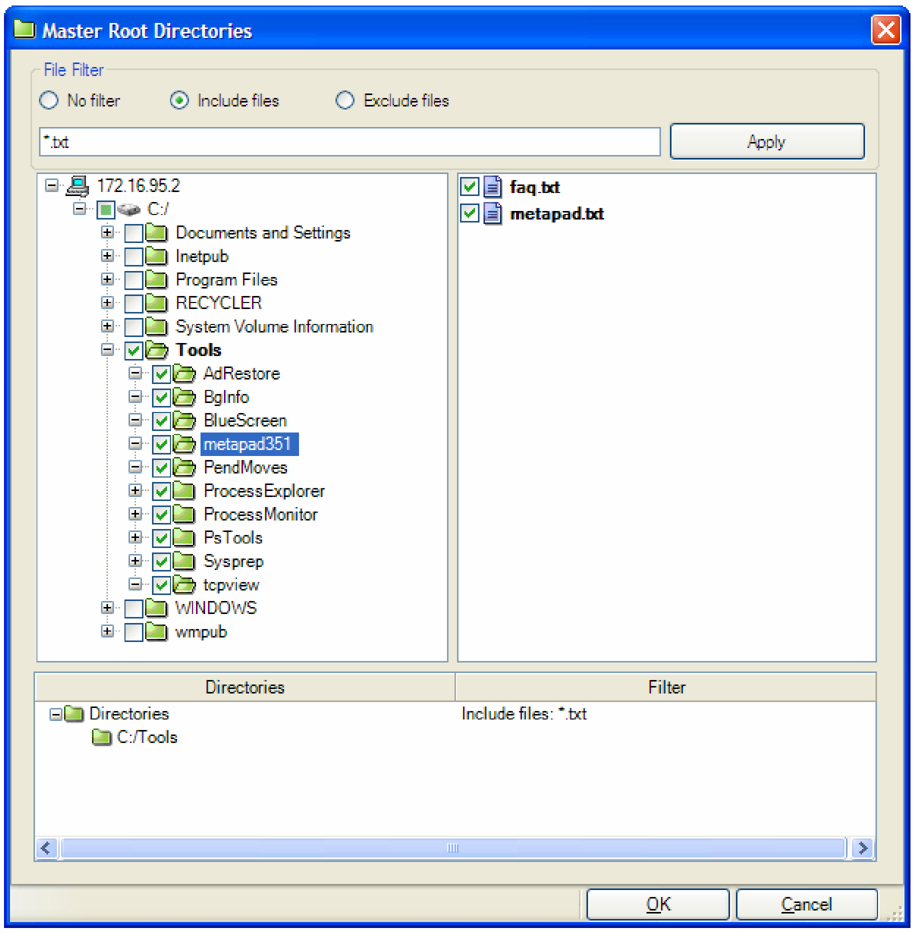

The Master root directories filters enables you to use a variety of filtering characters, such as characters, strings, wildcards, file names or extensions, etc. The following standard wildcards are available:

- An asterisk (*) selects any number of characters/digits of all types.
- A question mark (?) selects any single character or numeric digit.
- A pound sign (#) selects any numeric digit.
- An "at" sign (@) selects any single character.
- Entering other characters (one or many) selects for those specific characters.

A given filter selection applies to all files in all selected directories in the scenario.

The Filter options are as follows:

- **No filter** – all the directories and files you manually select will be replicated. This is the default option. See *Select Master Directories and Their Contents for Replication*, page 98.
- **Include files** - ONLY the selected files or file-types will be replicated. See Include Files, page 104.
- **Exclude files** - ONLY the selected files or file-types will be excluded from replication, and all others will be included. See *Exclude Files*, page 107.

Include Files

When using **Include files**, only the files or file-types entered into the Filter box are included in the replication scenario, and only if they are selected (checked). You need to manually select the directories in which these file are stored, and if you manually clear a file check box, you override the **Include files** option.

To include files

1. On the **Master Root Directories** dialog, manually select the directories on which you want to apply the filter:

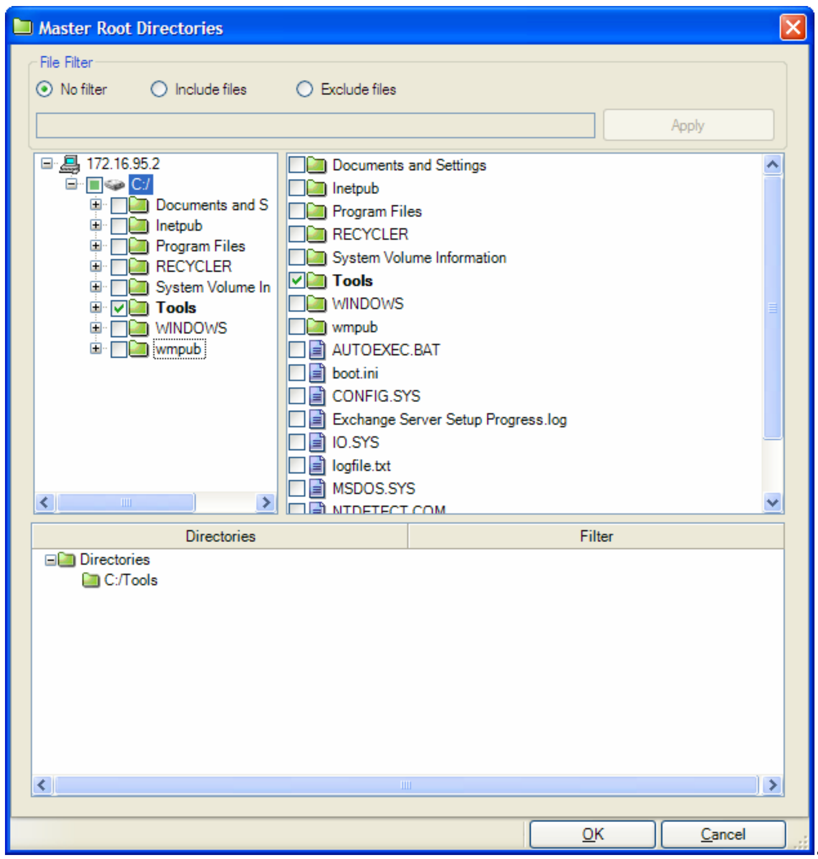

Note: Alternatively, you can manually select the directory check box AFTER you enter the filter parameters.

2. Click the **Include files** option button at the top of the **Master Root Directories** dialog. The Filter box is enabled with an asterisk (*) wildcard:

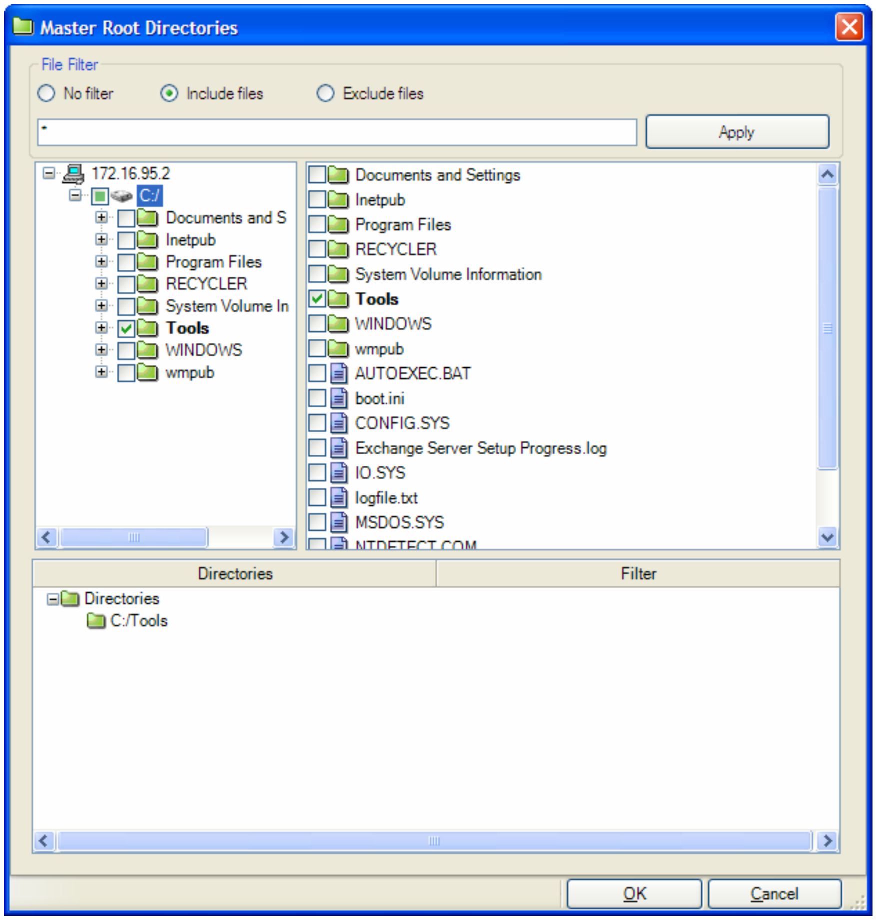

3. Enter the file types you want to include in the Filter box using the appropriate filtering characters. For example, include all files ending with the extensions *.txt *.exe. Separate the extensions using a space.

   Note: Do NOT use a comma or a semi-colon to separate extensions. If a file name includes blanks, enclose the complete file name between quotation marks ("").

4. Click the **Apply** button to filter the directories you selected according to the filter parameters:

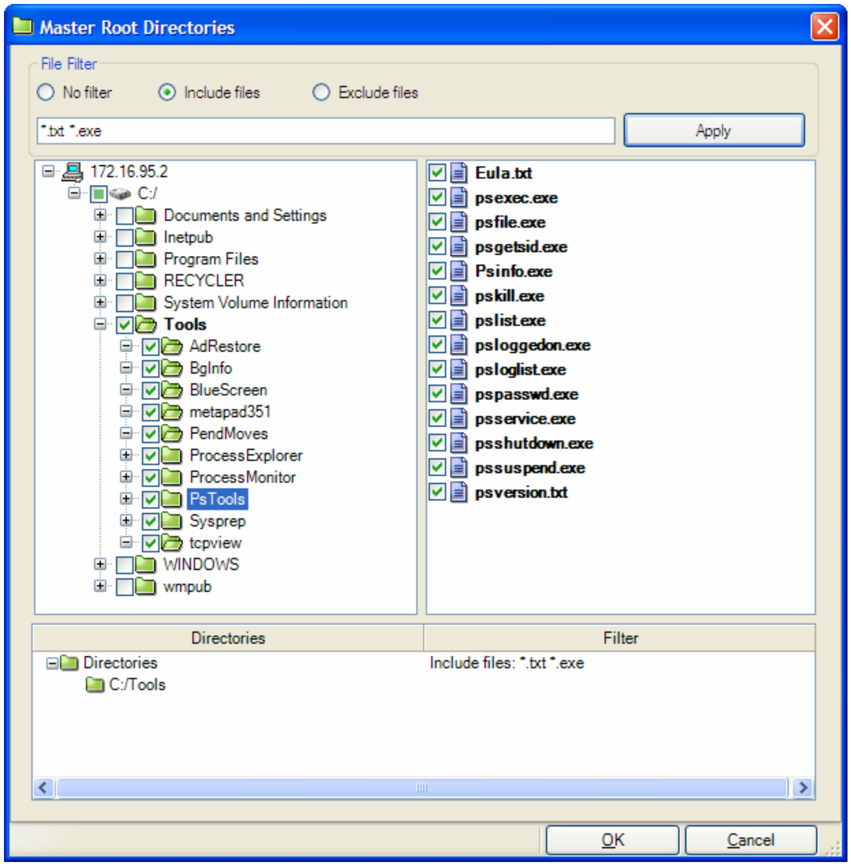

The only files that are displayed on the right area are those that meet the filtering criteria.

5. [Optional] You can manually select or clear individual directories and files. This action overrides the **Include files** option regarding the individual directory or file.

6. Click **OK** to save your root directory selection and close the **Master Root Directories** dialog.

Exclude Files

When using **Exclude files**, the replication scenario ignores any files that have been filtered out (excluded), and include all others.

To exclude files

1. On the **Master Root Directories** dialog, manually select the directories on which you want to apply the filter:

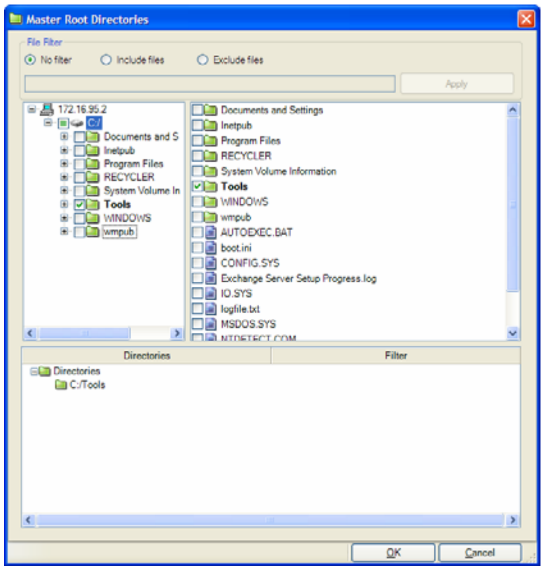

Note: Alternatively, you can manually select the directories AFTER you enter the filter parameters.

2. Click the **Exclude files** option button at the top of the **Master Root Directories** dialog. The Filter box is enabled with an asterisk (*) wildcard:

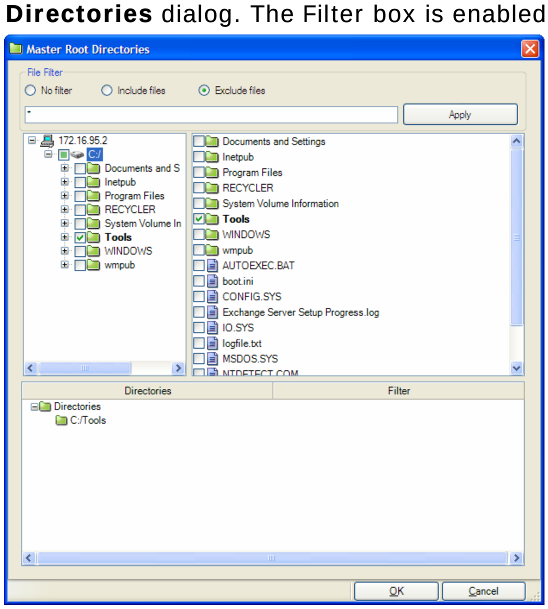

Enter the file types you want to exclude in the Filter box using the appropriate filtering characters. For example, exclude all files ending with the extensions *.bat *.dll. Separate the extensions using a space.

Note: Do NOT use a comma or semi-colon to separate extensions. If a file name includes blanks, enclose the complete file name between quotation marks ("").

3. Click the **Apply** button to filter the directories you selected according to the filter parameters:

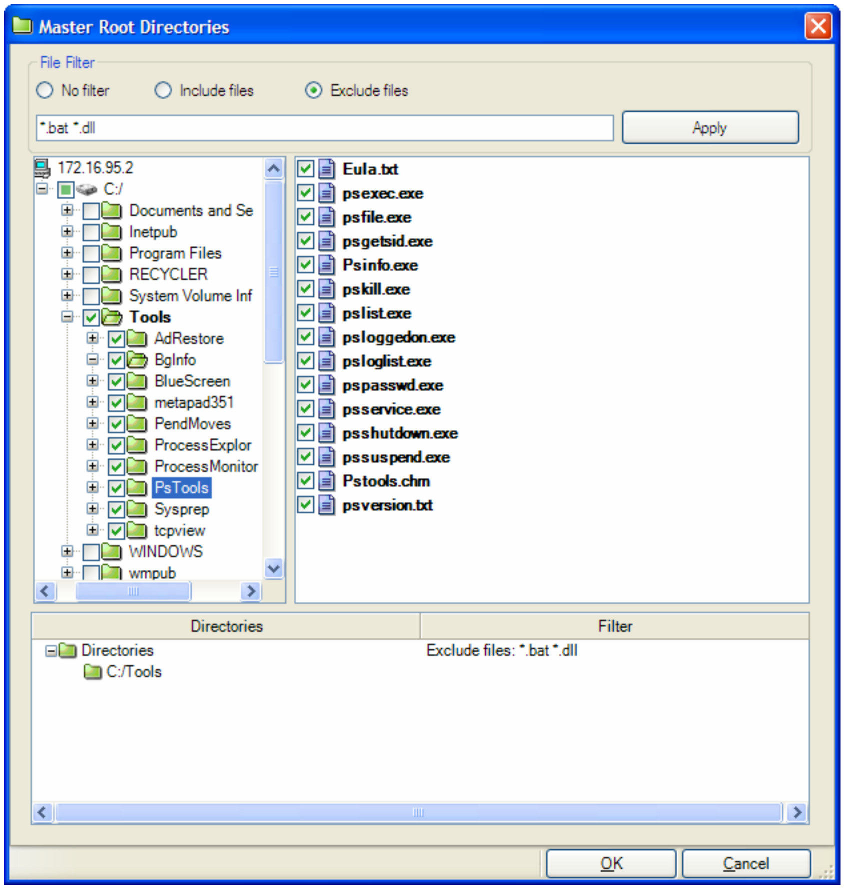

The excluded files are not displayed on the right area, and the displayed files are the ones that will be replicated.

4. [Optional] You can manually select or clear individual directories and files. This action overrides the **Include files** option regarding the individual directory or file.

5. Click **OK** to save your root directory selection and close the **Master Root Directories** dialog.

Auto-discover Database Files for all Databases

To facilitate easy directory selection for standard databases that are supported by CA XOsoft, database directories and files are identified in the scenario by using database APIs. Thus, CA XOsoft displays the structure of the database and makes the appropriate selections, which may then be easily modified if necessary.

The CA XOsoft auto-discovery function automatically discovers all database objects, related files and directories on your database or mail server — whether local or on a network. This function is currently available for all supported applications.

Note: Autodiscovery is possible only if both the CA XOsoft Engine and the database are installed and running on the Master server.

To use autodiscovery for selecting database files

1. On the scenario pane, select the scenario whose database you want to auto-discover, and verify it is NOT running.
2. On the Framework pane, open the **Root Directories** tab for the Master.

   The **Auto-discovered** database files icon appears in the Master Root Directories tab:

   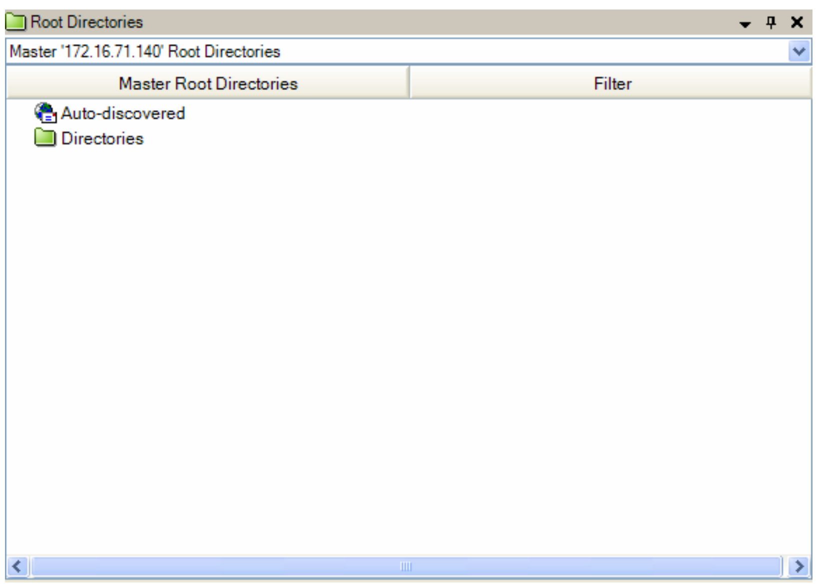

3. To start auto-discovery, double-click the **Auto-discovered** icon. Alternatively, select from the **Edit** menu the **Auto-discovery of Database Files** option.

Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

The **Autodiscovery** dialog appears:

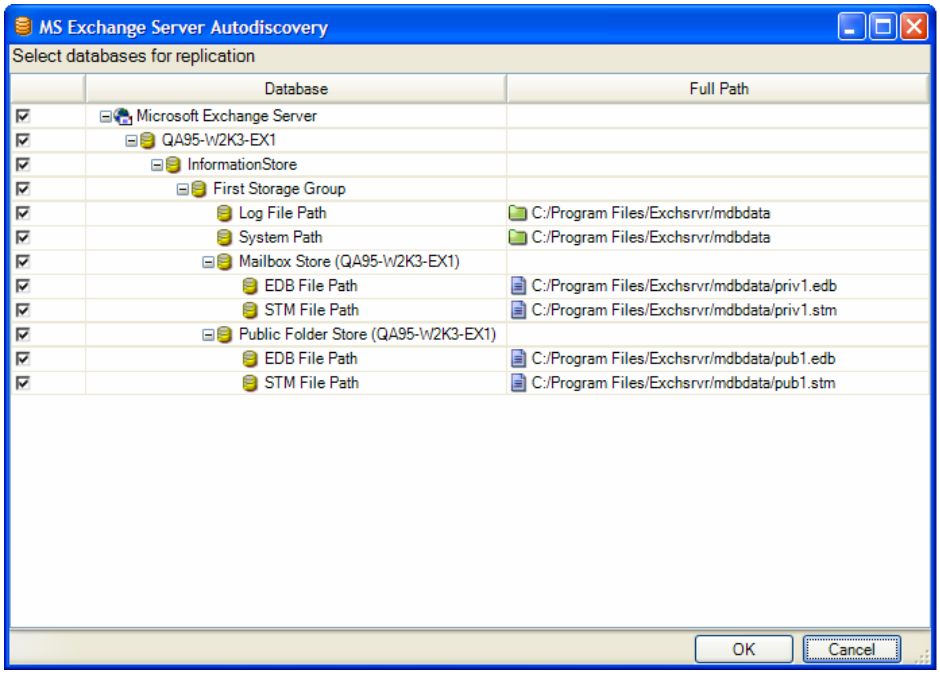

The **Autodiscovery** dialog displays all database directories and files that were auto-discovered.

4. Select the check boxes of the items you want to replicate, and clear the check boxes of the items you want to exclude from replication. Then, click **OK** to save your root directory selection and close the **Autodiscovery** dialog.

Select Replica Root Directories

You must select the Master directories before this function becomes available. For each Master root directory, you must define a Replica root directory on each related Replica.

Important! Special limitations apply to remote root directories, such as mapped drives or UNC paths (\\server\share). These path types are not supported as a source for real-time replication. However, they can be the target for data replicated in real-time, meaning they can be used to store data on the Replica. In this case, these root directories can even support ACL replication. If these path types are used as the source, you are limited to scheduled synchronization only. Open files without shared read-permissions are not synchronized. You should also take into account that using mapped drives/UNC paths as target directories will affect the performance and will slow down the synchronization and replication processes. Therefore, we do not recommend using this option.

Note: Browsing for a directory is possible only if the CA XOsoft Engine is installed and running on the selected server.

To select Replica root directories

1. In the Scenario pane, select the name of the Replica where you want store replicated data.
2. In the Properties pane, click the **Root Directories** tab at the bottom. The Replica Root Directories information appears in the pane:

   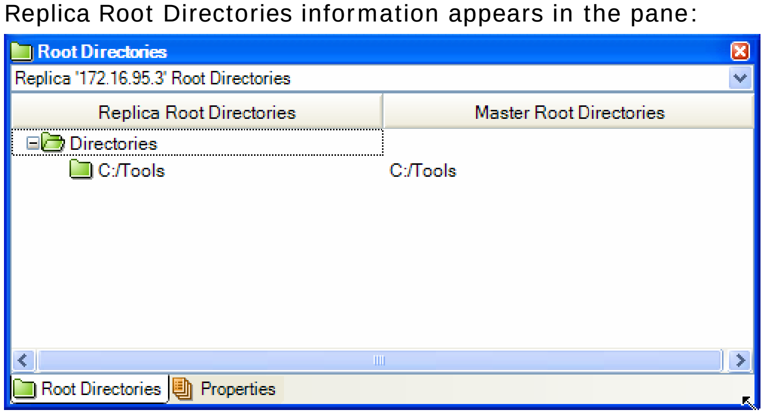

   Important! the Scenario Creation Wizard automatically configures the Replica root directories to be the same as the Master root directories. If you want to keep this configuration, ensure that your Replica server has the same drive letters as the Master server, and that the selected directories on the Replica do not contain data you want to save.

3. To change the default Replica root directories, right-click anywhere in the pane, and select **Browse and Select Directories**. Alternatively, double-click the name of the specified Replica root directory.

Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

CA XOsoft connects to the Replica server and opens the **Browse and Select Replica Directory** dialog appears:

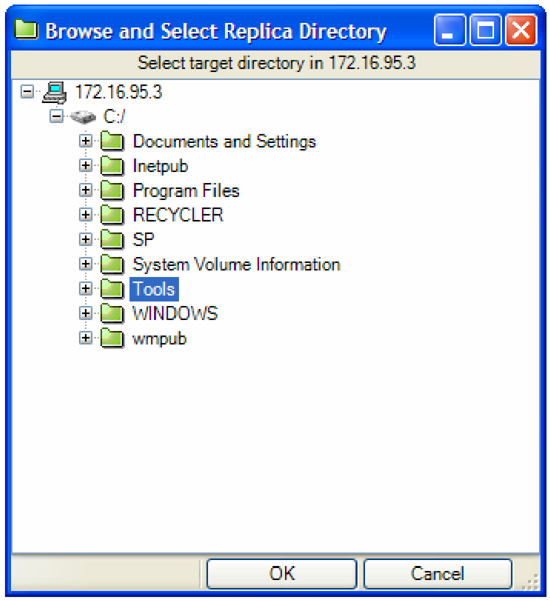

The **Browse and Select Replica Directory** dialog displays the Replica server's directory list.

4. Select a Replica directory to hold the corresponding Master directory. Repeat this for each Master directory.

5. In order for the replication process to succeed, verify that the user under which the CA XOsoft Engine is running has permission for each replication root directory.

   Note: The Replica root directory does not have to actually exist. You can enter the directory name by selecting the entry using the standard Windows conventions and CA XOsoft creates it when the replication starts.

6. Click **OK** to save your selection and close the **Browse and Select Replica Directory**.

Scenario Operations

The following sections describe scenario operations – saving, removing, exporting and importing.

Save Scenarios

There are two methods of saving scenarios, either per scenario or by a global saving of all scenarios.

To save scenarios

- On the Scenario pane, select the scenario and click the **Save** icon, or select the **Save** option from the **Scenario** menu.

  – or -

- Click the **Save All** icon or select **Save All** from the **Scenario** menu, to save all scenarios on the Manager.

Remove Scenarios

Important! Before removing a scenario, make sure you want to permanently delete it. There is no undo action.

To remove a scenario

1. On the Scenario pane, select the scenario and right-click it.

   A pop-up menu appears.

2. From the pop-up menu, select the **Remove** option.

   A message appears asking you to confirm the removal:

   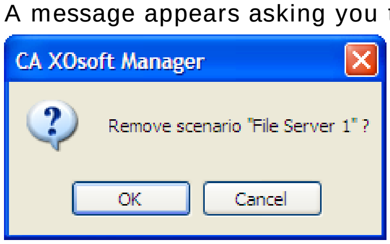

3. Click **OK**. The scenario is permanently removed.

Export Scenarios

You can export scenarios to other locations in order to reuse them.

To export a scenario

1. On the Scenario pane, select the scenario you want to export. Then, right click it and select **Export**, or select the **Export** option From the **Scenario** menu.

   The **Export scenario** dialog is opened:

   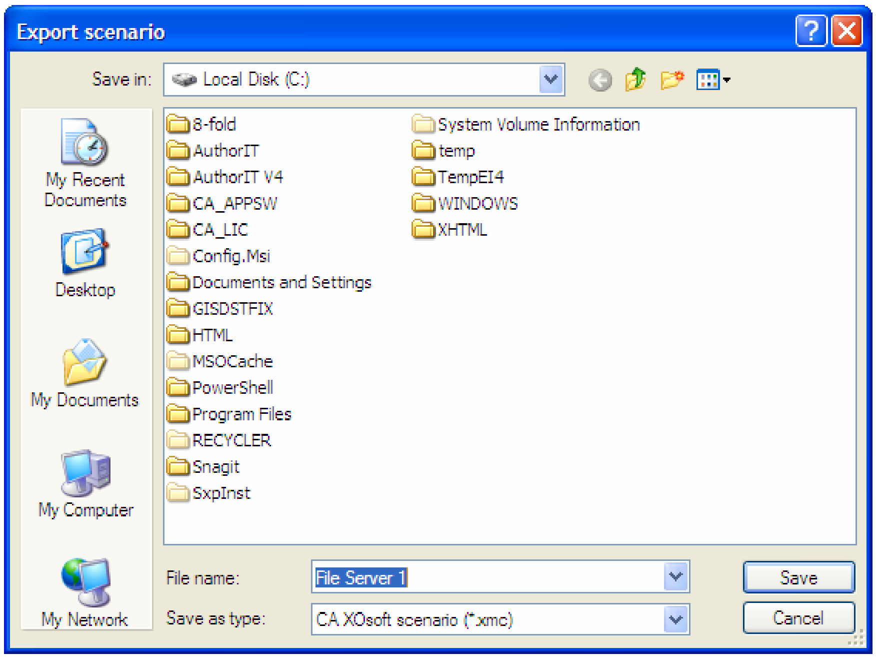

2. Name the scenario and save it by clicking the **Save** button.

   The scenario is saved as a *.xmc file.

Import Scenarios

You can import .xmc files that contain saved scenarios to your Manager. Use this option if you want to relocate scenarios from one workstation to another, or if you want to use older scenarios that were kept in your system.

To import a scenario

1. From the **Scenario** menu, select the **Import** option.

   An **Import scenario** dialog appears:

   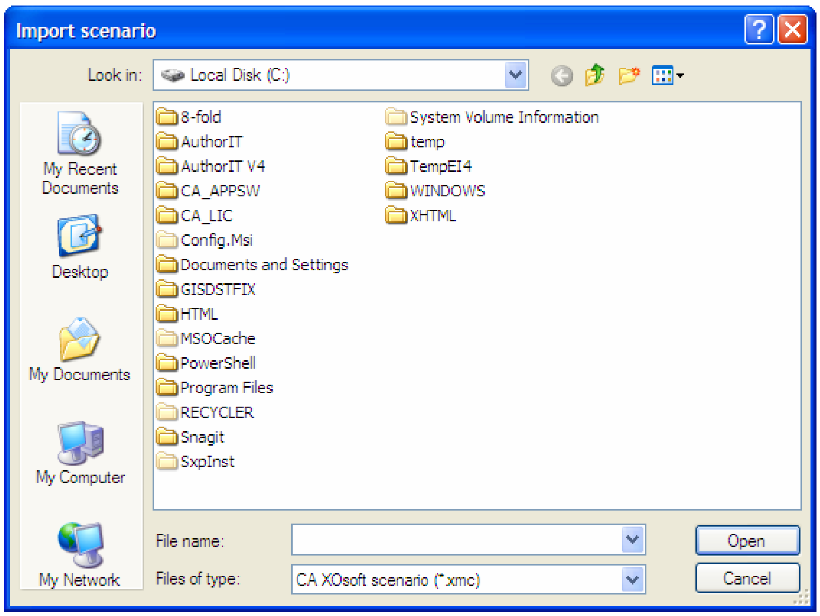

2. Locate the scenario you want to import and click **Open**.

   The scenario is imported to the Manager and appears on the Scenario pane.

   Note: If you are trying to import a scenario that has the same ID or name as an existing scenario, the following dialog appears:

   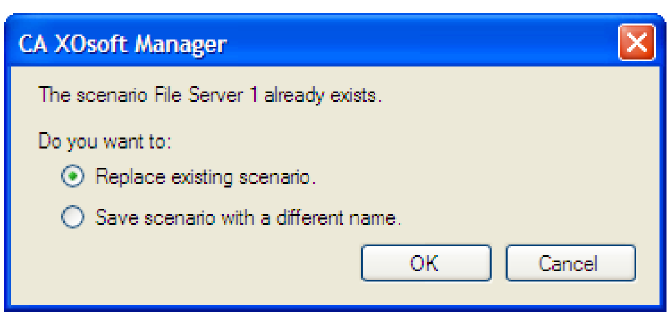

   Select your required option and click **OK**.

Host Maintenance

The Host Maintenance option enables you to perform maintenance procedures, such as rebooting a host or moving groups between MS Cluster nodes, without performing re-synchronization once these processes are completed. Usually, when the online replication process is critically interrupted, there is a need to compare the data between the source and target hosts and make it identical, to ensure the data integrity before the replication can continue. This resynchronization process consumes time and resources. The Host Maintenance option enables you to prepare your replicated system for planned maintenance procedures and avoid resynchronization.

The hosts that can be prepared for maintenance need to participate in running scenarios. The preparation is done on one host at a time, but this host can participate in multiple scenarios. In these scenarios the host can function both as the Master and the Replica. When a host participates in a scenario that is not running, the preparation that relates to this scenario will not occur. For example, a host can participate in both File Server and Exchange scenario. If before you start preparing the host, the File Server scenario is not running, only the Exchange services will be stopped during the preparation and the Server shares will remain intact.

When the selected host functions as the Master, during the preparation process either the DB services or File shares are stopped, depending on the scenario type. Then, all changes that occurred until that moment are passed on to the Replica. Once the Replica sends to the Master an acknowledgment that all changes were applied and the data integrity is ensured, the scenario is suspended and the host is ready for maintenance. When the selected host functions as the Replica, the changes that were sent to it are applied, and the Master stops sending new changes. The new changes are saved in the meantime in the Master's spool for future update. Then, the scenario is suspended and the host is declared as ready for maintenance.

Once the maintenance procedures are completed, CA XOsoft seamlessly resumes the real-time replication, avoiding any delay or disruption that data re-synchronization may cause.

Important! This option applies to Database and File Server applications. It supports both DR and HA scenarios. However, when using this option for File Server scenarios, and you have applications that are running locally on the host you want to reboot, you need to manually stop them before starting the host maintenance preparation, and manually restart them after the maintenance is completed.

To prepare your hosts for maintenance procedures

1. On the Scenario pane, verify that the scenarios whose host you want to reboot are running.

   Notes:

   - You do not have to run all the scenarios in which the host participates. The preparation will be done only on the parts that involve the running scenario, for example, Exchange services in the case of Exchange scenario.

   - The host maintenance preparation cannot be performed during synchronization. If a scenario currently synchronizes, wait until its completion.

2. Click the **Launch Host Maintenance** button, or select **Launch Host Maintenance** from the **Tools** menu.

   The **Host Maintenance** wizard appears:

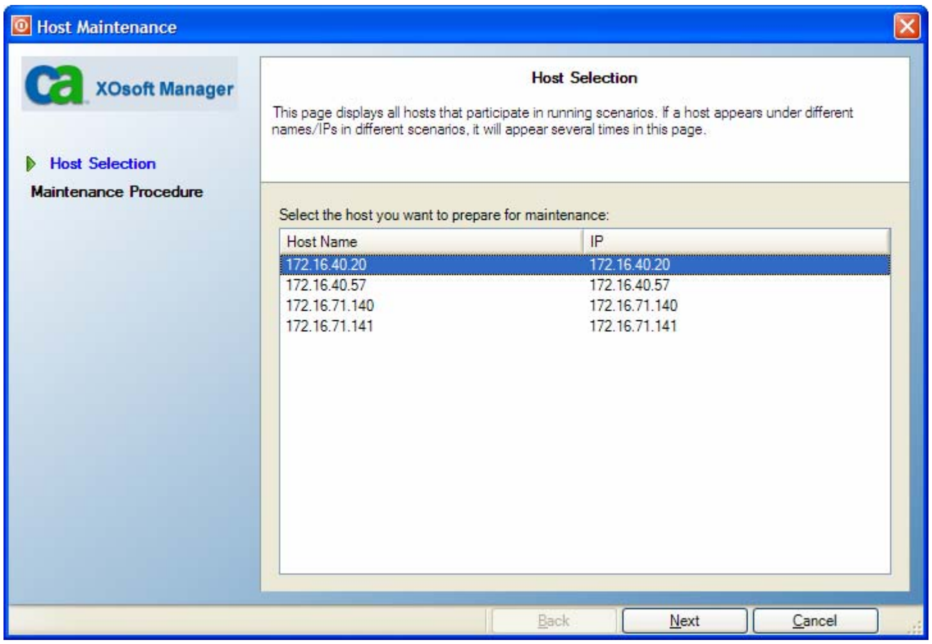

   The **Host Maintenance** wizard displays all hosts that participate in the running scenarios.

   Note: If the same host appears under different names/IPs in different scenarios, it will appear several times in this page.

3. Select the host you want to prepare for maintenance, and click **Next**.

The **Maintenance Procedure** page appears:

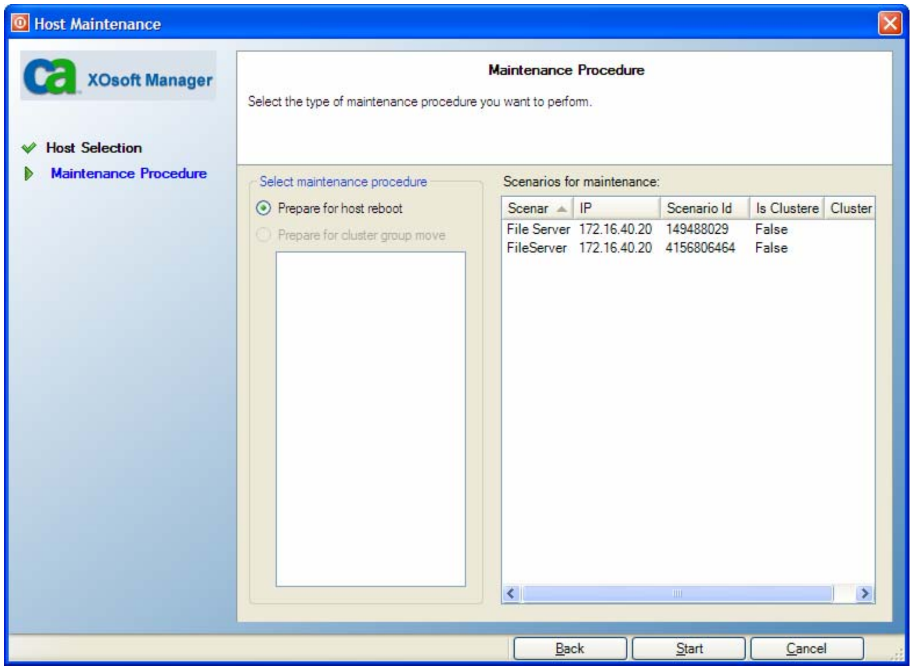

The **Maintenance Procedure** page displays the details of the scenarios in which the selected host participates.

4. On the **Select maintenance procedure** section on the left, select the operation you want to perform and click the **Start** button.

   On the Event pane, a message appears saying: **Preparing for reboot**. Then, another message appears saying: **Ready for reboot**.

   Note: If a message appears saying: **Not Ready for Reboot**, it means that the preparation did not succeed, and after you reboot the host, re-synchronization will be performed.

   Simultaneously, on the Scenario pane the scenario's state is changing to **Ready for HM**:

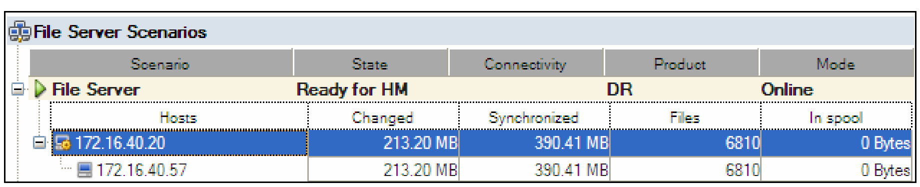

   Note: The scenario's state that appears on the Scenario pane refers only to the Master host's state. Therefore, if the host you are preparing for maintenance is functioning as the Replica, you will not see its changing status in the Scenario pane, only on the Event pane and the Host Maintenance Monitor.

5. To view the status of the selected host and the scenarios in which it participates, select from the **View** menu the **Active View**, **Host Maintenance Monitor** option, or click the **Host Maintenance Monitor**  button.

   The **Host Maintenance Monitor** view appears:

   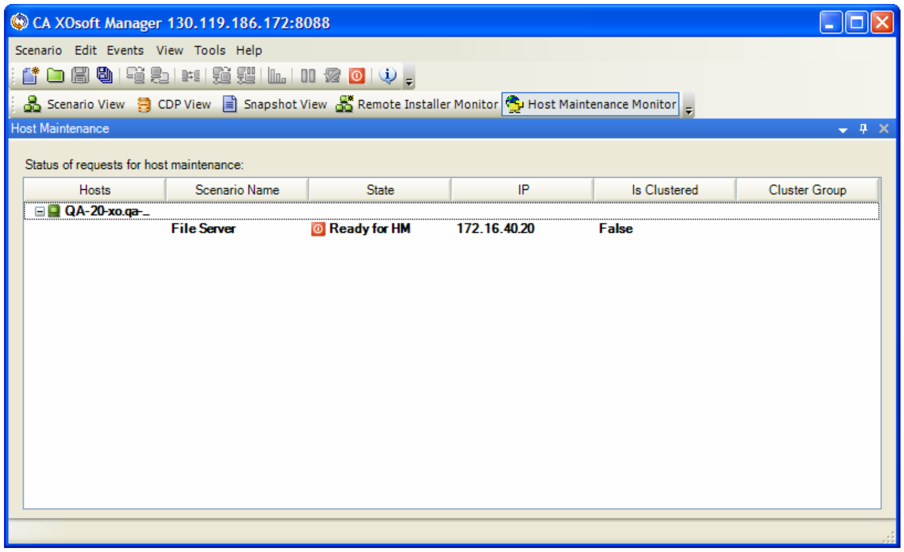

   The Host Maintenance Monitor displays all the requests for maintenance preparation. A maintenance request disappears from the Monitor when the involved scenario is either stopped or run. You cannot perform actions through the Monitor, which only displays information about the status of the current requests. The only action you can do is opening the Host Maintenance wizard by clicking anywhere in the screen and selecting **Launch Host Maintenance**.

   In this Monitor, the displayed host name is its fully qualified name, and not the name under which it appears in the scenarios. All the scenarios in which this host participates appear in the Monitor.

6. After you received the message informing you that the host is ready for reboot, you can reboot your host or switch groups between cluster nodes. Once you completed your maintenance procedures, the replication process automatically resumes, without performing re-synchronization.

   Note: If after preparing the host for maintenance, you decided not to reboot it and continue running its scenarios, you need to stop the scenarios and re-run them.

Chapter 7: Setting Scenario Properties

This chapter describes how to configure scenario properties, and provides the list of the scenario properties, their corresponding values, and an explanation of each property.

Configuring Scenario Properties

The scenario property values determine the entire scenario's default behavior concerning synchronization method, replication mode, event notification, reporting, and more.

Notes:

- The Properties pane and its tabs (Root Directories, Properties, Statistics) are context sensitive, and change whenever you select a different node from a scenario folder.
- To configure scenario properties, the scenario must be stopped.

Each scenario is identified by its specific Product type, Server type and unique Scenario ID. The values of these items cannot be modified.

To set scenario properties

1. On the Scenario pane, select the scenario its properties you want to configure. On the Framework pane on the left, the Scenario Properties list appears:

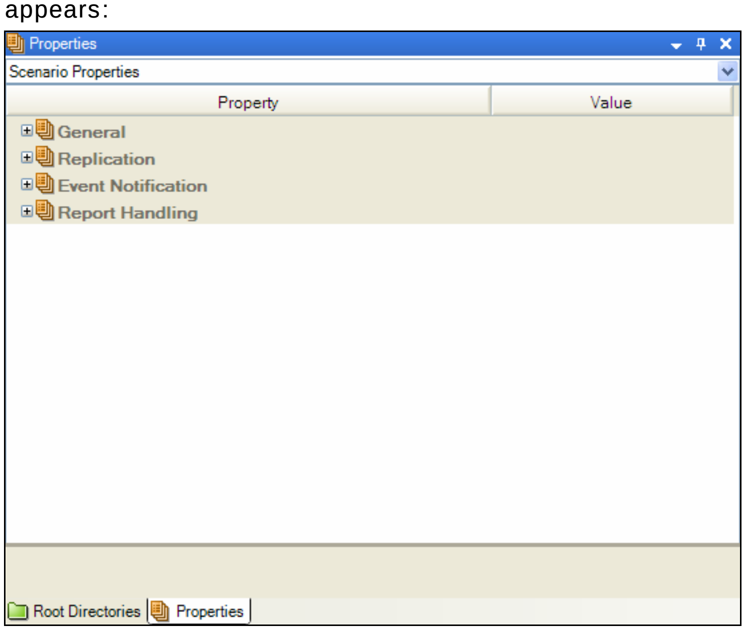

Note: A running scenario has a gray background, and scenarios that are not running have a white background.

2. If the scenario is running, click the **Stop** button on the toolbar. The scenario is stopped.
3. On the Scenario Properties list, open the desired group, select the required property, and select or enter the appropriate values. Some values can be manually entered in an edit box field, while others can be selected from a combo box or IP control by clicking the default value.
4. After you set the required properties, click the **Save** button on the Standard toolbar to save and apply your changes.

Understanding Scenario Properties

This section lists the scenario properties, corresponding values, and provides an explanation for each property.

General

The properties in this group cannot be changed. The Product Type and Server Type properties are set during the creation of a new scenario. The Scenario ID property is given automatically by the system. To change these properties you need to create a new scenario.

Product Type

Either DR (Disaster Recovery) or HA (High Availability).

Server Type

The type of application or database server that participates in the scenario.

Scenario ID

The unique ID of the scenario.

Replication

Mode

CA XOsoft supports two replication modes – Online and Scheduled.

Online

The Online replication mode activates real-time replication. Changes are replicated continuously, in real-time, using the XOMF driver.

The Online mode replicates all changes of files, even if files that are always open (as is the case of most database and mail servers). This mode maintains the order of file system operations. In this mode, the Engine records all I/O operations related to the root directories in journal files. The journal files are then sent to the Replicas where the operations that were recorded in the journal are replayed on the replicated files.

Scheduled

Scheduled replication mode is really nothing more than an automatically performed synchronization. The synchronization can be initiated either by a manual activation, or according to a pre-defined schedule, say every few hours or once a day. This replication mode does not differ in principle from a synchronization performed as part of initializing replication. Although there is no online replication in this mode, online changes made during synchronization are replicated.

When the **Scheduling** option is selected, two options are enabled:

- **By User Request**

  Synchronization is activated by a user running synchronization from CA XOsoft Manager, or CA XOsoft PowerShell.

- **Weekly Activity**

  In the Weekly Activity setting of the Scheduled mode, servers are synchronized at a scheduled, fixed time. When you select this option, you need to set the schedule for the recurring synchronization.

  For a detailed description of synchronization scheduling, see *Schedule Synchronization*, page 128.

Run after Reboot

If the Master is rebooted, CA XOsoft automatically re-synchronizes the Master and the Replica after the reboot.

Automatic Synchronization

Synchronization ensures that a set of folders and files on a Replica server involved in a scenario is identical to the set on the Master. There are two types of synchronization: File and Block synchronization.

Synchronization Type

- **File Synchronization**

  Compares files on the Master and Replica servers, and whenever they are different, copies the entire missing or modified files from Master to Replica.

  In order to synchronize your data for the first time, you need to choose the File Synchronization mode for each scenario. In subsequent cases, this synchronization type is best suited for File Server (a large number of small and medium sized files), preferably with the checked **Ignore Files of Same Size/Time** option. This may significantly decrease synchronization time.

- **Block Synchronization**

  Performs a block-by-block comparison of the Master and Replica files, and copies over only those blocks that are different. When differences exist between files, instead of requiring the transfer of the entire file, the block synchronization transfers only changes.

  Block synchronization is the method appropriate for database applications, such as MS Exchange, Oracle or SQL Server. You should use this method while clearing the **Ignore Files of Same Size/Time** option (unless database files are closed on the Master server).

Ignore Files of Same Size/Time

Skips comparison of files with the same path, name, size and modification time. This assumes that the files are identical.

Select this option when you are absolutely sure that files of this type are indeed identical. This option is best suited for File Server scenarios. It is not appropriate for database files for such applications as Exchange, SQL, or Oracle, since these databases modify files (leaving them open) without changing file modification time. You can use this option in database scenarios only when synchronized databases are un-mounted and files are closed on the Master server.

This option can reduce the overall synchronization time dramatically, but you must remember that this comes at the expense of content verification.

Optional Settings

Replicate NTFS Compress Attribute

(For Windows only) Replicates compress attribute of files or directories during synchronization and replication. .

Replicate NTFS ACL

(For Windows only) Replicates ACLs for files and directories during synchronization and replication.

Synchronize Windows Shares

If a directory has been set to allow sharing, then setting this Share option to On duplicates the property in the replicated directory. This occurs only during synchronization and on Windows OS.

Prevent Automatic Re-sync upon Error

An uncommon critical error on the Master can stop the replication continuance. In this case, setting this option to On prevents automatic re-synchronization. When this option is Off, re-synchronization starts automatically upon an error occurrence.

Event Notification

Notification

When an event occurs, you can set the system to run a script or send an email notification.

Notify by Email

Defines whether to send the details of an event by email to a specified address. If several events occur immediately one after the other, the system aggregates them and sends their details in one email.

Mail Server

Enter the mail server hostname or IP.

Email Address – To

Enter the receiver email address.

Email Address – From

Enter the sender email address.

Execute Script

Specifies a script for CA XOsoft to run whenever an event occurs.

Script Name (full path)

Enter the name and full path of the script that is invoked once an event occurs.

Arguments

Additional arguments to pass to the script, which is specified in the previous property. Any arguments entered here follow the argument sent automatically by CA XOsoft, which include the event details written in a notification file. Arguments entered here are static values.

Write to Event Log

Writes the events to the Windows event log.

Report Handling

Report Saving

Enter the report saving settings.

Report Directory

Specifies the location where the reports are saved.

Report Retention (days)

Specifies the number of days to retain replication reports. The default is Unlimited.

Notify by Email

Defines whether to send reports by email to the specified address.

- **Mail Server**

  Enter the mail server hostname or IP.

- **Email Address – To**

  Enter the receiver email address.

- **Email Address – From**

  Enter the sender email address.

Execute Script

Specify a script for CA XOsoft to run whenever it generates a report.

- **Script Name (full path)**

  Enter the name and full path of the script that is invoked once a report is generated.

- **Arguments**

  Additional arguments to pass to the script specified in the previous property. Any arguments entered here follow the argument sent automatically by CA XOsoft. This argument defines the full path of the generated report file and its type. Arguments entered here are static values.

Schedule Synchronization

When selecting a scheduled replication mode it means that synchronization will be performed automatically on a regular basis. Once you select this option the following flexible scheduling capabilities are offered:

- Synchronization on selected days of the week and for specific hours in a 24 hour cycle.
- Synchronization over selected periods (e.g., once every 36 hours) in a 7 day cycle.
- Exclusion of specific dates.

To open the schedule

1. On the Scenario Properties list, open the **Replication** group. On the **Mode** property, select the **Scheduling** value.

   The **Schedule** option appears.

2. On the **Schedule** option, select the **Weekly Activity** value. Then, on the **Schedule Settings** property, click the **Not Set** value:

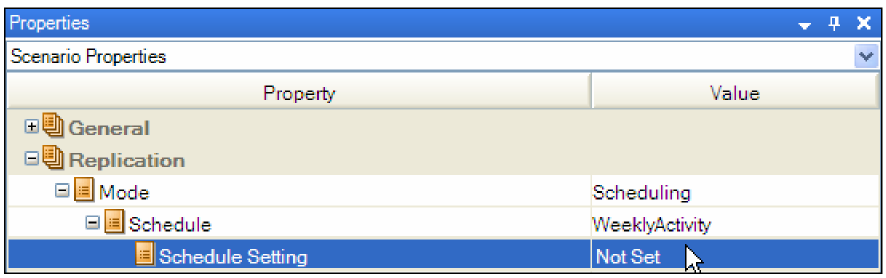

   The **Schedule Setting** dialog appears:

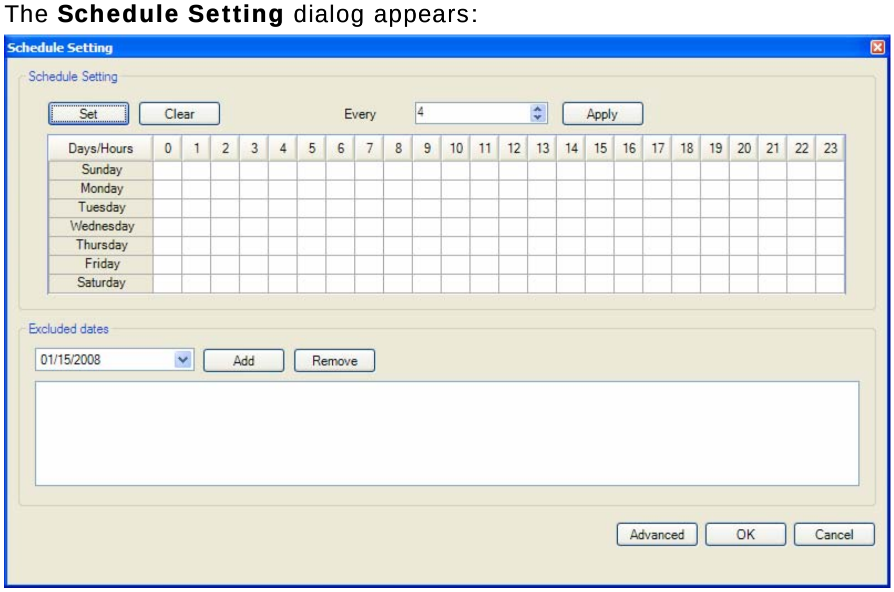

3. Set the schedule for automatic synchronization according to the guidelines described in the sections below.

   The following figure shows a sample scheduling setting in which the solid blue rectangles indicate the days and hours scheduled for synchronization. The excluded dates area lists the specific dates on which synchronization is not performed:

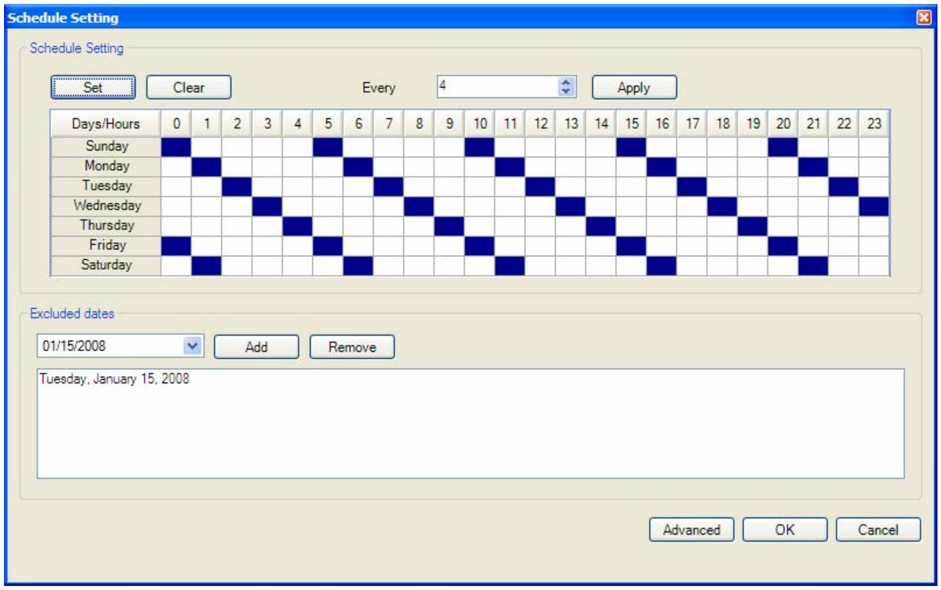

Set a Schedule for Automatic Synchronization

The following instructions describe how to set and clear hours and days in the **Schedule Setting** dialog for automatic synchronization.

To set a specific hour/day

- Select a single rectangle for a specific hour/day. Click the **Set** button to actually mark and enable that hour/day.

To set a specific hour for each day in the week

- Select a column and click the **Set** button.

To set each hour in a specific day in the week

- Select a row and click the **Set** button.

To set a repetitive cycle

- Enter a valid number of hours in the **Every** box, and click the **Apply** button.

Note: More than one rectangle can be simultaneously set by clicking and dragging the mouse. You can also use the **Ctrl** and **Shift** keys to set several dates at once.

To clear a setting

- Use the same technique of selecting, and click the **Clear** button.

Important! If synchronization is running, and the next scheduled synchronization hour comes up, the new synchronization stops the current one and starts again from the beginning.

Exclude Dates from Scheduled Synchronization

You can set specific dates that will be excluded from the automatic synchronization.

To exclude dates from automatic synchronization

- In the **Schedule Setting** dialog, on the **Excluded dates** section select or enter the excluded date in the **dates** box. Then, click the **Add** button.

  The selected date appears in the **Excluded dates** list:

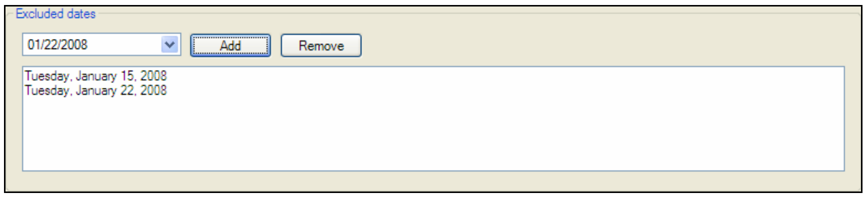

To remove an excluded date

- On the **Excluded dates** list, select the entry, and click the **Remove** button. Multiple entries can also be selected by dragging the mouse over them.

Set Advanced Schedules

The **Advanced Schedule Setting** page allows you to set times that are not on the hour:

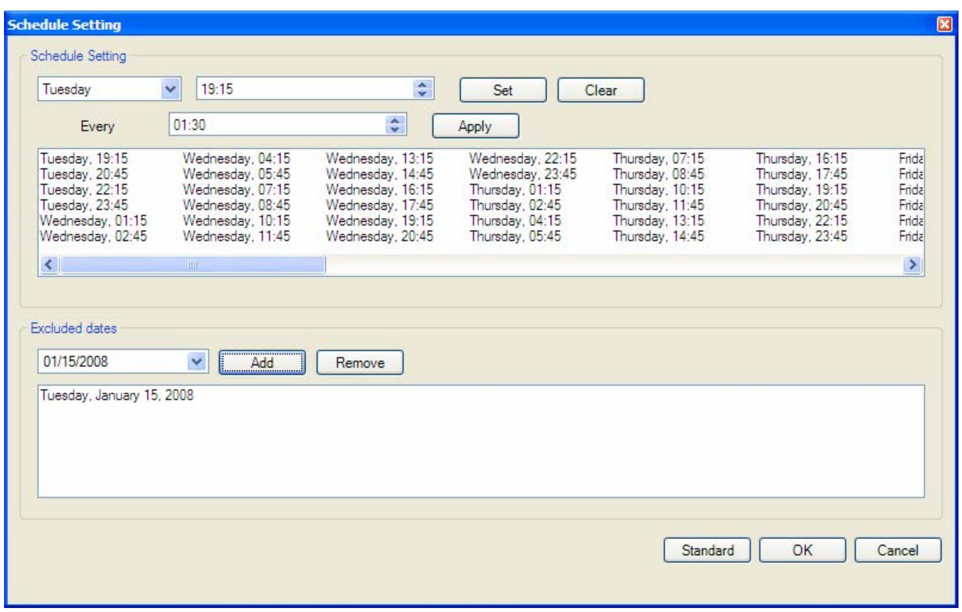

To open the Advanced Schedule Setting page

- In the **Schedule Setting** dialog, click the **Advanced** button at the bottom.

To return to the Standard Schedule Setting dialog

- In the **Advanced Schedule Setting** page, click the **Standard** button at the bottom.

Chapter 8: Setting Master and Replica Properties

This chapter describes how to configure the Master and Replica properties, and provides the list of their properties, the corresponding values, and an explanation of each property.

Configuring Master or Replica Server Properties

To configure Master or Replica properties, the scenario must be stopped.

To set Master or Replica properties

1. On the Scenario pane, select the Master or Replica whose properties you want to configure.

   On the Framework pane on the left, the Master/Replica Properties list appears:

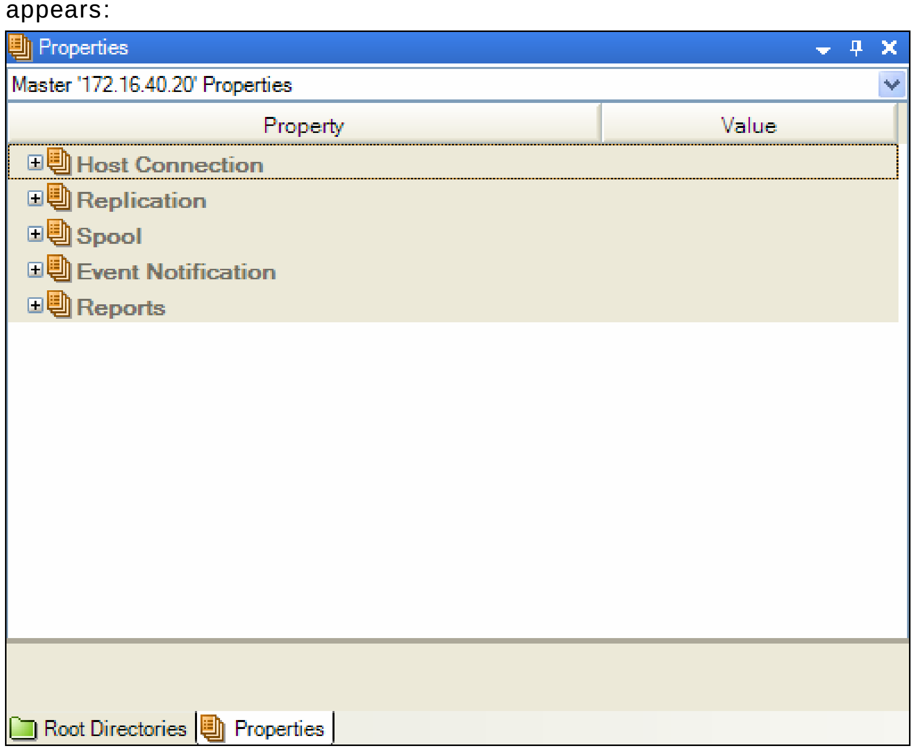

Note: A running scenario has a gray background, and scenarios that are not running have a white background.

2. If the scenario is running, click the **Stop** button on the toolbar. The scenario is stopped.
3. On the Master/Replica Properties list, open the desired group, select the required property, and select or enter the appropriate values. Some values can be manually entered in an edit box field, while other values can be selected from a combo box or IP control by clicking the default value.
4. Click the **Save** button on the toolbar to save and apply your changes.

Understanding Master Properties

This section lists the Master properties, corresponding values, and provides an explanation for each property.

Host Connection

IP Address

Enter the IP address of the Master host. If the Master name is changed, the IP address is updated. The Master can also be changed by entering another IP address in this field.

Port Number

Enter the number of the incoming port used for TCP communications. It can be changed to any unused port. Since the Engine uses only one port, make sure that the Engine uses the port specified here. The default port number is 25000.

Fully Qualified Name

Enter the fully qualified name of the Master host.

In HA scenario, this value is always synchronized with the value in the Switchover Properties pane.

Replication

Run Script before Synchronization

Triggers a script to run before each synchronization. The synchronization process does not start until this script run is completed.

Script Name

Enter the full name and path of the script.

Arguments

Arguments to pass to the script that is specified in the previous property. Arguments are static values.

Run Script after Synchronization

Triggers a script to run after each synchronization. The synchronization process does not wait for this script run to finish.

Script Name

Enter the full name and path of the script.

Arguments

Arguments to pass to the script that specified in the previous property. Arguments are static values.

Compress Data during Transfer

Compresses data before sending it to the Replica. This option optimizes bandwidth and transfer time. If the Master host is a busy production server, we recommend activating this feature on the first-level Replica that resides on the same LAN, and not on the Master.

- Compress data is resource consuming, and impacts server performance. If the typical file format being transferred does not compress much, this option is a waste of processing power and time. Although the transmission bandwidth can be lessened, the overall replication time is a function of compressibility and available power.
- Already compressed files such as .zip, .rar, .gz, .jpeg, etc., and any small file whose size is less than 512 bytes, are not compressed.

Run Script upon Trigger File Creation

[For File Server only] Defines whether special actions should be triggered via a script, when a specified trigger file appears.

Trigger File Name

The name of the file that triggers the script, which is specified in the next property. The script is triggered once the file creation event occurs.

Script to Run

Script Name

This script is invoked, when the trigger file that is specified in the previous property is created. Enter the full name and path of the script.

Arguments

Arguments to be passed to the script specified in the previous property. Arguments must be static values.

Spool

The CA XOsoft spool is a folder on a disk where data to be replicated is temporarily stored (that is, spooled). The spool stores changes captured during replication for some period of time before applying them to the Replica server. The spool parameters determine how much disk space is available for the spool. In most cases, the default values are sufficient. However, if you choose to change this value, it should be at least 10% of the total dataset size.

Max Spool Size

Enter the maximum spool size allowed. This disk space is used only if needed – it is not pre-allocated. The default is Unlimited. To enter a value of **Unlimited**, enter a zero.

Min Disk Free Size

Enter the free disk space threshold at which the system issues an error and stops replication.

Spool Directory

Enter the directory to be used to store the spool. The default directory is *INSTALLDIR/tmp* on Windows.

Important! If you change the spool location, remember to remove the new path from file level antivirus scans: both scheduled and real time.

Event Notification

Notification

When an event occurs, you can set the system to run a script or send an email notification.

Notify by Email

Defines whether to send the details of an event by email to a specified address. If several events occur immediately one after the other, the system aggregates them and sends their details in one email.

- **Mail Server**

  Enter the mail server hostname or IP.

- **Email Address – To**

  Enter the receiver email address.

- **Email Address – From**

  Enter the sender email address.

Execute Script

Specifies a script for CA XOsoft to run whenever it sends a report.

- **Script Name (full path)**

  Enter the name and full path of the script that is invoked once an event occurs.

- **Arguments**

  Additional arguments to pass to the script, which is specified in the previous property. Any arguments entered here follow the argument sent automatically by CA XOsoft, which include the event details written in a notification file. Arguments entered here are static values.

Write to Event Log

Writes the events to the Windows event log.

Reports

Generate Synchronization Report

Specifies whether to generate a synchronization report.

Generate Detailed Report

Specifies whether to generate a detailed synchronization report.

Generate Replication Report

Specifies whether to generate a replication report. Since replication is continuous, specify the frequency of the report generation in the property below.

Generation Frequency (hours)

Specifies how often to generate the replication report.

Generate Detailed Report

Specifies whether to generate a detailed replication report.

Report Handling

Notify by Email

Specify whether to send reports by email to the specified address.

- **Mail Server**

  Enter the mail server hostname or IP**.**

- **Email Address – To**

  Enter the receiver email address.

- **Email Address – From**

  Enter the sender email address.

Execute Script

Specify a script for CA XOsoft to run whenever it sends a report.

- **Script Name (full path)**

  Enter the name and full path of the script that is invoked once a report is generated.

- **Arguments**

  Additional arguments to pass to the script specified in the previous property. Any arguments entered here follow the argument sent automatically by CA XOsoft. This argument defines the full path of the generated report file and its type. Arguments entered here are static values.

Understanding Replica Properties

This section lists the Replica properties, corresponding values, and provides an explanation of each property.

Host Connection

IP Address

Enter the IP address of the Replica host. If the host name is changed, the IP address is updated. The host can also be changed by entering another IP address in this field.

Port Number

Enter the number of the incoming port used for TCP communications. It can be changed to any unused port. Since the Engine can use only one port, make sure that the Engine uses the port specified here. The default port number is 25000.

Fully Qualified Name

Enter the fully qualified name of the Master host.

In HA scenario, this value is always synchronized with the value in the Switchover Properties pane.

Replication

Run Script before Synchronization

Triggers a script to run before each synchronization. The synchronization process does not start until this script run is completed.

Script Name

Enter the full name and path of the script.

Arguments

Arguments to pass to the script that is specified in the previous property. Arguments are static values.

Run Script after Synchronization

Triggers a script to run after each synchronization. The synchronization process does not wait for this script's run to finish.

Script Name

Enter the full name and path of the script.

Arguments

Arguments to pass to the script that specified in the previous property. Arguments are static values.

Compress Data during Transfer

Compresses data before sending it to the Replica. This option optimizes bandwidth and transfer time.

- Compress data is resource consuming, and impacts server performance. If the typical file format being transferred does not compress much, this option is a waste of processing power and time. Although the transmission bandwidth can be lessened, the overall replication time is a function of compressibility and available power.
- Already compressed files such as .zip, .rar, .gz, .jpeg, etc., and any small file whose size is less than 512 bytes, are not compressed.

Keep Deleted Files during Synchronization

During synchronization, do not remove from the Replica files that were deleted from the Master. Best suited for cases in which several scenarios use the same Replica directories.

Keep Deleted Files during Replication

During replication, do not remove from the Replica files that were deleted from the Master.

Stop Database on Run

When set to On, if a database scenario (Exchange, SQL, Oracle) is running and the database is running on the Replica server, CA XOsoft stops the database services before running the scenario. [Does not apply to HA scenarios]

Retry if File is Busy

These options are relevant only for Windows servers. If changes were received for a busy file (opened as non-shared for read), these options define how many times and at what interval to attempt replacing this file with the one that contains the changes.

Number of Attempts

Enter the number of attempts to be made for replacing a modified file that is busy (and therefore cannot be replicated). If the file is not released before the last attempt is made, the change is lost and an error message is initiated.

Interval between Attempts (msec)

The time between an unsuccessful attempt and the next attempt.

Run Script upon Trigger File Creation

[For File Server only] Defines whether special actions should be triggered via a script, when a specified trigger file appears.

Trigger File Name

Enter the name of the file that triggers the script, which is specified in the next property. The script is triggered once the file creation event occurs.

Script to Run

Script Name

This script is invoked, when the trigger file that is specified in the previous property is created. Enter the full name and path of the script.

Arguments

Arguments to be passed to the script specified in the previous property. Arguments must be static values.

Note: You cannot change from here the size of the allowed incoming bandwidth. If you want to change the default, which is Unlimited, use the ws_rep.cfg file, and enter new bandwidth limit in Kbps.

Spool

The spool parameters determine how much disk space is available for the spool. In most cases, the default values are sufficient. However, if you choose to change this value, it should be at least 10% of the total dataset size.

Max Spool Size

Enter the maximum spool size allowed. This disk space is used only if needed – it is not pre-allocated. The default is Unlimited. To enter a value of **Unlimited**, enter a zero.

Min Disk Free Size

Enter the free disk space threshold at which the system issues an error and stops replication.

Spool Directory

Enter the directory to be used to store the spool. The default directory is *INSTALLDIR/tmp* on Windows.

Important! If you change the spool location, remember to remove the new path from file level antivirus scans: both scheduled and real time.

Recovery

Replication Delay

Data replication can be delayed in the parent server (Master or Replica) spool before sending it to this Replica. This is useful against data corruption or viruses. It enables stopping replication before corrupted or infected data is written to the Replica.

Delay Interval (min)

Enter the number of minutes of the replication delay.

Data Rewind

Keeps undo information needed to recover data from a certain action or point in time. It is useful in cases in which corrupted data on the Master was replicated to the Replica, and you want to restore the data to its previous state before the corruption occurred. Data Rewind is enabled for online replication only.

Retention Period (min)

I/O operations are recorded in the Rewind journal for this number of minutes. Then, they are discarded in first-in-first-out order.

Max Disk Size (MB)

Enter the maximum disk space allocated for the Rewind journal. Once this size is reached, old records are discarded in FIFO order.

Scheduled Tasks

Suspend

See *Schedule Replication Suspension*, page 73.

Replica Integrity Testing for Assured Recovery

See *Setting Assured Recovery Properties*, page 213.

Event Notification

Notification

When an event occurs, you can set the system to run a script or send an email notification.

Notify by Email

Defines whether to send the details of an event by email to a specified address. If several events occur immediately one after the other, the system aggregates them and sends their details in one email.

- **Mail Server**

  Enter the mail server hostname or IP.

- **Email Address – To**

  Enter the receiver email address.

- **Email Address – From**

  Enter the sender email address.

Execute Script

Specifies a script for CA XOsoft to run whenever it sends a report.

- **Script Name (full path)**

  Enter the name and full path of the script that is invoked once an event occurs.

- **Arguments**

  Additional arguments to pass to the script, which is specified in the previous property. Any arguments entered here follow the argument sent automatically by CA XOsoft, which include the event details written in a notification file. Arguments entered here are static values.

Write to Event Log

Writes the events to the Windows event log.

Reports

Generate Replication Report

Specifies whether to generate a replication report. Since replication is continuous, specify the frequency of the report generation in the property below.

Generation Frequency (hours)

Specifies how often to generate the replication report.

Generate Detailed Report

Specifies whether to generate a detailed replication report.

Generate Assured Recovery Report

Specifies whether to generate Assured Recovery report.

Report Handling

Notify by Email

Specify whether to send reports by email to the specified address.

- **Mail Server**

  Enter the mail server hostname or IP**.**

- **Email Address – To**

  Enter the receiver email address.

- **Email Address – From**

  Enter the sender email address.

Execute Script

Specify a script for CA XOsoft to run whenever it sends a report.

- **Script Name (full path)**

  Enter the name and full path of the script that is invoked once a report is generated.

- **Arguments**

  Additional arguments to pass to the script specified in the previous property. Any arguments entered here follow the argument sent automatically by CA XOsoft. This argument defines the full path of the generated report file and its type. Arguments entered here are static values.

Chapter 9: Data Recovery

This chapter describes how to restore lost data using CA XOsoft Manager, how to set bookmarks and how to rewind data.

The Data Recovery Process

When an event causes loss of Master data, the data can be restored from any Replica. The recovery process is in fact a synchronization process in the reverse direction - from a Replica to the Master.

CA XOsoft enables you to either recover all lost data from the Replica to the Master, or to recover data only from a certain event or point in time to a previous valid state. The latter option, called Data Rewind, is similar to an undo action. It is useful in cases in which corrupted data on the Master was replicated to the Replica, and you want to restore the data to its previous state before the corruption occurred. The rewind process is done using time-stamped checkpoints and user-defined bookmarks.

Important! You must stop replication in order to initiate recovery.

Recover Lost Data from Replica

Recover all lost data from a Replica

1. On the Manager, from the Scenario pane select the desired scenario and stop it.
2. [For database applications only] stop the database services on the Master host.
3. On the Manager, from the scenario folder select the Replica host:

   Note: If multiple Replica servers participate in the required scenario, select the Replica from which you want to recover data.

   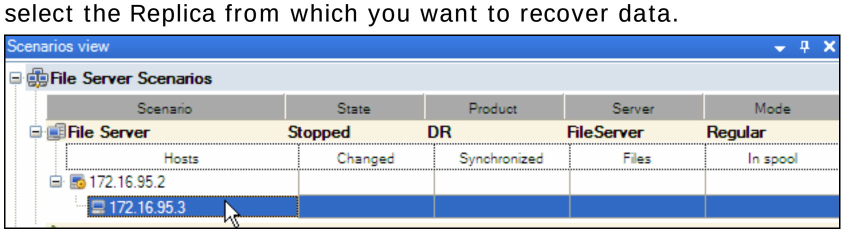

   The **Restore Data** option is enabled.

4. From the **Tools** menu, select **Restore Data**, or click the **Restore data** button on the Standard toolbar:

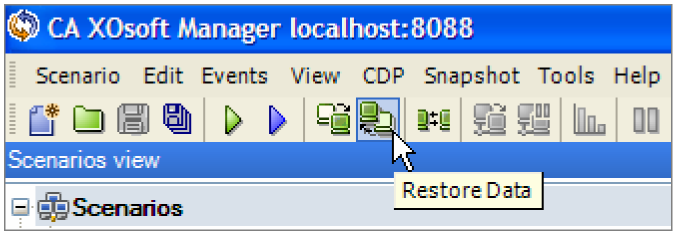

The **Recovery Method** page of the Restore Data wizard appears:

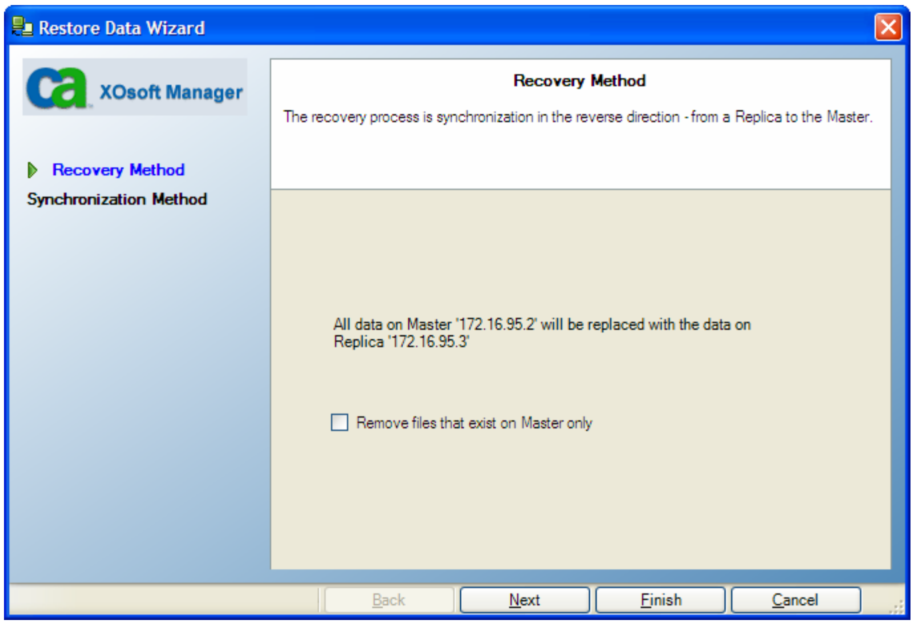

Note: If the Data Rewind property is set to On, another Restore Data dialog will appear (see step 4, page 150). In this case, select the first option – **Replace all data on Master with the data on Replica**.

5. Click **Next**. The **Synchronization Method** page appears:

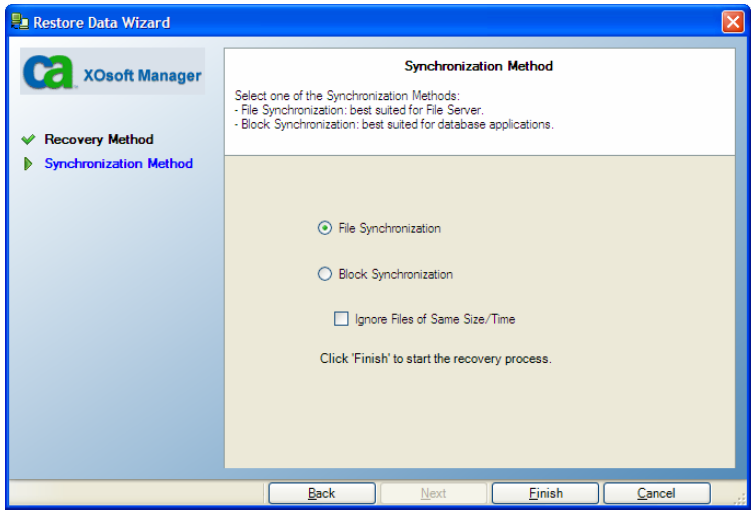

6. Make sure that the **File Synchronization** method is selected, and click **Finish.**

   Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

   Once you finished initiating the recovery process, CA XOsoft builds a temporary reverse tree using the selected Replica as the root, and the Master as the terminating node. After the Master recovery process ends, the temporary scenario is deleted, and you receive the following message in the Event pane: **Synchronization finished**.

7. By default, once a data recovery occurs a Synchronization Report is generated:

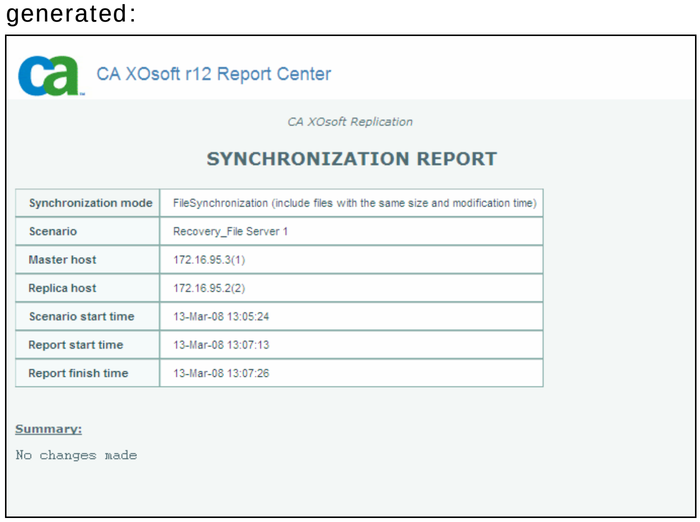

CA XOsoft r12 Report Center

CA XOsoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (include files with the same size and modification time)
Scenario	Recovery_File Server 1
Master host	172.16.95.3(1)
Replica host	172.16.95.2(2)
Scenario start time	13-Mar-08 13:05:24
Report start time	13-Mar-08 13:07:13
Report finish time	13-Mar-08 13:07:26

Summary:

No changes made

Now, the replication process can restart following the original scenario.

Setting Bookmarks

A *bookmark* is a checkpoint that is manually set to mark a state that you may want to rewind back to. We recommend setting a bookmark just before any activity that may cause data to become unstable. Bookmarks are set in real-time, and not for past events.

Notes:

- You can use this option only if you set in the Replica Properties list the **Recovery – Data Rewind** option to **On** (see page 141).
- You can not set bookmarks during the synchronization process.

To set a bookmark

1. When the required scenario is running, select from the **Tools** menu the **Set Rewind Bookmark** option.

   The **Rewind bookmark** dialog appears:

   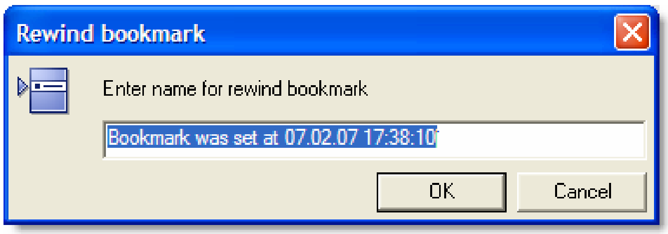

   The text that appears in the **Rewind bookmark** dialog will appear in the **Rewind Points Selection** dialog as the bookmark's name (see step 7, page 151). The default name includes date and time.

2. Accept the default name, or enter a new name for the bookmark. It is recommended to give a meaningful name that will later help you recognize the required bookmark. Then, click OK.

   The bookmark is set.

Data Rewind

The Data Rewind recovery method allows you to rewind data to a point in time before it was corrupted. The rewind process takes place on the Replica server before the reverse synchronization process starts. The Data Rewind method uses rewind points or bookmarks that enable you to reset the current data back to a previous state.

You can use this option only if you set the **Recovery – Data Rewind** option to **On** (see page 141):

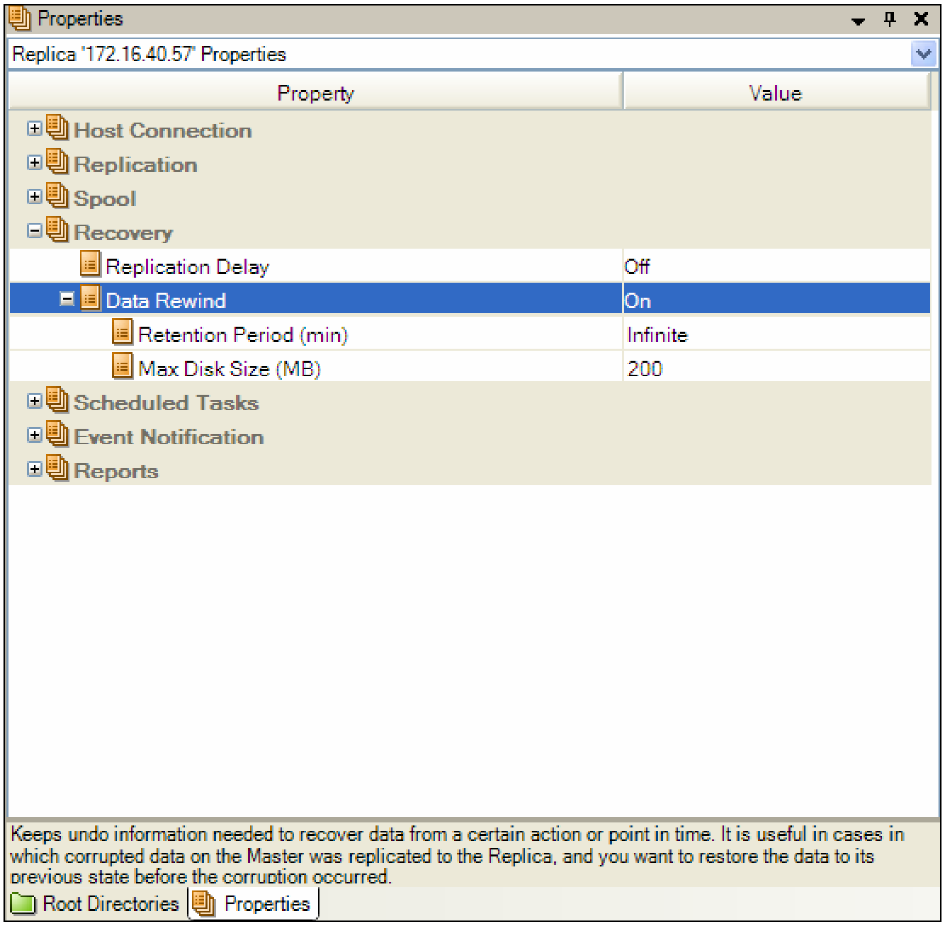

If this option is set to Off, the system will not register data rewind points.

Important! The data rewind process operates in one way only – there is no replay forward. After rewind, all data subsequent to the rewind point will be lost, since data after the rewind point will be overwritten with new data.

Note: The automatic registration of the rewind points starts only after the synchronization process is completed, and the message **All modifications during synchronization period are replicated** appears on the Event pane. Similarly, you can not manually set bookmarks during synchronization.

Recovering lost data using rewind points

1. On the Manager, from the Scenario pane select the desired scenario and stop it.
2. [For database applications only] stop the database services on the Master host.
3. On the Manager, from the scenario folder select the Replica host:

   Note: If multiple Replica servers participate in the required scenario, select the Replica from which you want to recover data.

   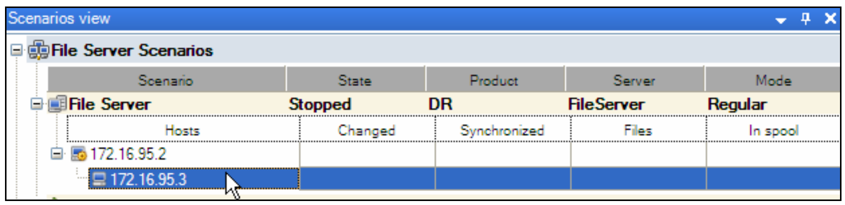

   The **Restore Data** option is enabled.

4. From the **Tools** menu, select **Restore Data**, or click the **Restore Data**  button.

   The **Recovery Method** page of the Restore Data Wizard appears:

   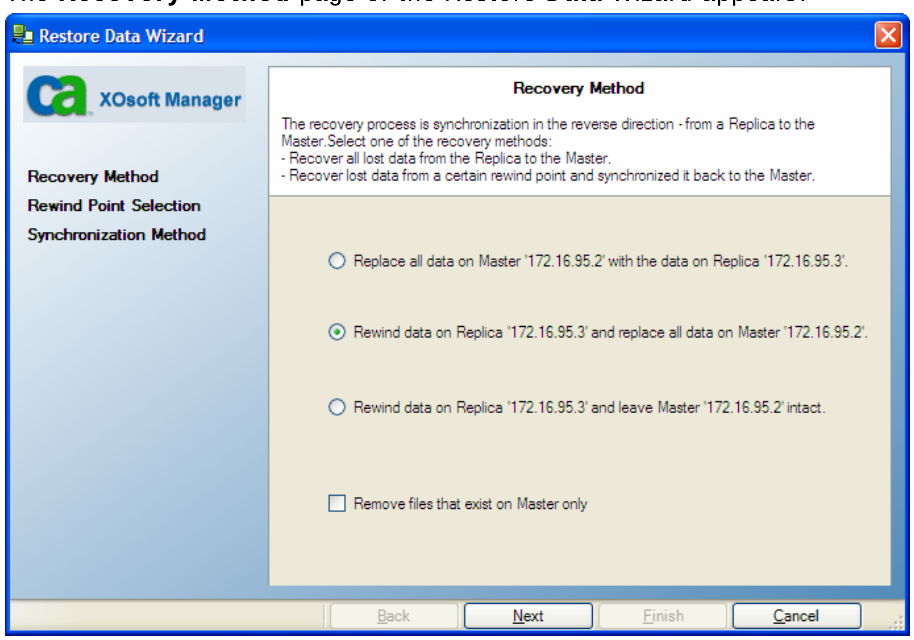

5. Select one of the Rewind data options, depending on whether you want the rewind data synchronized back to the Master (option 2) or left on the Replica only (option 3).

   Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

   After you select a Rewind data option, a Recovery scenario is automatically created. This Recovery scenario will run until the end of the rewind process.

6. Click **Next**. The **Rewind Point Selection** page is displayed:

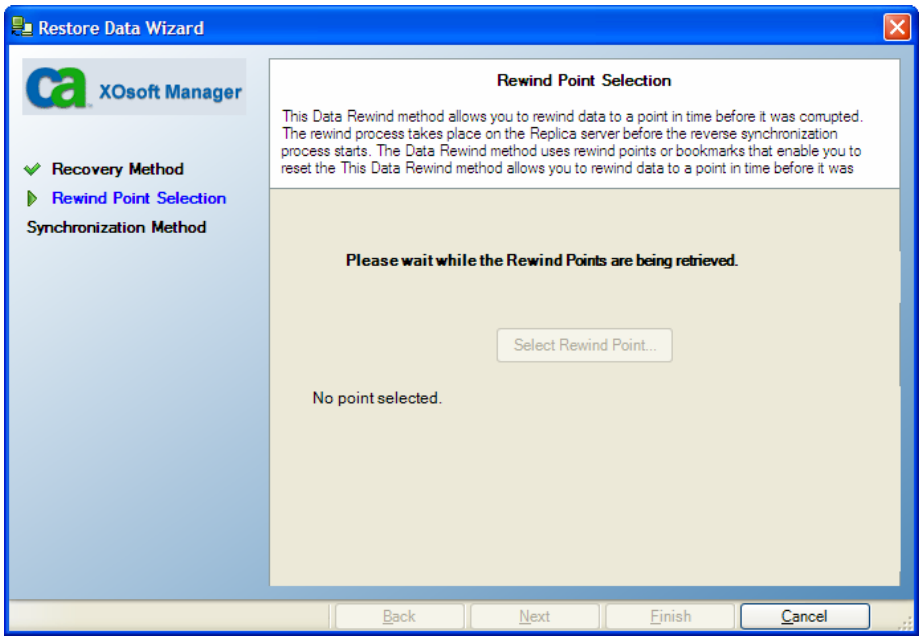

7. Wait until the **Select Rewind Point** button is enabled, and click it to view the existing rewind points.

The **Select Rewind Point** dialog appears:

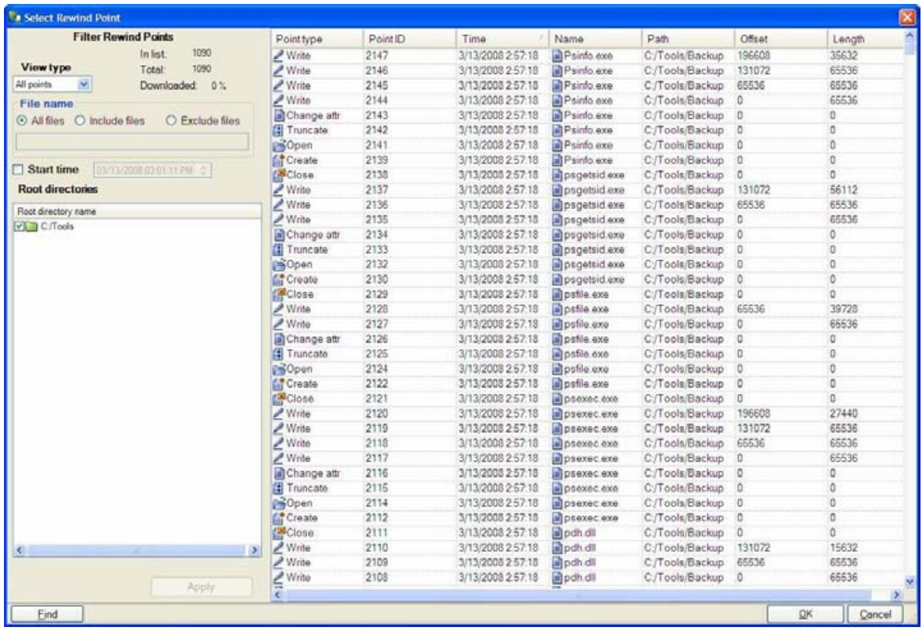

The **Select Rewind Point** dialog displays a list of all rewind points. These include modifications of folders and files that were automatically registered by the system and user-defined bookmarks.

The list can be filtered according to the rewind point type or other criteria, using the **Filter Rewind Points** pane on the left.

Note: If the **Select Rewind Points** dialog is empty, make sure that the **Data Rewind** property is enabled. For more information, see page 141.

8. Select the required rewind point, and click **OK.**

   Note: If you want to use a Bookmark as the rewind point, it is best practice to select the closest rewind point that indicates an actual event.

You return to the **Rewind Point Selection** page, now displaying information about the rewind point you selected:

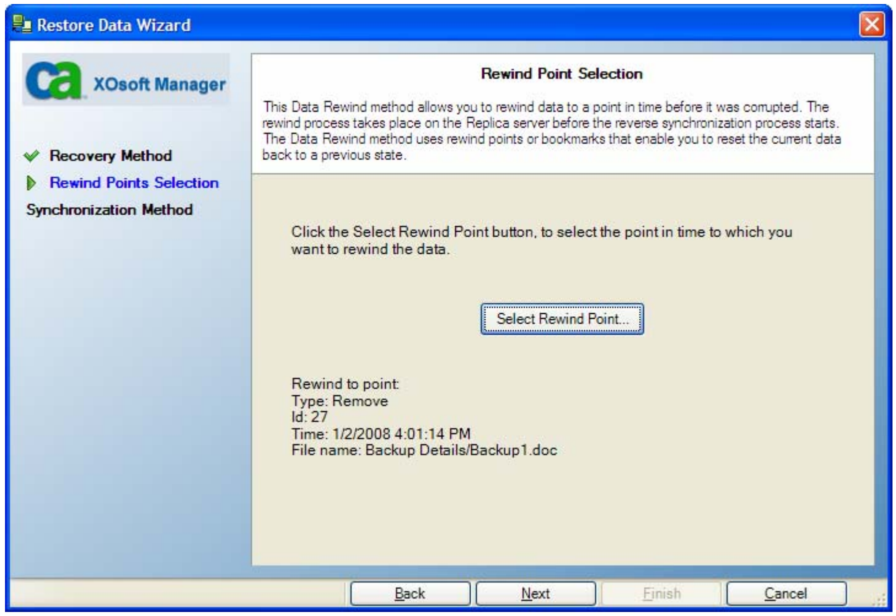

9. Click **Next**. The **Synchronization Method** page is displayed:

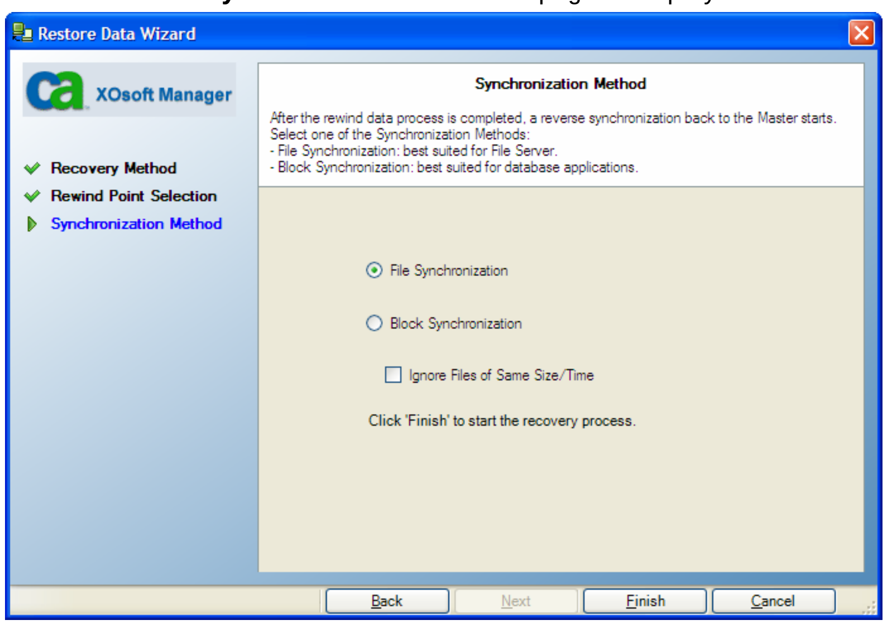

10. Select the **File Synchronization** method and click **Finish**.

Note: If the user credentials you used to log in to the Manager are different than the ones required for working with the Engine on the Replica, a **User credentials** dialog appears, asking you to enter log on account details for the selected Replica.

CA XOsoft rewinds the data to the point you selected. After the rewind process ends, you receive the following message in the Event pane: **Rewind process is completed successfully**.

If you chose to replace the data on the Master with the data on the Replica, CA XOsoft starts a synchronization process from the Replica to the Master. Once the process ends, the temporary Recovery scenario is stopped and then deleted.

11. By default, once a data recovery occurs a Synchronization Report is generated:

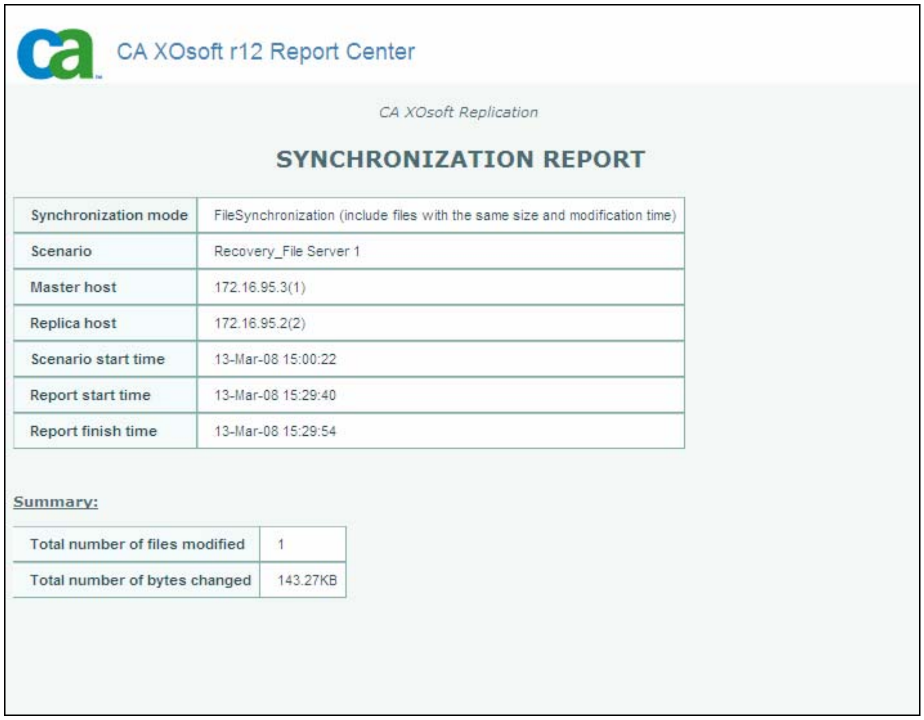

CA XOsoft r12 Report Center

CA XOsoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (include files with the same size and modification time)
Scenario	Recovery_File Server 1
Master host	172.16.95.3(1)
Replica host	172.16.95.2(2)
Scenario start time	13-Mar-08 15:00:22
Report start time	13-Mar-08 15:29:40
Report finish time	13-Mar-08 15:29:54

Summary:

Total number of files modified	1
Total number of bytes changed	143.27KB

Now, the Replication process can restart on the original scenario.

Chapter 10: Switching Over and Switching Back

This chapter explains the High Availability process and the switchover and switchback procedures, and describes the following operations: creating HA scenario, performing switchover and switchback, and recovering active server.

Understanding the High Availability System and the Switchover and Switchback Procedures

A High Availability scenario incorporates all the functionality and workflow of a Disaster Recovery replication scenario, but it adds three important new elements: pre-run verification, monitoring of the Master and the application running on it, and the switchover process itself.

- **Pre-run verification**

  During a switchover, there are many things that can go wrong – there might be problems with permissions, or with the application configuration, or even with the settings within the HA scenario itself. For this reason, when HA scenario is created and initiated, CA XOsoftHA performs an extensive list of checks. These checks are designed to determine, whether any of the common issues that are known to cause problems during switchover can be found. When such issues are found in the pre-run verification, errors and warnings are presented, prompting you to solve these issues before running the HA scenario.

- **Automatic monitoring of the Master and the application running on it**

  As soon as the scenario is initiated, the Replica checks the Master on a regular basis, by default every 30 seconds. There are three types of monitoring checks – a ping request that is sent to the Master in order to verify that the Master is accessible and alive; a database check that verifies that the appropriate services are running and the data is in good state; a user-defined check that can be tailored to monitor specific applications.

  If an error occurs with any part of the set, the entire check is considered to have failed. If all checks fail throughout a configured timeout period (by default 5 minutes), the Master server is considered to be down. Depending on the HA scenario configuration, this will cause CA XOsoftHA to send you an alert or to automatically initiate a switchover.

- **Switchover and switchback workflow**

  In an initial HA scenario, the Master is the active computer, and the Replica is the standby computer. The standby computer is continuously checking the state of the active one, to determine whether it is alive and to decide whether to assume the active role.

  A switchover can be triggered automatically or with the push of a button. The first time a switchover occurs, the Replica that was on standby becomes the active computer, and the Master reverts to a standby mode (assuming it is still operational). When the Master (now the 'standby') is ready, a switchback process can be initiated, either automatically or manually. Following the switchback, the Master again becomes active, and the Replica returns to its previous standby and monitoring role.

Creating High Availability Scenario

The creation of High Availability scenario is similar to the creation of Disaster Recovery scenario. In both you are using the step-by-step Scenario Creation wizard. However, during the creation of HA scenario there are some additional verification checks that CA XOsoftHA performs automatically, and you need to configure several High Availability properties.

Note: This section demonstrates the configuration of a generic Exchange Server High Availability scenario. For more detailed instructions involving scenarios tailored to specific applications such as Oracle or SQL, see the appropriate Operation Guide.

To create High Availability scenario

1. Open the CA XOsoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New**  button on the Standard toolbar.

The **Scenario Creation Wizard** appears:

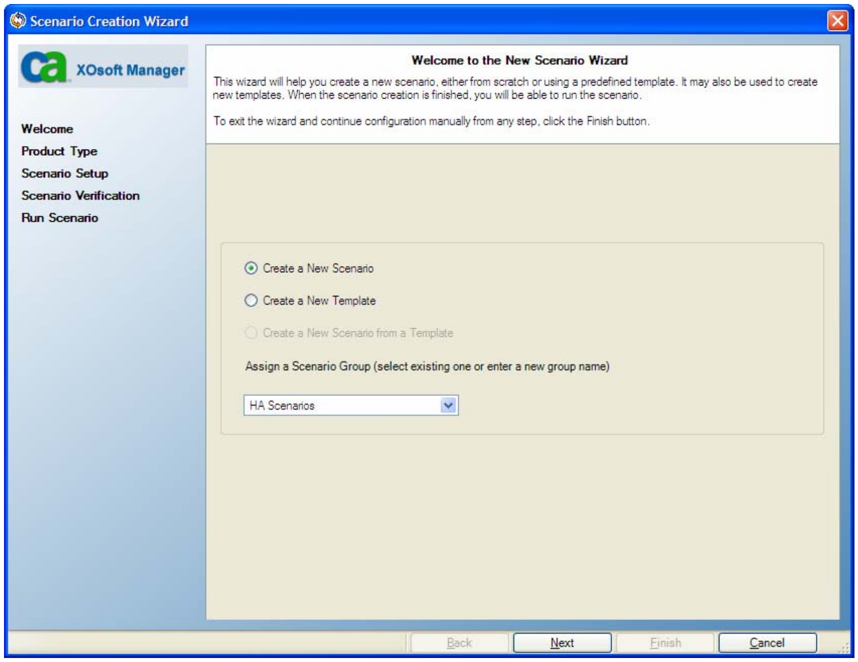

2. Select the required scenario options, as follows:
   - Select the **Create a New Scenario** option button.
   - From the **Group** drop-down list, select the group to which you want to assign the new scenario, or enter a name for a new scenario group.
3. Click **Next**. The **Select Server and Product Type** page appears:

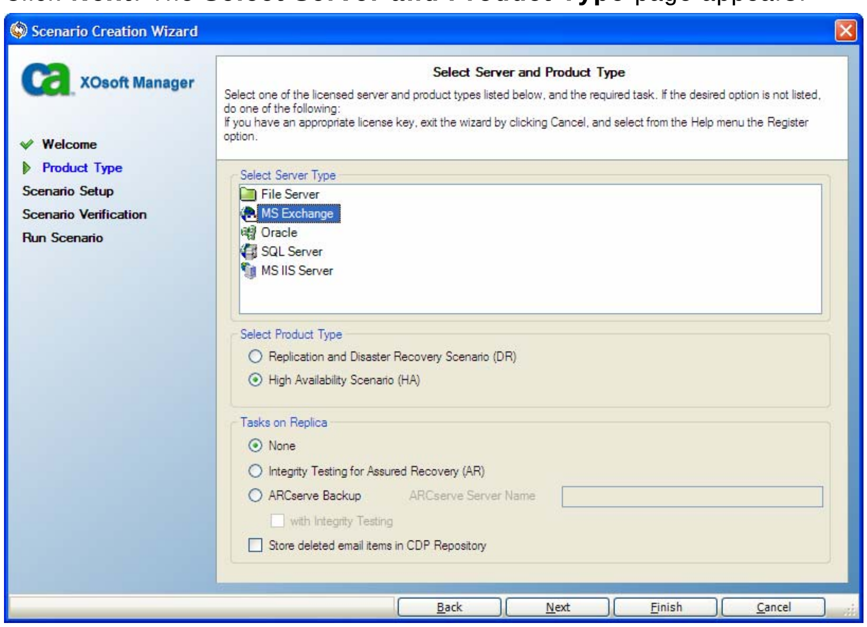

A list of available applications and scenario types is presented.

Note: The list of available applications depends on the licenses applied**.**

Select the required scenario options, as follows:

- From the **Select Server Type** list, select the type of server for which you want to create the scenario. In this example, we select **MS Exchange**.
- From the **Select Product Type** options, select **High Availability Scenario**.
- [Optional - a license is needed] From the **Tasks on Replica** options, select the tasks you want to implement in this scenario.

4. Click **Next**. The **Master and Replica Hosts** page appears:

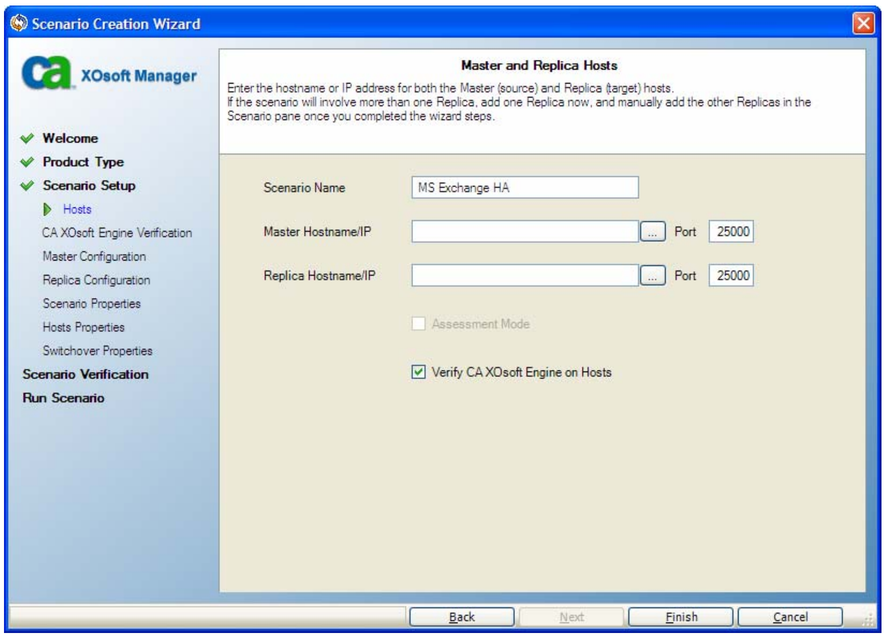

5. Enter the following information:

- In the **Scenario Name** box – accept the default name or enter a new name for the scenario. When entering a name, choose a unique name, since you cannot use the same name for more than one scenario.
- In the **Master** and **Replica Hostname/IP** boxes - enter the hostname or IP address of the Master (active) and Replica (standby) servers, or use the **Browse** buttons to find them.
- In the **Port** boxes: accept the default port no. (25000) or enter a new port numbers for the Master and Replica.

Important! The Replica server you enter here is the only Replica server that can participate in the switchover process. Even if you will later add to the scenario additional Replica servers, they will not be able to become the active server after switchover.

Notes:

– If either server is a MSCS cluster, enter the Exchange Virtual Server Name or IP address as the Master and/or Replica name (instead of the physical node's name/IP).

– If you want to include more than one Replica in the scenario, enter here the details of the first or most upstream Replica. After you finished the scenario creation, manually enter the other Replicas, as described on *Add Additional Replica Servers*, page 97.

- The **Assessment Mode** check box – verify that it is NOT selected.
- The **Verify CA XOsoft Engine on Hosts** check box - select this check box if you want the system to verify whether Engines are installed and running on the Master and Replica hosts you specified in this page. If Engines are not installed on the selected hosts, you can use this option to remotely install the Engines on one or both hosts. For more information about the **Host Verification** page, see page 46.

6. After you selected the desired options, click **Next**. The **Master Configuration** page appears:

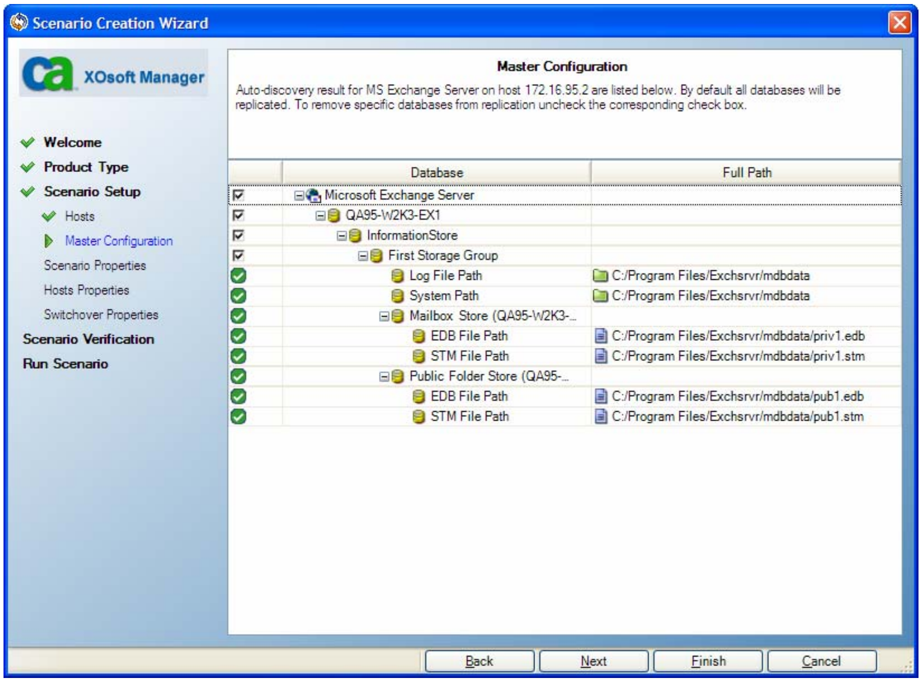

The CA XOsoft auto-discovery component automatically displays the Exchange databases that are on the Master server. These are the databases that can be replicated and protected.

7. By default, all the discovered databases are selected and all will be replicated. You can exclude any of these storage groups from replication by clearing their check boxes.

8. Click **Next**. The **Replica Configuration** page appears:

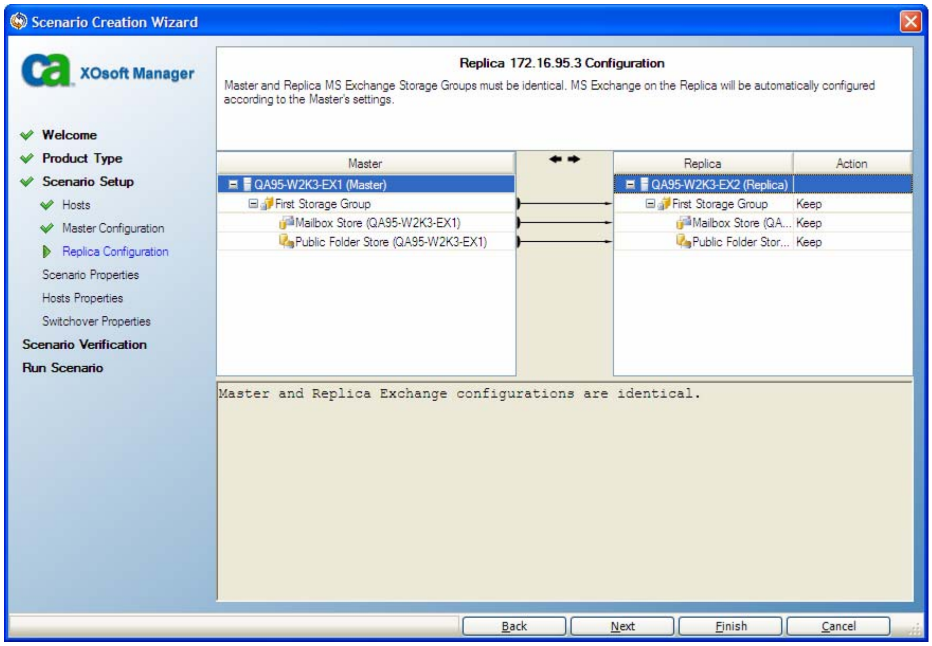

CA XOsoft auto-configuration component verifies that the Exchange Server configuration on the Master and Replica servers will be identical during the replication procedure. This means that if there are discrepancies, CA XOsoft will perform the required actions, including: deleting storage groups, public folders or mailbox stores from the Replica, create new ones and make modifications to existing ones. The actions that will be performed during the configuration process are indicated in the **Action** column on the right.

Note: The actions that can be performed in the auto-configuration process are as follows:

- **Create** - a new, storage group, public folder or mailbox store will be created.
- **Keep** - the existing storage items will remain the same.
- **Remove** - the existing storage items store will be deleted.
- **Update** - the existing storage items will remain, but its location will be changed

9. Review the changes that will occur during the automatic configuration on the Replica Exchange Server, and make sure you want them to be performed.

   Note: If a **Remove** action is indicated, make sure that you are ready to delete the specified storage item since it does not have an automatic backup. If you want to save it in a different location before deletion, click the **Finish** button to exit the wizard.

10. Click **Next** to start the Replica configuration process. The **Scenario Properties** page appears:

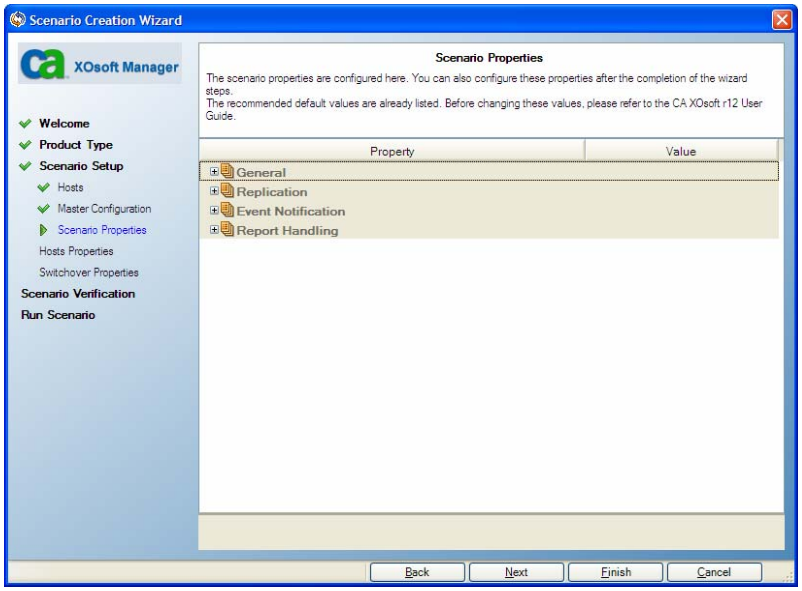

The **Scenario Properties** page enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

If you want to configure the scenario properties at this stage, refer to *Understanding Scenario* Properties, page 123. To configure the scenario properties at a later stage, refer to *Configuring Scenario Properties*, page 121.

11. Click **Next**. The **Master and Replica Properties** page appears:

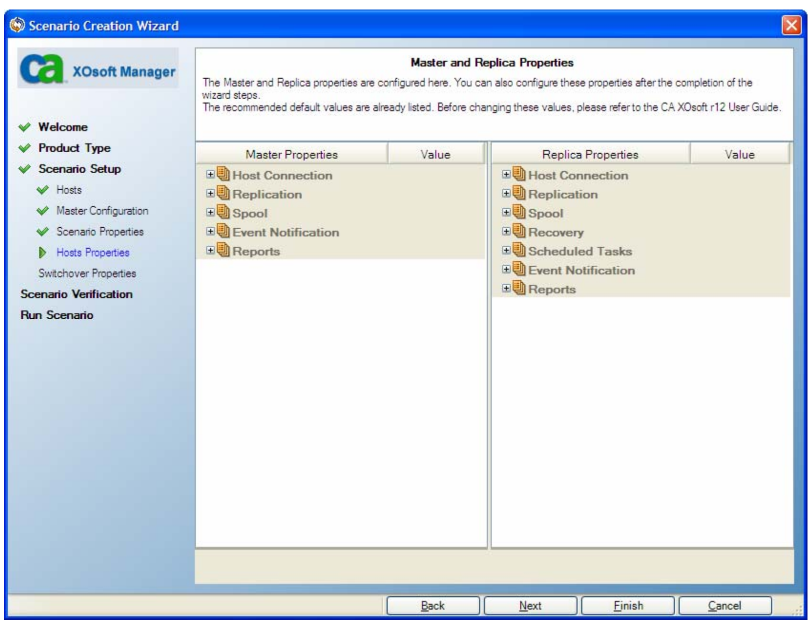

The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

If you want to configure the Master and Replica properties at this stage, refer to *Setting Master and Replica Properties*, page 133. To configure the Master and Replica properties at a later stage, refer to *Configuring Master or Replica Server Properties*, page 133.

Note: You can modify all the settings in this pane after the scenario is created. However, before changing any Spool properties (which can be configured here), review the *Spool* info on page 136 for configuration details.

12. Once you are satisfied with the Master and Replica properties, click **Next**.

    The **Switchover Properties** page appears:

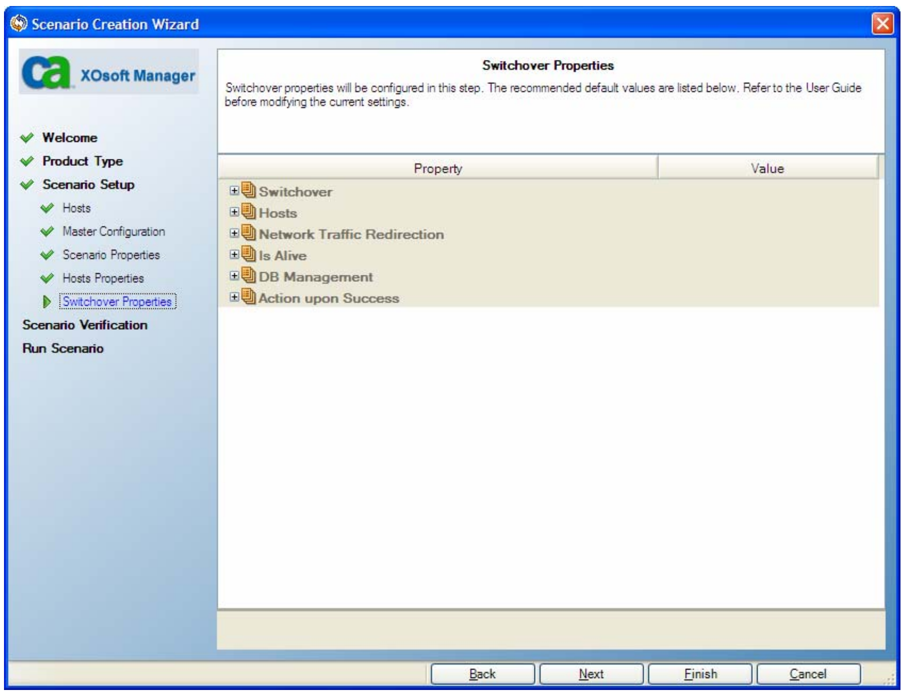

The **Switchover Properties** page allows you to modify switchover parameters. As with the prior steps, no changes are required.

If you want to configure the switchover properties at this stage, refer to *Understanding* High Availability Properties, page 179. To configure the switchover properties at a later stage, refer to *Configuring High Availability Properties,* page 177.

Note: If you plan to use automatic switchover, note that the **Is alive timeout** (sec) setting controls how long to wait after a failure is detected before triggering a switchover. For more information see *Is Alive*, page 184.

13. Click **Next**. The **Switchover and Reverse Replication Initiation** page appears:

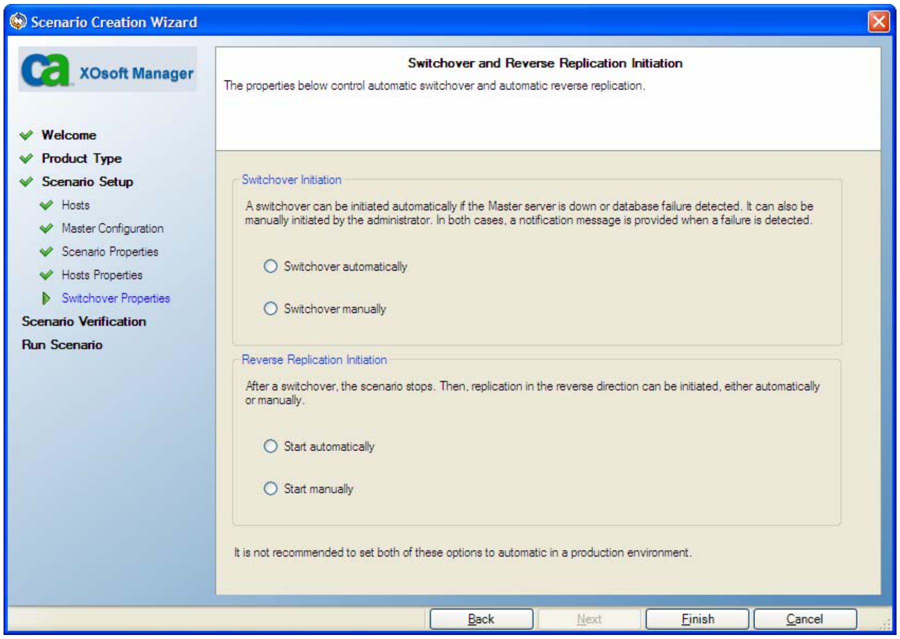

The **Switchover and Reverse Replication Initiation** page enables you to define how the switchover and the reverse (backward) replication scenario will be initiated: either automatically by CA XOsoftHA, or manually by you once a failure is detected.

Note: Automatic Switchover and Run Reverse Replication Scenario

Automatic switchover is in all ways identical to manual switchover performed by the administrator. Automatic switchover is triggered by a resource failure on the Master server rather than by an administrator who manually initiating the switchover. Server ping response, application service status and database connectivity are routinely tested. The timeout parameters are configurable, as described on the *Is Alive* section, page 184.

When **Reverse Replication Initiation** is set to **Start automatically**, CA XOsoftHA will automatically run a reverse scenario during switchover – thus saving the need for a resynchronization when running the "backward" scenario. Note that the reverse scenario can only run automatically if the Master server was online during switchover.

14. For each switchover property, select one of the available options:

- **Switchover**: either automatically or manually.
- **Reverse Replication**: either automatically or manually.

Important! We advise you against setting both of these options to automatic in a production environment. While these options are individually beneficial, it is best practice to set only one of these options to automatic.

15. Click **Next.** A notification message appears informing you that CA XOsoftHA verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful switchover.

    Note: Numerous checks are performed at this stage. Broadly, they fall into three categories:

    - **Consistency checks** – ensuring that the situation is as expected, that network resources point to the current Master (active) server, that the application is not active on the standby server, that there are not any conflicting configurations, etc. This is the largest category.
    - **Permission checks** – verifying that the Engine log on account has the authority to make changes, especially if DNS redirection is used.
    - **Application-specific tests** – for Exchange, SQL or Oracle, these tests include verifying that configurations of the active and standby server installations are compatible, ensuring that services that must run are running, checking that Active Directory configuration is correct, etc.

16. once the verification is completed the **Scenario Verification** page appears:

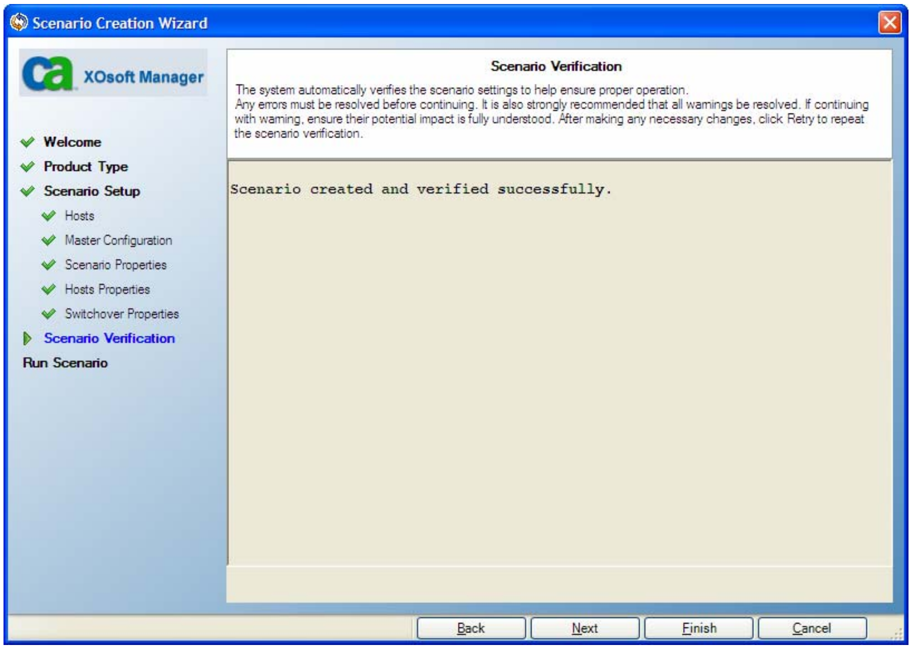

17. If the scenario was not set up correctly, or problems occurred in the participating hosts or the connection between the CA XOsoft components, the **Verification Results** dialog opens, listing all the errors and warnings detected:

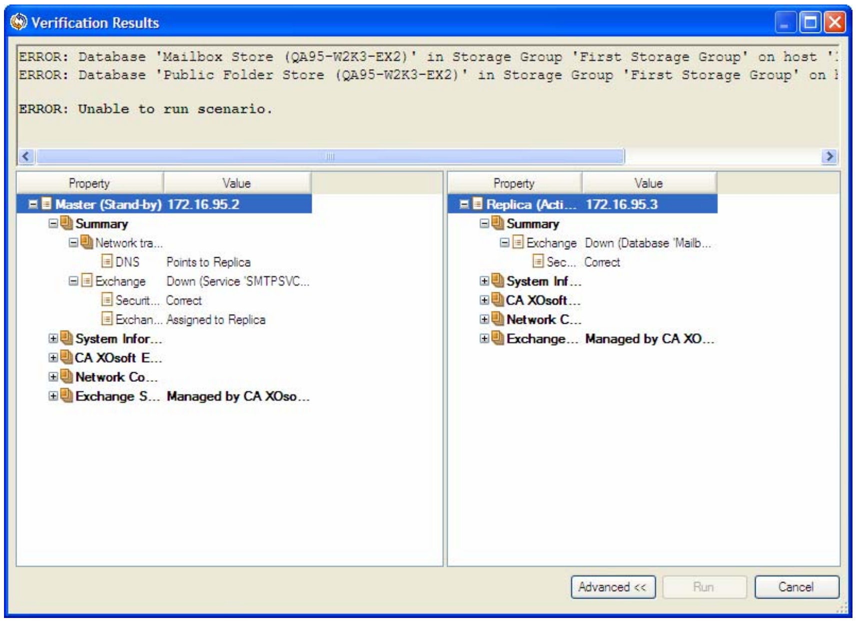

- If any errors are displayed, you cannot run the scenario. These errors must be corrected before you can start the synchronization, replication and HA processes.
- If only warnings are displayed, you can run the scenario by clicking the **Run** button. However, it is important that you consider the warning carefully since they indicate conditions that are known to potentially cause problems with replication or switchover.

Note: The **Advanced** button opens an additional pane with detailed information about the checks performed to help diagnose problems. It is intended to help you resolve any issues encountered in running the software. You can also contact Technical Support for further assistance.

18. When no errors or warnings are reported, once the verification is completed the **Scenario Verification** page appears:

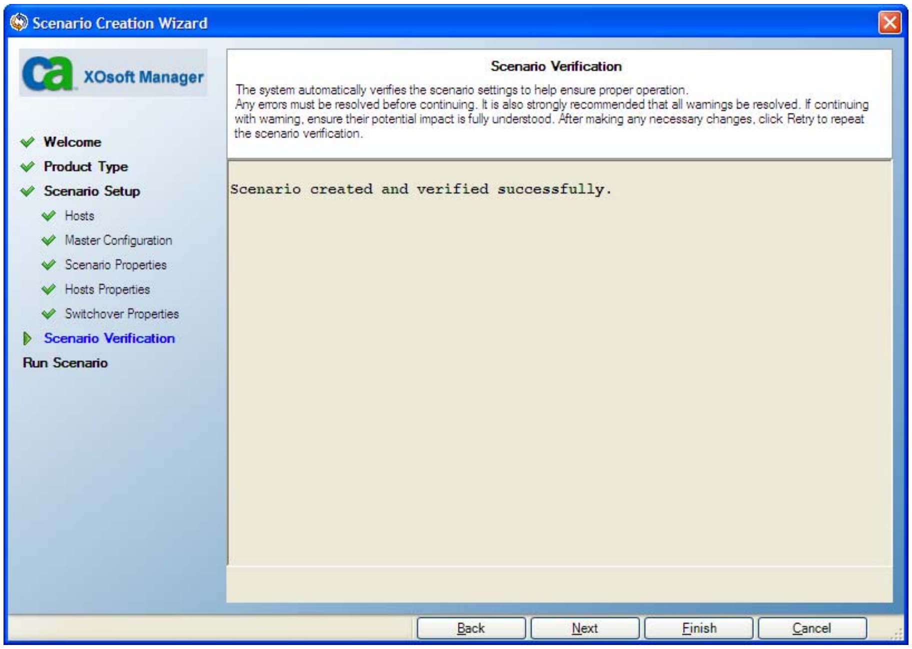

19. When the scenario is verified successfully, click **Next** to continue.

The **Scenario Run** page appears:

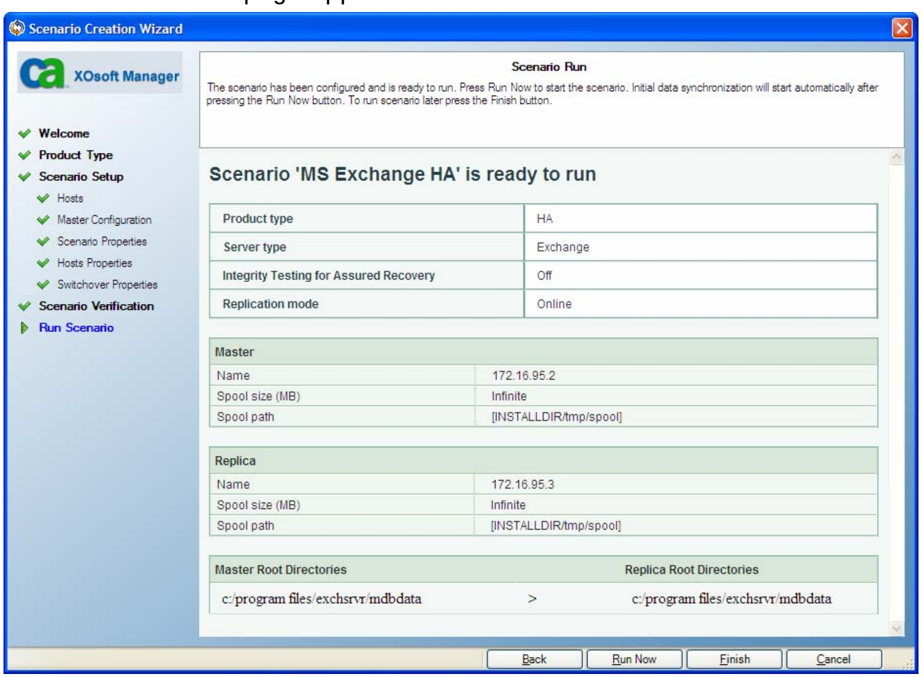

20. The scenario configuration is now completed and you are prompted to run it. Running the scenario starts the data synchronization process, following by replication and is alive checks. .
    - To finish the scenario creation and run it later, select **Finish**.
    - To run the scenario, click **Run Now**.
21. The synchronization process starts. Synchronization may take a while depending on database size and network bandwidth between the Master and Replica. You will receive the following message in the Event pane when synchronization is complete: **All modifications during synchronization are replicated**.

    From this point, real-time replication occurs and the High Availability solution is active:

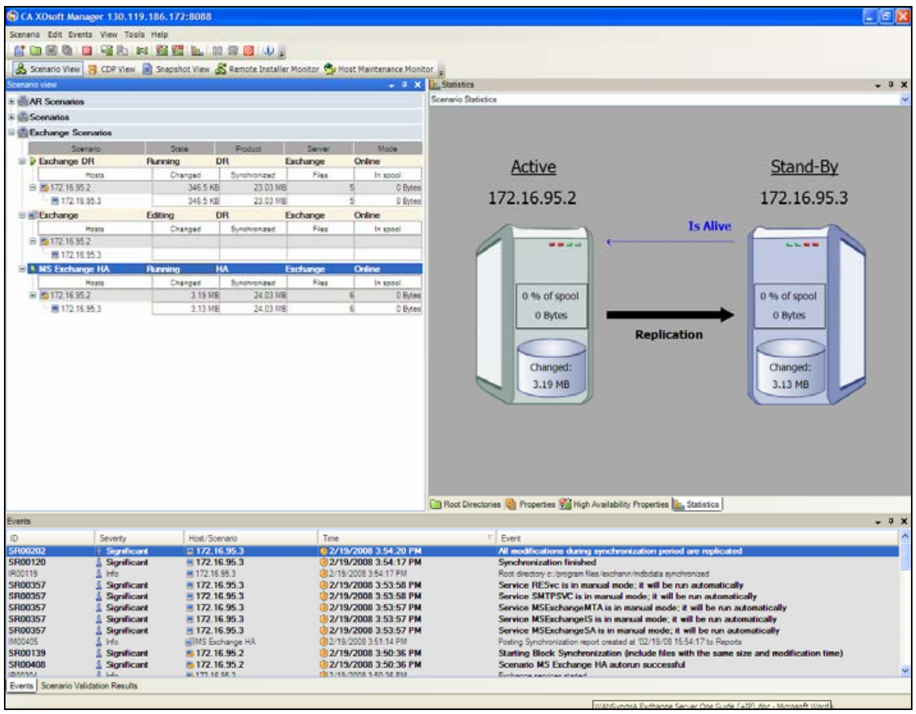

Switchover

Switchover is the process of changing roles between the Master and Replica, meaning making the Master server the standby server, and the Replica server the active server.

Switchover can be triggered automatically by CA XOsoftHA when it detects that the Master is unavailable. Alternatively, CA XOsoftHA can simply alert you to the problem, and then you can manually initiate switchover from the CA XOsoft Manager.

During the creation of HA scenario, you define how you want the switchover to be initiated. If you selected in the **Switchover and Reverse Replication Initiation** page the **Initiate Switchover manually** option (see page 163), you need to perform a manual switchover. However, if you selected the **Initiate Switchover automatically** option, you can still perform a manual switchover, even if the Master is alive. You can initiate switchover when, for example, you want to test your system, or you want to use the Replica server to continue the application service while some form of maintenance is performed on the Master server.

Once triggered, whether manually or automatically, the switchover process itself is fully automated.

To initiate switchover

1. Open the Manager and select the desired scenario from the Scenario pane. Verify it is running.
2. Click on the **Perform Switchover** button, or select from the **Tools** menu the **Perform Switchover** option:

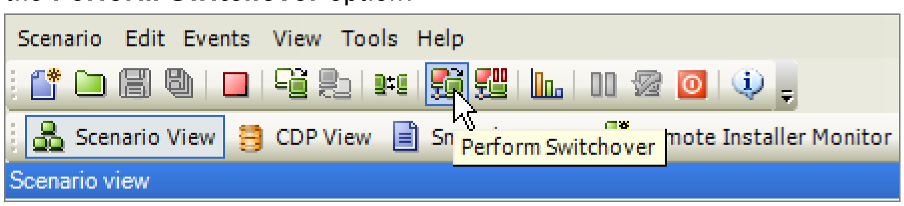

A confirmation message appears:

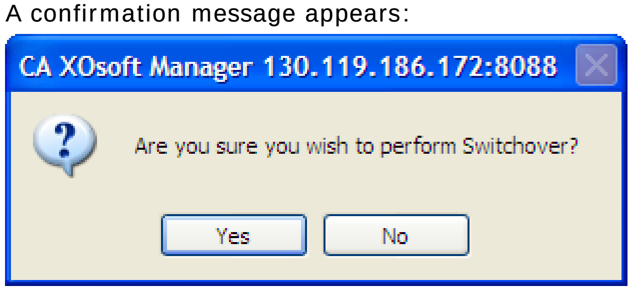

3. Click **OK** on the **Perform Switchover** confirmation message. This procedure initiates a switchover from the Master server to the Replica server:

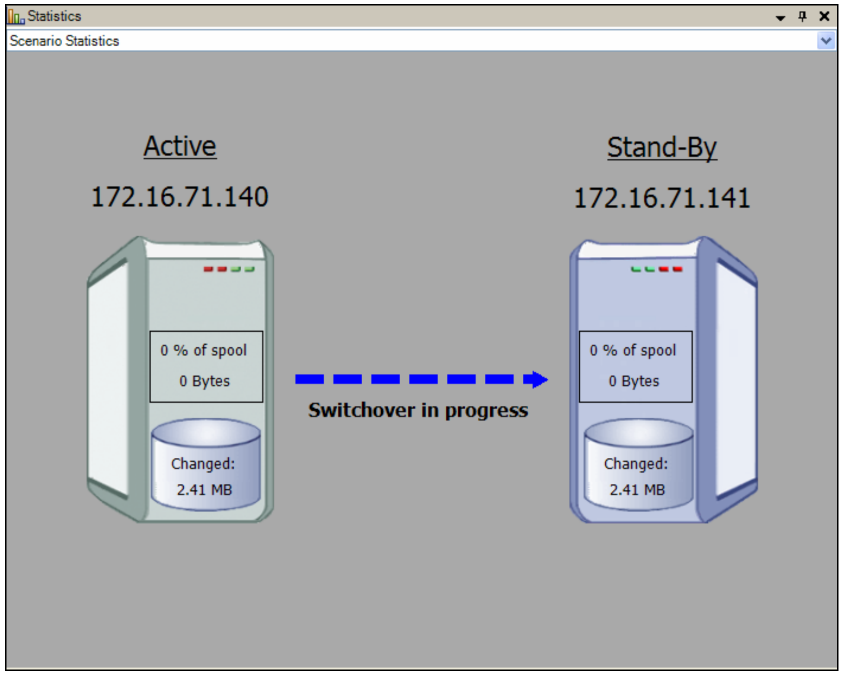

Detailed information about the switchover processes is located in the Events pane during switchover.

4. After the switchover is completed the scenario stops:

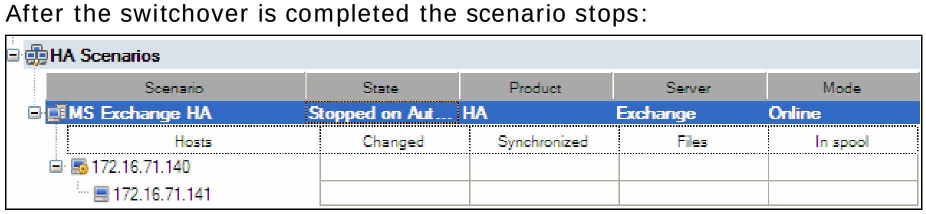

Note: The only case in which the scenario may continue to run after switchover is when **automatic reverse replication** is defined as **Start automatically**. For more information, see page *163.*

In the Event pane a message appears, informing you that **Switchover completed**, and then that the **Scenario has stopped**.

Now, the original Master becomes the Replica and the original Replica becomes the Master.

Switchback

After a switchover was initiated, whether manually or automatically, at some point you will want to reverse back the server roles, and make the original Master the active server again and the Replica the standby server. Before you switch back the roles between them, if you want the data on the active server, meaning the original Replica, to overwrite the data on the standby server, you need to run a reverse scenario (also called "backward scenario").

During the creation of the HA scenario, you defined how you want the reverse scenario to be initiated. If you selected in the **Switchover and Reverse Replication Initiation** page the **Initiate Reverse Replication automatically** option (see page 163), replication in the reverse direction (from Replica to Master) will automatically begin after a switchover, once the original Master will become available. However, if you selected the **Initiate Reverse Replication manually** option, you need to perform a manual switchback. If the manual option is selected and you will not initiate a manual switchback, a resynchronization of the data from Replica to Master will have to be performed, even after a test of clean switchover without an actual Master failure.

Note: After a switchover, in certain circumstances you may want to switch the Master and Replica roles without overwriting the original Master data with the Replica data. To perform this, use the **Recover Active Server** option, as described on page 174.

To initiate switchback

1. Ensure that both Master and Replica servers are available on the network and that the CA XOsoft Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. [Skip this step if the backward scenario is already running, and move to step 7, page 173.]

If the backward scenario is not running, select the **Run** button to start the scenario. CA XOsoftHA detects that a switchover has occurred, verify its state and configuration, and prompts you to approve the running of the backward scenario:

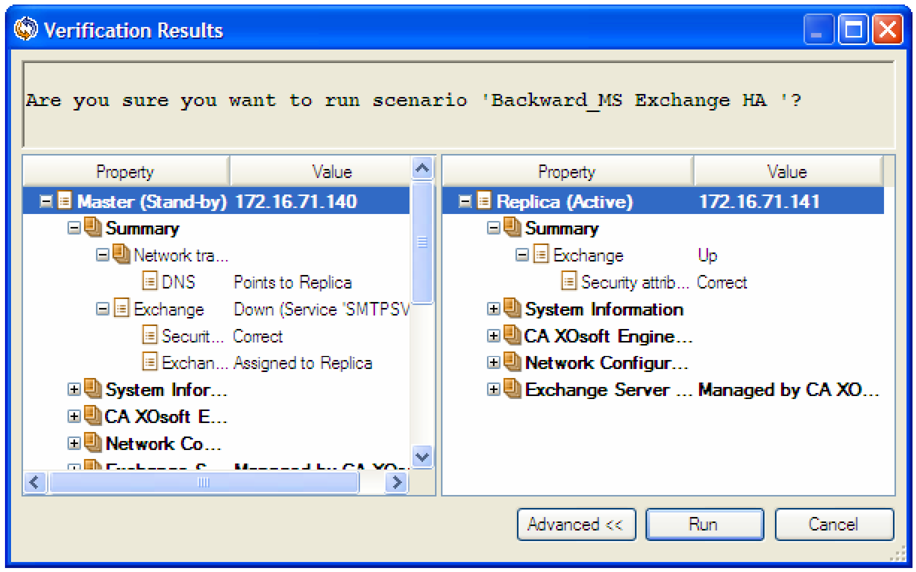

Note: The **Advanced** button opens an additional pane with detailed information about the hosts that participate in the scenario.

4. Click the **Run** button to start the backward scenario.

   The **Run** dialog appears:

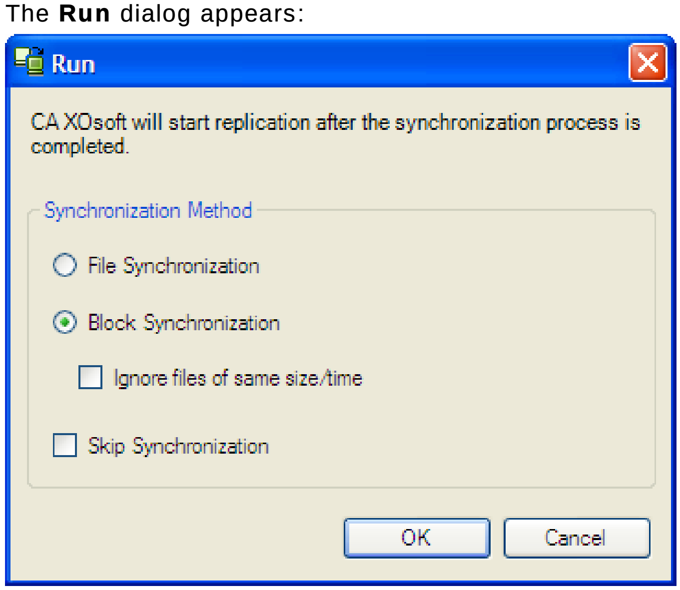

5. For MS Exchange, select **Block Synchronization**, and click **OK**. The resynchronization starts:

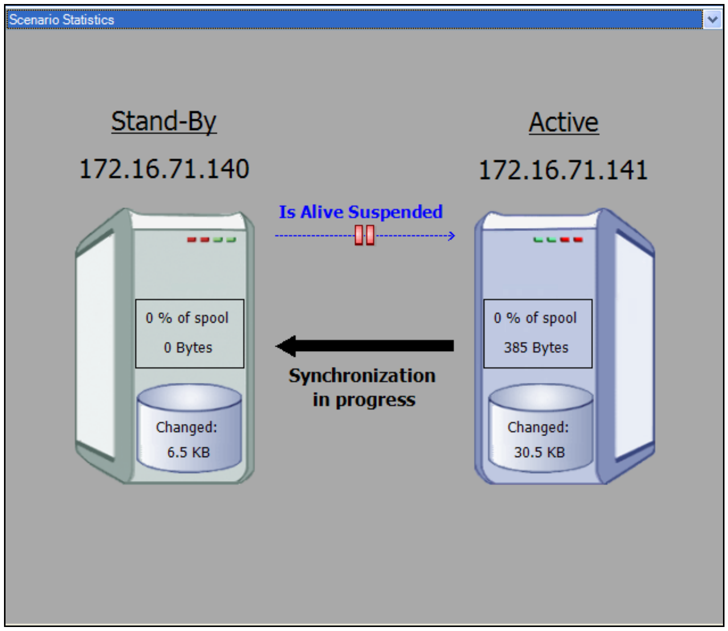

Wait until the resynchronization is completed.

6. Once the resynchronization is completed, you receive the following message in the Event pane: **All modifications during synchronization are replicated**. Then, replication from the active server to the standby server starts:

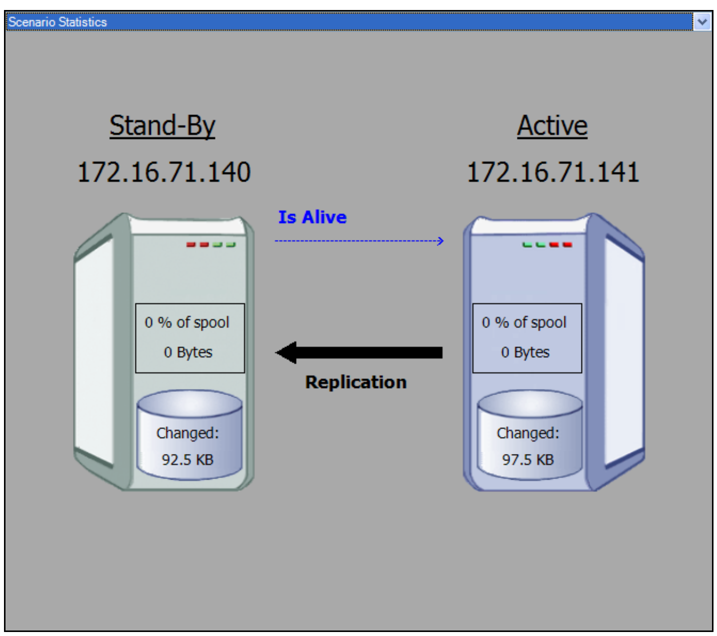

7. Now, you can reverse back the roles between the Master and Replica servers. To reverse the roles, while the backward scenario is running, click the **Perform Switchover**  button or select the **Perform Switchover** option from the **Tools** menu.

   A confirmation message appears:

   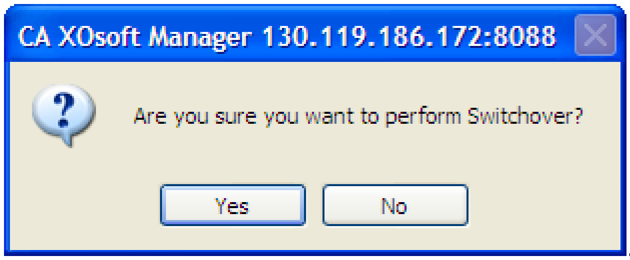

8. Click **Yes** on the **Perform Switchover** confirmation dialog. This procedure initiates a switchback from the original Replica server to the Master server:

   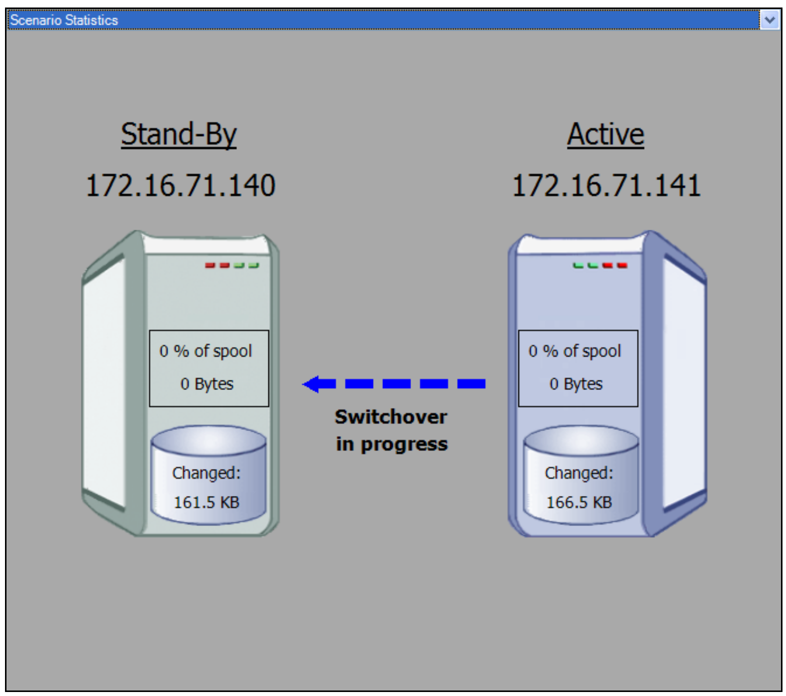

9. After the switchback is completed, and the server roles are reversed back, the scenario automatically stops:

   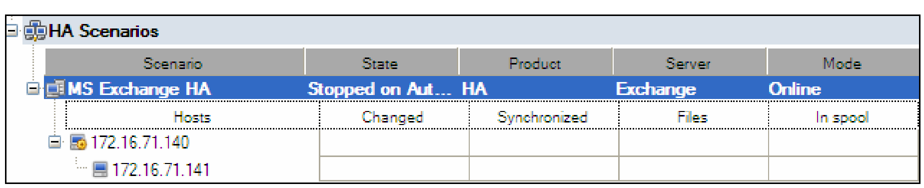

   Now, you can run again the scenario in its original (forward) state.

Recover Active Server

In certain circumstances, it may be necessary to intentionally make the Master or Replica server the active server without completing the synchronization process. This may happen when, for example, a switchover occurred but no data was changed on the Replica server, and you might even have newer data on the Master server. In this case, it is undesirable to synchronize data from the Replica to the Master server. CA XOsoft allows you to manually select the active server option through a process called Recover Active Server. To use this option, ensure that the scenario is stopped, and select *Recover Active Server* from the Tools menu.

Important! While this option is the right choice in certain situations, use it with caution. If used improperly, data loss can occur. Normally, CA XOsoft will not allow switchover from one host to another until all data is synchronized. It is designed this way so users are not redirected to an out-of-date data set, which then overwrites what may be a more current data set. When using **Recover Active Server**, CA XOsoft is redirecting users to one server or the other with no regard as to which server has the correct data set. Thus, as an administrator, you must manually ensure that the server you are making active has the most up-to-date data set.

To recover active server:

1. On the Scenario pane, select the scenario whose active server you want to recover and stop it.
2. From the **Tools** menu, select the **Recover Active Server** option.

   CA XOsoft verifies which host is currently the active one, and presents the results in the **Recover Active Server** dialog:

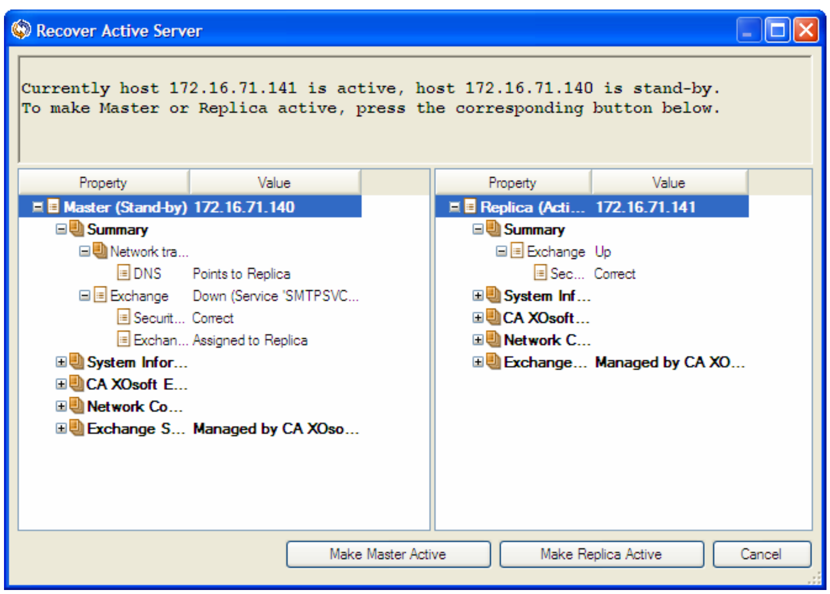

3. Click either the **Make Master Active** or **Make Replica Active** button, depending on which server you want to have the active role.

   Messages appear in the Event pane, informing you that one of the server becomes active while the other becomes inactive. Now, the host you selected becomes the active server, and users are directed to it.

Important! In a disaster situation, if a standard switchover occurs and users are redirected to the Replica server for any period of time, it is important to replicate all changes that occurred on the Replica back to the Master, before making the Master server active again. Using the **Recover Active Server** option in such a situation, may result in loss of data.

Chapter 11: Setting High Availability Properties

This chapter describes how to configure High Availability properties, and provides a list of the HA properties, their corresponding values, and an explanation of each property.

Note: These options are available only to those who are licensed for High Availability.

Some properties apply only to certain server types (e.g., Exchange, SQL, etc.). For more information, see the specific Operations Guide.

Configuring High Availability Properties

The HA property values determine the entire scenario's default behavior concerning network traffic redirection method, database management, and more.

Notes:

- The Properties pane and its tabs (Root Directories, Properties, Statistics) are context sensitive, and change whenever you select a different node from a scenario folder.
- To configure scenario properties, the scenario must be stopped.

To set scenario properties

1. On the Scenario pane, select the HA scenario whose properties you want to configure.
2. On the Framework pane on the left, select the High Availability Properties tab.

The High Availability Properties list appears:

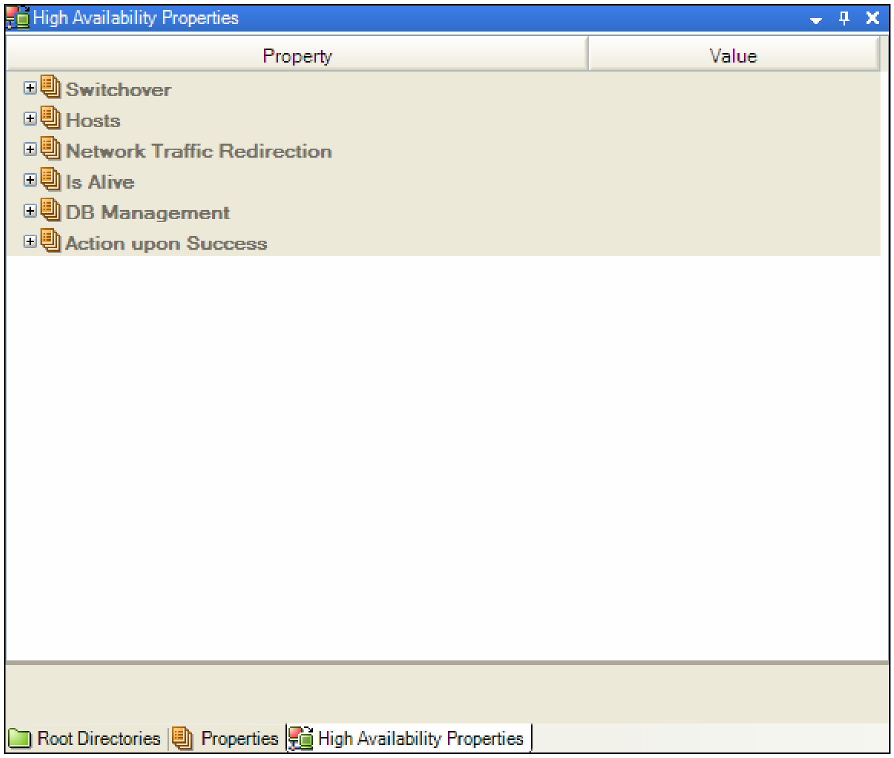

Note: A running scenario has a gray background, and scenarios that are not running have a white background.

3. If the scenario is running, click the **Stop**  button on the toolbar. The scenario is stopped.
4. On the Scenario Properties list, open the desired group, select the required property, and select or enter the appropriate values. Some values can be manually entered in an edit box field, while others can be selected from a combo box by clicking the default value.
5. After you set the required properties, click the **Save**  button on the Standard toolbar to save and apply your changes.

Understanding High Availability Properties

This section lists the High Availability properties, corresponding values, and provides an explanation of each property. It also explains the Active and Standby host concepts.

Active and Standby Hosts

In an initial scenario, the Master is the active computer, and the Replica is the standby computer. The standby computer is continuously checking the state of the active one, to decide whether to become active computer.

The first time a switchover occurs, the Replica that was on standby becomes the active computer, and the Master reverts to a standby mode (assuming it is still operational). When the Master (now the 'standby') is ready, a switchback process can be initiated where the Master again becomes active, and the Replica returns to its previous standby and monitoring role.

Switchover Properties

The system continuously checks for a switchover situation (see *Is Alive Group*) and informs the user according to the defined notification settings. When the system detects a switchover situation, the **Perform Switchover Automatically** option is checked to determine if there should be an automatic switchover, or only a notification of the situation. In the latter case, switchover may be triggered with the **Perform Switchover** menu command or toolbar button.

Switchover Hostname

Enter the hostname/IP address of the Replica host to which the Master switches. You can choose only one Replica.

Any time this Name/IP, or Master's Name/IP, is changed, all the switchover properties are reset to their default values.

Perform Switchover Automatically

When this option is On, switchover is initiated automatically if the Master server is down or database failure detected.

Run Reverse Replication Scenario after Switchover

After a switchover, this option determines whether replication in the reverse direction would begin automatically.

When set to On, CA XOsoft verifies that the data integrity is kept. If the data is found to be consistent, re-synchronization is avoided and the backward scenario is initiated. However, in DB scenarios when this option is set to On, the reverse replication starts in a suspended state. It is unsuspended only after the database on the Active server has passed all tests typically performed in the **Check DB** properties in the **Is Alive** properties.

Hosts

Master Fully Qualified Name

Enter the fully qualified name of the Master host.

Replica Fully Qualified Name

Enter the fully qualified name of the Replica host.

Redirection

There are four redirection methods:

- Move IP
- Redirect DNS
- Switch Computer Name
- User-defined scripts

Move IP:

During switchover, the switchover IP(s) are released on the active host and added to the standby host. This redirection method is applicable only when both the Master server and the Replica server are on the same IP subnet.

Choosing Off or On affects the available options in the **Check With** for a **Send ping request**. See *Is Alive group* below.

Add IP/Mask

Enter IPs for the active computer that will be moved to the standby computer during switchover. The Master IP address defined in the Master Properties (see page 134) must be different than the IPs entered here.

To add IP/Mask

1. Click the tree entry: **Click here to add new IP/Mask**.

   The **IP Address** dialog appears.

2. Enter the relevant IP/Mask data in the dialog and click **OK**.

   A new entry is added to the list, and a new row opens for another IP/Mask. Enter as many entries as you require.

Notes:

- The Master IP address on the Properties tab under **Host Connection** (the Master IP that was entered in the Scenario Creation Wizard), must NOT be one of the IPs included in this list.
- If the **Move IP** or the **Redirect DNS** property is set to On, CA XOsoft turns off the dynamic DNS registration for the Master. (The checkbox "Register this connection's addresses in DNS" in the Advanced TCP/IP Settings dialog is cleared).

Redirect DNS:

During the switchover, the A-record of the Master Server will be updated. This redirection option is applicable when the Master and the Replica are located on a different IP subnet or on the same subnet.

If the **Move IP** or the **Redirect DNS** property is set to On, CA XOsoft turns off the dynamic DNS registration for the Master. (The checkbox **Register this connection's addresses in DNS** in the **Advanced TCP/IP Settings** dialog is cleared).

DNS Server IPs

Enter IPs of DNS servers to update. CA XOsoft HA tries to update all servers listed. However, switchover is considered successful even if only one update is successful.

To enter the value, click the tree entry: **Click here to add new IP**.

DNS TTL

Enter number of seconds for DNS Time-To-Live. This value is changed in the DNS Server for the A-record that is updated.

Active Directory Integrated

Specify if DNS is an Active Directory Integrated. If the Master DNS is on a Windows platform and integrated with Active Directory, set this option to On.

DNS Key Filename (full path)

Enter the full path of the file containing the DNS secure key. This field appears only when **AD Integrated** is Off.

Master/Replica IPs in DNS

Enter the Master/Replica IPs in its DNS server.

The Master's DNS record is updated during every switchover: in the switchover from Master to Replica, the addresses in the Master's DNS record are replaced by the addresses of the Replica. In the switch back from Replica to Master, the addresses are restored to the original values.

To enter the value, click the tree entry: **Click here to add new IP**.

Switch Computer Name

[Not for Exchange] This redirection option is applicable when clients use NetBIOS name resolution for their connections with the Master. If the hostname and the NetBIOS name are not the same, this option cannot be used.

During the switchover, the Replica computer is renamed to the Master computer name and the Master computer is renamed to a temporary name (if the Master server is alive). During the switchback, the names are restored. Both the hostname and the NetBIOS name are changed.

Important! If you will be redirecting File Shares, in which clients connect to via the Master server name, Switch Computer Name must be enabled. For example, if the Master server's name is fs01 and clients connect to \\fs01\sharename or \\fs01.domain.com\sharename, then you must use the Switch Computer Name method. We also recommend enabling one other method as well. The most common method is to use both DNS Redirection and Switch Computer Name.

Master Computer Name

NetBIOS name of the Master computer. This name cannot be modified here.

Replica Computer Name

NetBIOS name of the Replica computer. This name cannot be modified here.

Reboot after Switchover and Switchback

After a switchover and a switchback, if this option is set to On, both Master and Replica computers are rebooted.

User Defined Scripts

This option allows the standard redirection methods to be enhanced or replaced by actions that are invoked by user-defined scripts.

Important! When using scripts, each script must reside in the same path and with the same name on both the Master and the Replica.

Active to Standby Redirection Script

- **Script Name**

  Enter the name and full path of script to be run on the active computer, if it is alive, in order to redirect clients to the standby computer or release network resources on active host.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values must be static.

Standby to Active Redirection Script

- **Script Name**

  Enter the name and full path of script to be run on the standby host, in order to redirect clients to it or add network resource.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values must be static.

Identify Network Traffic Direction Script

Required to fully support custom redirection methods. The custom scripts entered here are used to identify the active server. The Engine assumes that:

- If the script was executed on the host is returning 0, then the host is active (has all network resources on it or users directed to this host)
- If the script is returning a non-zero value, then the host is inactive (all or some of the network resources are absent, or users are not directed to this host).

- **Script Name**

  Enter the name and full path of script to be run. This script determines if the Forward or Backward scenario will run when the scenario is started. The script runs on both Master and Replica: the one that returns zero is active. If both return the same value, a conflict is reported.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values must be static.

Is Alive

CA XOsoft continuously checks to see if the active host is up (according to *Send ping request*, *Connect to DB*, *or User-Defined Script* method, see below). These checks are made in scheduled intervals according to the Heartbeat Frequency.

Checking the status is interpreted as follows:

- If there is an indication that the active computer is OK (alive), no new action is taken, and the system continues to check according to the scheduled intervals of the Heartbeat Frequency.
- If there is an indication that the active computer is not OK (is not alive), the active computer is checked again at the next Heartbeat timeout for a maximum period of the Is Alive timeout. If no indication of the active host being alive is found within the Is Alive timeout, CA XOsoft executes event notification. Simultaneously, it checks whether or not to perform a switchover, as defined by the *Perform switchover automatically* property (see page 163).

Important! When using scripts, each script must reside in the same path and with the same name on both the Master and the Replica.

Is Alive Timeout (sec)

If the standby host does not receive indication that the active host is alive during this interval (in seconds), switchover or notification is initiated. The checks are performed at the Heartbeat Frequency.

Default is 300 seconds.

Heartbeat Frequency (sec)

Interval (in seconds) for sending heartbeat requests (performing the checks below).

Default is 30 seconds.

Check Method

Send Ping Request

ICMP requests are sent automatically from the standby host to the active host, to check if the active host is up.

The available options depend on the value of the **Move IP** property (see page 180).

- If **Move IP** is **On**

  During the switchover, the IP is moved from the active computer to the standby. Therefore, the standby computer must check this IP continuously.

  In the **IP for Ping** property, enter IP Address to ping.

- If **Move IP** is **Off**

  During the switchover, the IP is not moved from the active to the standby computer. Therefore, define two IPs for ping:

  IP for ping from Master to Replica
  Enter IP address to ping. When the Replica computer is the active host, an ICMP request is made from the Master to the Replica. If no reply is received within two seconds, the Replica computer is considered non-operational.

  IP for ping from Replica to Master
  Enter IP address to send ping to. When the Master computer is the active host, an ICMP request is made from the Replica to the Master. If no reply is received within 2 seconds, then the Master computer is considered to be non-operational.

Connect to DB

[For database applications only] When this property is set to On, CA XOsoft connects to the active computer's database according to the Heartbeat Frequency, and checks to see if the database services are running and all databases are mounted.

User-Defined Scripts

Allows the standard check methods to be augmented or replaced by user-defined actions in scripts.

Check Script on Active Host

This script runs on the active server, and checks if it is alive.

- **Script Name**

  Enter the name and full path of script to run. CA XOsoft connects to the active computer once every Heartbeat timeout, and executes the script. If the return value equals zero, the active computer is OK (alive). If the return value is different than zero, the active server is not responding and switchover is required.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values must be static.

Check script on Standby host

This script runs on the standby server, and checks if the active server it is alive.

- **Script Name**

  Enter name and full path of script to be run. CA XOsoft connects to the active computer once every Heartbeat timeout, and executes the script. If the return value equals zero, the active computer is OK (alive). If the return value is different than zero, the active server is not responding and switchover is required.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values must be static.

DB Management Table/Application/Shares Management

Automatic

If you want CA XOsoft to manage shares on your File Server or services on your DB Server, set this option to On. CA XOsoft then:

1. Auto-discovers the list of application (DB) services that must be managed.
2. Starts DB services on the active host (if they are not already running), and stops them on the standby host (if they are running) when you start the CA XOsoft scenario.
3. Stops DB services on the active host and starts them on the standby host during the switchover.
4. Run DB services on the active host after reboot, if a CA XOsoft scenario was running previously.

User-Defined Scripts

Start DB

If set to On, runs a user-defined script to augment or replace the start of DB services. This action occurs during scenario run or switchover on the active host.

- **Script Name (full path)**

  Enter the name and full path of the script to be run.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values must be static.

Stop DB

If set to On, runs a user-defined script to augment or replace the stop of DB services. This action occurs during scenario run or switchover on the active host.

- **Script Name (full path)**

  Enter the name and full path of the script to be run.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values must be static.

Actions upon Success

Important! When using scripts, each script must reside in the same path and with the same name on both the Master and the Replica.

User-Defined Script

When set to On, runs a user-defined script. The actions invoked by the script will be performed following the completion of a successful switchover.

- **Script Name (full path)**

  Enter the name and full path of script. This script runs on the active server after the switchover completion.

- **Arguments**

  Arguments to be passed to the script specified in the previous property. Argument values are static.

Move IP Redirection

This section describes the necessary steps for adding Move IP redirection to the High Availability scenario. Use this method only when both servers are on the same IP subnet.

Move IP on the Master Server

Before configuring the scenario, add an additional IP address to the Master host, (which is denoted as *XO-IP* in the following steps). If you have already configured your scenario go ahead and add this IP address; special steps for this situation are listed below. This new IP address is used for CA XOsoft internal communication and replication. This is necessary since the current production IP address is not available on the Master after switchover -- it switches to the Replica server.

Note: In the following screenshot, the XO-IP IP address is 192.168.220.23 and the current production server IP address is 192.168.220.111.

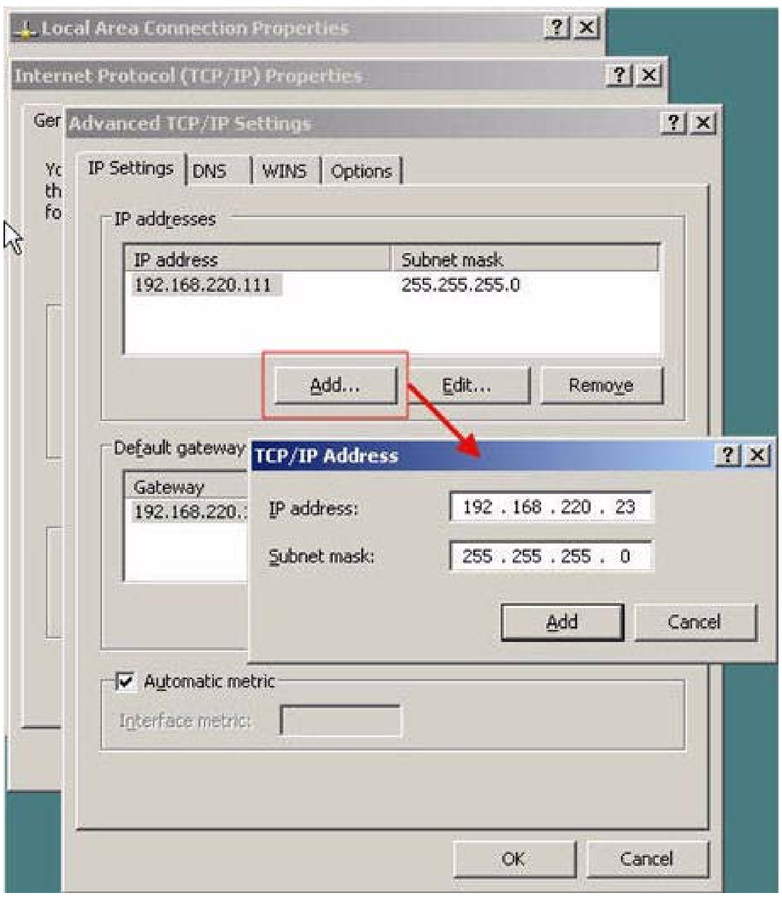

Move IP through the Manager

This section details Move IP redirection through the CA XOsoft Manager.

For New Scenarios

During the initial run of the Scenario Creation Wizard, enter the XO-IP and Replica IP addresses instead of the server names.

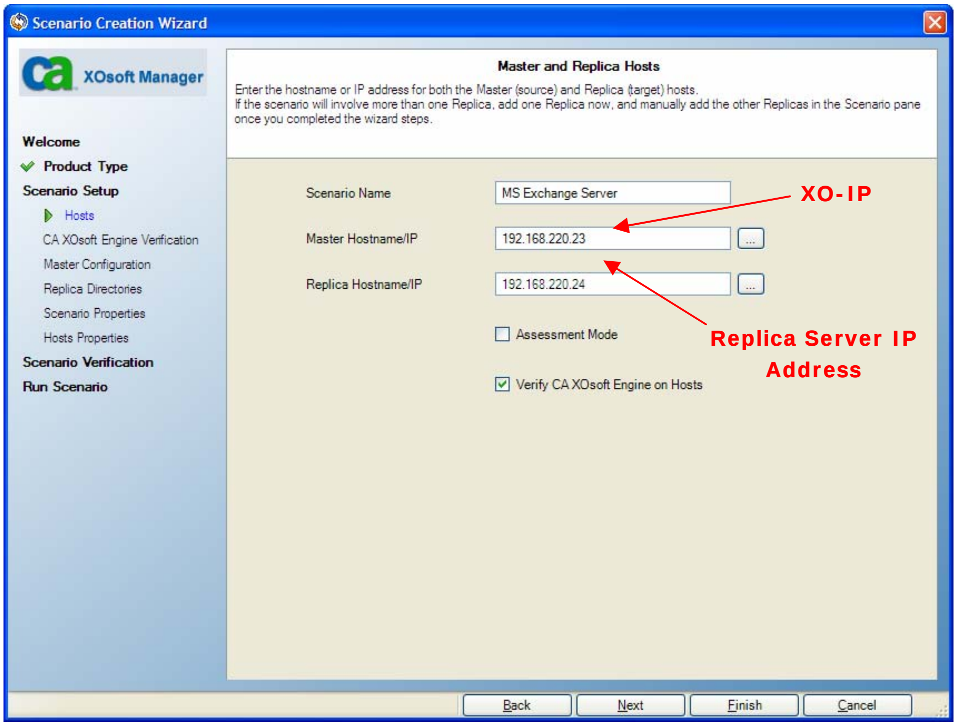

For Existing Scenarios

To use Move IP with existing scenarios:

1. On the Scenario pane, select the required Master host:

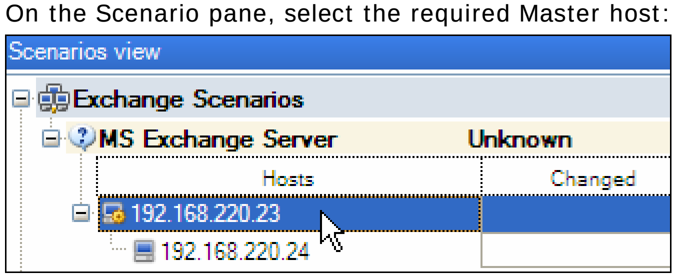

2. Right-click the Master and select **Rename** from the pop-up menu. Then, enter the XO-IP address.

3. On the Framework pane, select the **Switchover** tab and then select the Replica server as the switchover host.

4. Set the **Move IP** option to On. Ensure that the IP address under **Move IP**, **IP/Mask** matches the production server IP address: this is the IP address that will switch over. If you are moving more than one IP address you can add multiple production IP addresses by selecting **Click here to add new IP/Mask**:

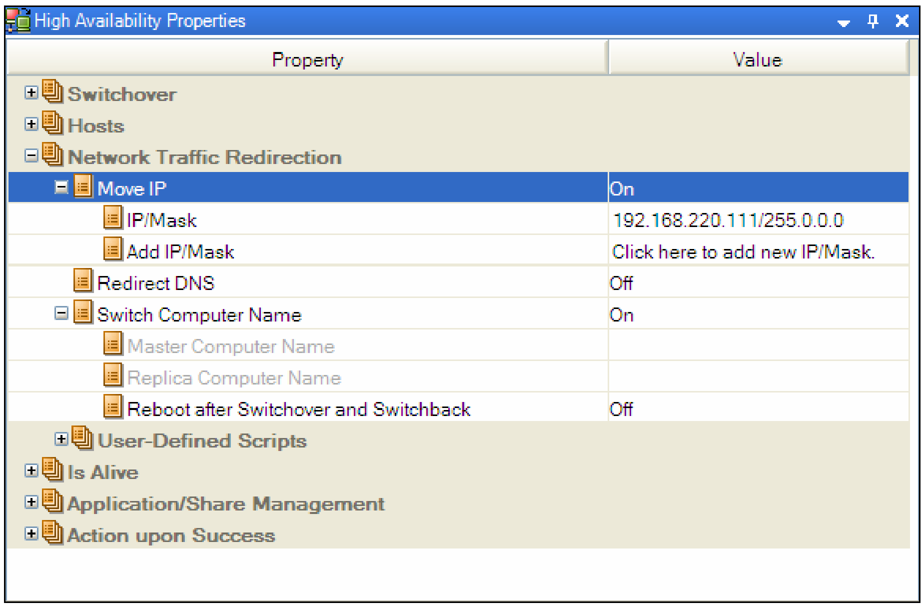

Cluster Move IP

Using Move IP redirection with a clustered Master (MSCS with shared storage) requires you to add an additional IP resource to the Master Exchange resource group. This section describes how to configure this redirection method.

Note: If both Master *and* Replica are clusters there are special configuration issues involved in the Move IP redirection process that are not detailed in this Guide. For a cluster-cluster scenario, use Redirect DNS or contact technical support to receive detailed instructions and guidance.

Using the Master Cluster

To use Cluster Move IP through the Master cluster:

1. Open the Cluster Administrator.

2. In the Master Cluster Exchange Resource Group, create a new IP resource and name it ***XO-IP***.

3. Bring this resource online and verify it is visible from the Replica via the ping command. This new IP address is used for CA XOsoftHA internal communication and replication. This is necessary since the current production IP address is not available on the Master cluster after switchover -- it switches to the Replica server.

Using the Manager

This section details Cluster Move IP redirection using the Manager.

For New Scenarios

During the initial run of the Wizard, enter the XO-IP and Replica IP addresses instead of the cluster virtual server names.

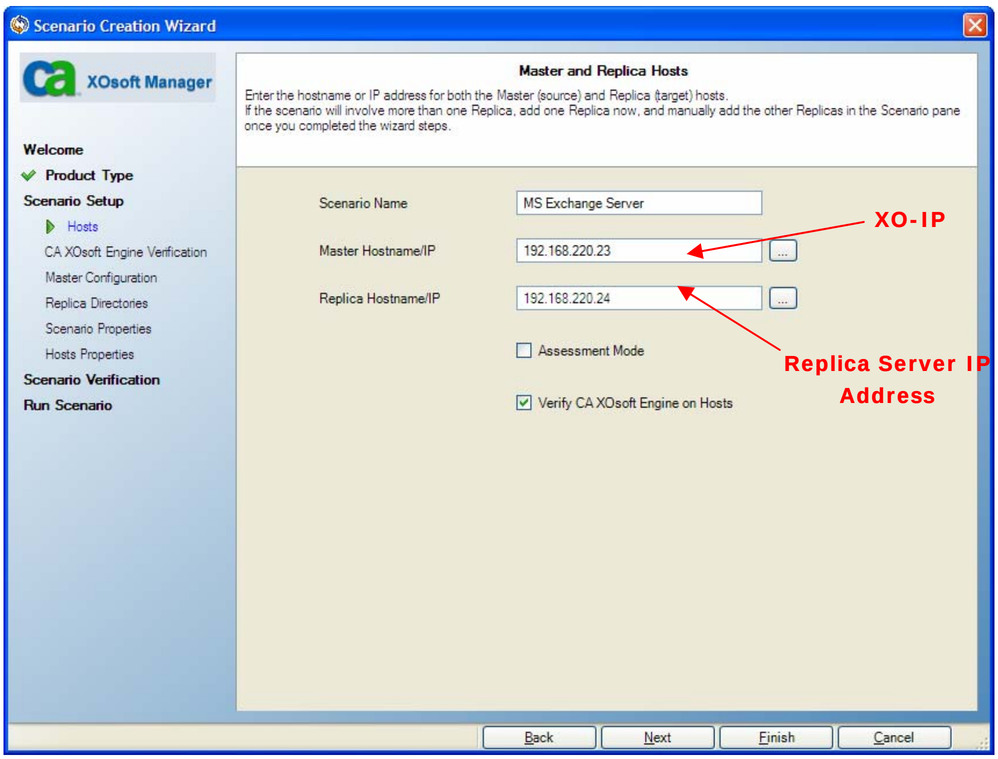

For Existing Scenarios

To use Cluster Move IP with existing scenarios

1. On the Scenario pane, select the required Master host:

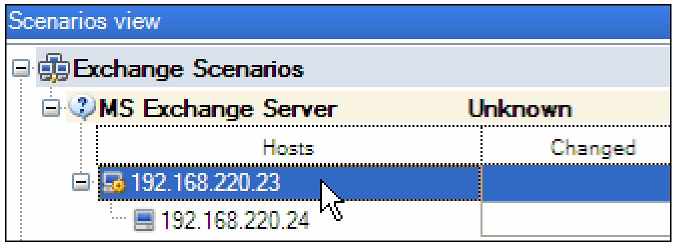

2. Right-click the Master and select **Rename** from the pop-up menu. Then, enter the XO-IP address.
3. On the Framework pane, select the **Switchover** tab and then select the Replica server as the switchover host.
4. Set the **Move IP** option to On. Ensure that the IP address under **Move IP, IP/Mask** matches the production server IP address: this is the IP address that will switch over. If you are moving more than one IP address you can add multiple production IP addresses by selecting **Click here to add new IP/Mask:**

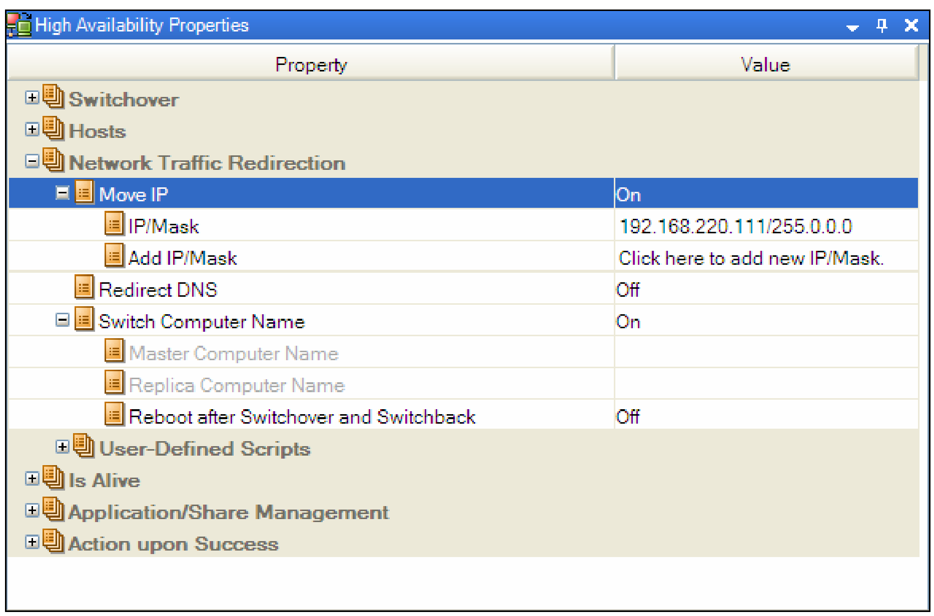

Chapter 12: Testing Assured Recovery and Managing VSS Snapshots

This chapter explains the Assured Recovery testing option, and describes the following operations: creating AR scenario, performing AR test in a scheduled and non-scheduled mode, and configuring the AR properties. In addition, this chapter describes how to set up VSS snapshot creation and how to manage snapshots.

About Assured Recovery

The Assured Recovery option enables you to perform a full transparent test of the recoverability of your data on the Replica server. The Replica server that is tested is the one that would take over the production server if it will be down. The Assured Recovery option is a true test of the actual server, applications and actions that will be required in the event the Replica server will have to switch, become the Active server, and carry out its functions.

This Assured Recovery test is executed by starting up database services, and performing whatever operations that are required to verify the integrity of the data. All this is done without any need to perform resynchronization, and without impacting either the availability of the production server, or the safety that the DR and HA systems are designed to provide.

During the test, the data changes that continue to take place on the Master are sent to the Replica, but they are not immediately applied. Instead, these changes are accumulated and stored in a spool, and only when the testing is completed they are applied to the Replica data. Since the spooling occurs on the Replica, if something happens to the Master during the testing process, none of the accumulated changes are lost.

Once the testing is finished, the Assured Recovery option stops the application services it started on the Replica. Then, the Replica server is automatically rewind to precisely the state that existed when the replication was paused and the test started. This way, the accumulated changes in the spool can be applied as if no testing has occurred. From this point on, the DR or HA scenario continues normally. In the case of an HA scenario, if a failure of the Master occurred during the testing, switchover will begin.

The Assured Recovery test can be fully automated and performed on a scheduled basis as often as needed. Upon completion, appropriate personnel can be alerted with the status of the test, and additional actions can be triggered on success, for example, taking a VSS snapshot of the Replica data or creating a backup. In addition, you can perform AR test in a non-scheduled mode when the need arises.

The Assured Recovery testing is tailored to all supported application and database servers. However, since the Assured Recovery option tests database services, it is less applicable for File and IIS Servers. You can still use the Assured Recovery option with these servers for special tasks. For example, you can automatically suspend replication on a regular basis during several hours each day, week or month, and run scripts in this interval, or you can use this suspension to take VSS snapshots on the Replica. Since there is no *application* per se, testing of the data with File and IIS Servers scenarios requires additional custom scripts.

The Assured Recovery option supports both DR and HA solutions. However, it is best suited for HA since in this case the Replica server necessarily contains the actual database servers, on which the test is performed, and not only data. If you are using AR test as a part of DR scenario, you must verify that the root directories path is the same on the Master and the Replica. In addition, the Replica should have database application installed, or share files if you test a File Server, and they need to be configured on the Master and the Replica in exactly the same way. Otherwise, the AR test will not produce meaningful results.

Creating Assured Recovery Testing Scenario

The Assured Recovery testing feature must be enabled during the creation of the scenario that will later use it. For this reason, you can not perform Replica testing within a DR or HA scenario that is already running, and was not configured to use the Assured Recovery option. To use Assured Recovery, it is necessary to create a new scenario with the **Integrity Testing for Assured Recovery** option turned to On.

Note: This section demonstrates the creation of an Assured Recovery testing scenario for Exchange Server HA.

To set Assured Recovery testing scenario

1. Open the CA XOsoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New** button on the Standard toolbar.

   The **Scenario Creation Wizard** appears:

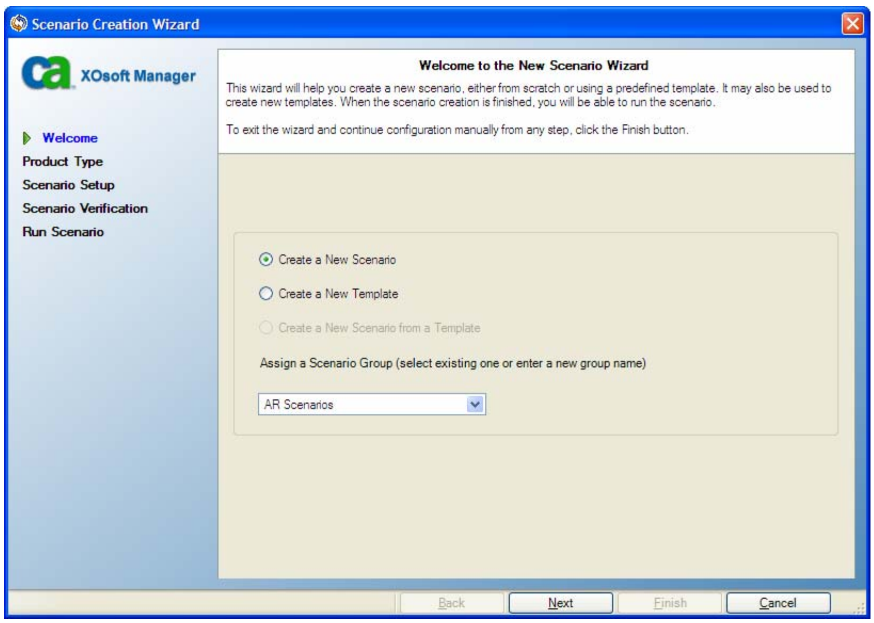

2. Select the required scenario options, as follows:
   - Select the **Create a New Scenario** option button.
   - From the **Group** drop-down list, select the group to which you want to assign the new scenario, or enter a name for a new scenario group.

3. Click **Next**. The **Select Server and Product Type** page appears:

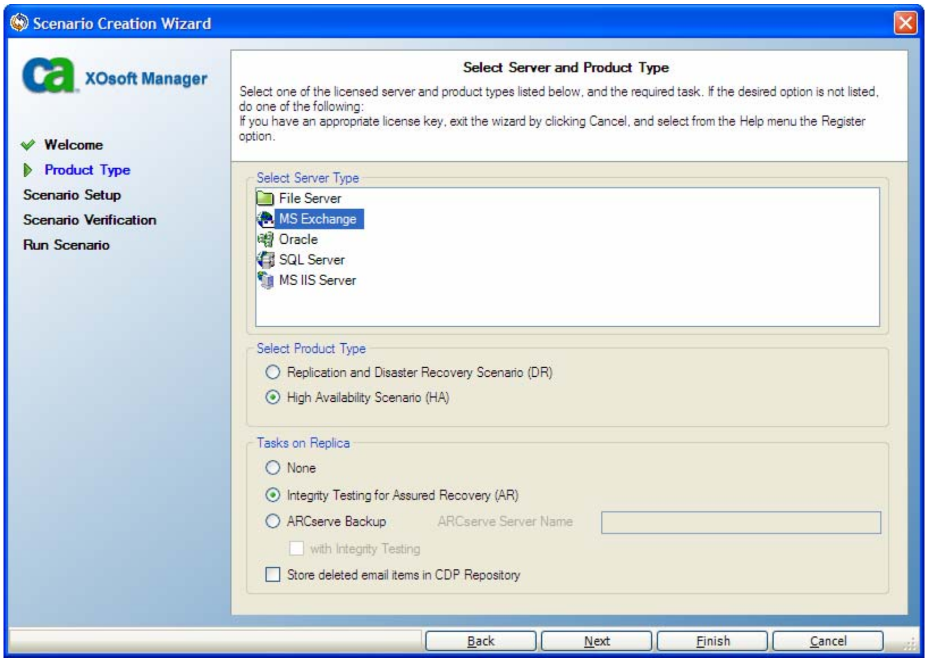

A list of available applications and scenario types is presented.

Note: The list of available applications depends on the licenses applied**.**

4. Select the required scenario options, as follows:

- From the **Select Server Type** list, select the type of server for which you want to create the AR scenario. For this example, we will use **MS Exchange**.
- From the **Select Product Type** options, select either **Replication and Disaster Recovery** or **High Availability Scenario**.

  Note: The Assured Recovery test is best suited for HA scenarios. If you select the DR option, you must verify that the root directories path is the same on the Master and the Replica. In addition, the Replica should have database application installed, or share files if you test a File Server. Otherwise, the AR test will NOT produce meaningful results.
- From the **Tasks on Replica** options, select the **Integrity Testing for Assured Recovery** option.

5. Click **Next**. The **Master and Replica Hosts** page appears:

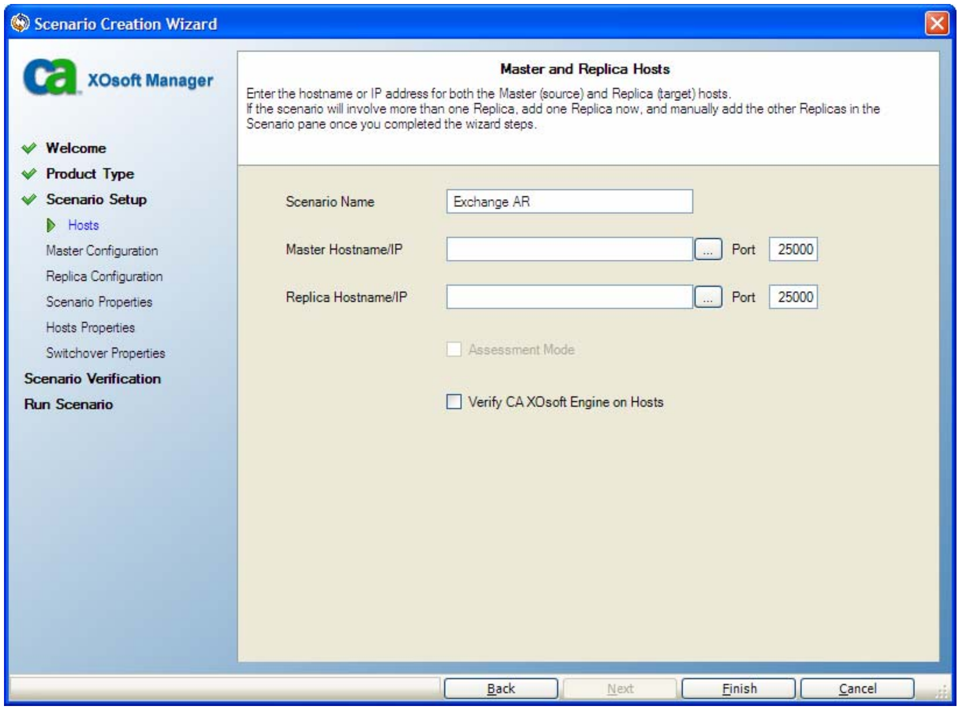

6. Enter the following information:

   - In the **Scenario Name** box – accept the default name or enter a new name for the scenario. When entering a name, choose a unique name, since you cannot use the same name for more than one scenario.
   - In the **Master** and **Replica Hostname/IP** boxes - enter the hostname or IP address of the Master (active) and Replica (standby) servers, or use the **Browse** buttons to find them.

Important! Only one Replica can be configured for AR testing in a single scenario. If at a later stage you will add a Replica to the scenario and try to configure it for the AR test, the following message will appear: **Only one scheduled task per scenario can be set. Replica integrity testing for Assured Recovery for host [Replica_name] is already switched on. Do you want to turn this option off now?** To switch the test to the second Replica, you will need to click **Yes**.

Note: If either server is a MSCS cluster, enter the Exchange Virtual Server Name or IP address as the Master and/or Replica name (instead of the physical node's name/IP).

- In the **Port** boxes: accept the default port no. (25000), or enter a new port numbers for the Master and Replica.
- The **Verify CA XOsoft Engine on Hosts** check box - select this check box if you want the system to verify whether Engines are installed and running on the Master and Replica hosts you specified in this page. If Engines are not installed on the selected hosts, you can use this option to remotely install the Engines on one or both hosts. For more information about the **Host Verification** page, see page 46.

7. After you entered or selected the desired options, click **Next**. The **Master Configuration** page appears:

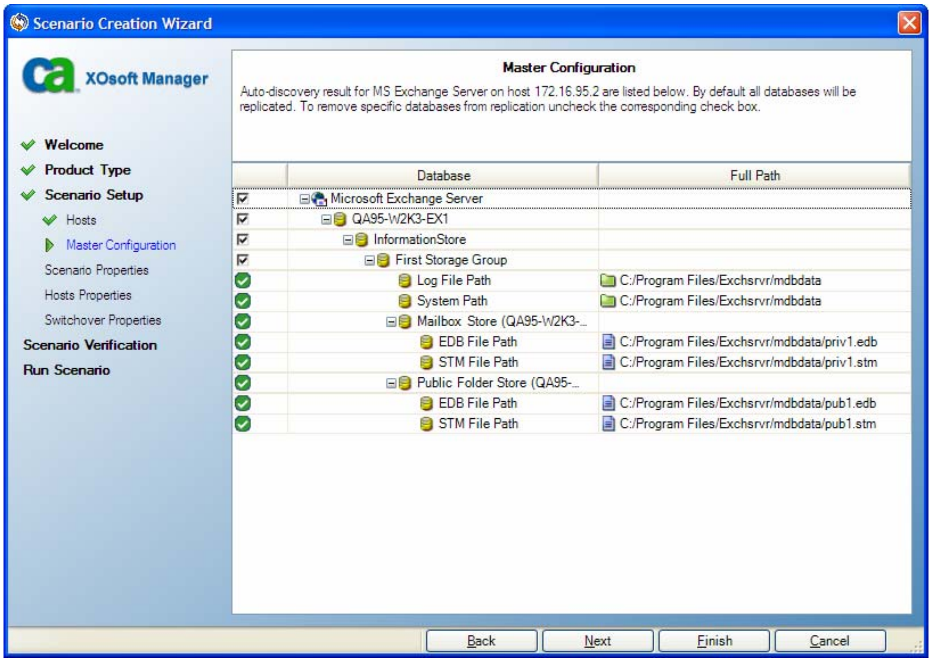

The CA XOsoft auto-discovery component automatically displays the Exchange databases that are on the Master server. These are the databases that can be replicated and protected.

8. By default, all the discovered databases are selected and all will be replicated. You can exclude any of these storage groups from replication by clearing their check boxes.

9. Click **Next**. The **Replica Configuration** page appears:

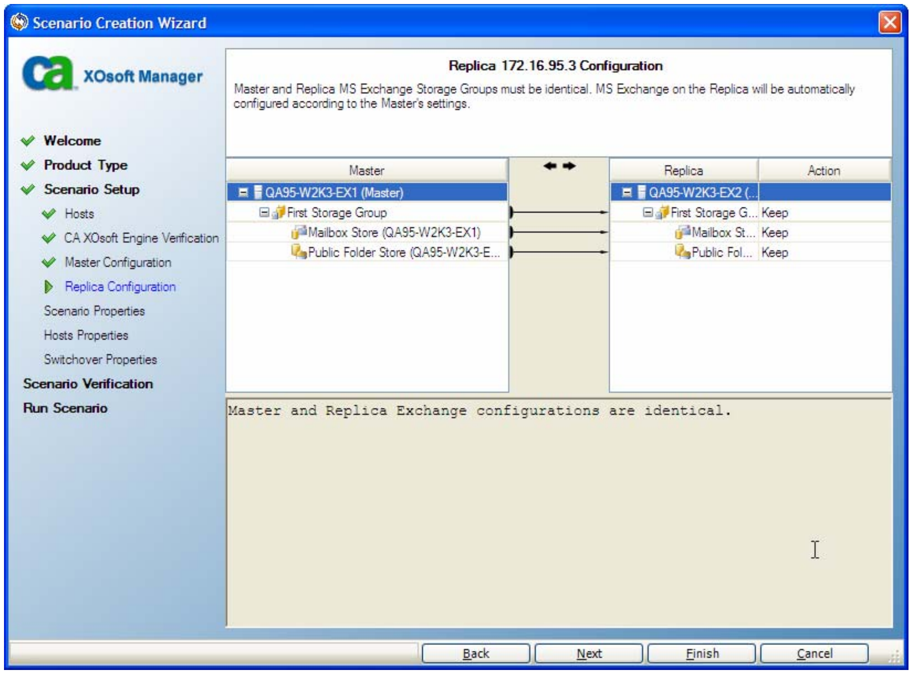

CA XOsoft auto-configuration component verifies that the Exchange Server configuration on the Master and Replica servers will be identical during the replication procedure. This means that if there are discrepancies, CA XOsoft will perform the required actions, including: deleting storage groups, public folders or mailbox stores from the Replica, create new ones and make modifications to existing ones. The actions that will be performed during the configuration process are indicated in the **Action** column on the right.

Note: For more information about the actions that can be performed in the Replica auto-configuration process, see page 160.

10. Review the changes that will occur during the automatic configuration on the Replica Exchange server, and make sure you want them to be performed.

Note: If a **Remove** action is indicated, make sure that you are ready to delete the specified storage item from the Replica server, since it does not have an automatic backup. If you want to save it in a different location before deletion, click the **Finish** button to exit the wizard.

11. Click **Next** to start the Replica configuration process. The **Scenario Properties** page appears:

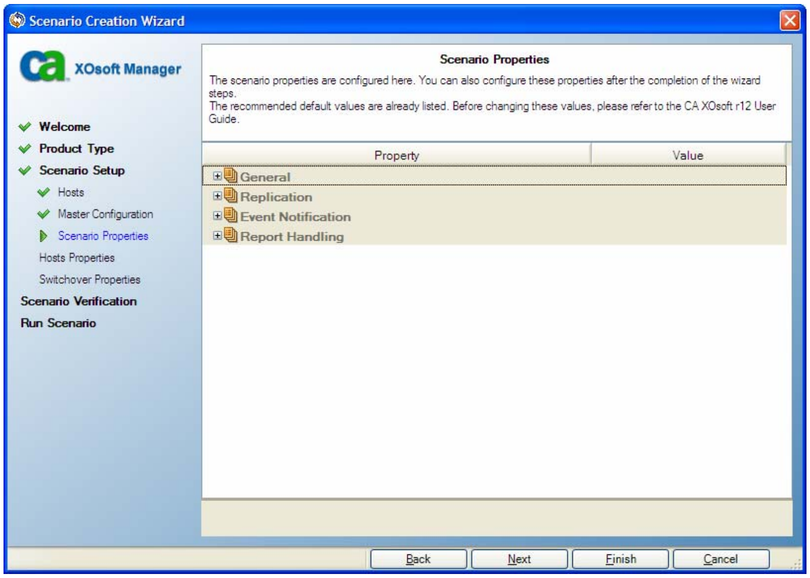

The **Scenario Properties** page enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

If you want to configure the scenario properties at this stage, refer to *Understanding Scenario* Properties, page 123. To configure the scenario properties at a later stage, refer to *Configuring Scenario Properties*, page 121.

12. Click **Next**. The **Master and Replica Properties** page appears:

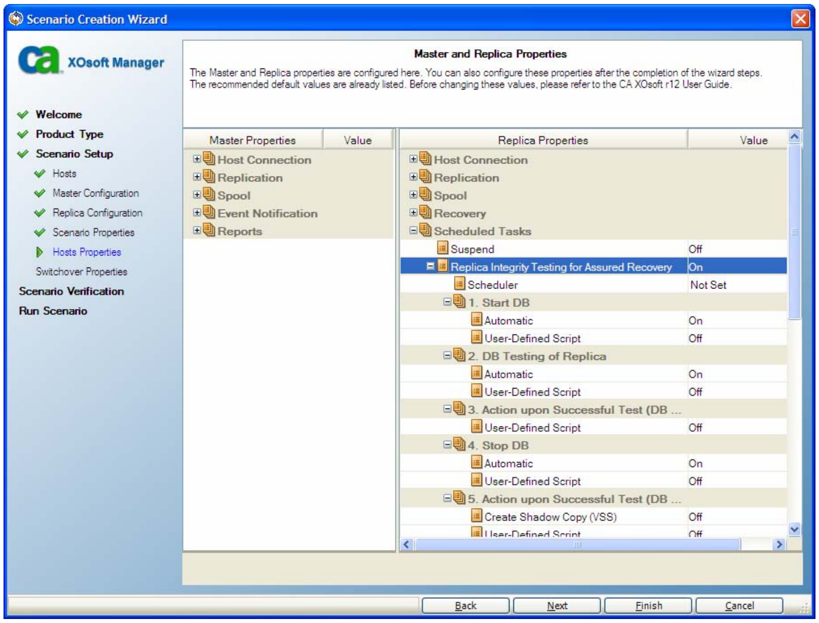

The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

13. To verify that the Assured Recovery option is active, under the **Replica Properties** list on the right, open the **Scheduled Tasks** group and ensure that the **Replica Integrity Testing for Assured Recovery** property is set to On. You can leave the default values of the other related properties, and change them later if needed. For more information about AR properties refer to *Understanding* Assured Recovery Properties, page 215.

    If you want to configure the Master and Replica properties at this stage, refer to *Setting Master and Replica Properties*, page 133. To configure the Master and Replica properties at a later stage, refer to *Configuring Master or Replica Server Properties*, page 133.

    Note: You can modify all the settings in this pane after the scenario is created. However, before changing any Spool properties (which can be configured here), review the *Spool* info on page 136 for configuration details.

14. Click **Next**. If you selected HA solution the **Switchover Properties** page appears:

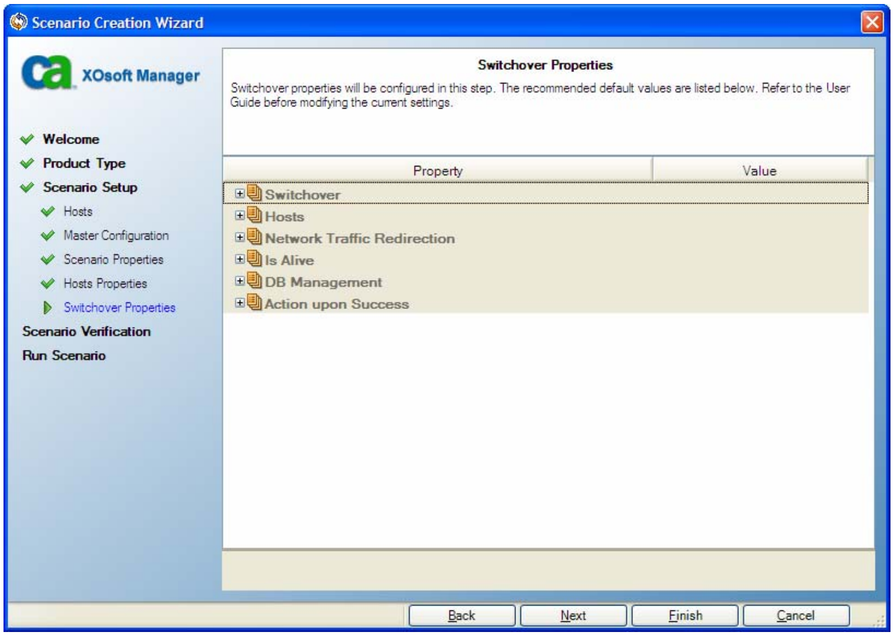

15. From this stage, set up the scenario as you would normally following the instructions on page 42 or in the appropriate Operation Guide. After the scenario creation is completed, run the scenario:

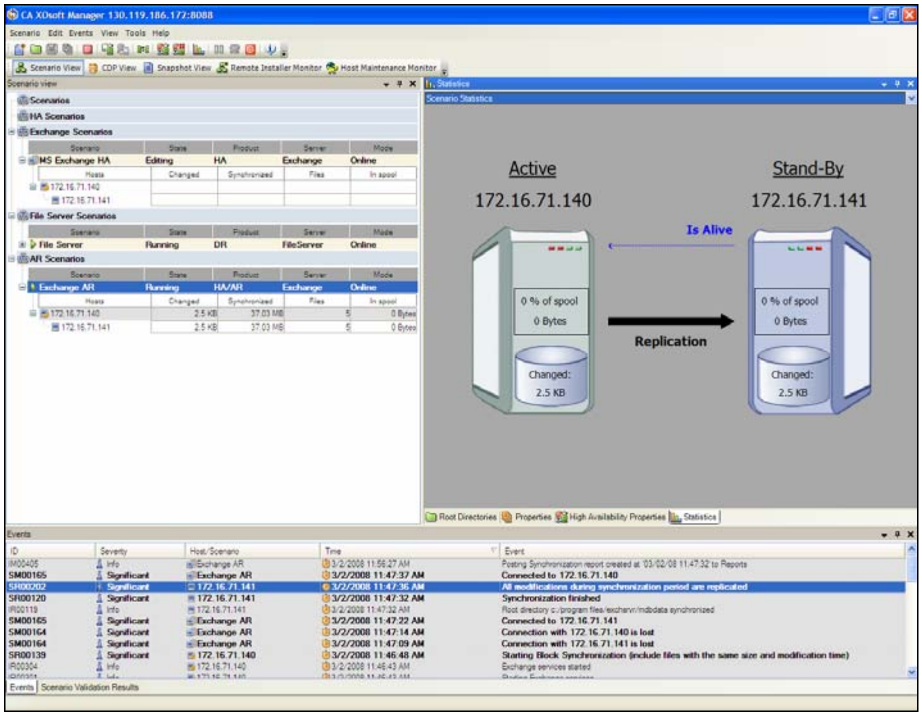

Once the initial synchronization is completed and the replication process is active, the AR test can be performed.

Performing Assured Recovery Test

The Assured Recovery test can be fully automated and performed on a scheduled basis as often as needed. Upon completion, appropriate personnel can be alerted with the status of the test, and additional actions can be triggered on success, for example, taking a VSS snapshot of the data or a running a backup. Alternatively, you can perform AR test in a non-scheduled mode, by initiating the test whenever the need arises.

In both modes, the AR test is performed in steps, according to the AR configuration settings. Some of the steps are transparent, and are executed automatically by CA XOsoft whenever an AR test is performed. Other steps are visible, and can be configured as to whether and how they will be performed.

The standard steps are as follows:

1. Initiating Assured Recovery test – the AR test can be initiated either on a scheduled basis, or in a non-scheduled mode by clicking the **Replica Integrity Testing** button on the Manager.
2. Suspending application of data changes on the tested Replica – this step is performed automatically by CA XOsoft at the beginning of each AR test.
3. Initiating a rewind component on the tested Replica - this step is performed automatically by CA XOsoft. It is aimed at capturing all the changes that are made to the Replica data during the test, so they can be later rewind back to the point when the replication was suspended.
4. Starting the database services – by default, this step is performed automatically by CA XOsoft. However, it can be switched of, replaced, or followed by a user-defined script.
5. Testing the database - the databases are verified, by default, using the same tests that are used to monitor the database in HA. These tests include verifying that all services have correctly started and that all databases have been successfully mounted. These tests can be switched off, replaced, or followed by a user-defined script.
6. Performing actions upon successful test while the database services are running - a user-defined script may be registered at this point to perform actions that are desired in the event of a successful test, but which also require that the application will be running.
7. Stopping the database services – by default, this step is performed automatically by CA XOsoft. However, it can be switched of, replaced, or followed by a user-defined script.
8. Performing additional actions upon successful test while the database services are stopped – this step is optional, and it may be used to perform actions that take advantage of the fact that the application passed validation tests and that it was stopped in a systematic order.

9. Rewinding AR Replica data and resuming replication - this step is performed automatically by CA XOsoft at the end of each AR test. It restores the Replica data to precisely the state it was before the test begun using the rewind technology. Then, it resumes replication.

Performing Assured Recovery Test in a Scheduled Mode

When you set the AR test to run in a scheduled mode, it means that an AR test will be performed automatically on a regular basis. Once you select this option the following flexible scheduling capabilities are offered:

- Testing on selected days of the week and for specific hours in a 24 hour cycle.
- Testing over selected periods (e.g., once every 36 hours) in a 7 day cycle.
- Exclusion of specific dates.

The AR test schedule can be set when the scenario is created or at a later stage.

Note: You can define only one scheduled task per scenario. If you attempt to configure AR testing while you already have a scheduled **Suspend** operation configured, the following message appears: **Only one scheduled task per scenario can be set. Suspend for host [Replica_name] is already switched on. Do you want to turn this option off now?** To switch the schedule option to the AR test, you need to click **Yes**.

To schedule the AR test

1. On the Scenario pane, select the Replica you want to test.
2. On the Framework pane on the left, select the properties tab.

   The Replica Properties list appears.
3. If the scenario is running, click the **Stop** button on the Standard toolbar.

   The scenario is stopped.

4. On the Replica properties list, open the **Scheduled Tasks** group. Then, under the **Replica Integrity Testing for Assured Recovery** group, select the **Scheduler** property, and click the **Not Set** value:

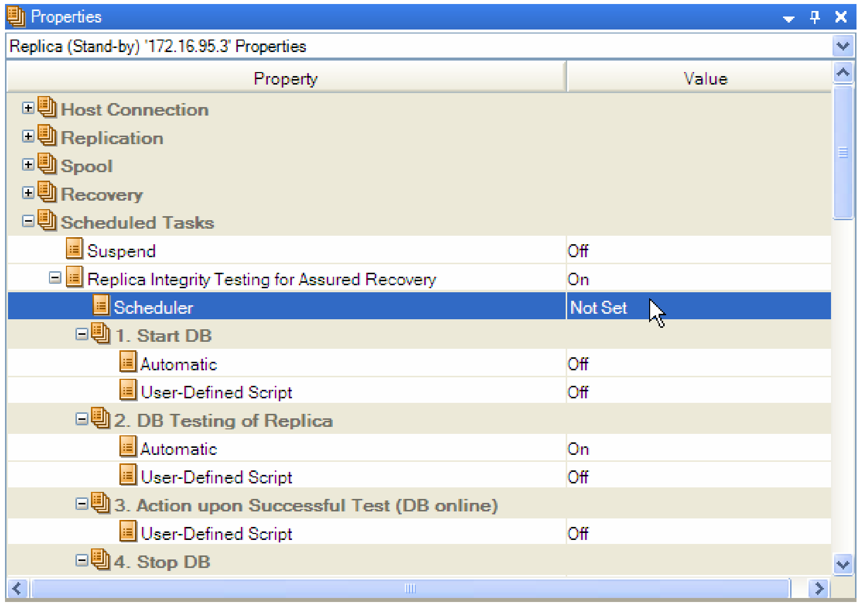

The **Assured Recovery hours** dialog appears:

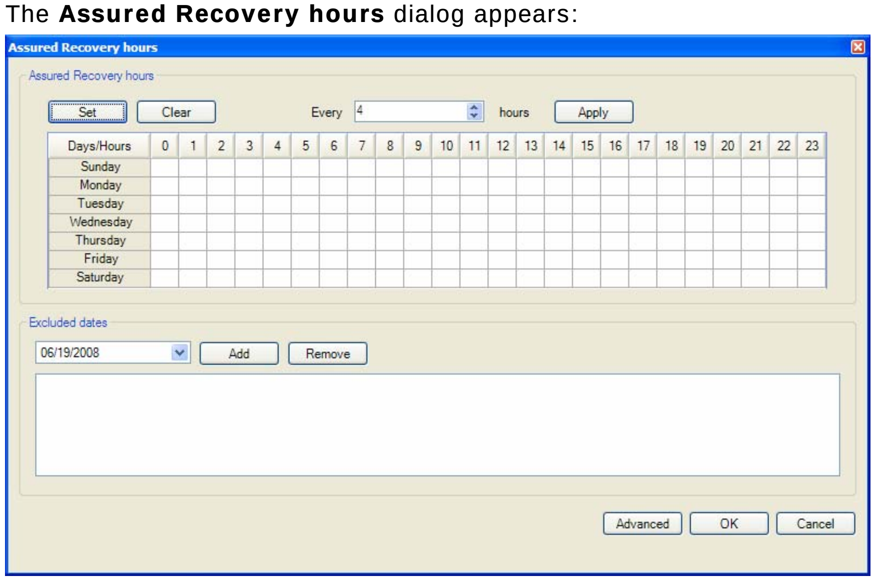

The **Assured Recovery hours** dialog is similar to the **Schedule Setting** dialog, which is used for scheduling automatic synchronization. For information about setting a schedule, see *Schedule Synchronization*, page 128.

5. Set the schedule for automatic AR testing in the **Assured Recovery hours** dialog, and click **OK** to save your schedule and close the dialog.

6. To activate the scheduler, click the **Save** button on the Standard toolbar and start the AR scenario.

   The Replica you selected for testing will be tested on a regular basis according to the schedule you set.

Performing Assured Recovery Test in a Non-Scheduled Mode

In a non-scheduled mode, you can test Assured Recovery either automatically or manually. When you are using the automatic method, all you need to do is initiate the AR test by a click of a button. Then, CA XOsoft automatically performs all the test steps, according to the AR configuration setting. Once the test is completed, the regular replication is resumed. There is only one difference between this method and a scheduled AR test. In a non-scheduled automatic mode you are the one who initiates the test whenever you need, not the Scheduler.

When you are using the manual method, you also need to initiate the AR test by a click of a button. However, unlike the automatic method, CA XOsoft will suspend the test after the first standard step – starting the database service. This will occur even when all standard steps are configured as Automatic.

Note: If the **Start DB** property is set to Off, and there is no user-defined script that replaces it, the only thing that CA XOsoft will do is suspending the application of changes to the Replica in preparation for the manual AR test.

Once the replication is suspended, you can perform tests or actions directly on the Replica without the need to later resynchronize the Master and Replica. You can use this option for manually testing applications or data on the Replica, or for performing tasks on the Replica instead of the Master, such as report generation, in order to reduce the Master workload.

When you finish the manual testing or operation, you need to manually stop the AR test suspension. This is done again by a click of a button. If other steps and actions were configured in the AR test, such as stopping the database services, they will be performed after you clicked the button for stopping the test and before the test will be declared as finished. When the test is considered finished, the replication will resume automatically.

Perform Assured Recovery Test Automatically

To perform Assured Recovery test automatically

1. On the Manager, verify that the AR scenario is running:

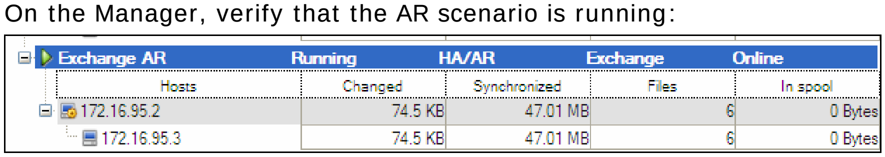

2. To start the AR testing, on the Scenario pane select the Replica you want to test. Then, click the **Replica Integrity Testing** button on the Standard toolbar, or right-click the Replica and select **Replica Integrity Testing** from the pop up menu:

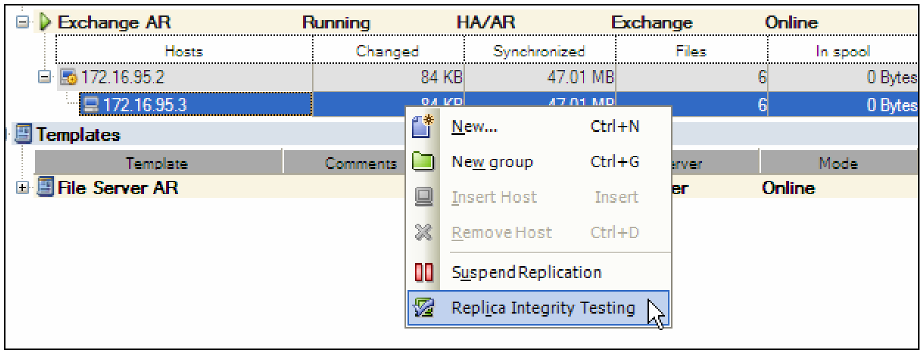

The **Replica Integrity Testing for Assured Recovery** dialog appears:

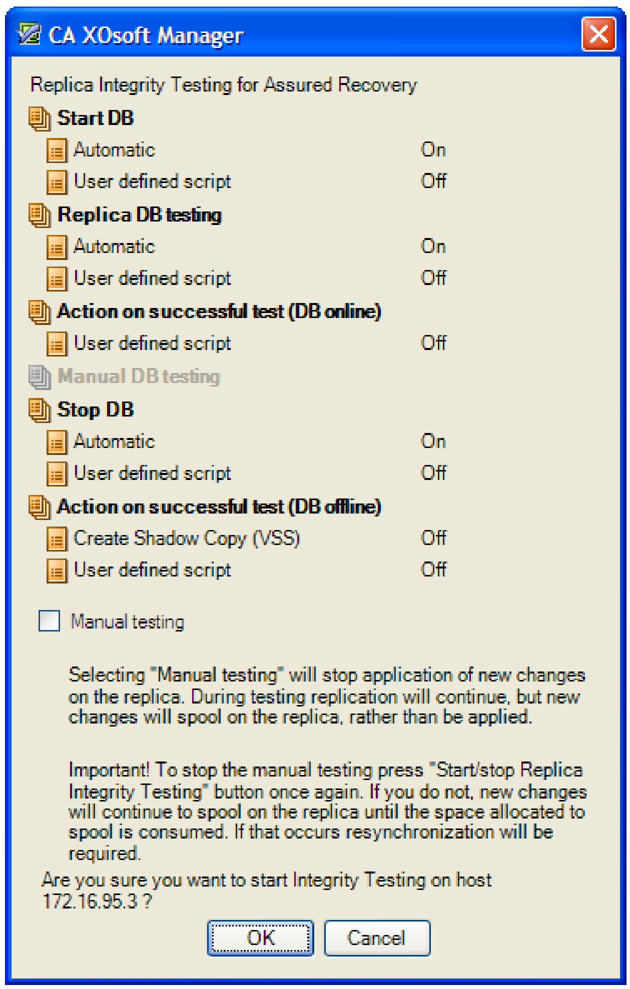

In this dialog, the configuration you set for the AR test is displayed.

3. To start the automatic AR test using the existing configuration, click **OK**.

   Notes:

   - To change the test configuration before running the test, click **Cancel,** and refer to *Setting Assured Recovery Properties*, page 213.
   - To manually perform the AR test, select the **Manual testing** check box, click **OK**, and refer to *Perform Assured Recovery Test* Manually, page 210.

4. After you initiate the AR testing, the **Replica Integrity Testing for Assured Recovery** dialog is closed. Then, before the test begins to run, CA XOsoft verifies that no synchronization, AR test or replication suspension are currently in progress on any of the hosts that participate in the current scenario.

5. Once the verification stage ends, the AR test begins:

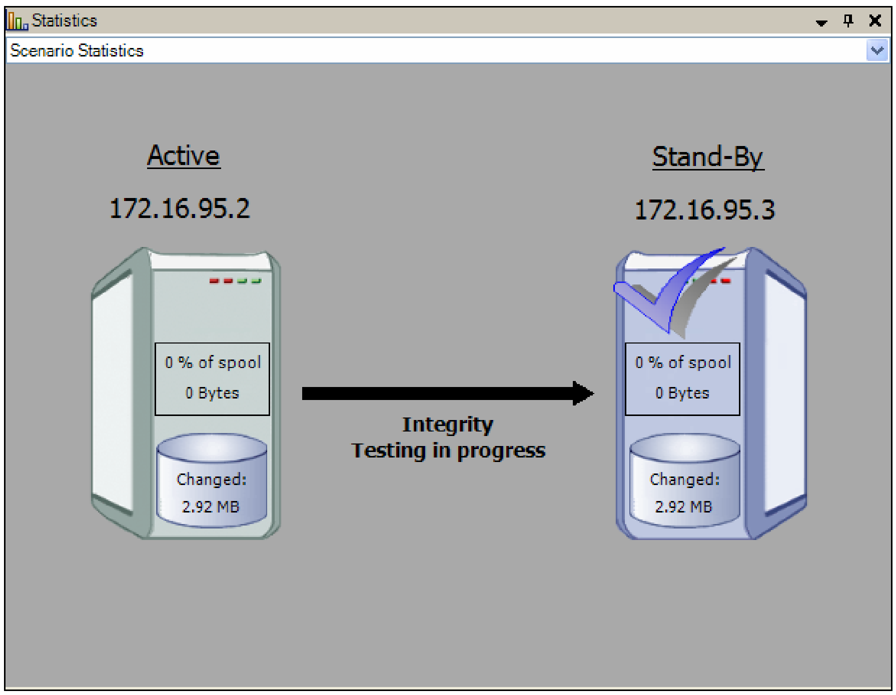

   The steps of the test are displayed as messages in the Event pane:

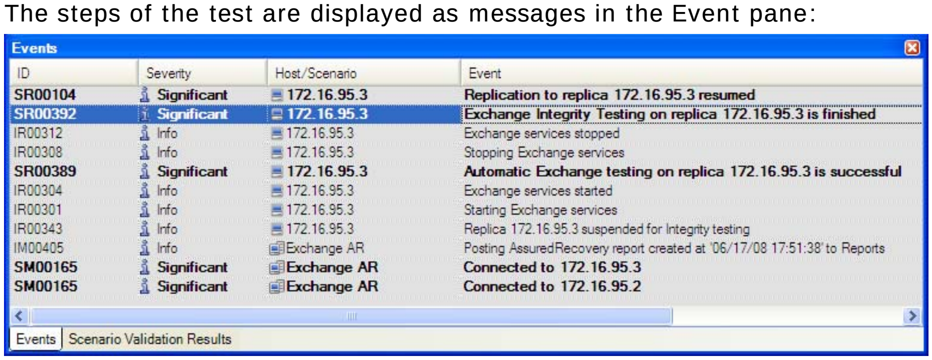

Events

ID	Severity	Host/Scenario	Event
SR00104	Significant	172.16.95.3	Replication to replica 172.16.95.3 resumed
SR00392	Significant	172.16.95.3	Exchange Integrity Testing on replica 172.16.95.3 is finished
IR00312	Info	172.16.95.3	Exchange services stopped
IR00308	Info	172.16.95.3	Stopping Exchange services
SR00389	Significant	172.16.95.3	Automatic Exchange testing on replica 172.16.95.3 is successful
IR00304	Info	172.16.95.3	Exchange services started
IR00301	Info	172.16.95.3	Starting Exchange services
IR00343	Info	172.16.95.3	Replica 172.16.95.3 suspended for Integrity testing
IM00405	Info	Exchange AR	Posting AssuredRecovery report created at '06/17/08 17:51:38' to Reports
SM00165	Significant	Exchange AR	Connected to 172.16.95.3
SM00165	Significant	Exchange AR	Connected to 172.16.95.2

Events | Scenario Validation Results

6. After the test is finished, the Replica is automatically restored to precisely the same state it was when the replication was suspended. This is done through the underlying rewind technology. Then, the changes that were accumulated in the spool are applied, and the replication is resumed and continues normally.
7. By default, once AR test is performed, an Assured Recovery Report is generated:

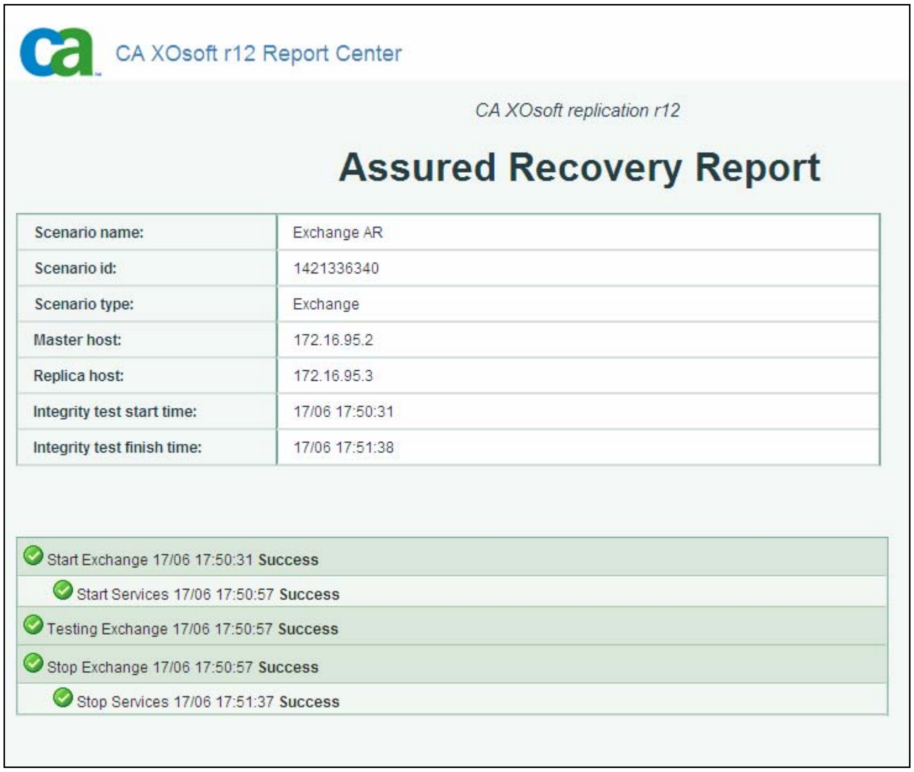

CA XOsoft r12 Report Center

CA XOsoft replication r12

Assured Recovery Report

Scenario name:	Exchange AR
Scenario id:	1421336340
Scenario type:	Exchange
Master host:	172.16.95.2
Replica host:	172.16.95.3
Integrity test start time:	17/06 17:50:31
Integrity test finish time:	17/06 17:51:38

Start Exchange 17/06 17:50:31 **Success**
Start Services 17/06 17:50:57 **Success**
Testing Exchange 17/06 17:50:57 **Success**
Stop Exchange 17/06 17:50:57 **Success**
Stop Services 17/06 17:51:37 **Success**

Notes:

- If the Assured Recovery Report was not generated, on the Replica Properties list under the **Reports** group, check the value of the **Generate Assured Recovery Report** property (see page 143).
- To view the report, refer to *Viewing a Report*, page 84.

All the tasks that were performed during the test are listed in the AR Report, along with their activation time and status.

Perform Assured Recovery Test Manually

To perform Assured Recovery test manually

1. On the Manager, verify that the AR scenario is running.
2. To start the AR testing, on the Scenario pane select the Replica you want to test. Then, click the **Replica Integrity Testing**  button on the Standard toolbar, or right-click the Replica and select **Replica Integrity Testing** from the pop up menu.

   The **Replica Integrity Testing for Assured Recovery** dialog appears:

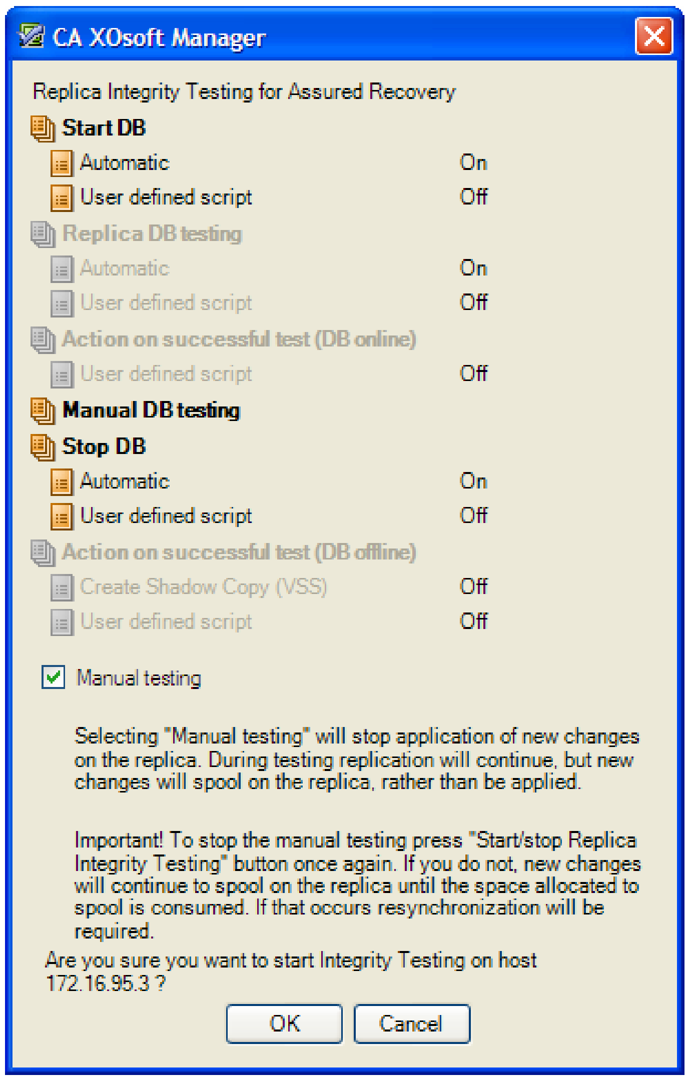

In this dialog, the configuration you set for the AR test is displayed.

3. To start the manual AR test using the existing configuration, select the **Manual testing** check box. Once this check box is selected, the dialog changes to reflect only the actions that will be performed in a manual mode.

   Notes:

   - To change the test configuration before running the test, click **Cancel** and refer to *Setting Assured Recovery Properties*, page 213.
   - To automatically perform the AR test, clear the **Manual testing** check box, click **OK** and refer to *Perform Assured Recovery Test* Automatically, page 206.

4. Click **OK** to close the dialog and start the manual testing.

   - If the **Start DB** property is set to On, or a user-defined script is set to replace it, these actions are performed and then the test is suspend.
   - If no action is set to take place at this step, the replication and test are suspended at this stage.

5. From this stage, the only automatic action that CA XOsoft performs, unless other actions are configured as Automatic, is suspension of updates on the Replica.

6. Once the replication is suspended, the following message appears in the Event pane: **Replica is ready for Manual Integrity Testing**.

   Now, you can start performing any test you want directly on the Replica host, including making changes to the database. Note that these changes will not be saved once the AR test will be finished, due to the rewind process.

   Important! Do not restart the tested Replica at this stage. If you do, all the changes that accumulated on the spool will be lost.

7. After you finished testing the Replica host, click again the **Replica Integrity Testing**  button to resume replication.

   Important! If you do not click the **Replica Integrity Testing** button a second time at the end of the test, changes will continue to spool up on the Replica host. Eventually, the spool on the Replica host overflows and the scenario is stopped.

   A confirmation message appears:

8. Click **Yes** to stop the AR test. If other steps and actions were configured in the AR test, such as stopping the database services, they will be performed before the test will be declared as finished. When the test is considered finished, the replication will be resumed automatically.
9. After the test is finished, the Replica is automatically restored to precisely the same state it was when the replication was suspended. Then, the changes that were accumulated in the spool are applied, and the replication is resumed and continues normally.
10. By default, once AR test is performed, an Assured Recovery Report is generated. (For more detail about the Assured Recovery Report, see page 209.)

Setting Assured Recovery Properties

The Assured Recovery property values determine the entire scenario's default behavior concerning AR testing: scheduling, starting and stopping database services, running user-defined scripts upon success, creating VSS snapshots of the Replica data, and more.

Configuring Assured Recovery Properties

To configure AR properties, the scenario must be stopped.

Note: The Properties pane and its tabs (Root Directories, Properties, Statistics) are context sensitive, and change whenever you select a different node from a scenario folder.

To set AR scenario properties

1. On the Scenario pane, select the Replica that you want to test and whose properties you want to configure.
2. On the Framework pane on the left, select the Properties tab.

   The Replica Properties list appears:

   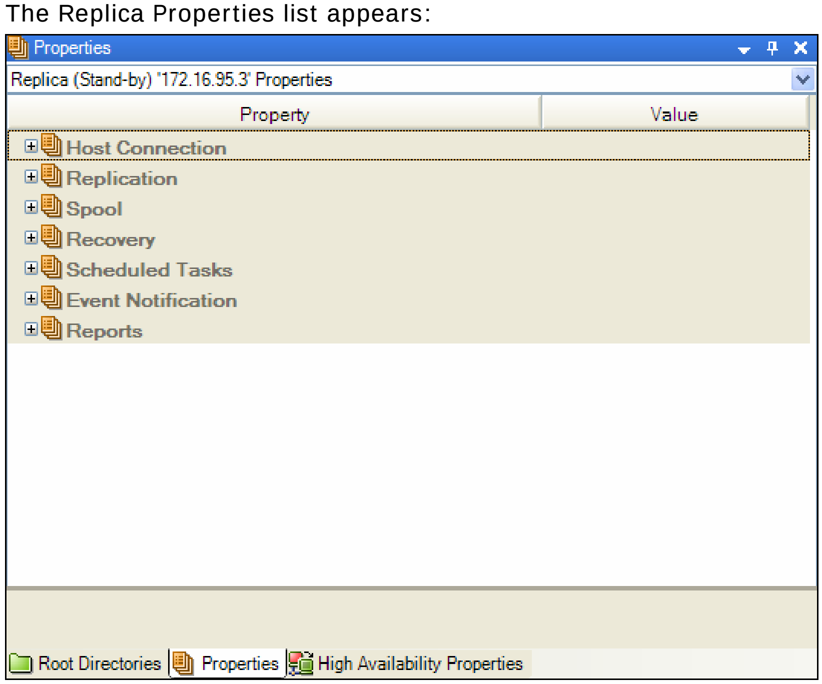

   Note: A running scenario has a gray background, and scenarios that are not running have a white background.

3. If the scenario is running, click the **Stop** button on the toolbar. The scenario is stopped.

4. On the Replica Properties list, open the **Scheduled Tasks** group to display the **Replica Integrity testing for Assured Recovery** properties:

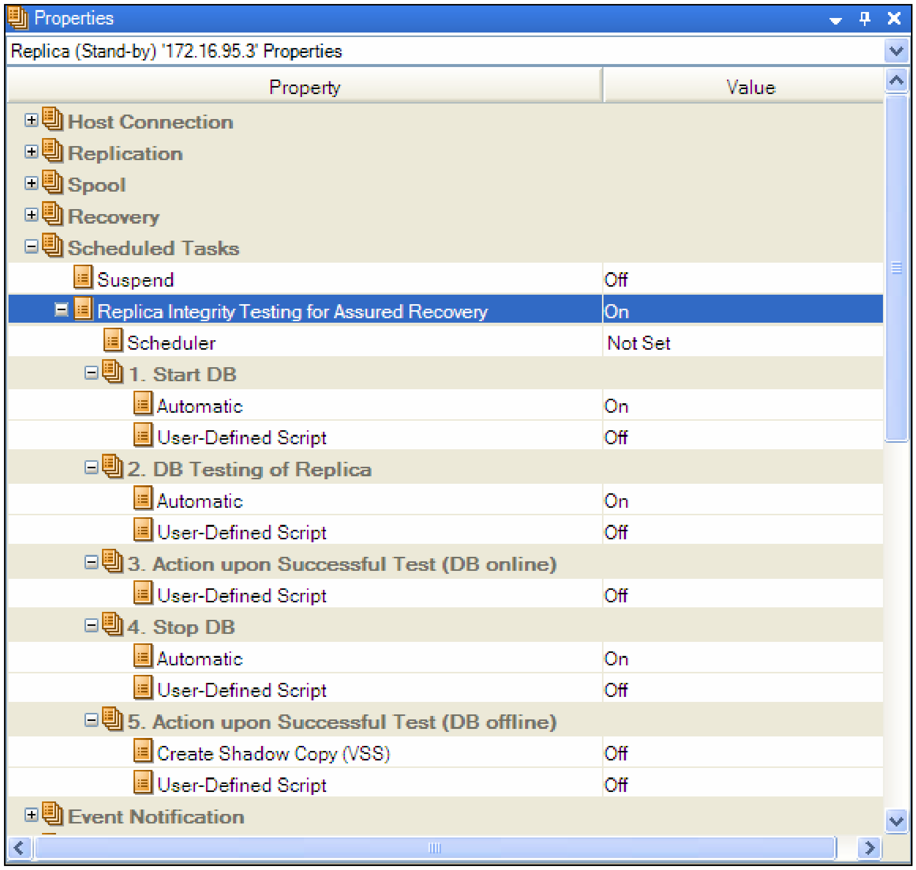

5. From the list, select the required property, and select or enter the appropriate values. Some values can be selected from a combo box while other values can be manually entered in an edit box field.

6. Click the **Save** button on the toolbar to save and apply your changes.

Understanding Assured Recovery Properties

This section lists the Assured Recovery properties, corresponding values, and provides an explanation of each property.

Scheduler

The Scheduler enables you to automatically run AR test according to a pre-defined schedule, say every few hours, once a day, or several times a month. To set the scheduler, see *Performing Assured Recovery Test in a Scheduled* Mode, page 204.

Start DB

This property defines the first step in the AR test: starting the database services on the Replica.

Automatic

By default, this property is set to On. To use script to replace the automatic initiation of database services, set this option to Off.

User-Defined Script

You can specify a script to augment or replace the standard step to start the database services.

To replace the standard step, set **Automatic** to Off and set **User-Defined Script** to On. Then, specify the full pathname of the script to be executed in the **Script Name** box.

To execute the script following the standard step, leave **Automatic** set to On.

Script Name (full path)

Enter the name and full path of the script that is invoked following the starting of database services or instead of it.

Arguments

Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

DB Testing of Replica

This property defines the second step in the AR test: verifying that all application services have started properly and that all databases or information stores have mounted successfully and are in a valid state.

Automatic

By default, this property is set to On. To use script to replace the automatic actions performed during this database validation stage, set this option to Off.

User-Defined Script

You can specify a script to augment or replace the actions performed during this database validation stage.

To replace the standard step, set **Automatic** to Off and set **User-Defined Script** to On. Then, specify the full pathname of the script to be executed in the **Script Name** box.

To execute the script following the standard step, leave **Automatic** set to On.

Script Name (full path)

Enter the name and full path of the script that is invoked following the database validation step or instead of it.

Arguments

Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

Actions on Successful Test (DB Online)

After the Replica is successfully tested, the application data is in a known, valid state. You may want to make use of this fact, for example, to ensure that a backup is performed at this point on validated data. If the action you want to perform requires that the application is running and the databases or information stores mounted, then it should be registered through a script here, in this step, by specifying the script details in the **User-Defined Script** boxes. This section has no default actions.

User-Defined Script

Script Name (full path)

Enter the name and full path of the script that is invoked when the application is still running and the databases or information stores are mounted.

Arguments

Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

Stop DB

This property defines the third and final step in a standard AR test: stopping the database services once the testing is complete.

Automatic

By default, this property is set to On. To use script to replace the automatic stopping of database services, set this option to Off.

User-Defined Script

You can specify a script to augment or replace the standard step to stop the database services.

To replace the standard step, set **Automatic** to Off and set **User-Defined Script** to On. Then, specify the full pathname of the script to be executed in the **Script Name** box.

To execute the script following the standard step, leave **Automatic** set to On.

Script Name (full path)

Enter the name and full path of the script that is invoked following the stopping of database services or instead of it.

Arguments

Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

Actions on Successful Test (DB Offline)

As noted in **Actions on Successful Test (DB Online)** above , at this stage the application data is in a known, valid state and you may want to copy it or perform a backup or a snapshot. If the action does not require the application to be running, it should be registered through a script here, in this step, by specifying the full pathname of a script in the **User-Defined Script** section.

Note: On Windows Server 2003 and up, you can generate VSS snapshots automatically (see *Configuring VSS Snapshot Creation*, page 218).

User-Defined Script

Script Name (full path)

Enter the name and full path of the script that is invoked after the AR test is completed successfully.

Arguments

Additional arguments to pass to the script, which is specified in the previous property. Arguments entered here are static values.

Using VSS Snapshots

CA XOsoft enables you to easily use Microsoft's Volume Shadow Copy Service (VSS) to create, view and manage VSS snapshots of the Replica data.

Important! You can use VSS only on Windows Server 2003 and up (not on earlier versions).

You can set up automatic creation of VSS snapshots in association with two operations: during replication suspension and after the Assured Recovery test is completed successfully. In addition, when CA XOsoft is integrated with ARCserve, a VSS snapshot is automatically created upon each ARCserve backup. All these snapshots are displayed in CA XOsoft Snapshots Management window, which allows you to monitor and manage them.

Configuring VSS Snapshot Creation

By default, CA XOsoft does not automatically create VSS snapshots. In order to activate this option, you need to set to On the **Create Shadow Copy (VSS)** property of the required Replica. This property is associated with two operations – replication suspension and AR test. Since you cannot set both operations on a scheduled mode for the same Replica, you need to configure the **Create Shadow Copy (VSS)** property with regards to one of these operations.

Setting Up Snapshot Creation

To set up a snapshot creation

1. On the Scenario pane, select the Replica for which you want to create VSS snapshots.
2. On the Framework pane on the left, select the Properties tab.

   The Replica Properties list appears:

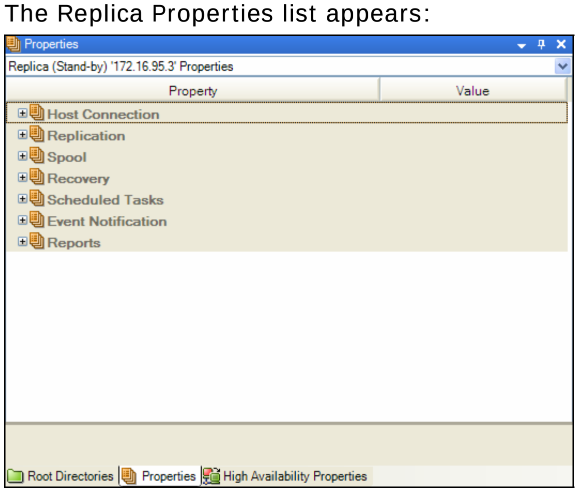

3. If the scenario is running, click the **Stop**  button on the toolbar. The scenario is stopped.

4. On the Replica Properties list, open the **Scheduled Tasks** group to display the **Suspend** and the **Replica Integrity testing for Assured Recovery** properties:

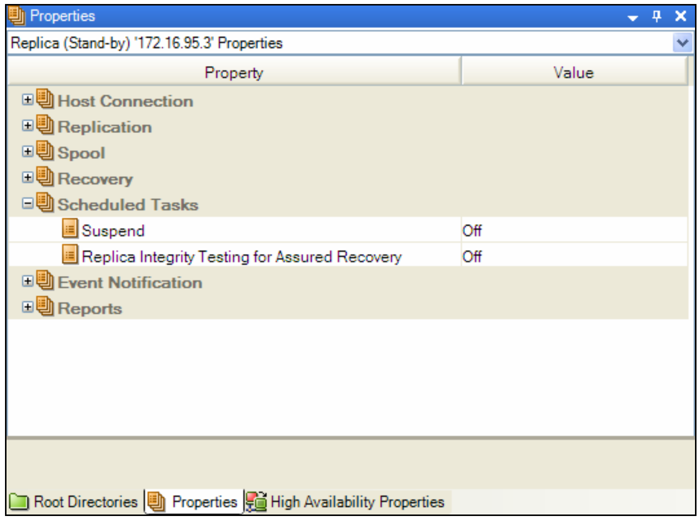

5. On either the **Suspend** or the **Replica Integrity testing for Assured Recovery** property, set the value to On.

   The **Create Shadow Copy (VSS)** property appears along with its related properties:

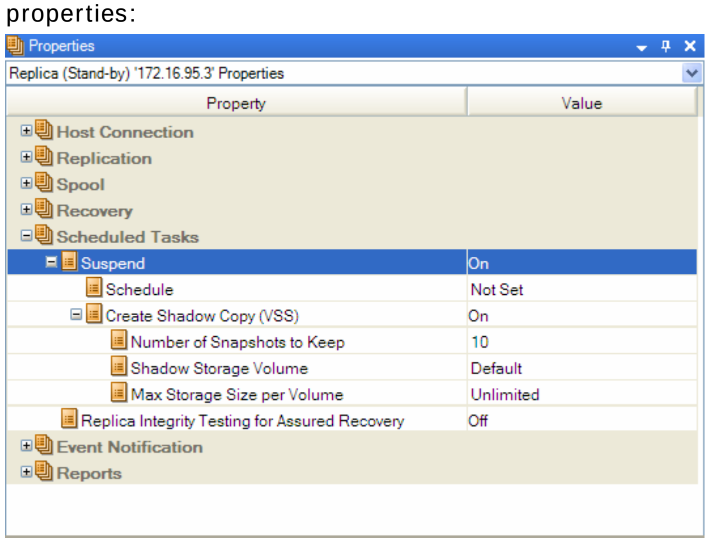

   Note: If you set to On the **Replica Integrity testing for Assured Recovery** property, the **Create Shadow Copy (VSS)** property appears under the **Actions on Successful Test (DB Offline)** group.

6. To activate the automatic creation of snapshots, Set the **Create Shadow Copy (VSS)** property value to On.

7. Set the other VSS properties, according to the information provided in *Understanding VSS Snapshot Properties*, page 220.

8. Click the **Save** button on the Standard toolbar to save and apply your changes, and start the scenario.

Now, after an AR test or during suspension, a VSS snapshot will be created automatically. The creation of the snapshot is indicated in the Event pane, as the following example shows:

Events

ID	Host/Scenario	Event
SR00104	**172.16.95.3**	**Replication to replica 172.16.95.3 resumed**
SR00392	**172.16.95.3**	**Exchange Integrity Testing on replica 172.16.95.3 is finished**
IR00199	172.16.95.3	Shadow Copy Id for volume C:\ is {27565e61-2635-4df2-89df-3a94a01ccb48}
IR00197	172.16.95.3	Shadow Copy is built successfully
IR00175	172.16.95.3	Building Shadow Copy
IR00312	172.16.95.3	Exchange services stopped
IM00405	Exchange AR	Posting AssuredRecovery report created at '06/29/08 17:02:03' to Reports
IR00308	172.16.95.3	Stopping Exchange services
SR00389	**172.16.95.3**	**Automatic Exchange testing on replica 172.16.95.3 is successful**

Once a snapshot is created, you can view and manage it through the Snapshots Management window, as described on page 221.

Understanding VSS Snapshot Properties

This section lists the VSS Snapshot properties, corresponding values, and provides an explanation of each property.

Create Shadow Copy (VSS)

To create VSS snapshots automatically during replication suspension or after successful AR test, set this option to On.

Preferred Number of Snapshots to Keep

Enter the number of snapshots you prefer to save and monitor. Once this number is reached, the oldest snapshots are replaced with newer ones. However, if the oldest snapshot is mounted or locked for backup, it is not deleted. Then, the new snapshot is added to the snapshot list even if the number is exceeded. The default no. is 10 snapshots.

Universal Shadow Storage Volume

Specify the volume on which the snapshots will be stored. Note that this property cannot be set for each scenario separately. The storage location of the first VSS snapshot that is created in the system, applies to all other succeeding snapshots.

Max Storage Size per Volume

Enter the maximum storage allowed per volume used by snapshots (MB).

Viewing and Managing Snapshots

CA XOsoft provides you with a special window for managing your VSS snapshots.

Viewing Snapshots

To open the Snapshots Management window

- On the Manager, click the **Snapshot View** 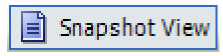

button on the Viewing toolbar.

  The **Snapshots Management** window appears:

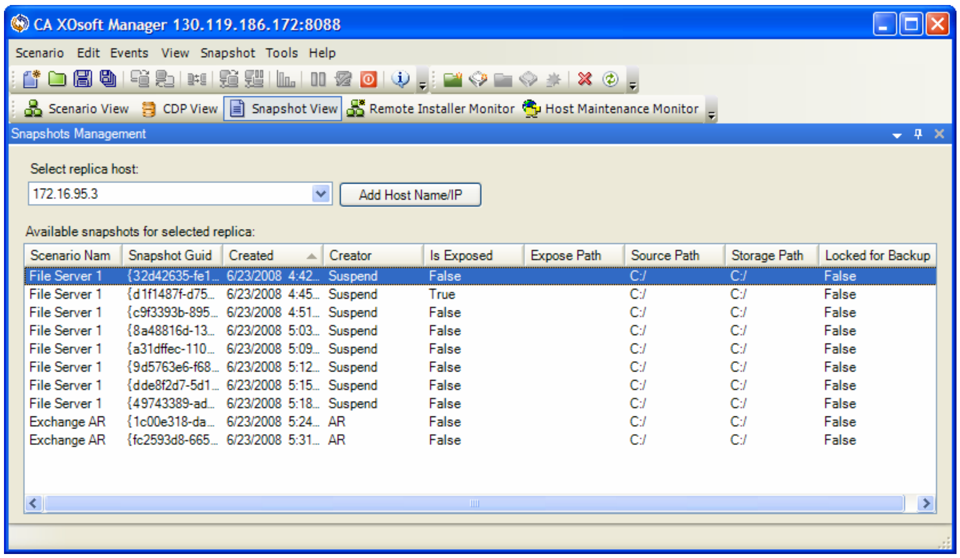

In this window, the VSS snapshots that were created for each existing Replica are displayed, according to the selected Replica.

You can change the Replica whose snapshots are displayed by using the **Select replica host** drop-down list. The Replica hosts that appear on the list are all the Replica hosts that participate in existing scenarios.

If a Replica that had snapshots participated in a scenario that was removed from the Manager, it does not appear on the list. To display snapshots of a Replica that no longer appear on the list, you can add it manually by using the **Add Host Name/IP** 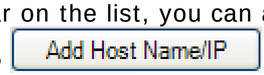

button.

The following information is provided for each snapshot:

- **Scenario name** – the scenario in which the snapshot was created.
- **Snapshot Guid** – the unique ID that identifies the snapshot.
- **Created** – the date and time of the snapshot creation.

- **Creator** – the type of operation that is associated with the creation of the snapshot. Two types are available: Suspend and AR.
- **Is Exposed** – indicates whether the snapshot was exposed ("True") or not ("False").
- **Expose Path** – where the snapshot was exposed.
- **Source Path** – the volume/directory that the snapshot captured.
- **Storage Path** – where the snapshot was stored.
- **Locked for Backup** – this column refers to snapshots that were taken as a part of Arcserve backup. If the backup is not complete yet, you cannot manage the snapshot, and the value that appears is "True". If the backup is complete, or if the snapshot is not associated with ARCserve, the value is "False".

Once the snapshots are displayed, you can start managing them.

Managing Snapshots

To manage snapshots

- On the **Snapshots Management** window, select the snapshot you want to manage. Then, open the **Snapshot** menu and select the required option, or right-click and select the required option from the pop-up menu:

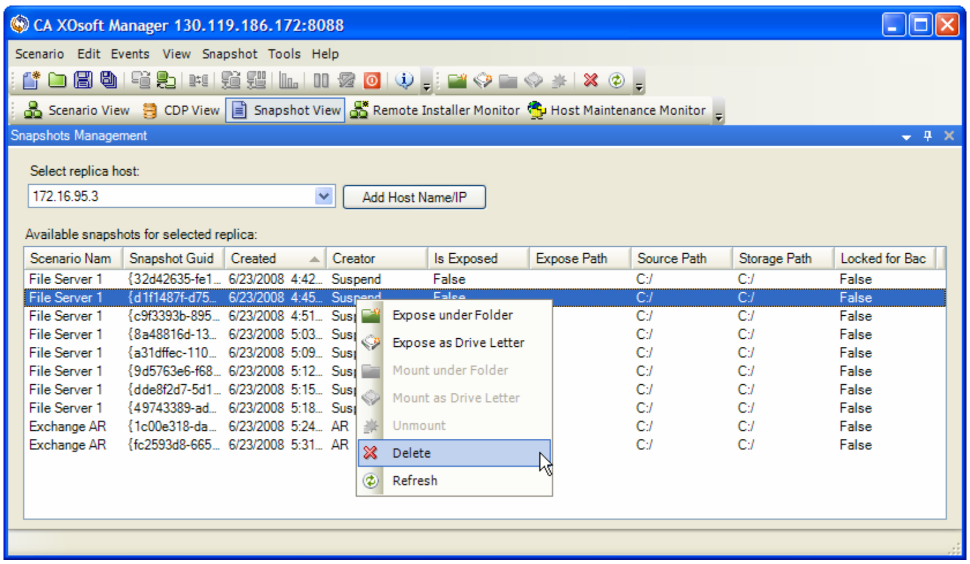

The available actions are:

- **Expose under Folder** – expose a snapshot as a local read-only folder by mounting it on an unused folder.
- **Expose as Drive Letter** - expose a snapshot as a local read-only volume by mounting it on an unused drive letter.

Notes:

- An exposed snapshot remains exposed through subsequent boots. Dismounting an exposed snapshot releases it without losing the snapshot itself.
- The **Expose** and **Mount** actions produce the same result - mounting a snapshot to a certain path. The difference between them is that when you want to mount a snapshot for the first time, you cannot use the **Mount** action directly and you need to use the **Expose** action. The **Expose** action both exposes and mounts the snapshot. Then, you can use the **Unmount** and **Moun**t actions.

- **Mount under Folder** – mount an exposed snapshot on an unused folder.
- **Mount as Drive Letter** – mount an exposed snapshot on an unused drive letter.
- **Unmount** - release an exposed snapshot without losing the snapshot itself. The snapshot is still exposed but it does not use a mount point.
- **Delete** – delete a snapshot. You can delete several snapshots at once by using the **Ctrl** key.
- **Refresh** – refresh the snapshot list to display the most up-to-date snapshots.

Chapter 13: Using the CDP Repository

This chapter provides instructions for creating, managing and using the CDP Repository module.

Understanding the CDP Repository

The CDP Repository module provides the ability to store deleted Outlook items, to search for certain items according to different criteria, and to retrieve them upon end-users requests. Thus, it helps the enterprise to better protect, manage and use its Exchange server environment.

The CDP Repository module uses CA XOsoft data replication and recovery capabilities, and introduces the ability to restore and retrieve a single or numerous deleted messages upon end-user requests, without administrative intervention. The CDP Repository uses the data that was replicated from the Master and stored on the Replica servers. Therefore, the bulk of the data processing is done outside the production servers, avoiding performance overload.

When scanning the replicated data, the CDP Repository use of indexing allows it to capture and process only changes that occur in this data. This means that when a request arrives, only deleted messages are scanned, and consequently the retrieval process is rapid and efficient. The deleted messages are kept in the CDP Repository even if the scenario that initially replicated them is deleted from the Manager.

There are two kinds of users who use the CDP Repository:

- **End-user**: A user who has an account in the enterprise domain and a mailbox in the replicated Exchange Store. Using the Web-based E-mail Retrieval, each end-user can directly access his or her own deleted messages, search and sort them, and retrieve the ones that are needed. The end-users have access only to their mailbox. They do not need to have any knowledge of the underlying system, only an access to the easy-to-use E-mail Retrieval.
- **Administrator**: A user who configures and manages the CDP Repository via the CA XOsoft Manager. The administrator configures through the Manager the CDP Database and the Exchange scenario properties that define when to run the data extraction operation and where to store the data extracted. The administrator is not expected to be a DBA in order to manage the CDP Repository.

The administrator is responsible for defining the types of Outlook items that can be retrieved. These can include: e-mail messages, appointments, contacts, tasks, journal entries, notes, and attachments. The administrator also set a retention policy that determines for how long the deleted messages will be kept, according to size and type, and the maximum disk space allocated for the deleted messages. Once this size is reached, new deleted messages are not inserted into the CDP database.

The administrator can configure all the relevant CDP settings. However, unless the administrator will explicitly logon as a DB owner to the SQL Server (not through the CA XOsoft Manager) he will not be able to view the content of the users' mailboxes. That way, an organization can set a stricter privacy policy, by which the CA XOsoft administrator has no direct access to the users' mailboxes.

Important! The CDP Repository module can be used with both DR and HA solutions. It is activated solely with Exchange scenario.

CDP Repository Components

The CA XOsoft CDP Repository consists of five components:

- **CDP Storage –** a storage area that resides in an instance of SQL Server 2005 and contains the entire deleted message data. The deleted messages can be stored in one or several databases. The SQL configuration is done through the CDP Admin, and multiple Exchange servers can use the same repository. Besides SQL Server 2005, this component does not need any additional installation.
- **CDP Web Server –** a component that receives end-user requests regarding deleted messages, passes queries on to the CDP Storage, receives from it the requested information, and passes it back to the user via the E-mail Retrieval component.
- **CDP Support** – a component that supports the CDP Repository functions and activities. It extracts deleted messages from database files and feeds them to the SQL Server. This component is installed as an additional component during the Engine installation.
- **CDP Admin –** a User Interface that resides in the Manager, which enables administrators to configure and deploy the CDP Storage retention and quota policies. It is installed as part of the Manager installation.
- **E-mail Retrieval –** an end-user web-based GUI, which enables users to search for deleted Outlook items and retrieve them. It can be opened from any workstation with a Web browser and a connection to the CDP Web Server machine, without additional installation.

The following illustration shows a typical deployment of CDP Repository components:

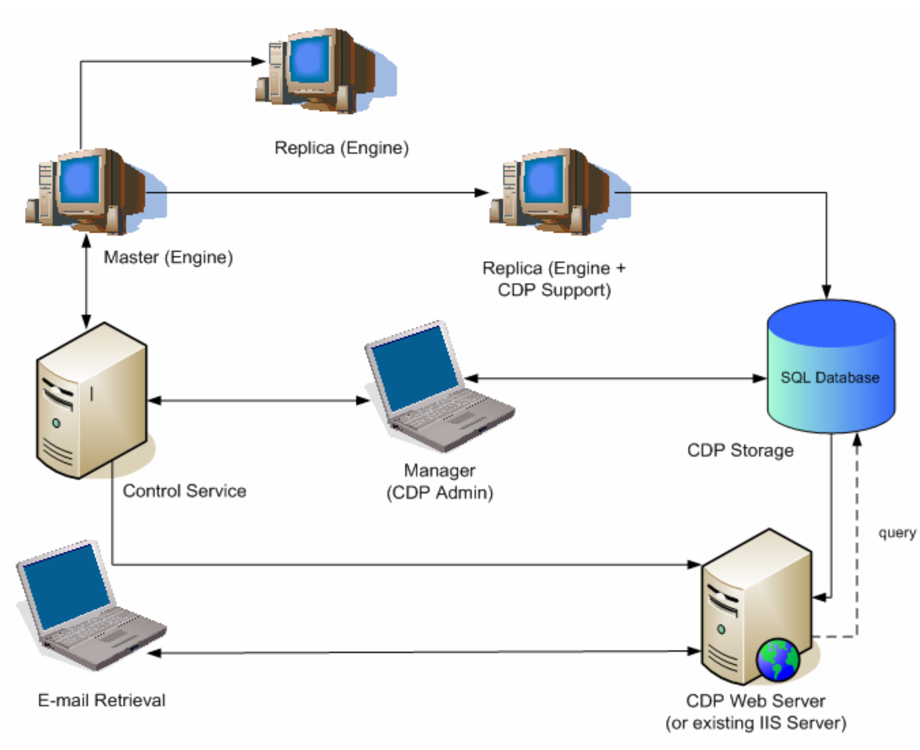

Configuring the CDP Repository

There are three steps in configuring the CDP Repository:

1. Defining the CDP Database.
1. Creating an Exchange scenario with the CDP option.
2. Running the scenario.

The CDP Database configuration needs to be done before the scenario creation.

Defining the CDP Database

The CDP Database configuration consists of three steps:

1. Selecting the SQL instance that will contain the new CDP repository.
2. Defining the name and the path of the repository.
3. Setting the retention and quota policy for the entire database and for each message type folder.

To define the CDP database

1. On the Manager, open the CDP View by selecting the **CDP View** 

option button on the Viewing toolbar. Alternatively, select from the **View** menu the **Active View** - **CDP View** option.

   The **CDP View** appears:

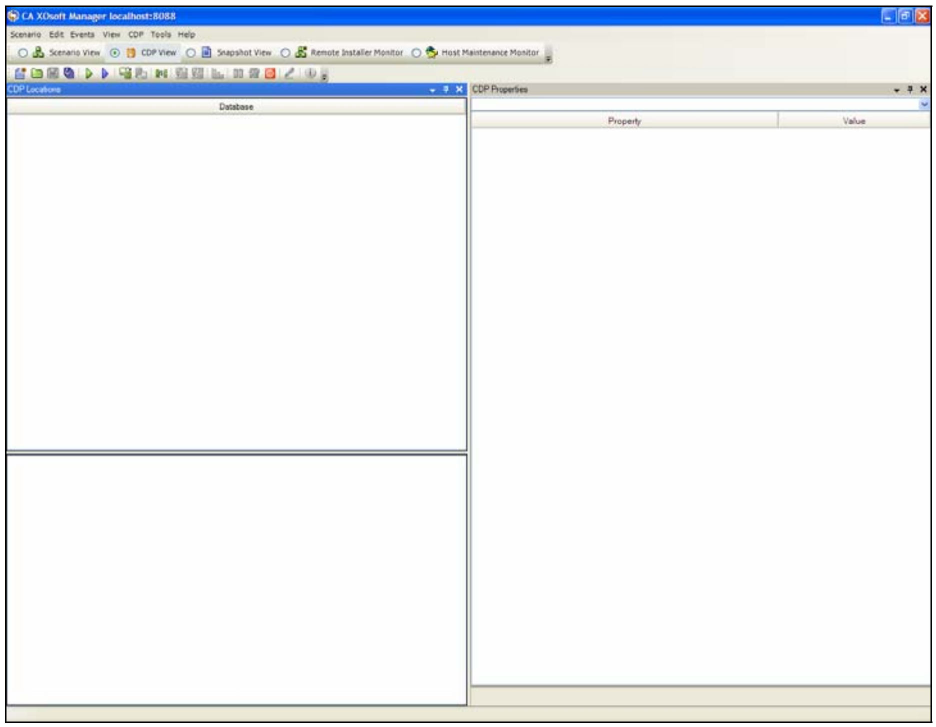

2. To create a CDP Database, select from the **CDP** menu the **Add Database** option.

The **Select SQL Instance** dialog appears:

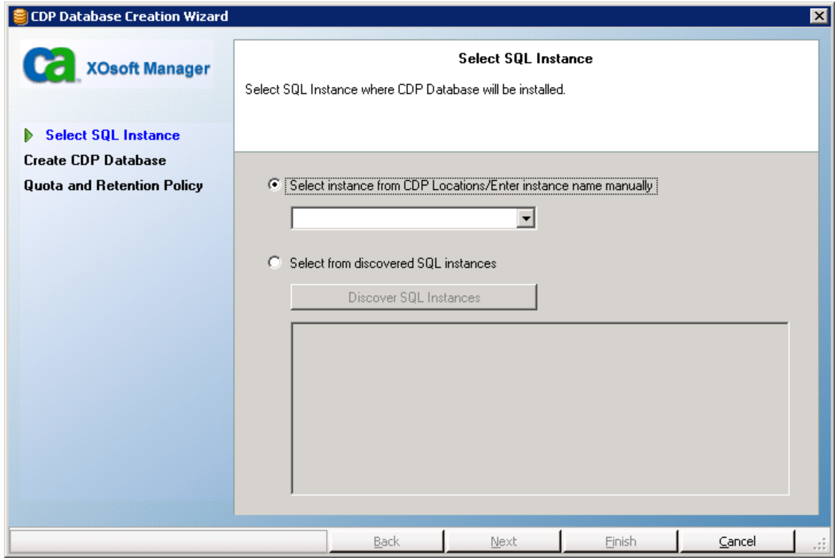

In this page, you select the SQL instance where the CDP database will be created.

3. When creating a CDP database for the first time, do one of the following:

- Manually enter the SQL instance name - select the first option button, and enter the name in the empty box.
- Select the instance name from a list of existing instances - select the second option button and click the **Discover SQL Instances** button. CA XOsoft automatically discovers the existing SQL instances, and displays them:

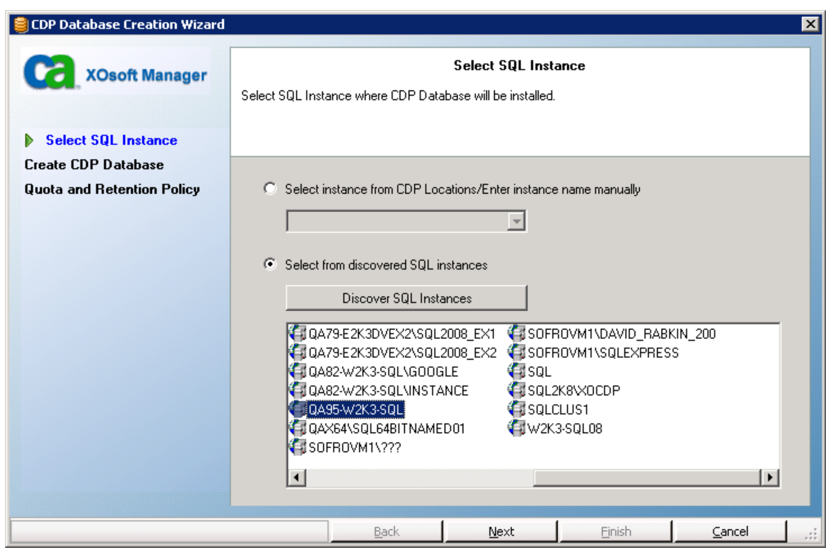

Select from the list the required SQL instance.

4. After defining the SQL instance, click **Next**. The **Import or create CDP Database** page appears:

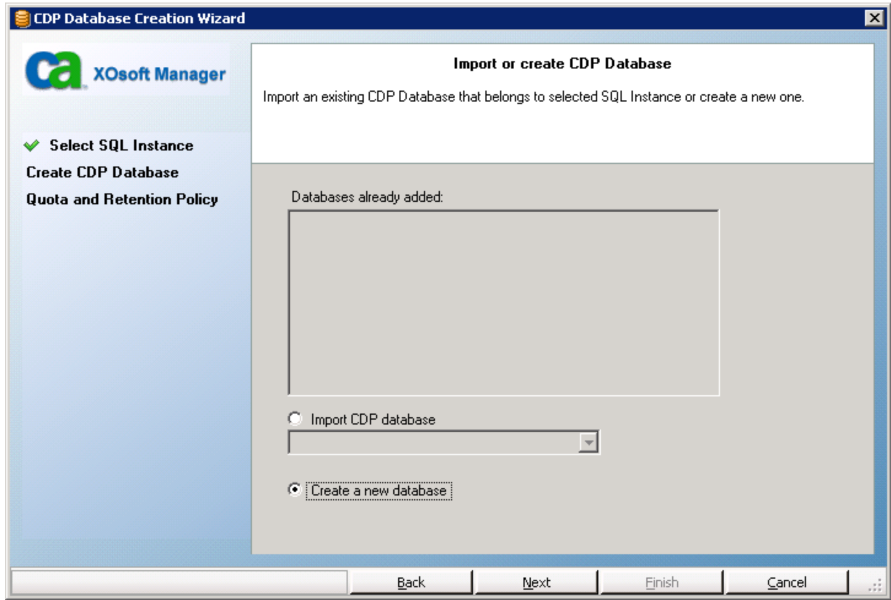

To create a new CDP database, select the **Create a new database** option button. If you already created a CDP database and you want to re-use it, select the **Import CDP database** option button and import the database.

5. Click **Next**. The **Create CDP Database** page appears:

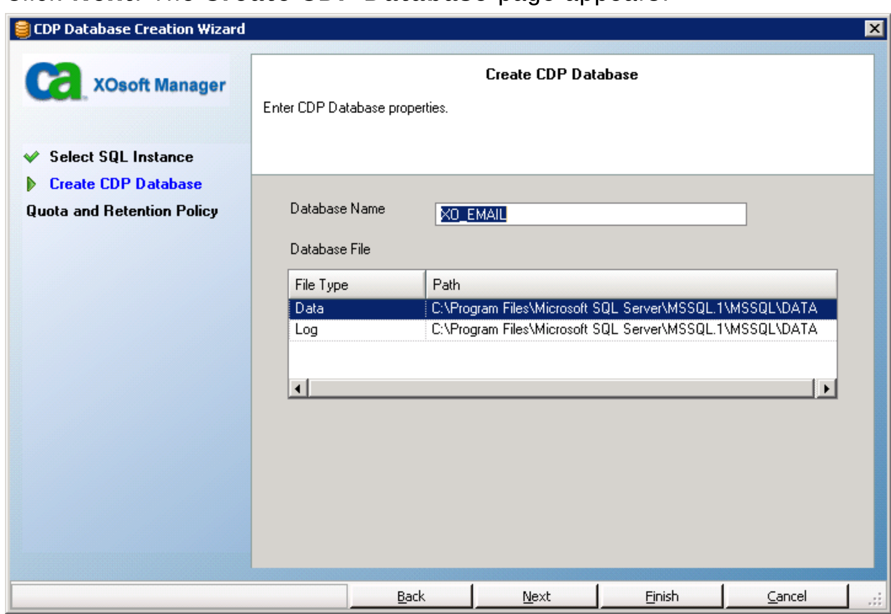

In this page, you can see the default CDP Database name and its storing path.

Note: If you click **Cancel** at this stage, the following message appears:

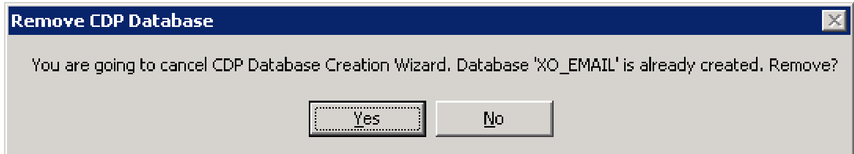

Since the CDP Database is already created, you have the option of keeping it and configuring it at a later stage, or removing it from the SQL Instance. Select the required option. If you click **Finish**, the new CDP Database will be kept automatically.

6. In the **Create CDP Database** page, keep the default database values or change them. Then, click **Next**.

   The **Quota and Aging** page appears:

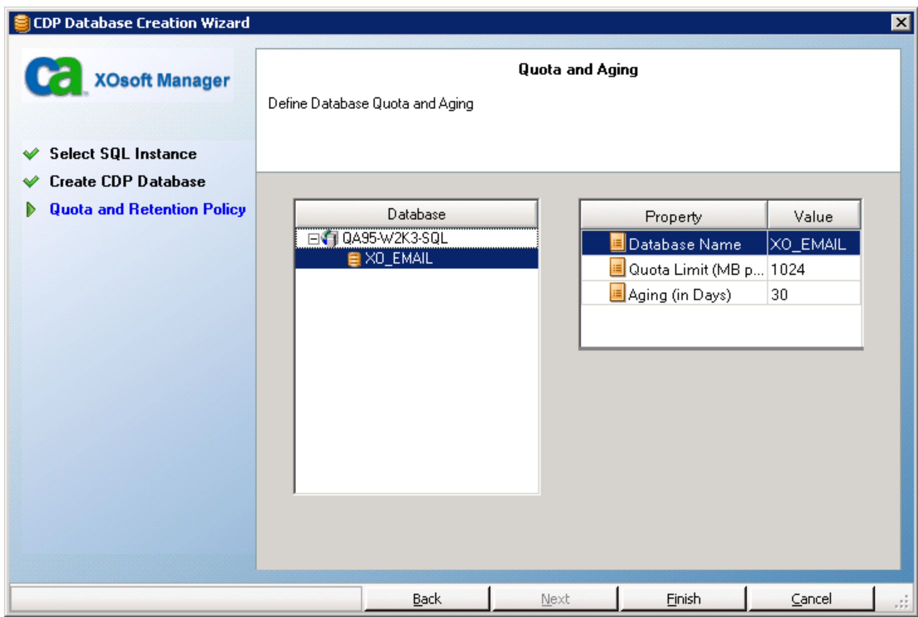

   In this page, you define for the entire database its quota limit, meaning the maximum disk space allocated for the deleted messages. Once this size is reached, new messages will not be inserted into the database. You can also define an aging period for the database, meaning for how many days the deleted messages will be kept in the CDP database.

7. To change the default quota limit of the entire database, from the **Property** area on the right, select the **Quota Limit** property, and double-click the value field to enter a new size in MB.

8. To change the default aging period of the entire database, from the **Property** area on the right, select the **Aging** property, and double-click the value field to enter a new number in days.

9. After you defined the quota and aging policy for the entire CDP database, select **Finish**.

The CDP database is created and configured. It appears on the CDP View window:

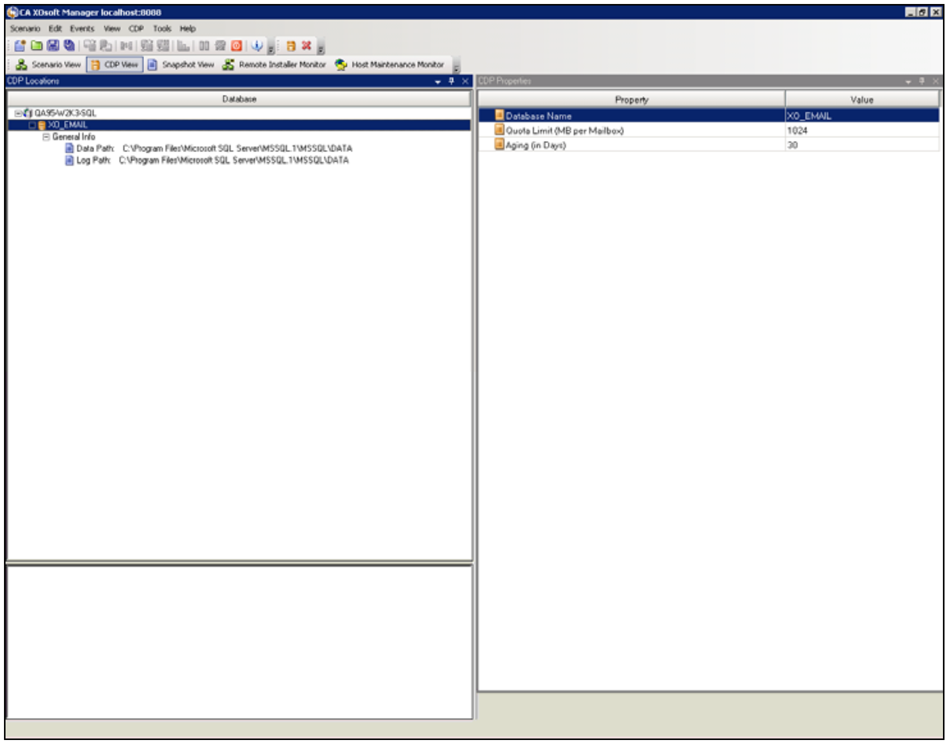

10. You can change the database configuration from here as well.

    To configure the CDP database from the CDP View, select the database on the Database pane on the left, and change its property values on the Properties pane on the right.

Creating an Exchange Scenario with the CDP Option

To activate the CDP Repository module, you need to create an Exchange scenario with the CDP option enabled.

To create an Exchange scenario with CDP

1. On the Manager, click the **New** button on the Standard toolbar.

   The **Scenario Creation Wizard** appears:

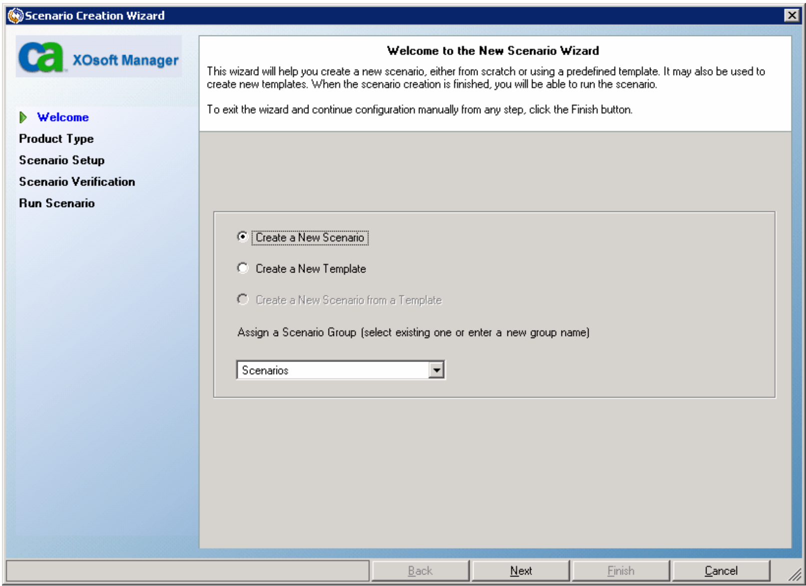

2. Select the **Create a New Scenario** option button, and click **Next**.

   The **Select Server and Product Type** page appears:

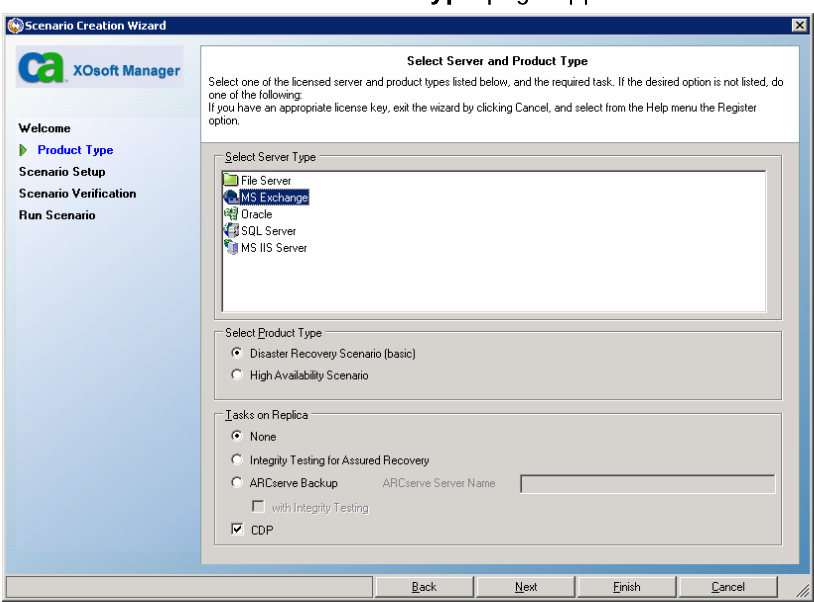

3. Select the required scenario options, as follows:
   - From the S**elect Server Type** list, select **MS Exchange**.
   - From the **Select Product Type** options, select either Disaster Recovery or High Availability Scenario.
   - From the **Tasks on Replica** options, select **CDP**.
4. Click **Next**. The **Master and Replica Hosts** page appears:

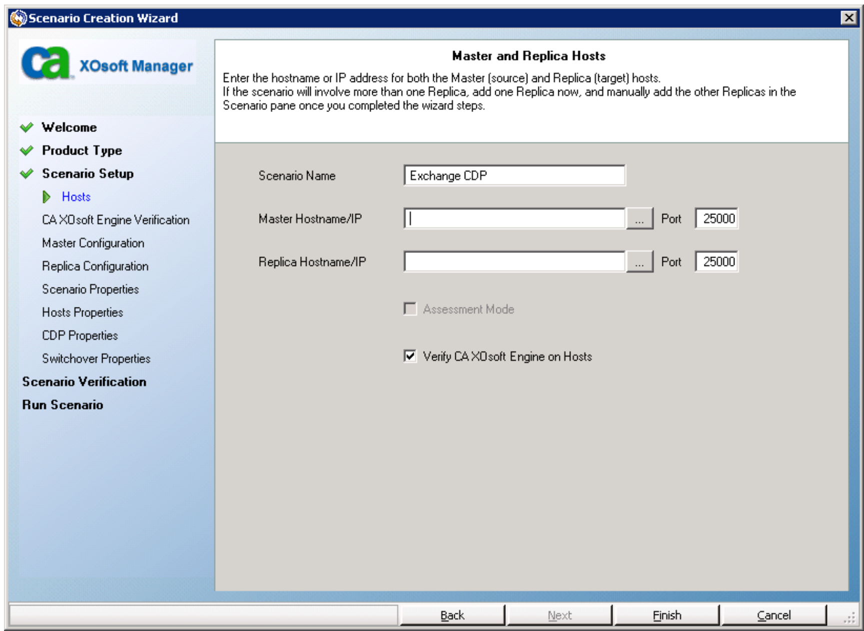

5. Enter the required information, as follows:
   - **Scenario name** – accept the default name or enter a new name for the scenario.
   - **Master Hostname/IP** and **Replica Hostname/IP** - enter the name or IP of the Master and Replica hosts, or use the **Browse** button to find it.

   Note: If either server is a MSCS cluster, enter the Exchange Virtual Server Name or IP address as the Master and/or Replica name (instead of the physical node's name or IP).

   - In the **Port** boxes: accept the default port no. (25000) or enter a new port numbers for the Master and Replica.
   - **Assessment Mode** – make sure this check box is NOT selected.
   - **Verify CA XOsoft Engine on Host** – select this option if you want the system to verify whether Engines are installed and running on the Master and Replica hosts you specified in this page. If you select this option, the **Hosts Verification** page appears after you click **Next**. For more information about host verification, see page 45.

6. Click **Next.** The **Master Configuration** page appears:

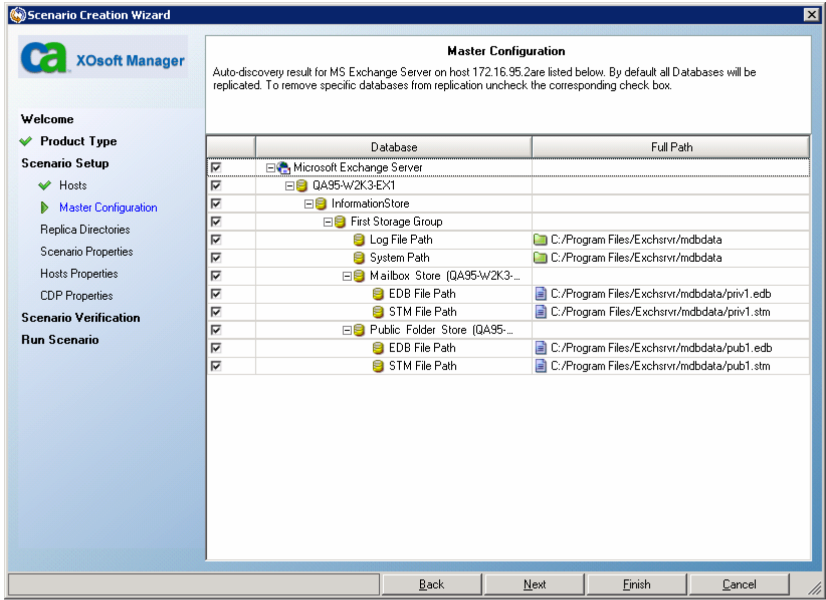

CA XOsoft auto-discovery component automatically discovers all Exchange databases on your Master server. These are the databases that can be replicated and protected.

7. By default, all the discovered databases are selected and all will be replicated. You can exclude some of these storage groups from replication by clearing their check boxes.

8. After defining the data to be replicated, click **Next**.

   The **Replica Root Directories** page appears:

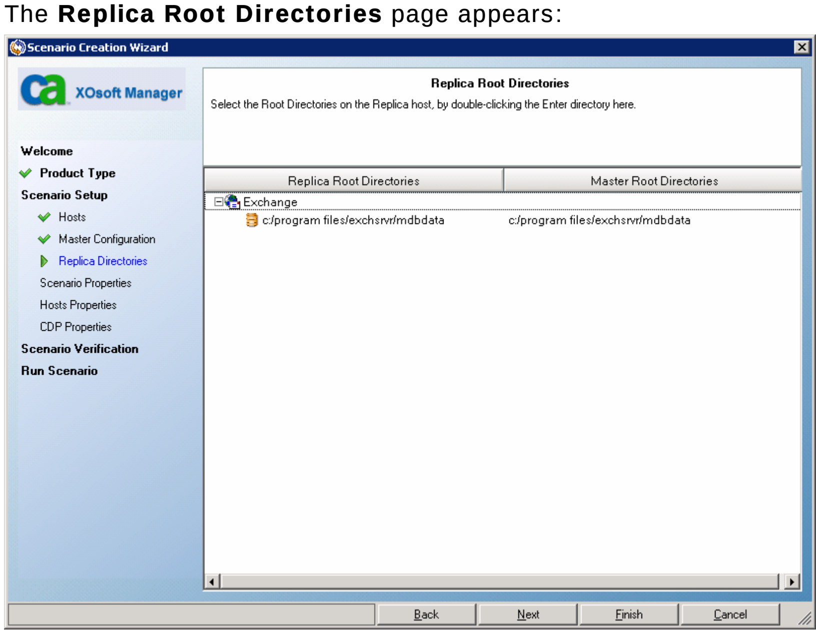

CA XOsoft auto-configuration component verifies that the Exchange Server configuration on the Master and Replica servers will be identical during the replication procedure.

9. After defining the storage location of the replicated data, click **Next**.

   The **Scenario Properties** page appears:

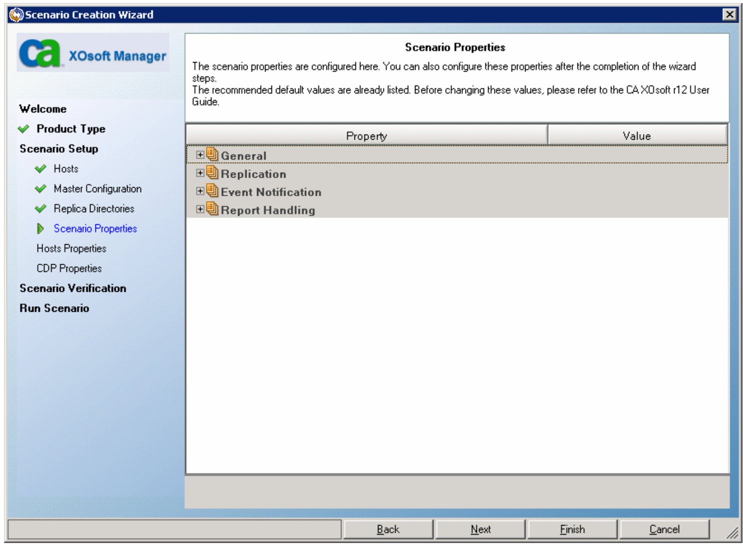

   The **Scenario Properties** page enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

   If you want to configure the scenario properties at this stage, refer to *Understanding Scenario* Properties, page 123. To configure the scenario properties at a later stage, refer to *Configuring Scenario Properties*, page 121.

10. Click Next. The **Master and Replica Properties** page appears:

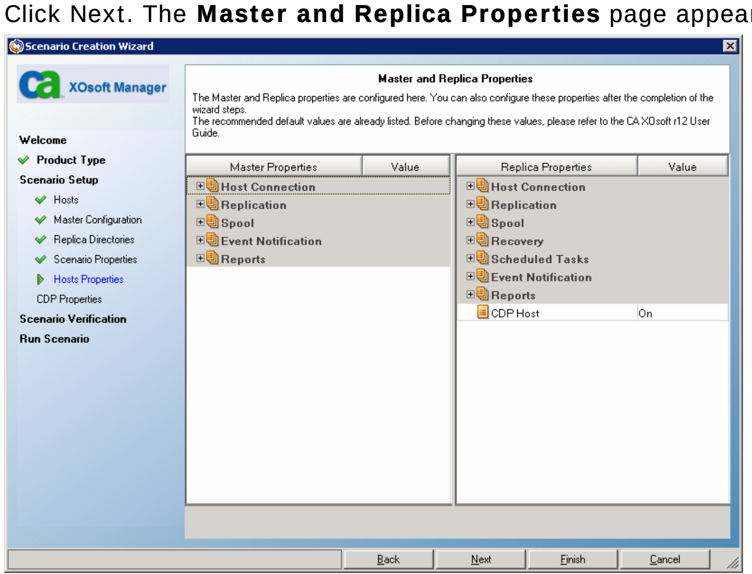

The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient. All you need to verify at this stage is that on the **Replica Properties** list, the **CDP Host** property is **On**.

If you want to configure the Master and Replica properties at this stage, refer to *Setting Master and Replica Properties*, page 133. To configure the scenario properties later, refer to *Configuring Master or Replica Server Properties*, page 133.

Note: You can modify all the settings in this pane after the scenario is created. However, before changing any Spool properties (which can be configured here), review the *Spool* info on page 136 for configuration details.

11. Once you are satisfied with the Master and Replica properties, click **Next**.

    The **CDP Properties** page appears:

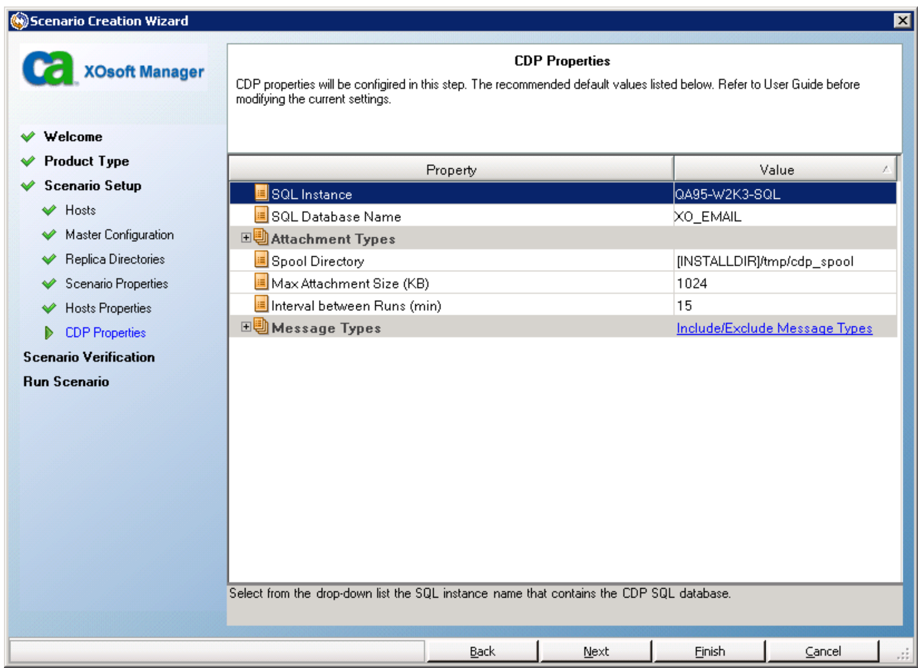

The **CDP Properties** page enables you to configure the properties that are related to the CDP Repository.

12. Set the CDP properties as follows:

    - **SQL Instance** – select from the drop-down list the SQL instance name that contains the CDP SQL database.
    - **SQL Database Name** – select from the drop-down list the SQL database name that you defined for the CDP Repository.

- **Attachment Types** – enter the type of attachments to be stored. By default, all attachment types are excluded from storage. To include an attachment type, enter the type in the value box, for example **doc**, and press the **Enter** key.

  Each included type should appear in a new row, as the following picture illustrates:

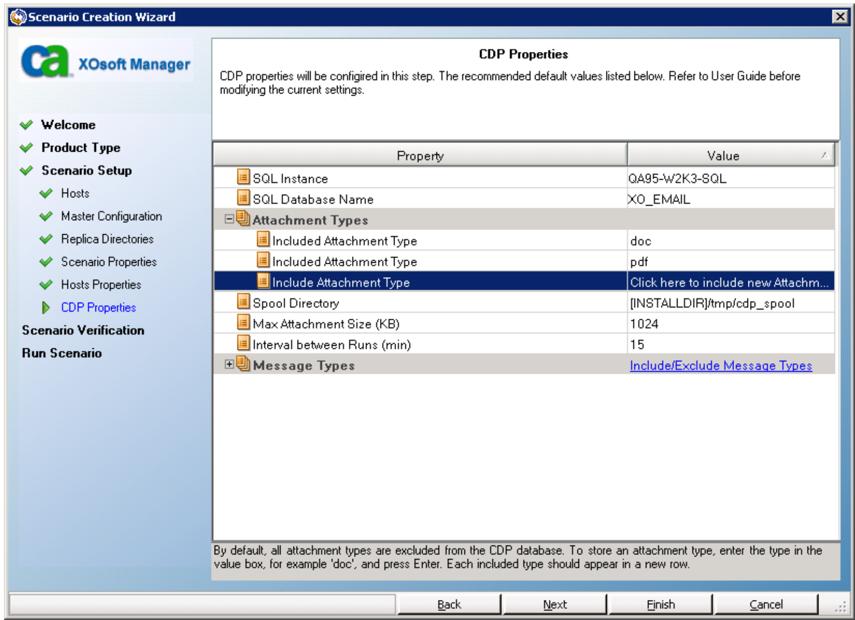

- **Spool Directory** – enter the directory to be used to store the spool for the CDP Repository on the Replica.
- **Interval between Runs (min)** – enter the number of minutes between each check the system will perform in the CDP Repository, in search for new deleted messages.
- **Max Attachment Size** – enter the attachment size allowed for storage. Deleted attachments that exceed this size will not be stored in the CDP Repository.

- **Message Types** - specify whether deleted message of this type will be stored in the CDP database. By default, all message types are stored. To exclude a type from being stored, click the **Include/Exclude Message Types** link. The **Message Types** dialog appears:

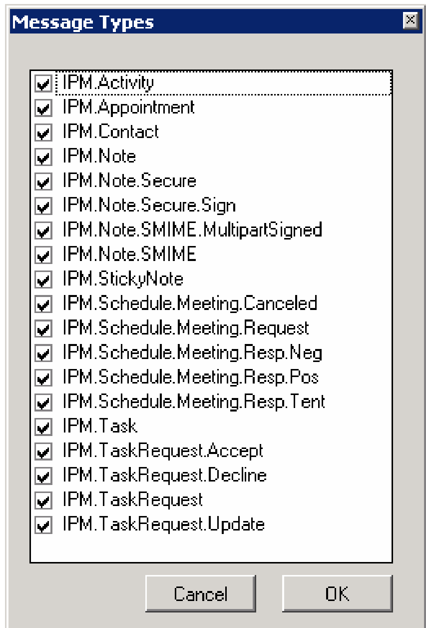

To exclude message types from storage, clear their check boxes.

13. Once you set the CDP Repository properties, click **Next**.

    CA XOsoft verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful replication and data recovery process. Once the verification is completed, the **Scenario Verification** page appears:

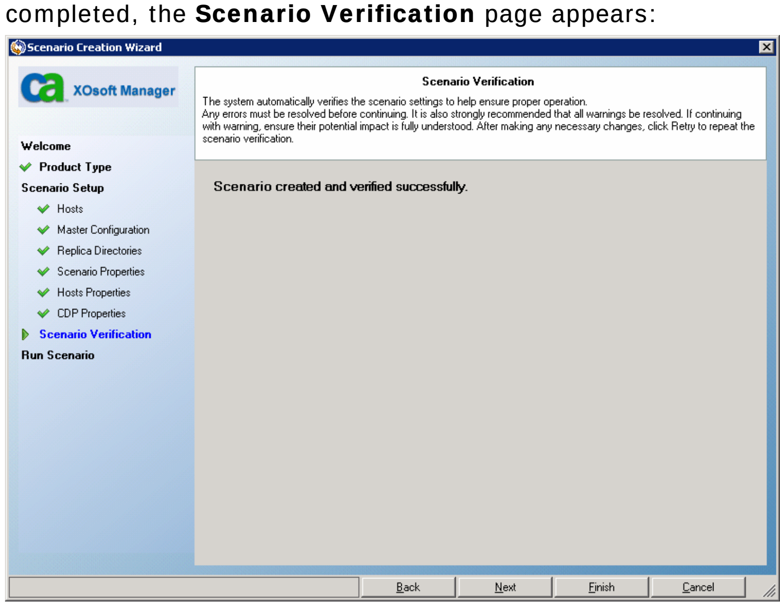

Note: Although the software allows you to continue with warnings, it is not recommended to do so. Resolve any warning situations before continuing to ensure proper operation of the application.

14. If the scenario is verified successfully, click **Next**.

    The **Scenario Run** page appears:

15. After the scenario is verified, you are prompted to run it. Running the scenario starts the data synchronization process.

    - To finish the scenario creation and run it later, select **Finish**.
    - To run the scenario now, click **Run Now**.

      The synchronization process starts, and the CDP Repository is starting to be filled with the enterprise deleted messages.

16. By default, once a synchronization occurs, a synchronization report is generated:

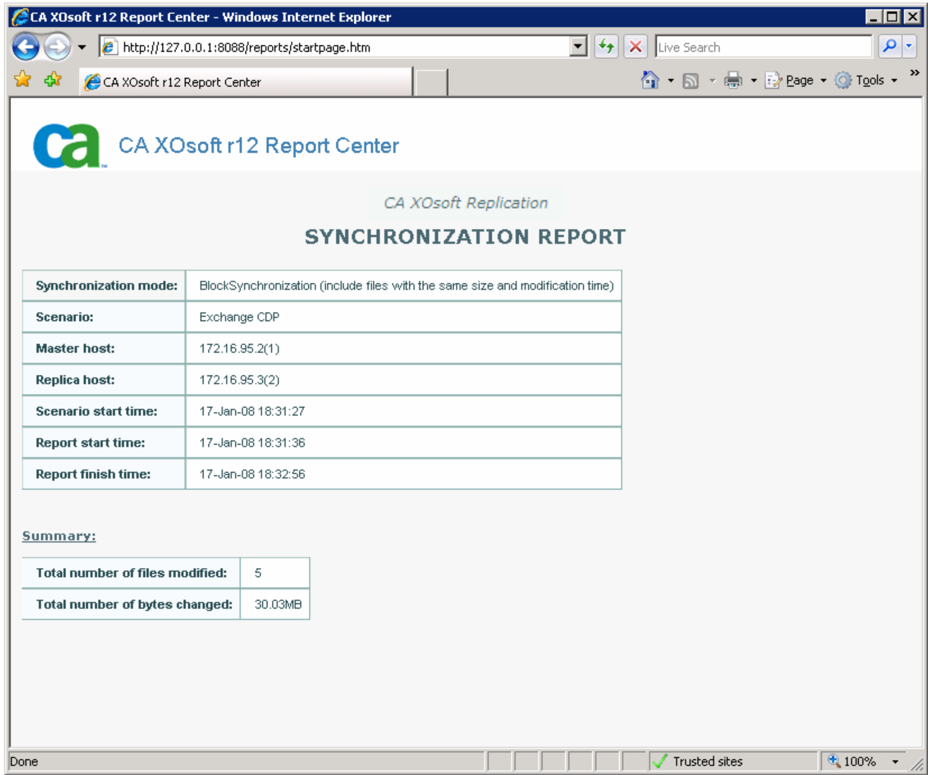

Note: For more information about opening a report, see page 63.

Setting CDP Scenario Properties

After you defined an Exchange scenario with the CDP option, you can set or modify its properties.

Note: To configure scenario properties, the scenario must be stopped.

To set CDP scenario properties

1. If you are not in the Scenario View, open it by selecting the **Scenario View** 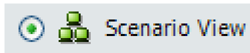

option button on the Viewing toolbar.
2. On the Scenario pane, select the Exchange scenario its CDP properties you want to configure. On the Framework pane on the left, the Scenario Properties list appears:

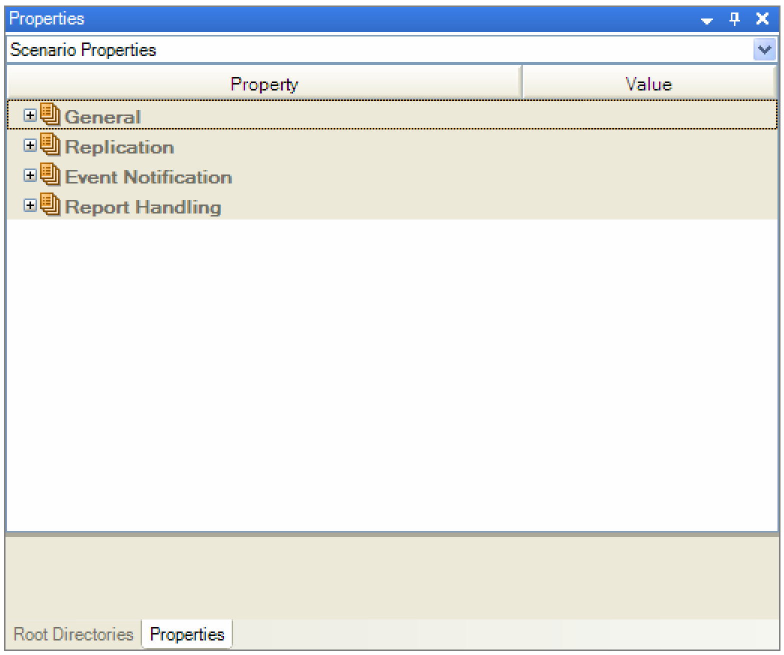

Note: A running scenario has a gray background, and scenarios that are not running have a white background.

3. From the **Properties** drop-down list on the Framework pane, select **CDP Properties**.

   The **CDP Properties** list appears:

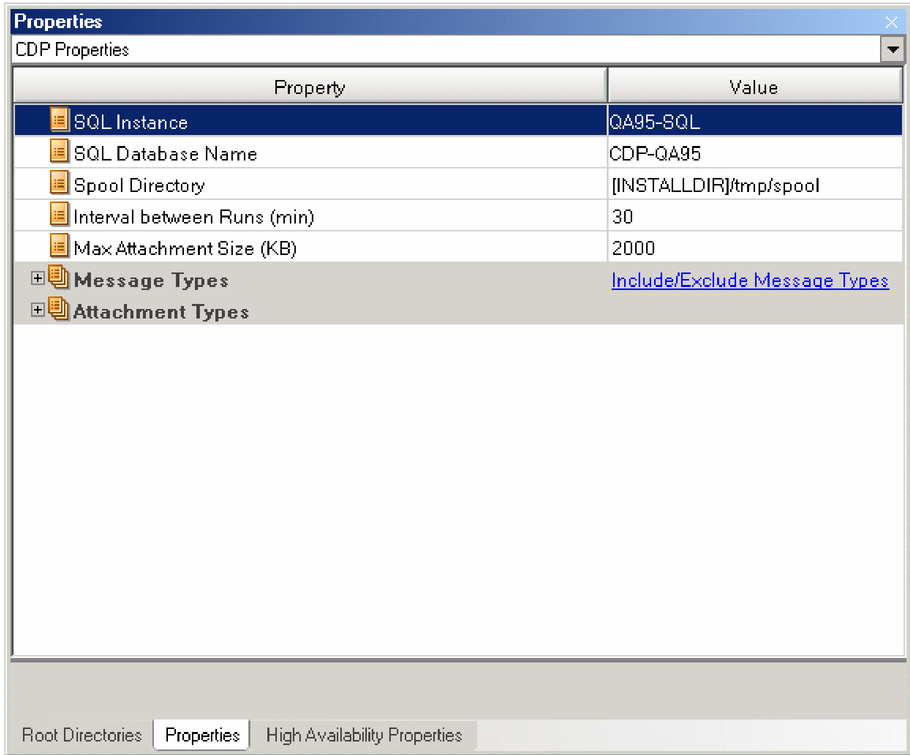

4. If the scenario is running, click the **Stop**  button on the toolbar. The scenario is stopped.

5. On the CDP Properties list, open the desired group, select the required property, and select or enter the appropriate values. Some values can be manually entered in an edit box field, while others can be selected from a down list.

   Note: For a description of each property, see page 237.

6. After you set the required properties, click the **Save**  button on the Standard toolbar to save and apply your changes.

Retrieving Deleted Outlook Items Using the E-mail Retrieval

The E-mail Retrieval enables you to search for deleted Outlook items and retrieve them. The types of Outlook items that can be retrieved are defined by the Administrator, and they can include: e-mail messages, appointments, contacts, tasks, journal entries, notes, and attachments.

Search for Deleted Items and Retrieve Them

To search and retrieve a deleted item:

1. To open the Email Retrieval window, select **Start**, **Programs**, **CA**, **XOsoft**, **Email Retrieva**l.

   Note: You can also enter the following default parameters on the **Address** box of your Internet Explore: (If you are opening the Management Center from the machine where the Control Service is installed): http://localhost:8088/start_page.aspx

   The Email Retrieval window appears.

2. On the Email Retrieval, enter any text you want to search for in one of the available fields: **Subject**, **From** or **Sent To**:

   Note: to search in the message body or attachments, enter the text in the **Subject** field and use check boxes that underneath it.

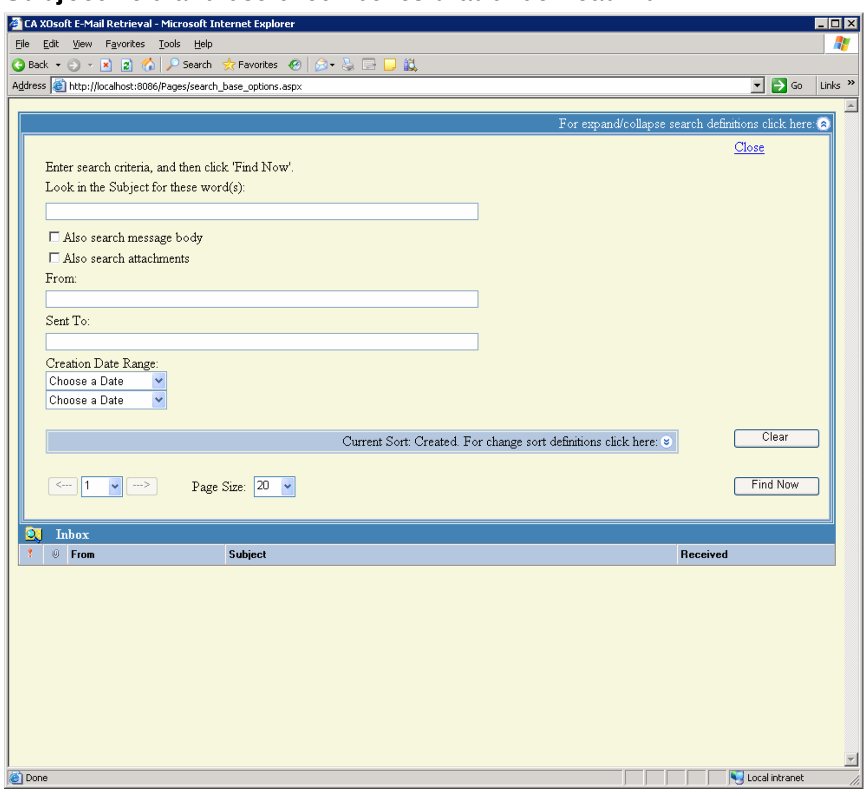

3. [Optional] Define the date range in which you want to conduct the search by using the **Creation Date Range** drop-down menus.

4. After you entered the criteria, click the **Find Now** button.

   The search results appear:

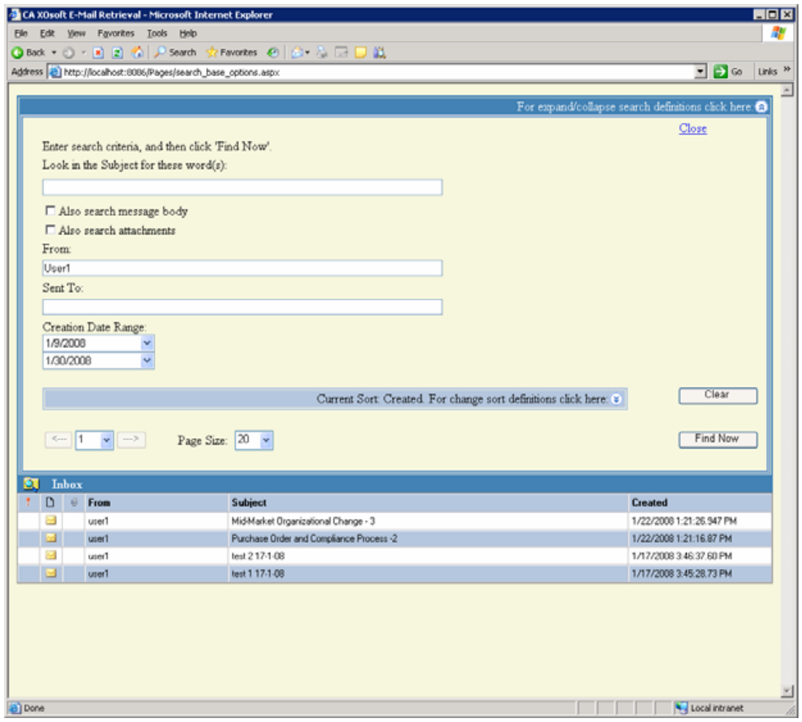

5. To view a deleted message, double-click it.

   The message opens:

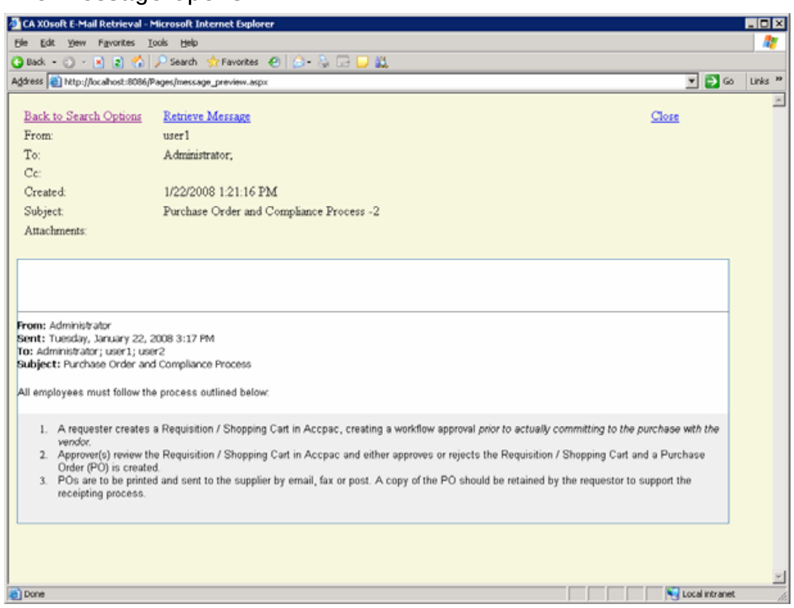

6. To retrieve the message, click the **Retrieve Message** link at the top of the page.

   A standard **File Download** dialog appears:

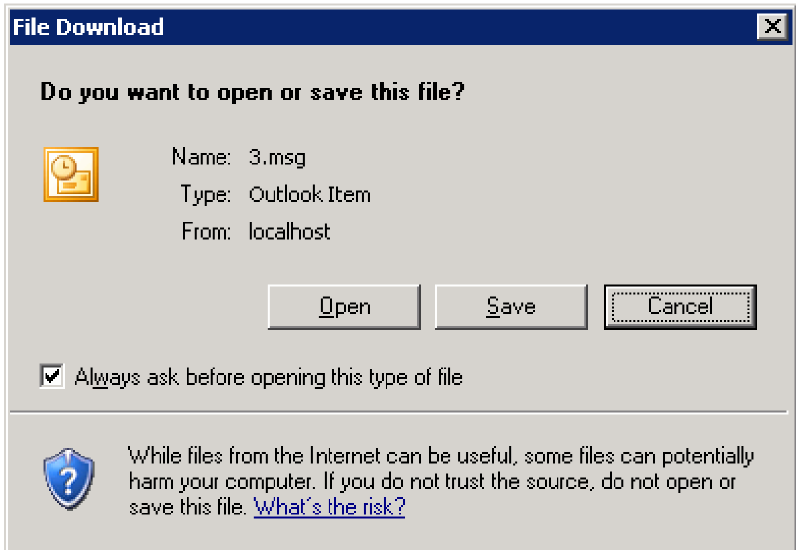

7. To open the message as an Outlook item, click **Open**:

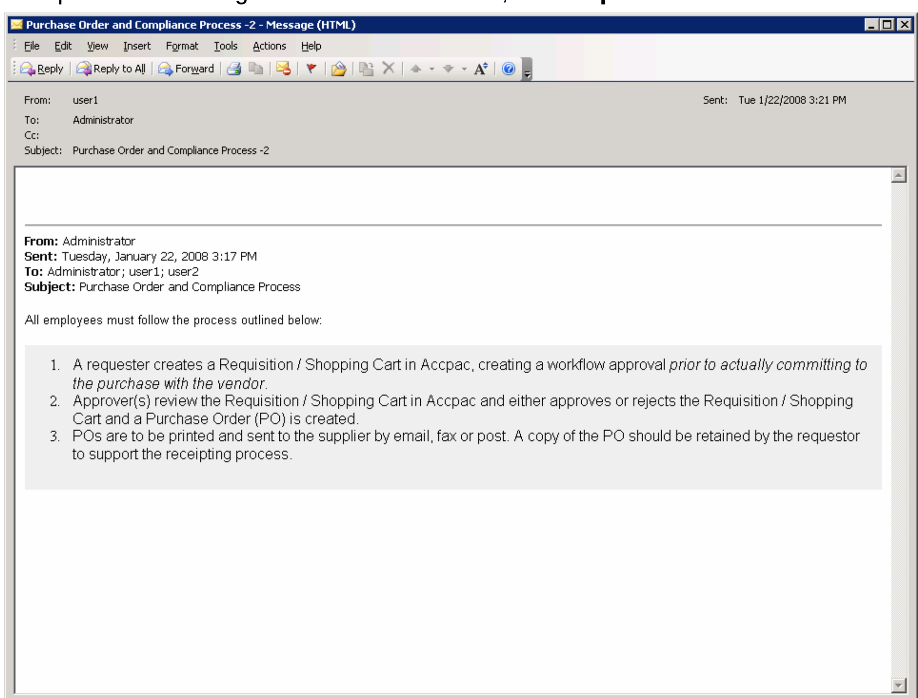

8. You can now save the message or perform any other required activity with it.

Index

A

B

C

F

G

H

I

K

L

M

N

O

P

Q

R

W