

CA XOssoft™

User Guide
r12.5



This documentation and any related computer software help programs (hereinafter referred to as the "Documentation") is for the end user's informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and protected by the copyright laws of the United States and international treaties.

Notwithstanding the foregoing, licensed users may print a reasonable number of copies of the Documentation for their own internal use, and may make one copy of the related software as reasonably required for back-up and disaster recovery purposes, provided that all CA copyright notices and legends are affixed to each reproduced copy. Only authorized employees, consultants, or agents of the user who are bound by the provisions of the license for the Product are permitted to have access to such copies.

The right to print copies of the Documentation and to make a copy of the related software is limited to the period during which the applicable license for the Product remains in full force and effect. Should the license terminate for any reason, it shall be the user's responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

EXCEPT AS OTHERWISE STATED IN THE APPLICABLE LICENSE AGREEMENT, TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO THE END USER OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED OF SUCH LOSS OR DAMAGE.

The use of any product referenced in the Documentation is governed by the end user's applicable license agreement.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Copyright © 2009 CA. All rights reserved.

Contact CA

Contact Technical Support

For your convenience, CA provides one site where you can access the information you need for your Home Office, Small Business, and Enterprise CA products. At <http://ca.com/support>, you can access the following:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Provide Feedback

If you have comments or questions about CA product documentation, you can send a message to techpubs@ca.com.

If you would like to provide feedback about CA product documentation, complete our short [customer survey](#), which is also available on the CA support website, found at <http://ca.com/support>.

Note: CA XOsoft is sold in Japan under the names, CA ARCserve Replication and CA ARCserve High Availability.

Contents

Chapter 1: Introduction	13
About this Guide	13
Terms and Abbreviations	14
Related Documentation	16
CA XOssoft Replication and High Availability Overview	17
CA XOssoft Replication	17
CA XOssoft High Availability	18
Supported Application and Database Servers	19
CA XOssoft Solutions	19
How Synchronization Works	19
How Replication Works	22
How Recovery Works	23
How Data Rewind Works	24
How Replication Suspension Works	24
How High Availability Works	24
CA XOssoft Components	25
CA XOssoft Control Service	25
CA XOssoft Engine	26
Management Center	27
CDP Repository	28
CA XOssoft PowerShell	28
CA XOssoft Deployment	29
Chapter 2: Exploring CA XOssoft Manager	31
Logging Into the Management Center	31
Exploring the CA XOssoft Manager Screen	34
Manager Screen Areas	35
Viewing and Arranging the Manager Screen	37
Viewing Options	37
Customizing the Scenario Pane	37
Rearranging Panes	38
Toolbars	40
Standard Toolbar	40
Viewing Toolbar	42
License Registration	43

Chapter 3: Creating Replication Scenarios **45**

Creating a Scenario Group	45
Creating a Scenario	46
Using Templates	60
Create a Template	60
Create a New Scenario using an Existing Template	63

Chapter 4: Running the Replication Process **67**

Initiate Replication	67
Run Mode	70
Stop Replication	71
Change Configuration when a Scenario is Running	72
Synchronization	73
Close and Open the CA XOssoft Manager during Replication	74
Suspend Replication	75
Suspend Replication Manually	76
Resume Replication after Manual Suspension	77
Schedule Replication Suspension	77

Chapter 5: Monitoring Replication **81**

The CA XOssoft Overview Page	81
The CA XOssoft Manager	82
Monitor Multiple Scenarios	82
State Information	83
Live Statistics	83
Events	90
View Events in a Separate Window	90
View Incoming Events	91
Copy Events	91
Filter Events	92
Reports	93
Viewing a Report	94
Deleting Reports	96
Synchronization Report	97
Replication Report	99
Backward Scenario Report	100
Difference Report	101

Chapter 6: Editing and Managing Scenarios and Hosts **103**

Define the Master and Replica Servers	104
---	-----

Add Additional Replica Servers	105
Select Master Directories and Their Contents for Replication	105
Edit Directory Names	108
Remove Master Root Directories	109
Filter Master Directory Files	109
Include Files	111
Exclude Files	114
Synchronize Registry Keys	117
Activate the Registry Synchronization Option	118
Select Registry Keys for Synchronization	119
Select the Name and Storage Location of the Synchronized Registry Keys	121
Auto-discover Database Files for all Databases	123
Select Replica Root Directories	125
Propagating Master Root Directories to Multiple Replica Hosts	127
Scenario Operations	129
Save Scenarios	129
Remove Scenarios	129
Export Scenarios	130
Import Scenarios	131
Host Maintenance	132

Chapter 7: Setting Scenario Properties **137**

Configuring Scenario Properties	137
Understanding Scenario Properties	138
General	139
Replication	139
Event Notification	143
Report Handling	144
Schedule Synchronization	145
Set a Schedule for Automatic Synchronization	146
Exclude Dates from Scheduled Synchronization	147
Set Advanced Schedules	148
Protecting Your System State	149
How to Configure System State Protection	150
Modify Scenario System State Protection	157
Restore System State Data	159
Command Line Enhancements for System State Protection	161
Additional System State Information	161

Chapter 8: Setting Master and Replica Properties **163**

Configuring Master or Replica Server Properties	164
---	-----

Understanding Master Properties	165
Host Connection	165
Replication	165
Spool	167
Event Notification	167
Reports	168
Understanding Replica Properties	170
Host Connection	170
Replication	170
Spool	173
Recovery	173
Scheduled Tasks	174
Event Notification	174
Reports	175
Schedule the Bandwidth Limit	176
Propagating Property Values	178

Chapter 9: Data Recovery **181**

The Data Recovery Process	181
Recover Lost Data from Replica	182
Setting Bookmarks	186
Data Rewind	187

Chapter 10: Switching Over and Switching Back **193**

Understanding the High Availability System and the Switchover and Switchback Procedures	193
Creating High Availability Scenario	195
Switchover	208
Switchback	211
Recover Active Server Using the Manager	215
Recover Active Server from Outside the Manager	216
Manually Recover a Failed Server - Move IP Address	217
Manually Recover a Failed Server-Switch Computer Name	218
Manually Recover Failed Server-IP and Switch Name	219

Chapter 11: Setting High Availability Properties **221**

Configuring High Availability Properties	221
Understanding High Availability Properties	223
Active and Standby Hosts	223
Switchover Properties	224
Hosts	224
Redirection	225

Is Alive	229
DB Management/Application/Shares Management	231
Actions upon Success	232
Move IP Redirection	233
Adding IP Address to the Master Host	233
Configuring the Move IP Method through the Manager	235
Cluster Move IP	239

Chapter 12: Data Replication and Role Switching of Control Services **243**

Understanding the Control Service Scenario	243
Opening the Manager for Using the HA Control Service Scenario	245
Creating High Availability Scenario for the Control Service	247
Switching the Roles of the Active and Standby Control Services	258
Manually Initiating a Control Service Switchover	259
The Switchover and Backward Scenario Processes	260
Switching Back the Control Service Roles	263

Chapter 13: Testing Assured Recovery and Managing VSS Snapshots **265**

About Assured Recovery	265
Creating Assured Recovery Testing Scenario	267
Performing Assured Recovery Test	275
Performing Assured Recovery Test in a Scheduled Mode	276
Performing Assured Recovery Test in a Non-Scheduled Mode	278
Setting Assured Recovery Properties	284
Configuring Assured Recovery Properties	284
Understanding Assured Recovery Properties	286
Using VSS Snapshots	289
Configuring VSS Snapshot Creation	289
Viewing and Managing Snapshots	294

Chapter 14: Using the CDP Repository **297**

Understanding the CDP Repository	297
CDP Repository Components	298
Configuring the CDP Repository	299
Defining the CDP Database	300
Creating an Exchange Scenario with the CDP Option	305
Setting CDP Scenario Properties	313
Configuring CDP Scenario Properties	314
Understanding CDP Scenario Properties	315
Configuring CDP Replica and Report Properties	318
Understanding the CDP Repository Statistics and Reports	319

CDP Repository Statistics	320
CDP Repository Reports	323
Retrieving Deleted Outlook Items Using the E-mail Retrieval	327
Logging Into the Email Retrieval	327
Search for Deleted Items and Retrieve Them	329
Chapter 15: Using the Content Distribution Solution	333
Understanding the Content Distribution Solution	333
Creating a Content Distribution Scenario	336
Chapter 16: Managing Users	347
How Delegated Security Works	347
Access Rights Considerations	348
Prerequisite Tasks	348
Create a User Group	348
Initial Group Selection	351
Set Up a User Group	352
How to Manage Users	352
Delegation of Rights	353
Set User Rights	354
Setting the Super User Group	355
Index	357

Documentation Changes

The following documentation updates have been made since the last release of this documentation:

- [Filter Events](#) (see page 92) - Describes how to filter the events that will be displayed in the Event pane according to different criteria.
- [Synchronize Registry Keys](#) (see page 117) - Describes how to synchronize the Master and the Replica registry keys, and how to define the synchronization frequency.
- [Propagate Root Directories](#) (see page 127) - Describes how to propagate the root directories you set for a Master to multiple Replica hosts at once.
- [Protect Your System State](#) (see page 149) - Describes how to create snapshots of your system and boot files on a scheduled basis, and how to store and restore them from the Replica.
- [Schedule the Bandwidth Limit](#) (see page 176) - Describes how to control the size of the allowed incoming bandwidth on the Replica host, using the Bandwidth Scheduler.
- [Propagate Property Values](#) (see page 178) - Describes how to apply the values you set for one scenario to multiple scenarios at once.
- [Create and Use the Control Service Scenario](#) (see page 243) - Describes how to replicate the Control Service data, and how to switch the roles of two Control Services when the active Control Service is down.
- [Understand and Set CDP Repository Properties](#) (see page 313) - Describes how to configure the CDP Repository properties, and provides the list of properties, corresponding values, and an explanation of each property.
- [View and Understand CDP Repository Statistics](#) (see page 320) - Describes how to display the new CDP Repository statistics, and what type of information appears in the two statistics tables: [Email CDP Repository Statistics Table](#) (see page 321), and [Email CDP Repository Statistics Summary Table](#) (see page 322).
- [Generate and Understand CDP Repository Reports](#) (see page 323) - Describes how to generate and view the new CDP Repository Reports, and what type of information appears in the two CDP Repository Reports: [Email CDP Repository Report](#) (see page 324), and [Detailed Email CDP Repository Report](#) (see page 326).
- [Create and Use the Content Distribution Scenario](#) (see page 333) - Describes how to create a CD scenario, and how to implement the Content Distribution solution.
- [Delegate and Manage Security Rights](#) (see page 347) - Describes how to delegate security rights, and how to manage users and users groups when using ACL-based Control Service.



Chapter 1: Introduction

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

This section contains the following topics:

[About this Guide](#) (see page 13)

[Terms and Abbreviations](#) (see page 14)

[Related Documentation](#) (see page 16)

[CA XOssoft Replication and High Availability Overview](#) (see page 17)

[CA XOssoft Solutions](#) (see page 19)

[CA XOssoft Components](#) (see page 25)

[CA XOssoft Deployment](#) (see page 29)

About this Guide

This Guide contains all of the necessary information for configuring and running the CA XOssoft 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.

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

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 Operation Guide. You can find the most up-to-date Operations Guides for each application on CA Support site. (For more information about viewing CA XOssoft Operation Guides, refer to [Related Documentation](#) (see page 16).)

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 XOssoft Installation Guide*
- *CA XOssoft PowerShell Guide*
- *CA XOssoft SQL Server Operation Guide*
- *CA XOssoft Exchange 200x Operation Guide*
- *CA XOssoft HA Exchange 200x Operation Guide*
- *CA XOssoft HA SQL Server Operation Guide*
- *CA XOssoft HA Oracle Server Operation Guide*
- *CA XOssoft HA File Server Operation Guide*
- *CA XOssoft HA IIS Server Operation Guide*
- *CA XOssoft HA for Blackberry Enterprise Server Operation Guide*
- *CA XOssoft UNIX and Linux Operation Guide*

To view CA XOssoft Operation Guides

1. Open CA Support site at <https://support.ca.com/> and log in.
2. Click the **Documentation** link on the left.
The Documentation page appears.
3. From **Select a Product**, select CA XOssoft.
4. From **Select a Release**, select the required release number.
5. From **Select a Language**, select the required language.
6. Click the **Go** button.

The Operation Guide list appears.

White papers offering a broader introduction to Disaster Recovery and business continuity solutions using CA XOssoft and CA XOssoft HA are also available on the CA XOssoft website: www.caxossoft.com

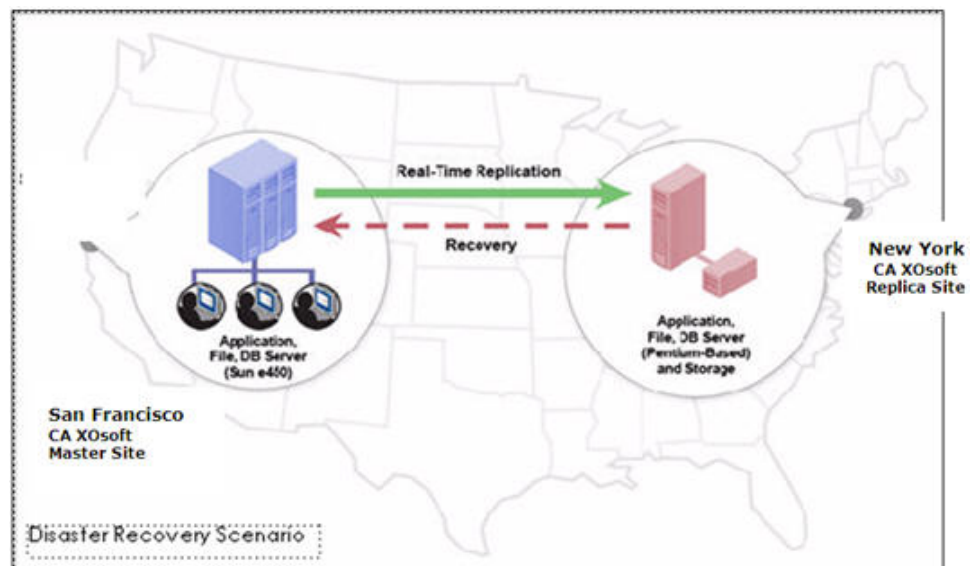
CA XOssoft Replication and High Availability Overview

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

CA XOssoft Replication

CA XOssoft 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 XOssoft 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.

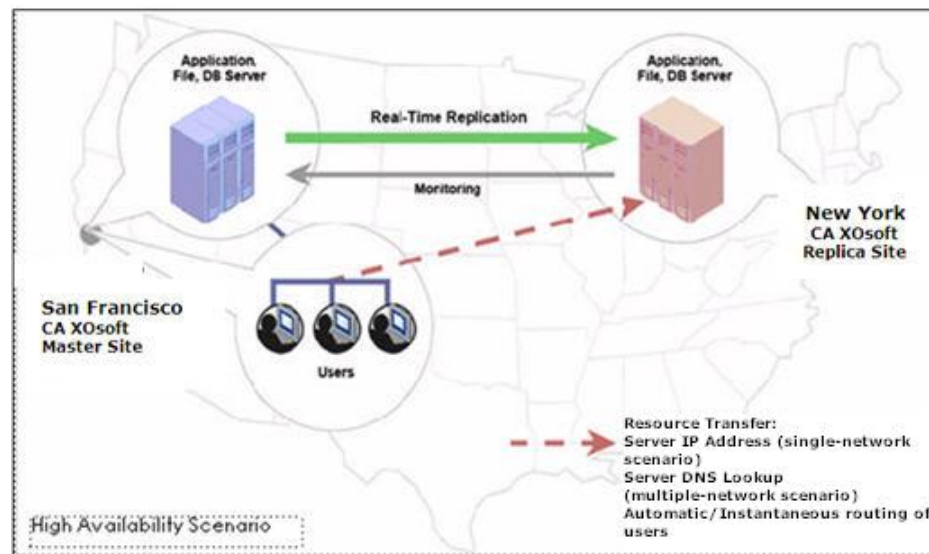
The replication process requires first setting up a valid replication scenario. Then, synchronization between the Master and Replicas data is performed, and once the Master and Replica servers have an identical set of data, the replication process starts. Once the replication process has started, the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



CA XOssoft High Availability

CA XOssoft High Availability (CA XOssoft HA) is based on asynchronous real-time replication and automated switchover and switchover. 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 XOssoft HA 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 XOssoft 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 XOssoft.

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 SharePoint
- Microsoft SQL Server
- Microsoft IIS Server
- Microsoft Hyper-V
- Oracle
- File Server
- VMware vCenter Server
- CA XOsoft Control Service

Note: BlackBerry Enterprise Servers can be protected using CA XOsoft for Microsoft SQL Server or File Server, depending upon your configuration. For details, see the CA XOsoft HA for Windows BlackBerry Enterprise Server Operation Guide.

For an up-to-date list of supported platforms and applications, see the *CA XOsoft r12.5 Supported Configurations* article in <https://support.ca.com/>

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.

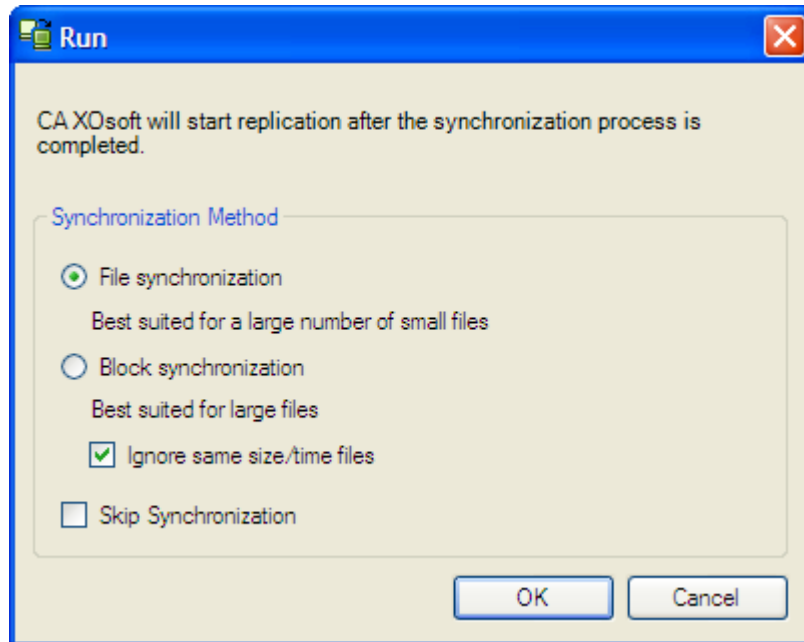
In File synchronization, the Replica server sends its snapshot to the Master server, which uses it to drive the information and content for the data comparison. After the comparison is performed, the CA XOsoft Engine on the Master server sends a sequence of commands to the Replica. These commands consist 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 XOssoft 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 XOssoft 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 XOssoft 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 XOssoft 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 XOssoft 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 XOssoft 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 XOssoft 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 XOssoft 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 XOssoft Components

CA XOssoft is comprised of the following components:

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

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

CA XOssoft Control Service

The CA XOssoft Control Service functions as the single-point-of-control of the CA XOssoft 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 XOssoft 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. We recommend installing the Control Service on a standalone host. If this is not possible, you can install the Control Service on either the Master or Replica servers. However, if the server is down, the connection with the Control Service is lost and scenarios will be unmanageable.

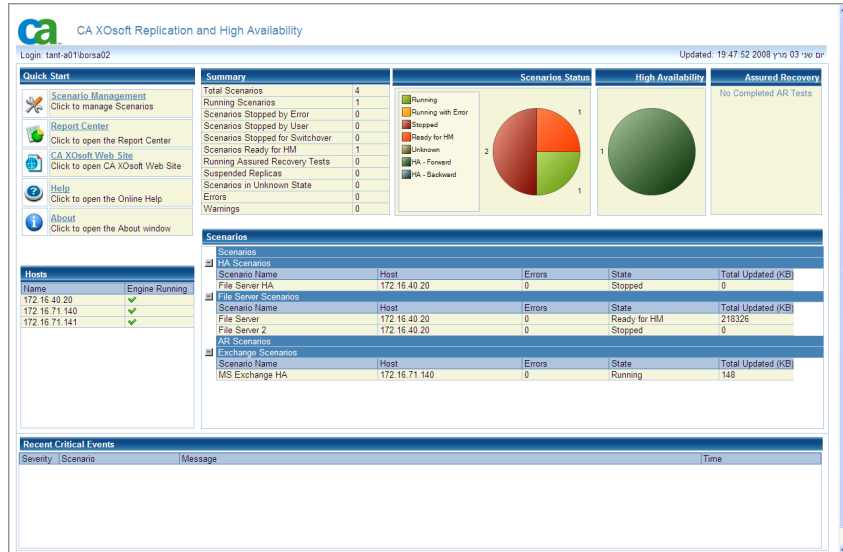
CA XOssoft Engine

The CA XOssoft 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.

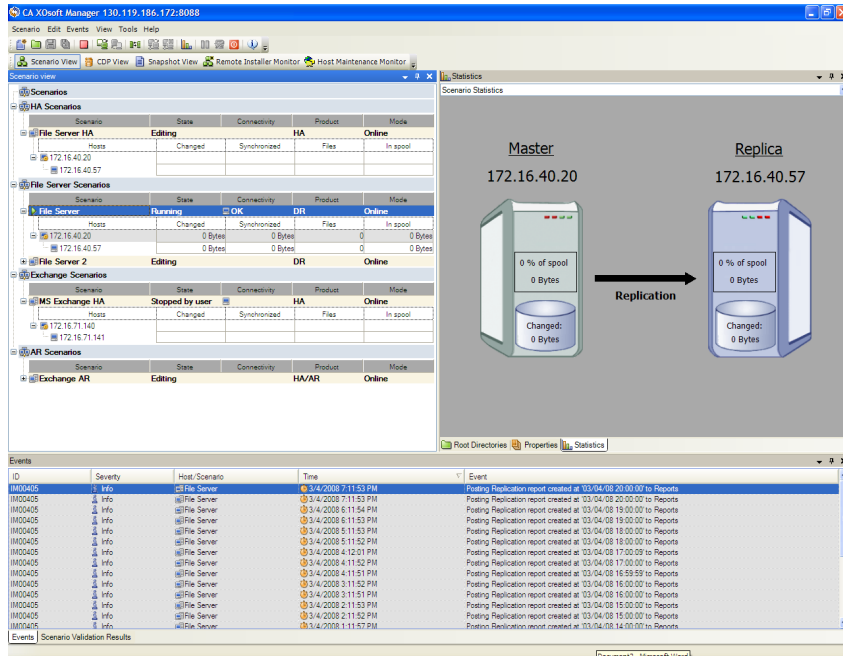
Management Center

The CA XOssoft 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.

CA XOsft r12 Report Center

Report Center Home Page

Updated: 18-11-28 2008 03:03

Available Reports per Scenario

Scenario Name	Synchronization	Difference	Replication	Assessment Mode	Assured Recovery	Total Reports
File Server 1	1	1	0	0	0	2
Exchange HA	2	0	0	0	0	2
Backward_Exchange HA	2	0	0	0	0	2

HA Scenarios

Scenario Name	Synchronization	Difference	Replication	Assessment Mode	Assured Recovery	Total Reports
Exchange AR	1	0	0	0	0	1
MS Exchange HA	1	0	0	0	0	1

Exchange Scenarios

Scenario Name	Synchronization	Difference	Replication	Assessment Mode	Assured Recovery	Total Reports
Exchange DR	5	0	0	0	0	5
MS Exchange HA	1	0	0	0	0	1

Reports

Drag a column header here to group by that column

Host	Date	Time	Type	Summary	Detailed	Size (bytes)
172.16.95.3	02/16/08	17:40:46	Verification			1490
172.16.95.3	02/16/08	17:22:22	Synchronization			1496

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, refer to [Using the CDP Repository](#) (see page 297).

CA XOsft PowerShell

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

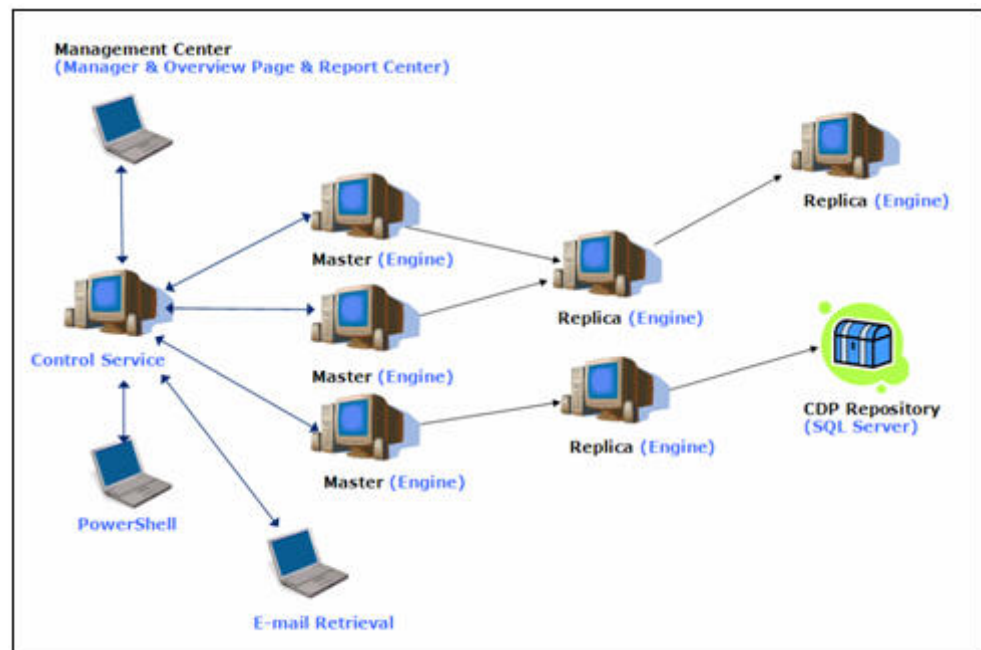
The CA XOssoft PowerShell is a command-line shell and scripting environment that allows users to configure a replication scenario and control and monitor the replication process. All the scenarios that are managed by the CA XOssoft PowerShell look and operate exactly as the ones that are managed by the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

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

CA XOssoft Deployment

The deployment of CA XOssoft 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 XOssoft different components on a Windows platform. For information regarding an efficient deployment of CA XOssoft components, see the *CA XOssoft Installation Guide*.

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



Chapter 2: Exploring CA XOsoft Manager

This section 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.

This section contains the following topics:

[Logging Into the Management Center](#) (see page 31)

[Exploring the CA XOsoft Manager Screen](#) (see page 34)

[Viewing and Arranging the Manager Screen](#) (see page 37)

[Toolbars](#) (see page 40)

[License Registration](#) (see page 43)

Logging Into 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

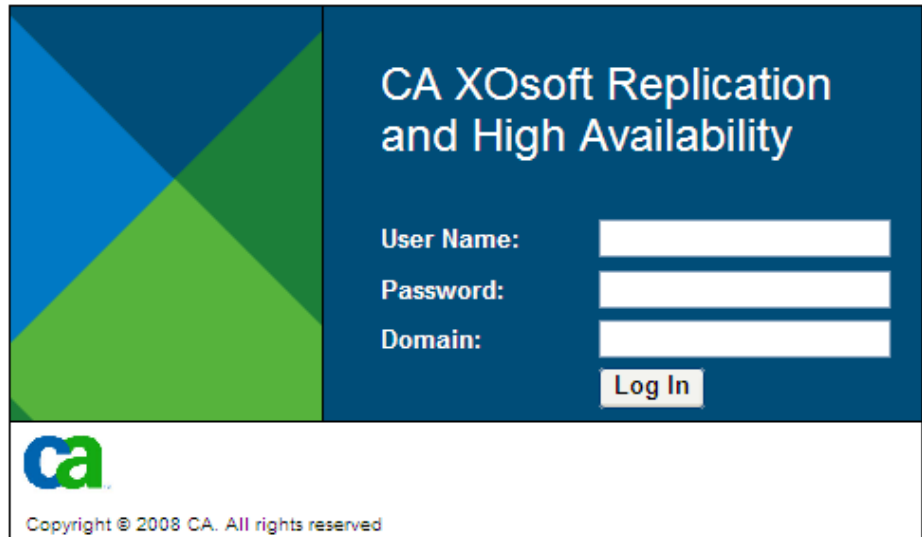
To open CA XOsoft Manager

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.

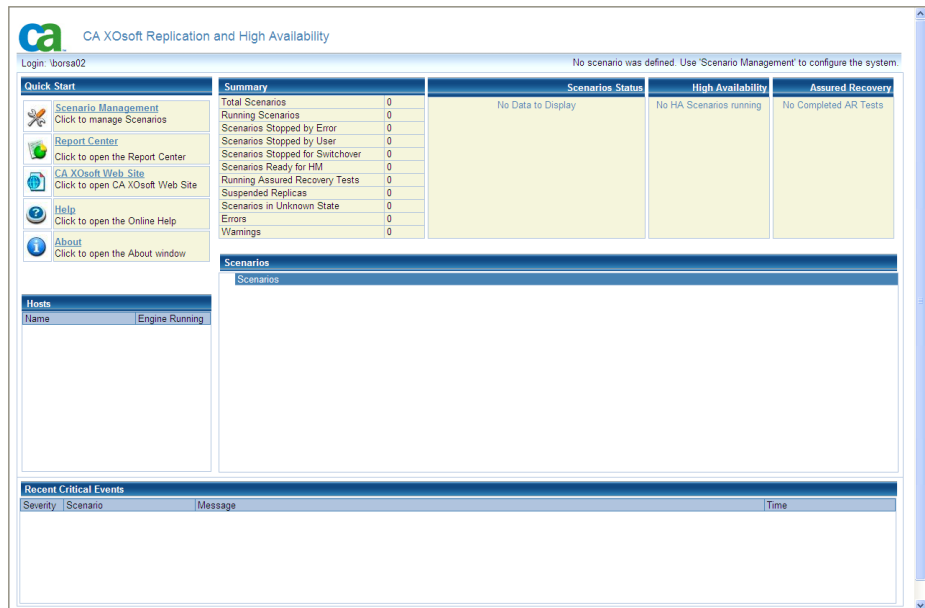


The login dialog features a blue background with a green and blue geometric pattern on the left. The title "CA XOsoft Replication and High Availability" is displayed in white. Below the title are three input fields for "User Name:", "Password:", and "Domain:". A "Log In" button is positioned at the bottom right. At the bottom of the dialog, the CA logo and the text "Copyright © 2008 CA. All rights reserved" are visible.

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

Important! To log in to the Management Center, you must be a member of the Administrators Group on the Local machine where the Control Service is installed.

The **Overview page** appears.



The Overview page displays the CA XOsoft Replication and High Availability interface. It includes a "Quick Start" sidebar with links to Scenario Management, Report Center, CA XOsoft Web Site, Help, and About. The main area features a "Summary" table, a "Scenarios Status" table, and a "Recent Critical Events" table.


Summary	Scenarios Status	High Availability	Assured Recovery
Total Scenarios	No Data to Display	No HA Scenarios running	No Completed AR Tests
Running Scenarios			
Scenarios Stopped by Error			
Scenarios Stopped by User			
Scenarios Stopped for Switchover			
Scenarios Ready for HM			
Running Assured Recovery Tests			
Suspended Replicas			
Scenarios in Unknown State			
Errors			
Warnings			

Recent Critical Events			
Severity	Scenario	Message	Time

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.

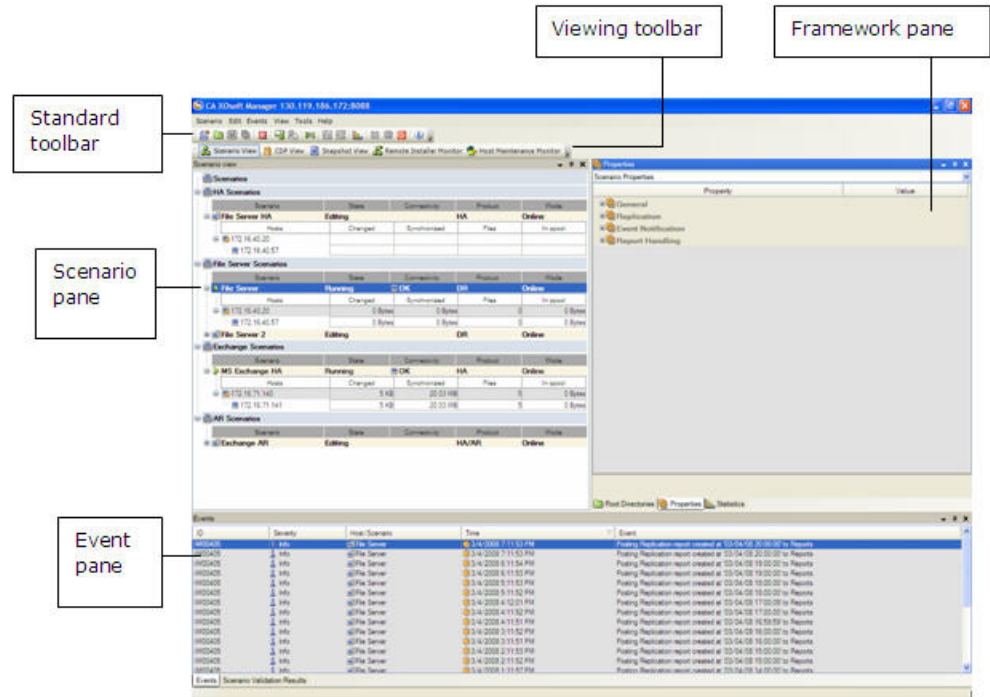


Important! Multiple administrators can simultaneously access CA XOssoft Manager, and they can make any changes anytime they need, depending on their privileges. The last update will be effective as the latest state of the scenario. Therefore, when multiple administrators are working with the Manager on the same time, it is important to be aware that one administrator can unintentionally overwrite the changes another administrator just did. We recommend taking internal measures to prevent the occurrence of this event.

Exploring the CA XOssoft Manager Screen

After logging in to the application, CA XOssoft 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. If active scenarios exist, they are shown on the left side of the Manager screen.




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

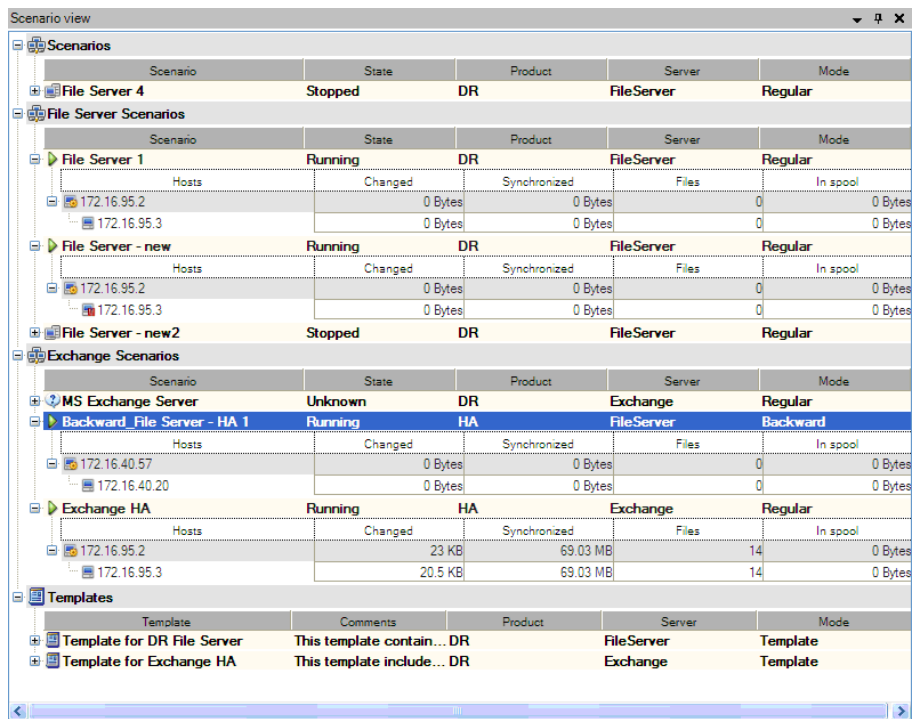
Manager Screen Areas

The Manager screen is divided into several areas:

- The application's name (CA XOssoft 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.

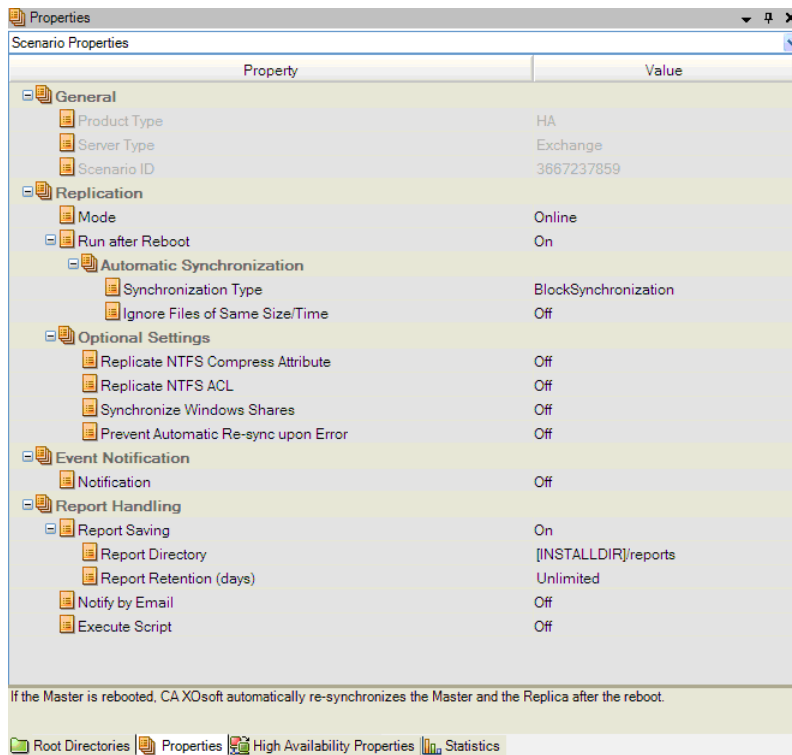


- Scenarios, including their replication tree, appear on the left pane.

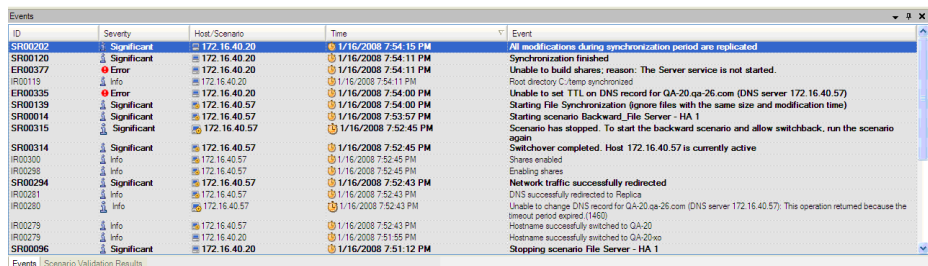


- The Framework pane appears on the right. In this pane, configurable property lists are displayed - the Scenario, Master, Replica, HA and Template properties. The displayed list depends on the item selected in the Scenario pane, or on the item chosen from the pane's drop-down list. The Framework pane also shows two, three or four tabs, depending on the selected CA XOssoft 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.

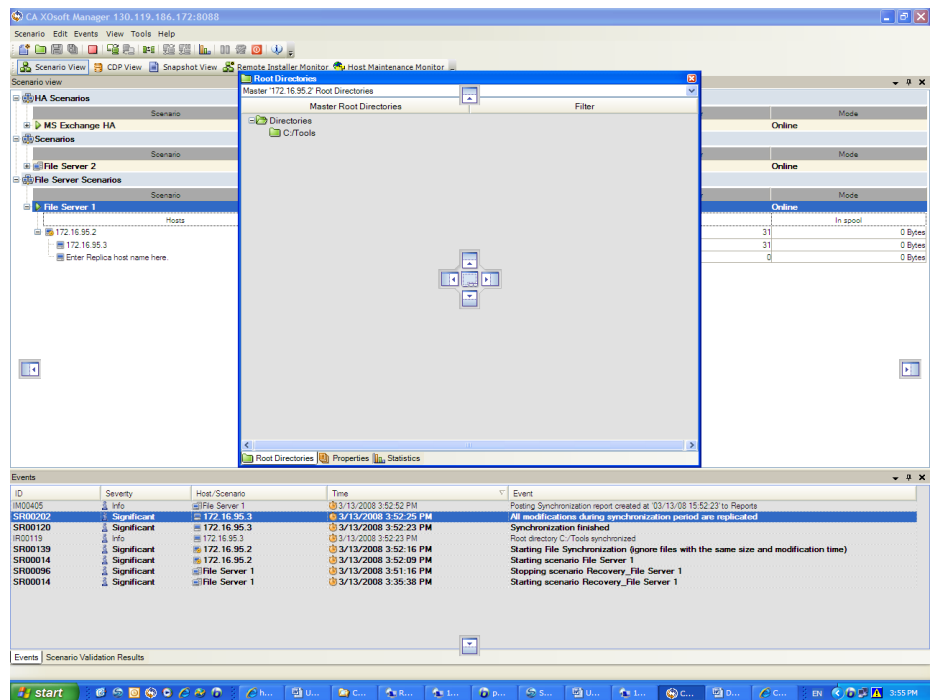
Rearranging Panes

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 XOssoft 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 XOssoft Manager. The following list contains a brief description of each toolbar option:

**New**

Create a new scenario using the Scenario Creation Wizard. Refer to [Creating a Scenario](#) (see page 46).

**Group**

Create a new scenario group. Refer to [Creating a Scenario Group](#) (see page 45).

**Save**

Save a selected scenario. Refer to [Save Scenarios](#) (see page 129).

**Save All**

Save at once all existing scenarios. Refer to [Save Scenarios](#) (see page 129).

**Run**

Run the selected scenario to start the replication process. Refer to [Initiate Replication](#) (see page 67).

**Run (Assessment mode)**

Run the selected scenario in Assessment Mode. Refer to [How Replication Works](#) (see page 22).

**Synchronize**

Activate the synchronization process (whether replication is running or not). Refer to [Synchronization](#) (see page 73).

**Restore Data**

Recover lost or corrupted Master data from any Replica by activating a synchronization process in the reverse direction. Refer to [Data Recovery](#) (see page 181).



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. Refer to [Difference Report](#) (see page 101).



Perform Switchover

[For HA only] Switch the active and passive roles between the Master and Replica servers following their current status. Refer to [Switchover](#) (see page 208).



Suspend Is Alive Check

[For HA only] Suspend the Is Alive check that verifies that the active server is operational. Refer to [Is Alive](#) (see page 229).



Refresh Statistics

Update the scenario state information and live statistics display. Refer to [Refresh Statistics Display Manually](#) (see page 89).



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. Refer to [Suspend Replication](#) (see page 75).



Replica Integrity Testing

Perform Assured Recovery test in a non-scheduled mode. Refer to [Performing Assured Recovery Test in a Non-Scheduled Mode](#) (see page 278).



Launch Host Maintenance

Prepare a node in your replicated system to planned maintenance procedures, while avoiding resynchronization once these procedures are completed. Refer to [Host Maintenance](#) (see page 132).




Help Topics


Open CA XOssoft Online Help.

Viewing Toolbar


The Viewing toolbar buttons provide quick access to the different windows and monitors in the CA XOssoft Manager. The following list contains a brief description of each toolbar option:



The **Scenario View** gives you access to the main [Manager screen](#) (see page 34), which enables you to create, monitor, and manage replication scenarios.




The **CDP View** gives you access to the [CDP View window](#) (see page 300), which enables you to define, change and manage the CDP database and its properties.



The **Snapshot View** gives you access to the [VSS Management window](#) (see page 294), which enables you to view and manage VSS snapshots.



The **Remote Installer Monitor** gives you access to the Remote Installer view, which enables you to view the installation status of Engines, you installed with the Remote Installer. For more information about the Remote Installer, refer to *CA XOssoft Installation Guide*.



The **Host Maintenance Monitor** gives you access to the [Host Maintenance Monitor view](#) (see page 132), which enables you to view the status of current requests for maintenance preparation.

License Registration


CA XOssoft 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 XOssoft product using your license key. To register the product, you need to open CA XOssoft 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.

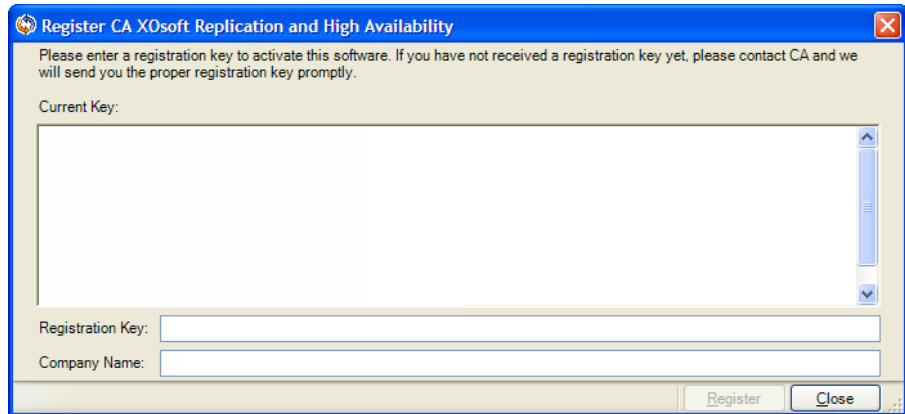
To register CA XOssoft using the license key

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 opens.



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 start working with the CA XOsoft Manager according to your license permissions.

Chapter 3: Creating Replication Scenarios

This section describes how to create the following: a scenario group, a scenario using both the Scenario Creation Wizard and a template, and a scenario template.

This section contains the following topics:

[Creating a Scenario Group](#) (see page 45)

[Creating a Scenario](#) (see page 46)

[Using Templates](#) (see page 60)

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 46).
- 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.

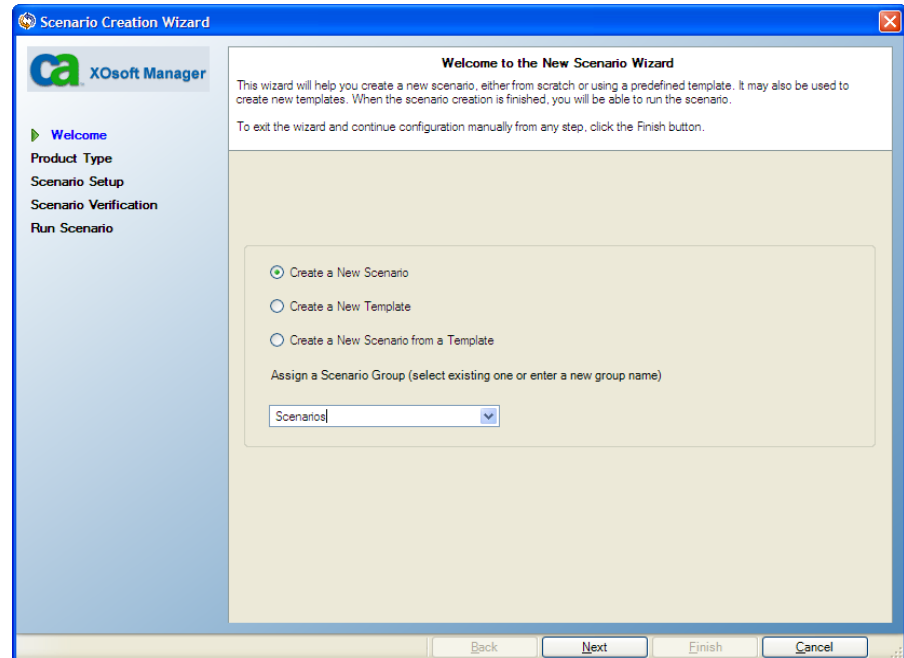
Each scenario defines a replication tree that set the flow of information from the Master server to any number of designated Replicas. It establishes the data recovery procedure, and, if applicable, the switchover parameters. You can configure, add or remove servers from a scenario and select or modify 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 Operation Guide.

To create a new scenario

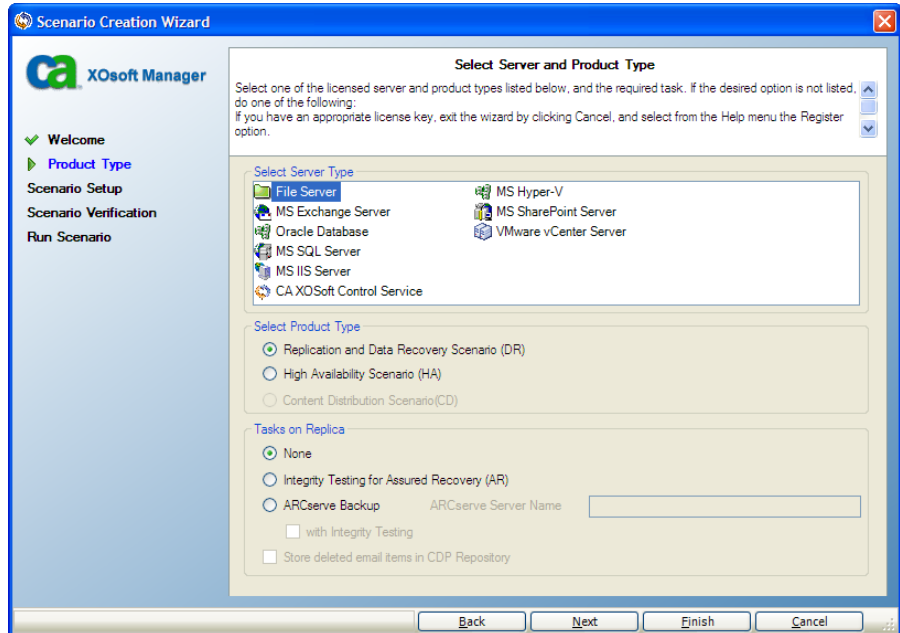
1. Open the CA XOssoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New**  button on the Standard toolbar.

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.

- 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** for DR, or select **High Availability Scenario** for HA.
- [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 **Integrity Testing for Assured Recovery** and how it can help ensure success through automatic testing, refer to the [About Assured Recovery](#) (see page 265).
- To learn more about **ARCserve Backup** and CA XOsoft r12.5 integration see the *CA ARCserve Backup for Windows CA XOsoft Integration Guide*.
- To learn more about the CDP Repository, see [Using the CDP Repository](#) (see page 297).

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.

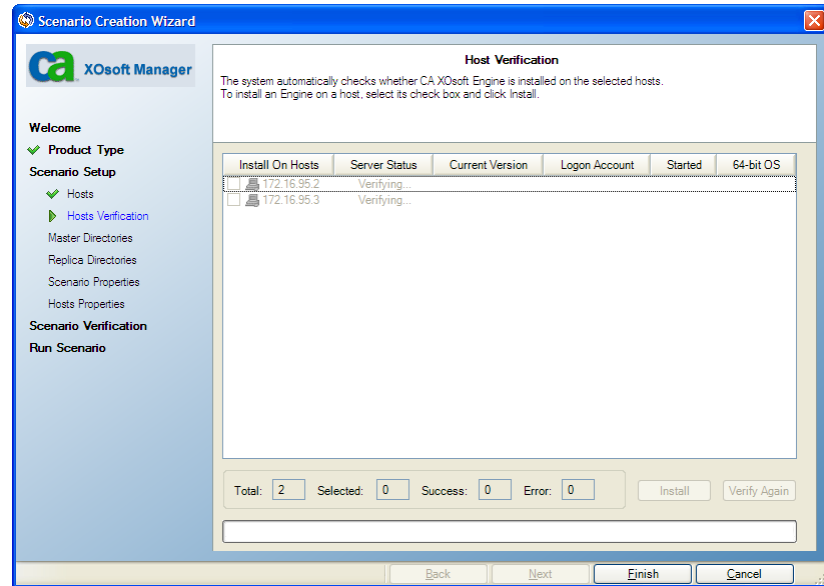
Note: If you want to include more than one Replica in the scenario, enter here the details of the first or most upstream Replica. After you finished the scenario creation, manually enter the other Replicas, as described on [Add Additional Replica Servers](#) (see page 105).

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 XOssoft 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.

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

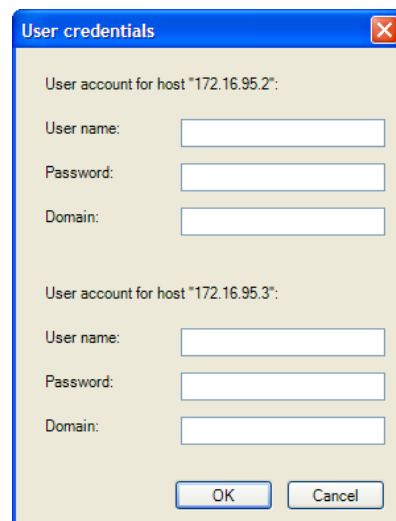
Note: If you did not select the **Verify CA XOssoft Engine on Hosts** option, move to step 11, page 50.

The **Host Verification** page appears:



CA XOssoft 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:



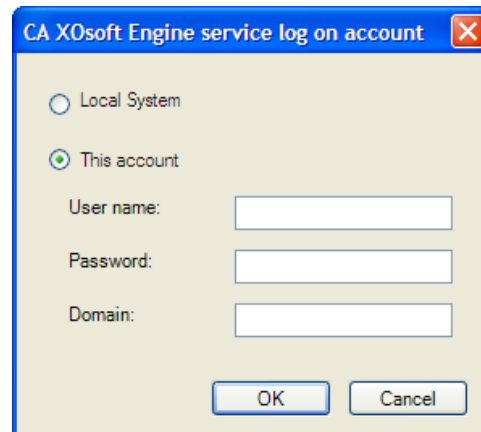
Note: You must enter exactly the same **User Credentials** you used for logging into the remote host.

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 XOssoft Engine service log on account** dialog appears:



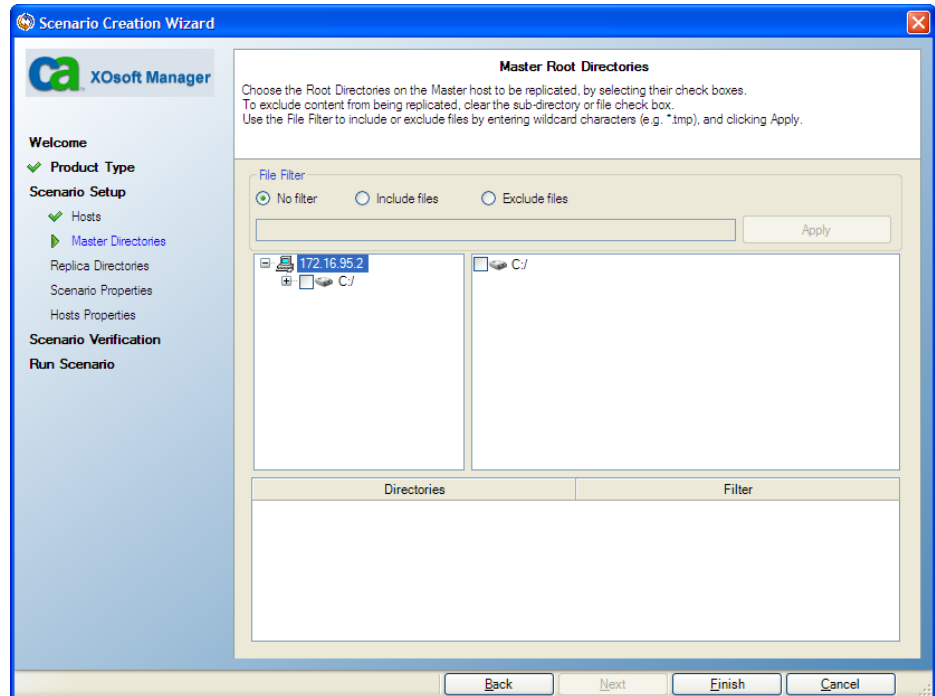
Enter the required information according to following guidelines:

- 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
<input checked="" type="checkbox"/>  172.16.95.2	Complete	5.0.0
<input checked="" type="checkbox"/>  172.16.95.3	Complete	5.0.0

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



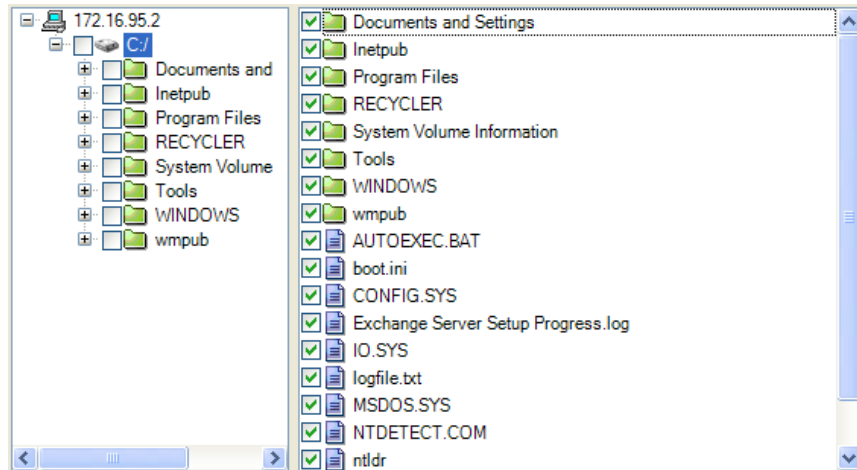
CA XOsft displays the directories and files that are on the Master server. These directories and files are the data that can be replicated and protected. CA XOsft automatically aggregates data that has a common path into one directory.

Notes:

- When you select root directories for the Master or Replica servers, the total character length of root directory plus subdirectory names should not exceed 1024 characters.
- After you finish creating the scenario through the wizard, you can also select registry keys for synchronization, as described in [Synchronize Registry Keys](#) (see page 117).


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:

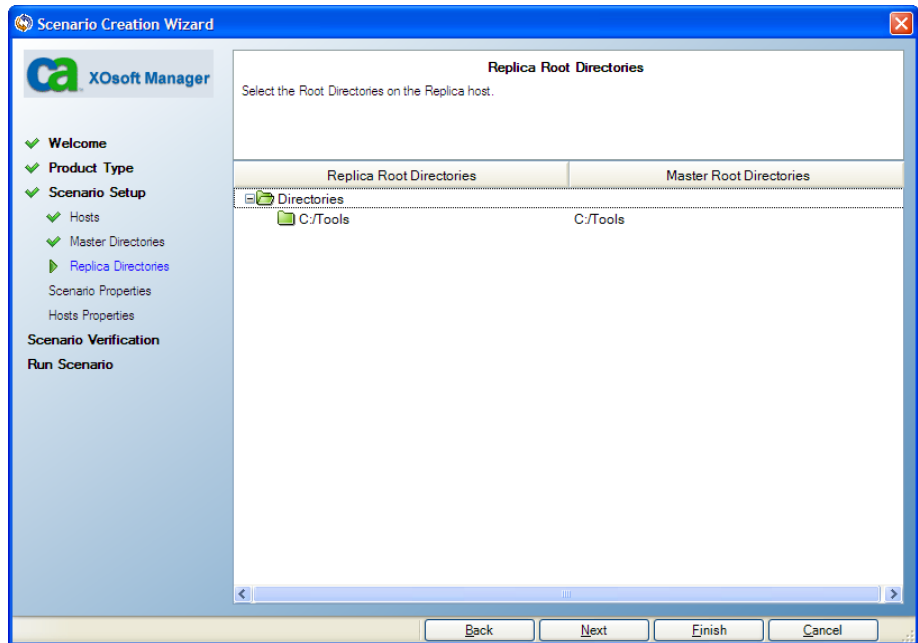
- You can use wildcards to filter the data that will be replicated by selecting the appropriate option button on the **File Filter** section, and entering the expression in the box below (for more information, see [Filter Master Directory Files](#) (see page 109)):



- 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**.

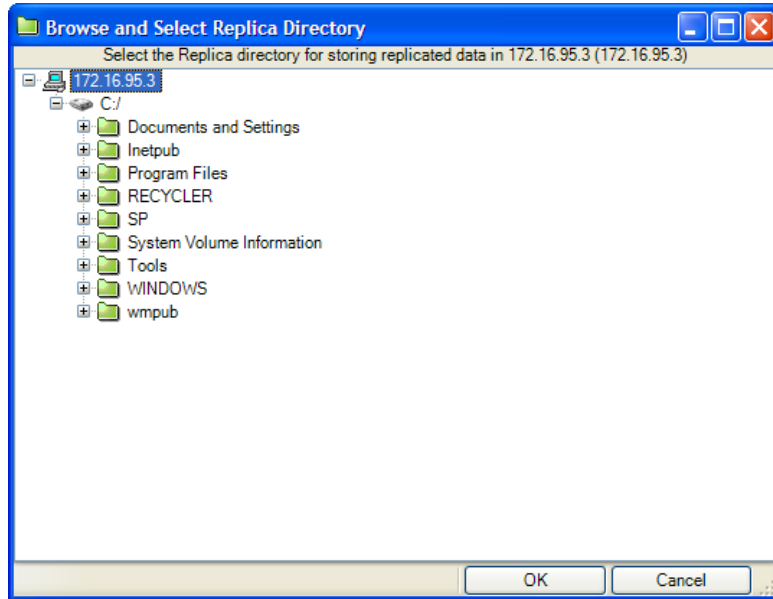
The **Replica Root Directories** page is displayed:



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](#) (see page 125).


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



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

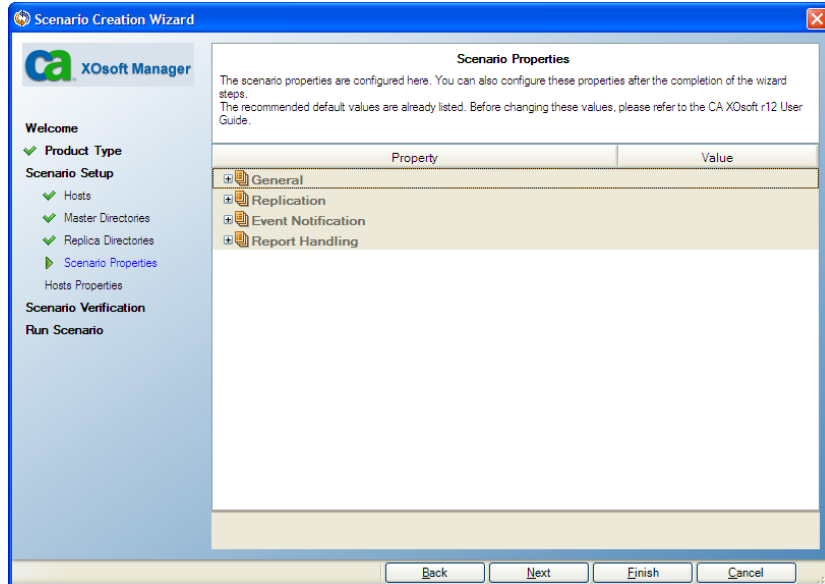
You return to the **Replica Root Directories** page.

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 XOssoft creates it automatically:



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

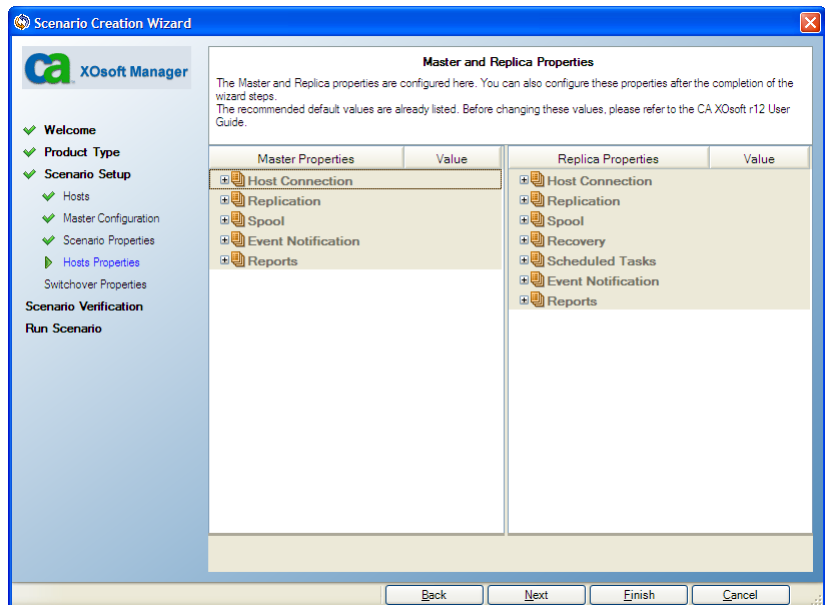
The **Scenario Properties** page appears:



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](#) (see page 138). To configure the scenario properties at a later stage, refer to [Configuring Scenario Properties](#) (see page 137).

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



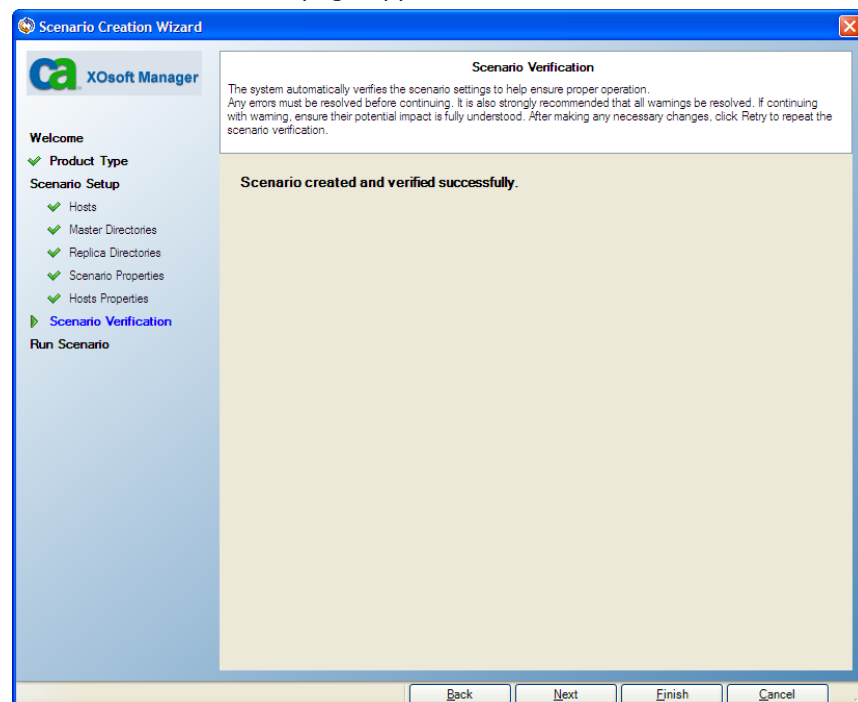
The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

If you want to configure the Master and Replica properties at this stage, refer to [Setting Master and Replica Properties](#) (see page 163). To configure the Master and Replica properties at a later stage, refer to [Configuring Master or Replica Server Properties](#) (see page 164).

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 information](#) (see page 167) for configuration details.

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

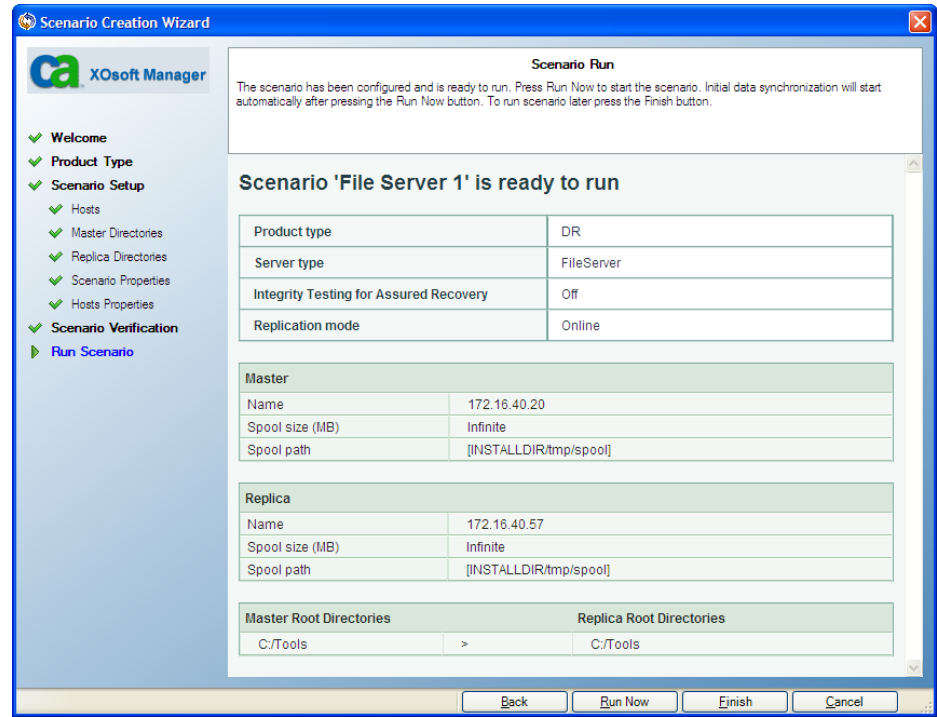
CA XOssoft verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful Disaster Recovery. Once the verification is completed, the **Scenario Verification** page appears:



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, refer to [Viewing a Report](#) (see page 94).



ca CA XOssoft Report Center Report Center Home Page

CA XOssoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server 1
Master host	130.119.185.152(1)
Replica host	130.119.185.153(2)
Scenario start time	18-Mar-09 12:18:09
Report start time	18-Mar-09 12:21:02
Report finish time	18-Mar-09 12:21:21

Summary:

Total number of files modified	54
Total number of bytes changed	12.76MB
Total number of directories created	13
Total number of files with different security attributes	1

Using Templates

Templates are a powerful facility for customizing CA XOssoft 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

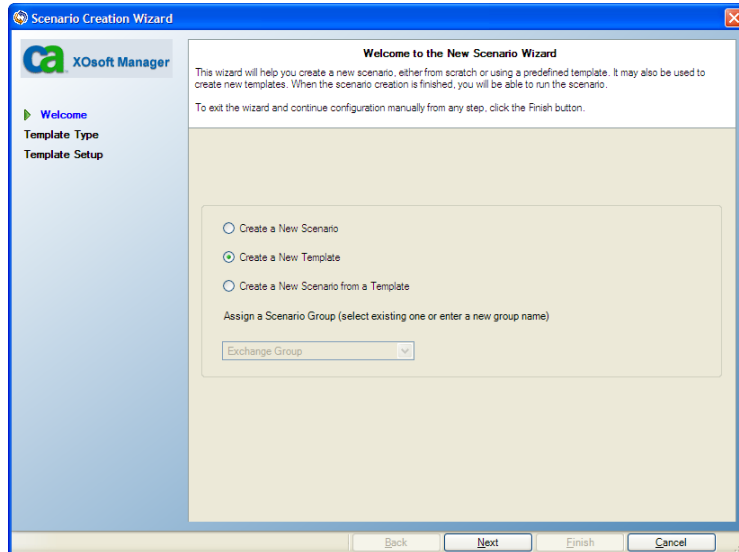
Templates are simple to create and use, and their creation process is basically similar to the creation of a new scenario. However, because a template is not associated with any actual servers, some values cannot be entered, such as the host names or IP addresses of the Master and Replica servers. Also, 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

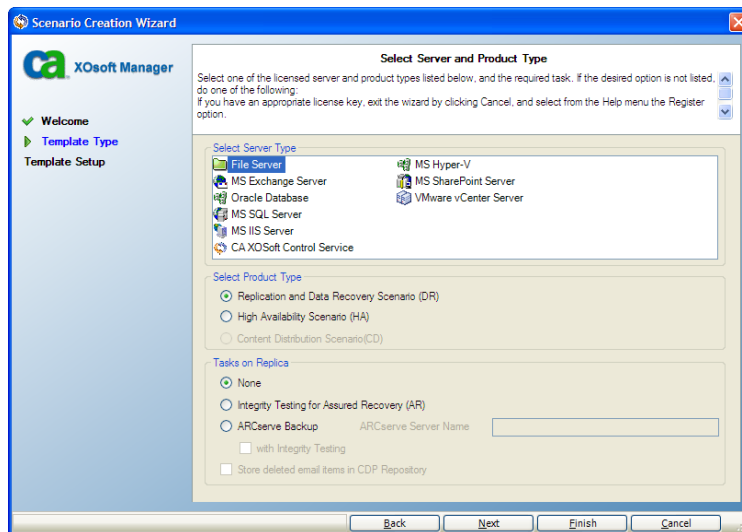
1. Open the Scenario Creation Wizard by clicking the **New**  button on the Standard toolbar, or selecting **New** from the **Scenario** menu.

The **Scenario Creation Wizard** opens.



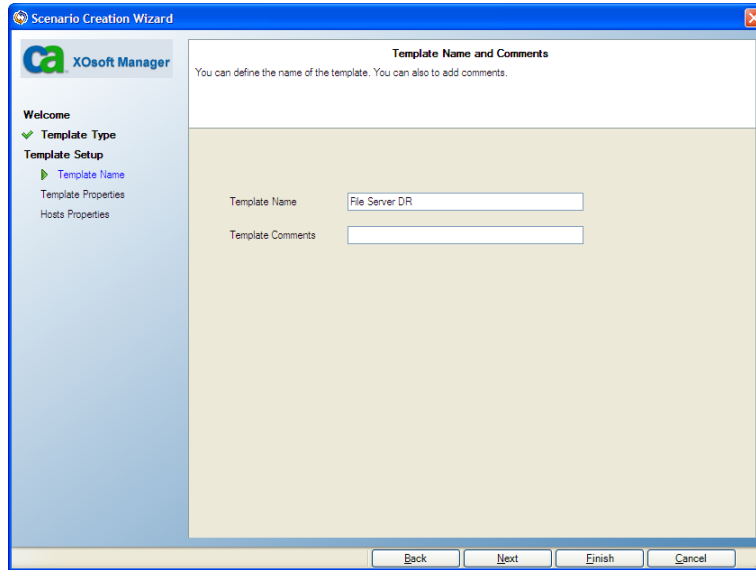
2. Select the **Create a New Template** option button, and click **Next**.

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



3. Select the required scenario options, and click **Next**.

The **Template Name and Comments** page opens.



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 are created using this template. You can also add description of the template in the **Template Comments** box. This description appears 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 46). Continue defining the template by following the wizard instructions.

Create a New Scenario using an Existing Template

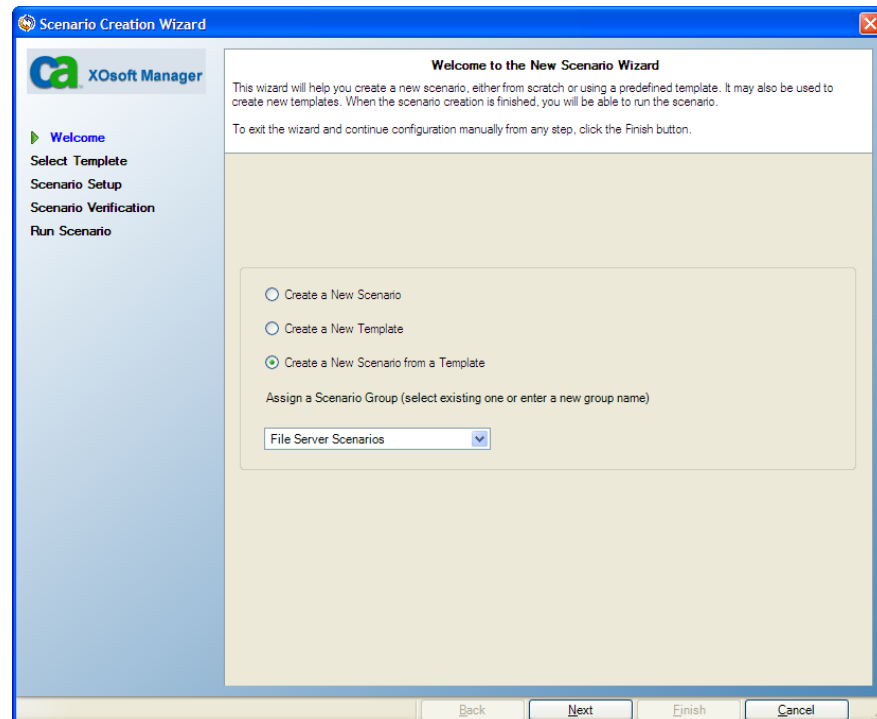
Creating a new scenario using a template saves you the need to separately configure each required scenario. When you are using one of the existing templates, a new scenario is created with all parameter values taken from that template.

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

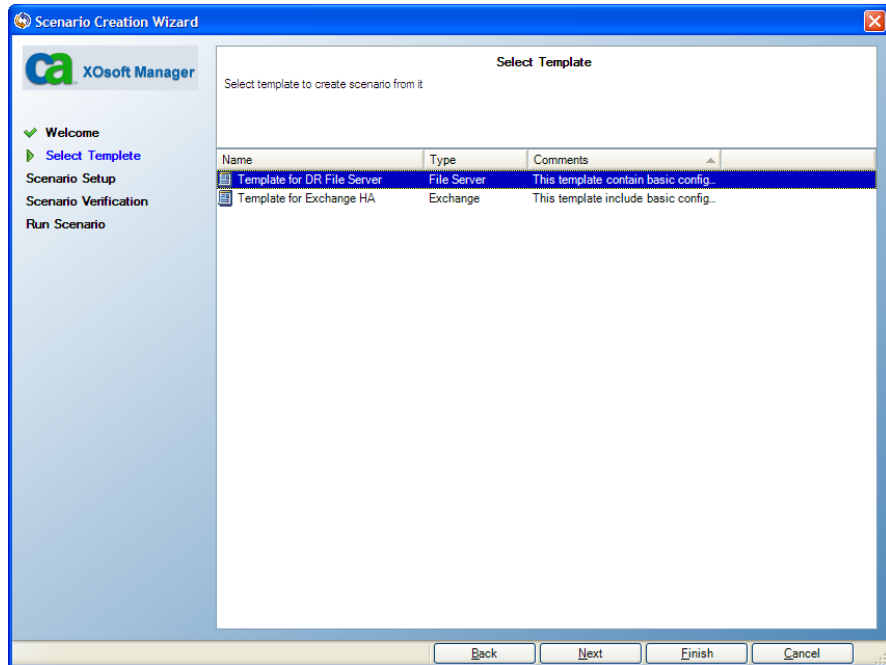
1. Open the Scenario Creation Wizard by clicking the **New**  button on the Standard toolbar, or selecting **New** from the **Scenario** menu.

The **Scenario Creation Wizard** opens.



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

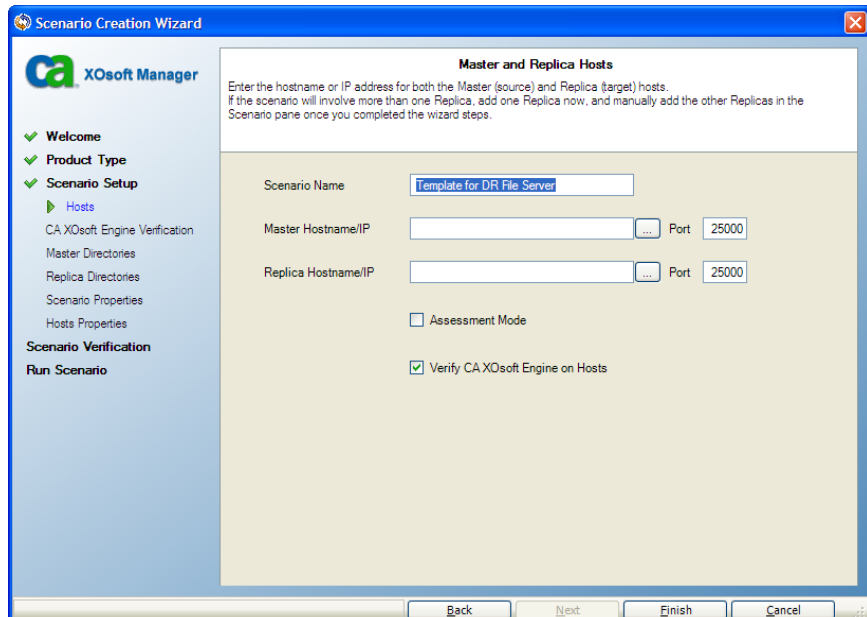
The **Select Template** page opens.



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 opens.



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 46). Continue defining the new scenario by following the wizard instructions.

Chapter 4: Running the Replication Process

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

This section contains the following topics:

[Initiate Replication](#) (see page 67)

[Stop Replication](#) (see page 71)

[Change Configuration when a Scenario is Running](#) (see page 72)

[Synchronization](#) (see page 73)

[Close and Open the CA XOssoft Manager during Replication](#) (see page 74)


[Suspend Replication](#) (see page 75)

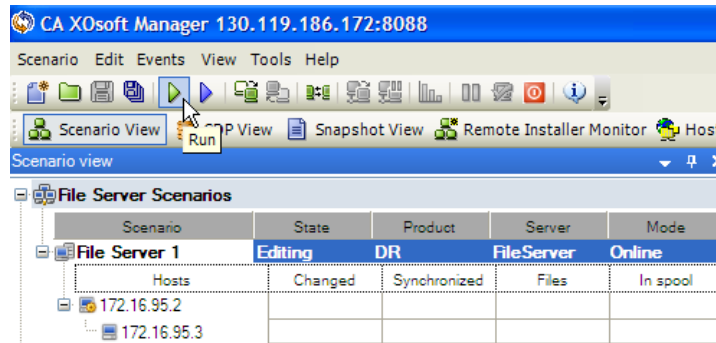
Initiate Replication

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 XOssoft 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 XOssoft 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.

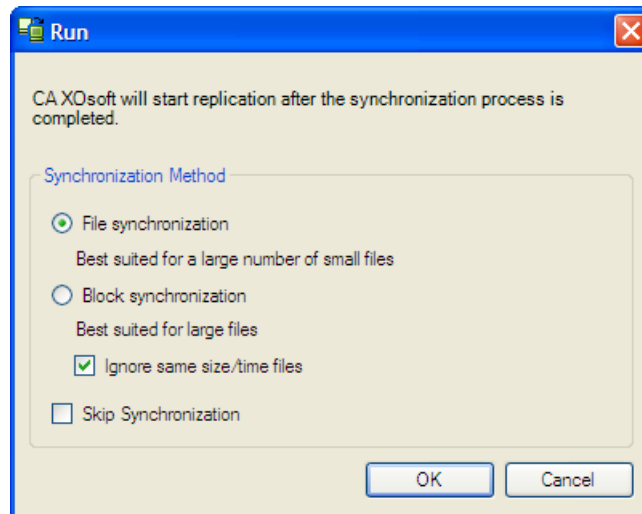


CA XOsoft verifies the scenario before running it.

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 opens.



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

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

- 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**:

Scenario	State	Product	Server	Mode
File Server 1	Running	DR	FileServer	Online
Hosts		Changed	Synchronized	Files
172.16.95.2		0 Bytes	0 Bytes	0
172.16.95.3		0 Bytes	0 Bytes	0

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

- By default, once synchronization occurs, a Synchronization Report is generated. To view the report, refer to [Viewing a Report](#) (see page 94).

Note: You can also generate a [Replication Report](#) (see page 99) on a regular basis to monitor the replication process on each participating server.

Run Mode

While replication is running and the running scenario is selected, the CA XOssoft 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.


The screenshot displays the CA XOssoft Manager interface. The main window is titled "CA XOssoft Manager 130.119.186.172:8088". The "Scenario view" pane on the left shows a tree structure with "File Server" selected and its state set to "Running". The "Scenario Statistics" pane on the right shows a replication diagram between a Master server (172.16.40.20) and a Replica server (172.16.40.57). Both servers show "0 % of spool" and "0 Bytes" of data, with a "Changed" amount of 101.38 MB. The "Events" pane at the bottom shows a log of events, including "Posting Replication report created at 10/28/08 17:00:00 to Reports" and "Connected to 172.16.40.20".

ID	Severity	Host/Scenario	Time	Event
M000405	Info	File Server	2/28/2008 4:11:43 PM	Posting Replication report created at 10/28/08 17:00:00 to Reports
M000405	Info	File Server	2/28/2008 4:11:42 PM	Posting Replication report created at 10/28/08 17:00:00 to Reports
SM00165	Significant	File Server	2/28/2008 3:50:10 PM	Connected to 172.16.40.20

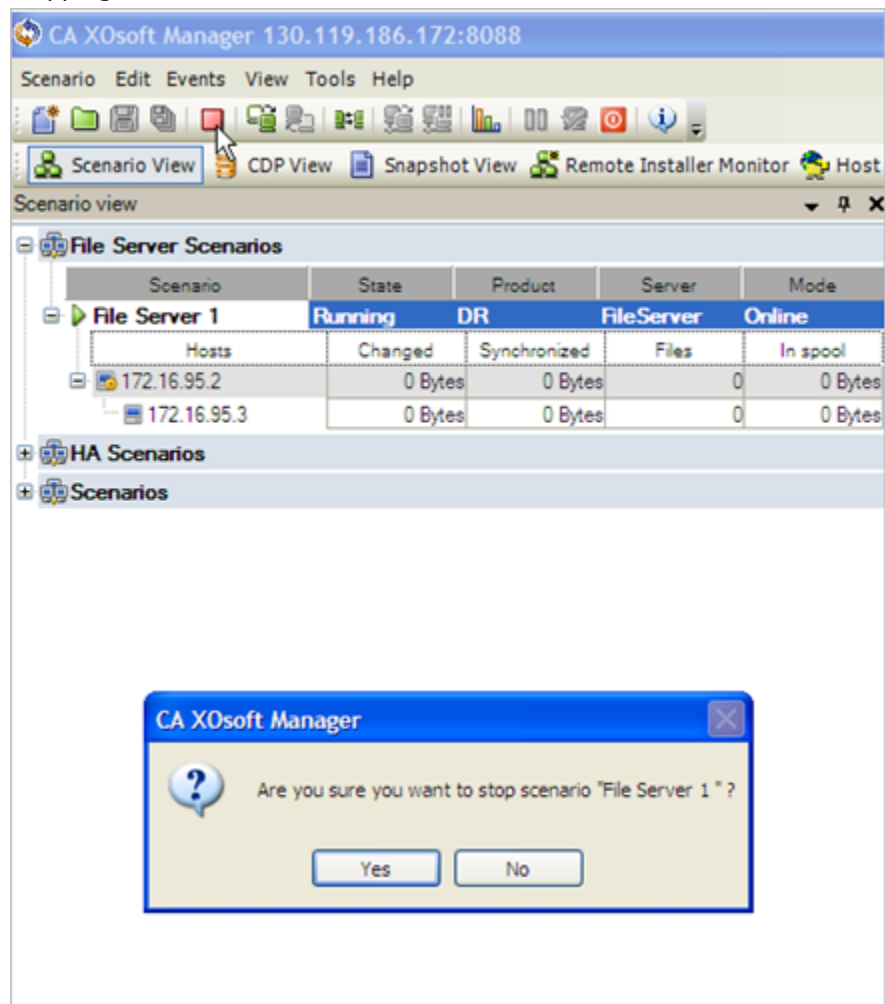
Note: When working with clusters you should be aware that CA XOssoft does not support renaming a Cluster Group while the Engine is running. To ensure the proper operation of clusters with CA XOssoft, stop the Engine before renaming a Cluster Group.

Stop Replication

To stop replication

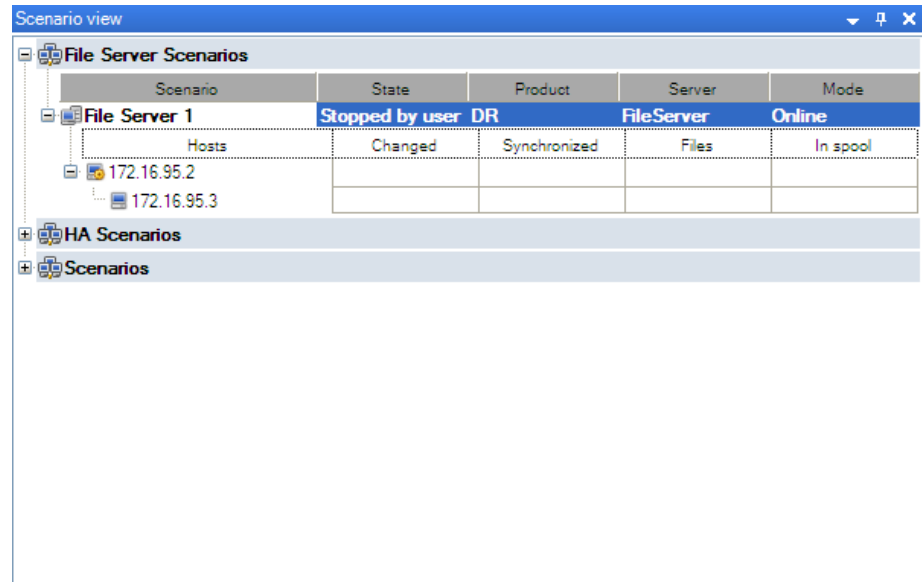
1. From the Scenario pane, select the scenario you want to stop.
2. To stop the scenario, click the **Stop**  button on the Standard toolbar.

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.

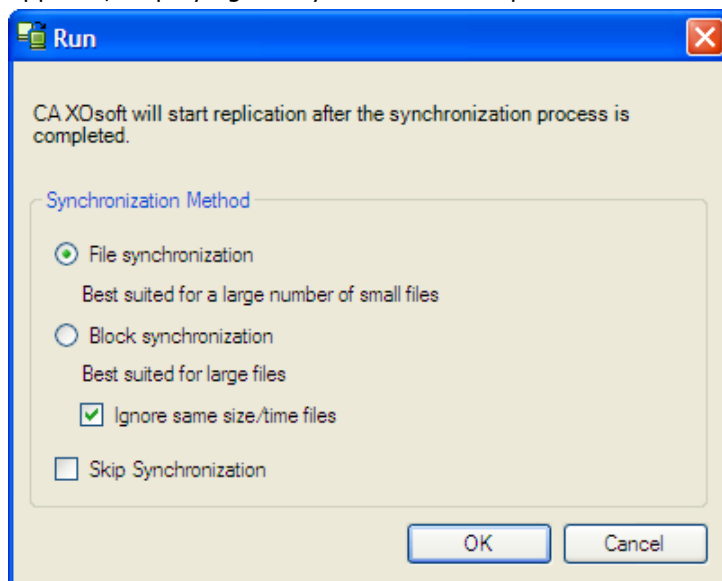
Synchronization

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.



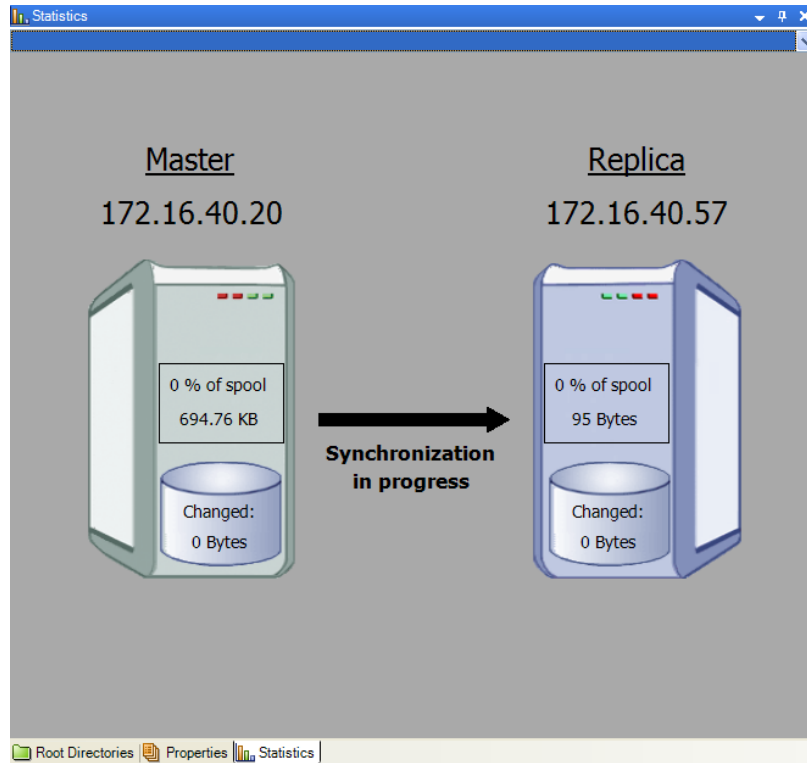
3. Choose the desired options. For more information about the synchronization methods refer to [How Synchronization Works](#) (see page 19) and [Synchronization Type](#) (see page 139).

Note: You can also set synchronization to run automatically at pre-scheduled hours on specific days, and exclude it from specific dates. Refer to [Schedule Synchronization](#) (see page 145).

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**.



Close and Open the CA XOssoft Manager during Replication

After the scenario has been defined, and replication has started, the CA XOssoft 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 XOssoft 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* sections.)

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](#) (see page 265).

Notes:


- 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.
- In a scenario that has more than one Replica host, you can only suspend one Replica at a time.

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. To bring the replica databases online without resynchronizing please use Assured Recovery.



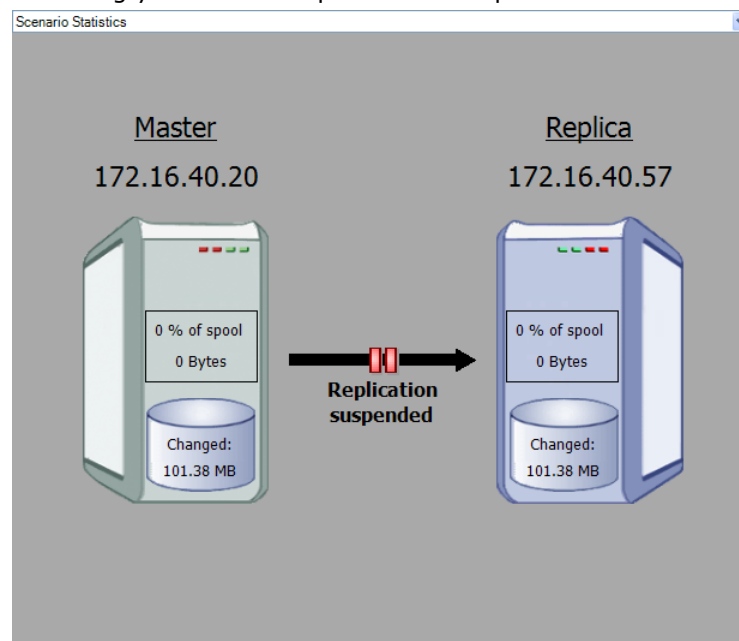
2. Click **Yes** to suspend the replication.

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.





- 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

- 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:



- Click **Yes** to resume replication.

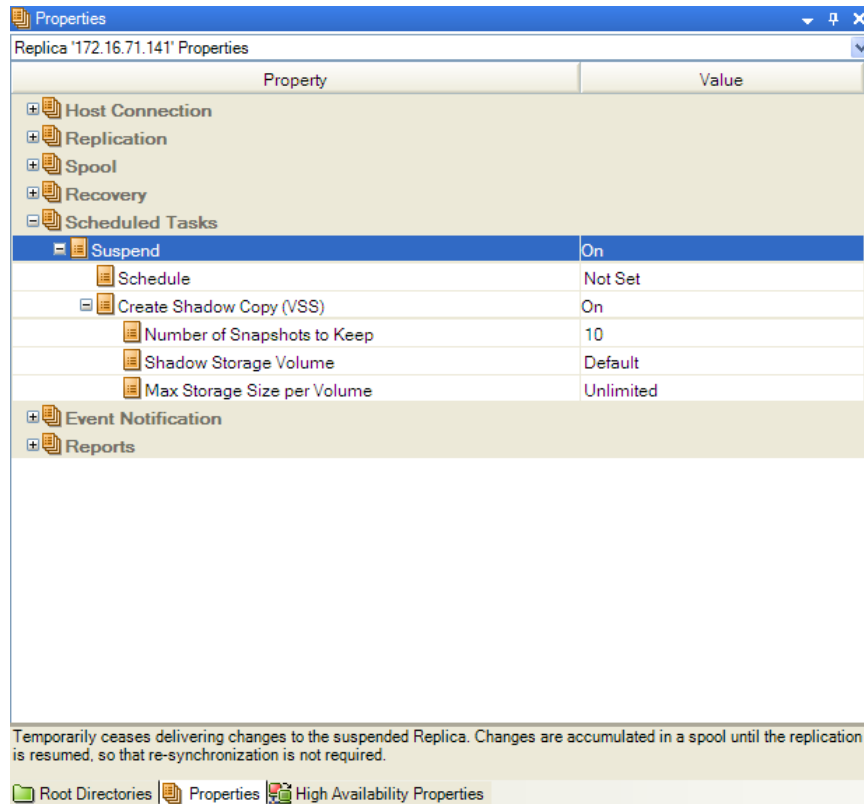
Once the replication is resumed, the red icon disappears from the Replica on the Scenario pane, and the suspension symbol disappears from the Scenario Statistics pane.

Schedule Replication Suspension

To schedule automatic replication suspension

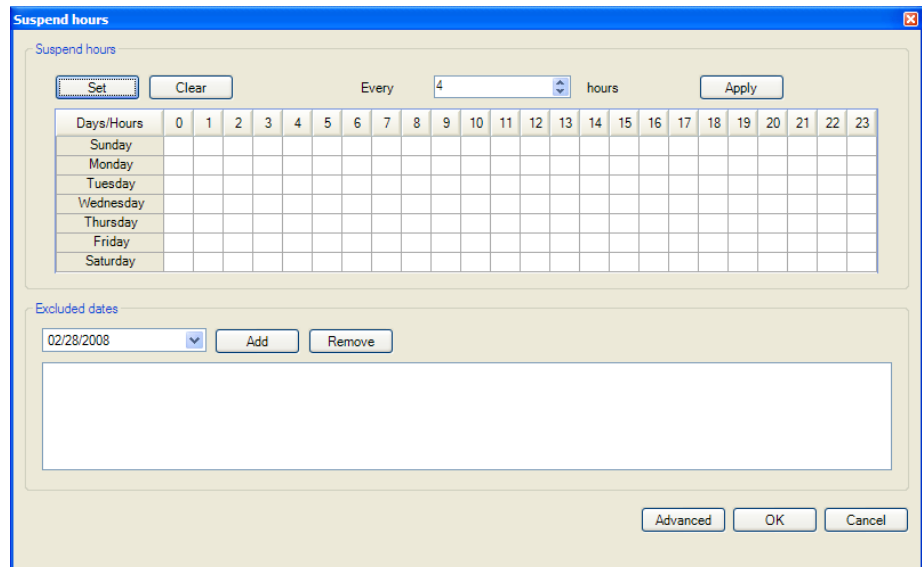
- On the Scenario pane, select the Replica you want to suspend and stop the scenario in which it participates.
- On the Framework pane, select the Properties tab to open the Replica Properties list.

- On the Replica Properties list, open the **Scheduled Tasks** group. On the **Suspension** property, set the value to **On**.



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

The **Suspend hours** dialog opens.



The **Suspend hours** dialog is similar to the **Schedule Setting** dialog, which is used for scheduling automatic synchronization. For information about setting a schedule, refer to [Schedule Synchronization](#) (see page 145).

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 section describes the various monitoring tools of CA XOssoft that enable you to easily control and monitor your replication environment.

This section contains the following topics:

[The CA XOssoft Overview Page](#) (see page 81)

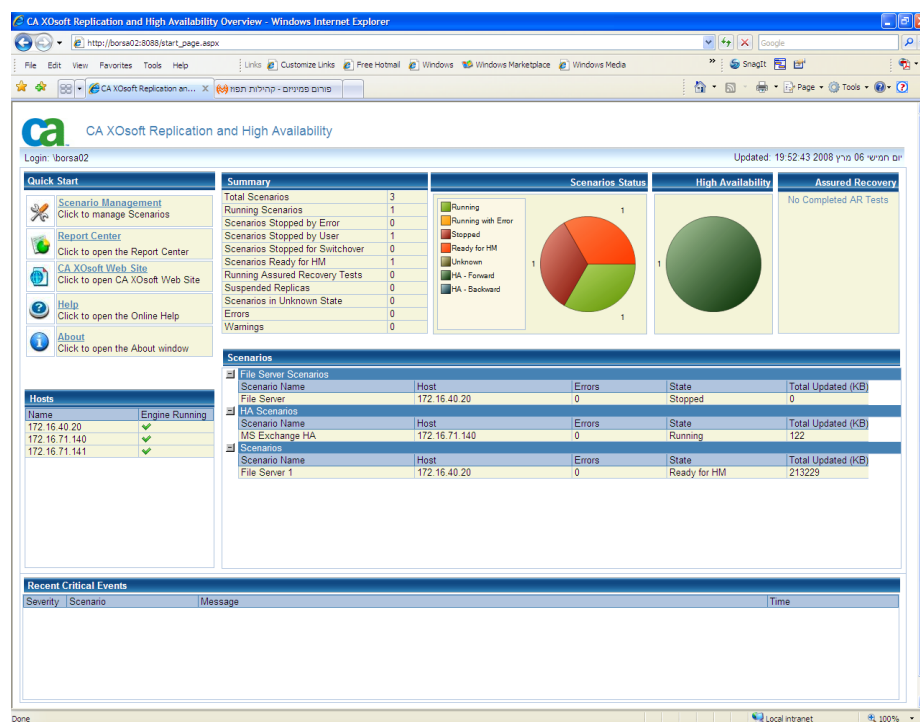
[The CA XOssoft Manager](#) (see page 82)

[Events](#) (see page 90)

[Reports](#) (see page 93)

The CA XOssoft Overview Page

The CA XOssoft 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 XOssoft Manager

The CA XOssoft 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.

The screenshot shows the 'Scenario view' window with a tree view on the left and a table of details on the right. The tree view includes 'File Server Scenarios', 'HA Scenarios', and 'Scenarios'. The table below shows the details for each scenario.





Scenario	State	Connectivity	Product	Mode
File Server	Stopped by user	DR	DR	Online
Hosts	Changed	Synchronized	Files	In spool
172.16.40.20				
172.16.40.57				
MS Exchange HA	Running	OK	HA	Online
Hosts	Changed	Synchronized	Files	In spool
172.16.71.140	112 KB	26.03 MB	6	0 Bytes
172.16.71.141	111 KB	26.03 MB	6	0 Bytes
File Server 1	Running	OK	DR	Online
Hosts	Changed	Synchronized	Files	In spool
172.16.40.20	97.47 MB	0 Bytes	0	33.68 MB
172.16.40.57	46.62 MB	0 Bytes	0	16.43 MB



You can customize the way the columns are displayed. For more information, refer to Customizing the [Scenario Pane](#) (see page 37).

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.

Scenario	State
File Server	Connecting...
Hosts	Changed
 172.16.40.20	
 172.16.40.57	


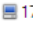
Live Statistics

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:

File Server Scenarios					
Scenario	State	Connectivity	Product	Mode	
File Server	Running	OK	DR	Online	
Hosts	Changed	Synchronized	Files	In spool	
 172.16.40.20	45.02 MB	390.41 MB	6810	23.34 MB	
 172.16.40.57	8.68 MB	390.41 MB	6810	16.89 MB	

- **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.

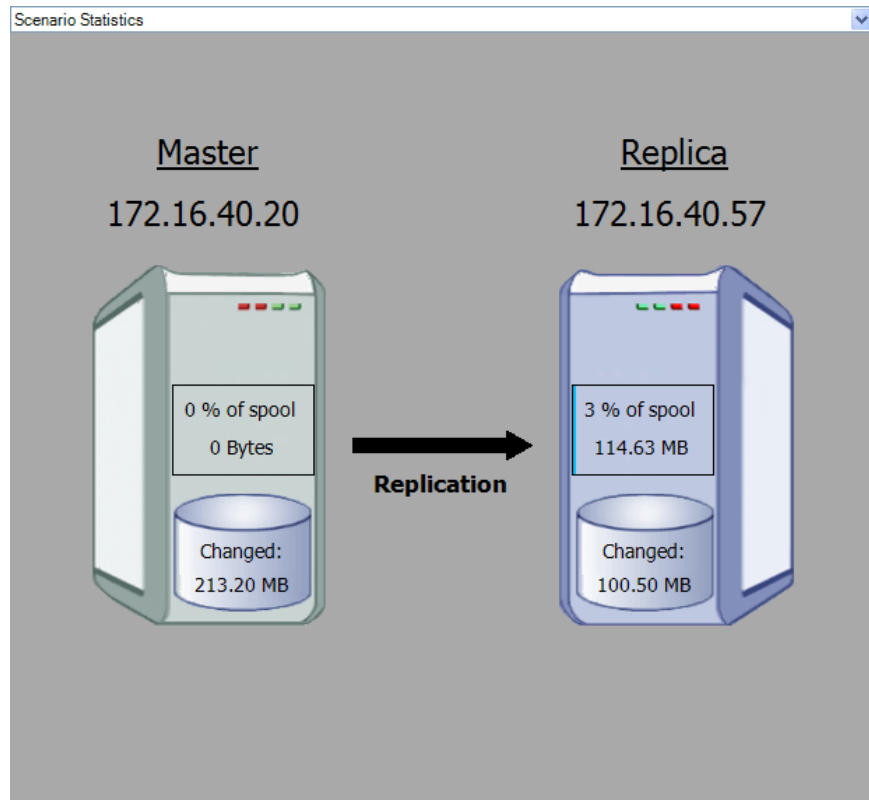
Statistics Tab

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 XSoft 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 sent and replicated from this host; synchronization method; and synchronization progress.

The screenshot shows the 'Statistics' window for Master '172.16.40.20'. The window displays various metrics related to the replication process, including the state, start time, version, spool space usage, online file changes, transferred bytes, and synchronization progress.

Master '172.16.40.20' Statistics

State:	Running
Start of replication:	03/03/08 20:08:49
Version:	5.0.1.153

Spool space:

Size	% of threshold
0Bytes	0%

Online file changes per root directory:

Root Directory	Size	Folders Created	Changed	Removed	Renamed
C:/Tools	213.21MB	263	3243	2	2

Transferred bytes to Replicas:

Host	Size	Current File	Progress
172.16.40.57	610.24MB		100%

Last synchronization statistics: File synchronization

Synchronization Progress:

- 172.16.40.20 -> 172.16.40.57

C:/Tools:

Total File Synchronization Progress	File Transferring Progress
100% of 6810 Files	100% of 6810 Files

The window also shows a taskbar at the bottom with 'Root Directories', 'Properties', and 'Statistics' tabs.

When running synchronization for very large files, additional tables appear, displaying in details the synchronization progress for each file in each root directory.

Notes:

- The definition of a large file depends on the value of the **BDMaxFileSizeToSendWholly** property. This property is stored on the Engine machine, in the INSTALLDIR\Engine\ws_rep.cfg file. The default is 10MB. When a synchronized file is smaller than this value, it will not appear on the table.
- The appearance of the detailed Synchronization Progress table also depends on the value of the **UseNewSynchStatistics** property. This property determines whether the detailed Synchronization Progress table will be displayed when there are large files. When the value of this property is **True**, the table will be displayed. The default value is **True**, and the property is also stored in the ws_rep.cfg file.

Statistics
Master '130.119.176.139' Statistics

State	Running
Start of replication	3/12/2009 4:31:05 PM
Version	12.5.1.205

Spool space:

Size	% of threshold
0Bytes	0%

Online file changes per root directory:

Root Directory	Size	Folders Created	Changed	Removed	Renamed
C:/test/.../m1 for 2 repls	0Bytes	0	0	0	0
C:/test/.../m2 for 2 repls	0Bytes	0	0	0	0
C:/test/.../m3 for 2 repls	0Bytes	0	0	0	0
Total	0Bytes	0	0	0	0

Transferred bytes to Replicas:

Host	Sent	Current File	Progress
130.119.176.115	34.65KB		0%

Last synchronization statistics: File synchronization

Synchronization Progress:

- 130.119.176.139 -> 130.119.176.115
C:/test/.../m1 for 2 repls:

Stage	File Info	Size	Already Checked	To be Sent	Saved	Starting Time
Finished	60M.txt	15MB	100% (15MB)	0Bytes	100% (15MB)	3/12/2009 16:31:08
Finished	50M.txt	75MB	100% (75MB)	0Bytes	100% (75MB)	3/12/2009 16:31:08
Finished	10m.txt	10MB	100% (10MB)	0Bytes	100% (10MB)	3/12/2009 16:31:08
Finished	Total Data	100.11MB	100% (100.11MB)	0Bytes	100% (100.11MB)	3/12/2009 16:31:08
	Total Files	6	6	0	6	3/12/2009 16:31:08

C:/test/.../m2 for 2 repls:

Stage	File Info	Size	Already Checked	To be Sent	Saved	Starting Time
Finished	60M.txt	15MB	100% (15MB)	0Bytes	100% (15MB)	3/12/2009 16:31:08
Finished	50M.txt	75MB	100% (75MB)	0Bytes	100% (75MB)	3/12/2009 16:31:08
Finished	20M.txt	20MB	100% (20MB)	0Bytes	100% (20MB)	3/12/2009 16:31:08
Finished	10m.txt	10MB	100% (10MB)	0Bytes	100% (10MB)	3/12/2009 16:31:08
Finished	Total Data	120MB	100% (120MB)	0Bytes	100% (120MB)	3/12/2009 16:31:08
	Total Files	4	4	0	4	3/12/2009 16:31:08

The Synchronization Progress tables contain the following information for each synchronized file: synchronization state; file name; file size; amount and percentage of data that was already compared between the Master and Replica; amount of data that needs to be sent from the Master to the Replica; amount and percentage of data that is the same on the Master and Replica, and therefore is not sent to the Replica; synchronization starting date and time; summary of the synchronization progress of each root directory.

Note: When the File Synchronization method is running, the files are synchronized in their entirety. Therefore, the **Already Checked** column can contain only two values - 0% or 100% - for an individual file.

Each synchronized root directory is represented by a separate Synchronization Progress table, and each Synchronization Progress table can display statistics of 10 files at the most. When a root directory contains more than 10 files, the 10 largest files will be presented in the table.

- **Statistics per Replica** - a table that containing the following information: state; replication starting date; CA XOssoft 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.

The screenshot shows a window titled 'Statistics' for 'Replica '172.16.40.57''. The window contains the following information:

State:	Running
Start of replication:	03/03/08 20:08:51
Version:	5.0.1.153

Spool space:

Size	% of threshold
20MB	0%

Online file changes per root directory:

Root Directory	Size	Folders Created	Changed	Removed	Renamed
C:/Tools	199.17MB	225	2296	2	2

The window also shows a taskbar at the bottom with icons for 'Root Directories', 'Properties', and 'Statistics'.


Refresh Statistics Display Automatically

The CA XOssoft 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 opens.




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

1. Click the **Refresh Statistics**  icon on the Standard toolbar.
2. Press the **F5** key.
3. 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 XOssoft 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 XOssoft 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.

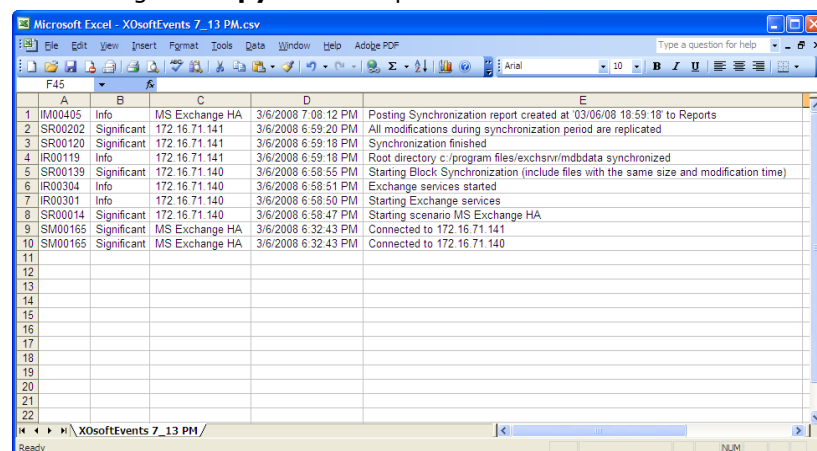
To copy the text of the events for use in other programs

- On the Event pane, select any number of events using the **Ctrl** key.
- 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 XOssoft 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.



The screenshot shows a Microsoft Excel window titled "Microsoft Excel - XOssoftEvents_7_13 PM.csv". The spreadsheet contains the following data:

	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
11					
12					
13					
14					
15					
16					
17					
18					
19					
20					
21					
22					

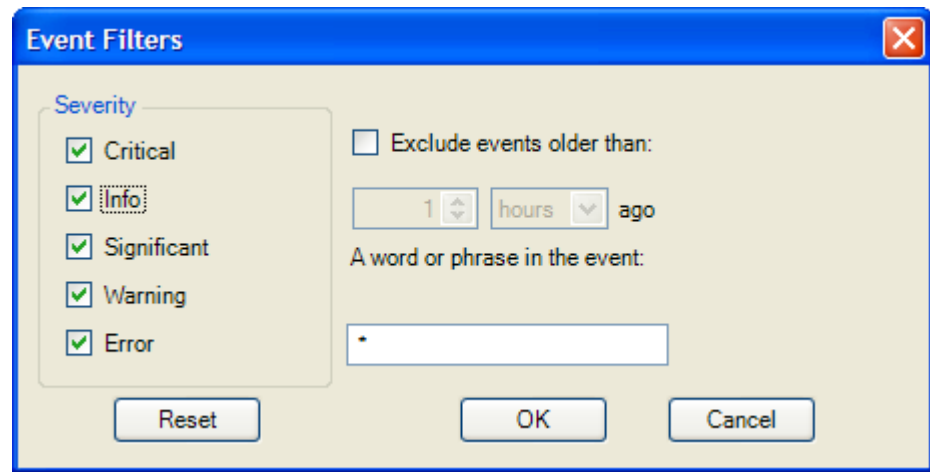
Filter Events

You can filter the events that will be displayed in the Event pane according to different criteria.

To filter displayed events

1. On the Event pane, right-click and select **Event Filters** from the pop-up menu, or select the **Event Filters** option from the **Events** menu.

The **Event Filters** dialog opens.



2. Use one of the following criteria to filter the events that will be displayed in the Event pane:
 - **Severity** - clear the severity level check boxes that you do not want to display, or select the severity level check boxes that you want to display.
 - **Date** - select the **Exclude events older than** check box, and then select the time unit (**hours/days/months**) and the number of units.
 - **Text** - in the **A word or phrase in the event** box, enter the word or the phrase that you want the displayed events to contain. You can use an asterisk (*) to select any number of characters/digits of all types.
3. To apply the criteria you selected and close the dialog, click **OK**.

Only the events that meet the criteria you defined are now displayed in the Event pane.
4. To clear the existing criteria and display all events, on the **Event Filters** dialog click the **Reset** button, and then **OK**.

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, refer to the [Report Handling section](#) (see page 144) in [Understanding Scenario Properties](#) (see page 138).
- For defining the automatic generation of synchronization and replication reports for the Master, refer to the [Reports section](#) (see page 168) in [Understanding Master Properties](#) (see page 165).
- For defining the automatic generation of replication reports for the Replica, refer to the [Reports section](#) (see page 175) in [Understanding Replica Properties](#) (see page 170).

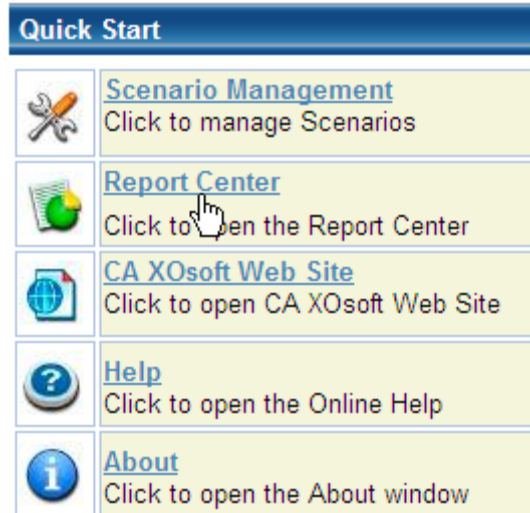
The default storage directory of the generated reports is:
[ProgramFilesFolder]\CA\XOsoft\Manager\reports

Important: A report cannot be created if the scenario name contains special characters (i.e. \?:"<>|,).

Viewing a Report

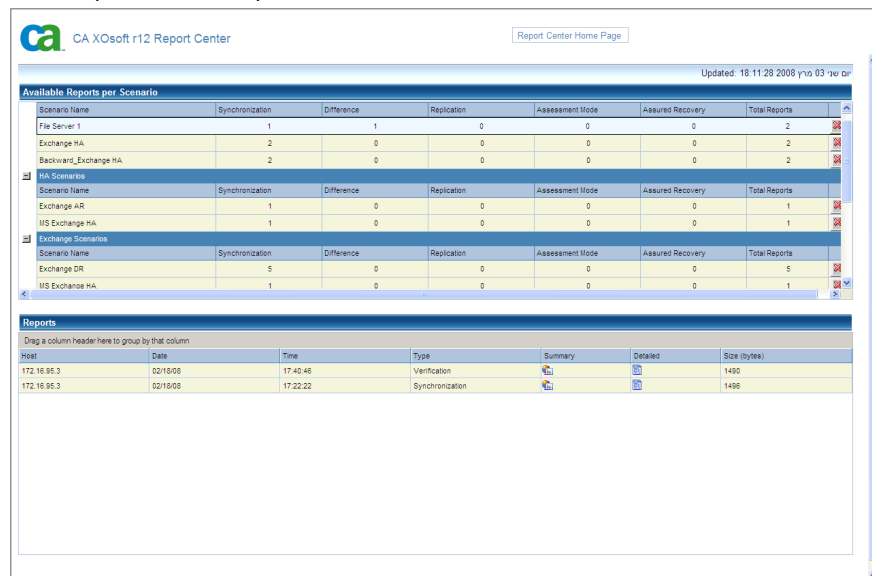
To view a report

- 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.




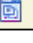
- From the **Tools** menu, select the **Reports** option and then **Show Scenario Reports**.

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.

Reports						
Drag a column header here to group by that column						
Host	Date	Time	Type	Summary	Detailed	Size (bytes)
172.16.95.3	01/17/08	14:50:37	Synchronization			1513

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 opens.


CA XOsoft Report Center

[Report Center Home Page](#)

CA XOsoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server 1
Master host	130.119.185.152(1)
Replica host	130.119.185.153(2)
Scenario start time	18-Mar-09 12:18:09
Report start time	18-Mar-09 12:21:02
Report finish time	18-Mar-09 12:21:21


Summary:

Total number of files modified	54
Total number of bytes changed	12.76MB
Total number of directories created	13
Total number of files with different security attributes	1

Deleting Reports

The reports that are displayed in the Report Center are stored for the period that is defined in their scenario properties, under the [Report Handling group](#) (see page 144). 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.

A confirmation message appears, asking you whether you want to delete the row.


3. Click **OK** on the confirmation message.

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 XOssoft 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, and the number of synchronized files that have different Windows security properties.



[Report Center Home Page](#)

CA XOssoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server 1
Master host	130.119.185.152(1)
Replica host	130.119.185.153(2)
Scenario start time	18-Mar-09 12:18:09
Report start time	18-Mar-09 12:21:02
Report finish time	18-Mar-09 12:21:21


Summary:

Total number of files modified	54
Total number of bytes changed	12.76MB
Total number of directories created	13
Total number of files with different security attributes	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.


CA XOssoft Report Center

[Report Center Home Page](#)

CA XOssoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	File Server 1
Master host	130.119.185.152(1)
Replica host	130.119.185.153(2)
Scenario start time	18-Mar-09 16:29:14
Report start time	18-Mar-09 16:29:23
Report finish time	18-Mar-09 16:29:30

EVENT	BYTES	TIME STAMP	FILE NAME
Create	39Bytes	01-Feb-07 12:22:52	C:/Tools/Bginfo/bginfo.bat
Create	6.84KB	28-Jul-06 08:32:44	C:/Tools/BlueScreen/Eula.txt
Remove			C:/Tools/BlueScreen/Eula12.txt
Create	70.37KB	16-Aug-06 15:39:10	C:/Tools/ProcessExplorer/proccexp.chm
Create	3.46MB	01-Nov-06 13:07:34	C:/Tools/ProcessExplorer/proccexp.exe
Remove			C:/Tools/ProcessMonitor/Procmon12.exe
Create Directory	0Bytes	01-Feb-07 18:47:28	C:/Tools/Sysprep/sysprep/386
Create Directory	0Bytes	01-Feb-07 18:47:28	C:/Tools/Sysprep/sysprep/386/Soem\$
Remove			C:/Tools/tcpview/Copy of TCPVIEW.HLP

Summary:

Total number of files/directories removed	3
Total number of files modified	4
Total number of bytes changed	3.53MB
Total number of directories created	2

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.

CA XOssoft Report Center

Report Center Home Page

CA XOssoft Replication

REPLICATION REPORT

Scenario	File Server 1
Host	130.119.185.153(2)
Scenario start time	18-Mar-09 16:47:03
Report start time	18-Mar-09 16:47:03
Report finish time	18-Mar-09 17:00:07

Summary:
Directory C:/Tools:

replicated bytes	38.27MB
changed files	162
removed files	0
renamed files	0
last change	18-Mar-09 16:49:06

Total number of bytes changed 38.27MB

Total statistics:
Directory C:/Tools:

replicated bytes	38.27MB
changed files	162
removed files	0
renamed files	0
last change	18-Mar-09 16:49:06

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

- Generating [Replication Reports for the Master host](#) (see page 168).
- Generating [Replication Reports for the Replica hosts](#) (see page 175).

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.

CA XOsoft Report Center Report Center Home Page

Updated: Thursday, March 19, 2009 7:00:25 PM

Available Reports per Scenario							
Scenario Name	Synchronization	Difference	Replication	Assessment Mode	Assured Recovery	CDP	Total Reports
Control Service HA	1	0	0	0	0	0	1
Exchange Server HA	3	0	0	0	0	0	3
Backward Exchange Server HA	1	0	0	0	0	0	1

Reports							
Drag a column header here to group by that column							
Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)
130.119.185.152	Changes found	Today	18:19:59	Synchronization			2403

2. On the **Reports** table, select the Synchronization Report you want to view.
3. The Synchronization Report for the Backward scenario opens.

CA XOsoft Report Center Report Center Home Page

CA XOsoft High Availability

SYNCHRONIZATION REPORT

Synchronization mode	BlockSynchronization (include files with the same size and modification time)
Scenario	Backward Exchange Server HA
Master host	130.119.185.153(1)
Replica host	130.119.185.152(2)
Scenario start time	19-Mar-09 18:19:30
Report start time	19-Mar-09 18:19:39
Report finish time	19-Mar-09 18:19:59

Summary:

Total number of files modified	5
Total number of bytes changed	20.03MB

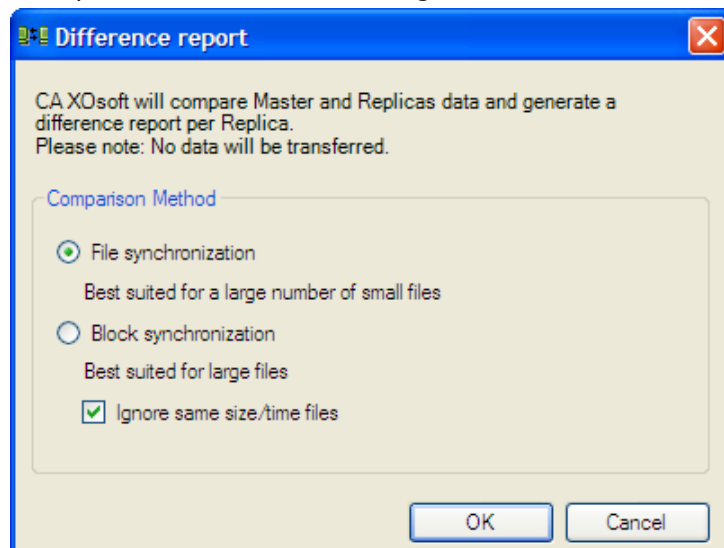
Difference Report

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, according to the instructions specified in [How Synchronization Works](#) (see page 19). Then, click **OK**.

At the end of the process, a Difference Report is generated for each replica, and all the reports are sent to the Report Center.

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

ca CA XOssoft Report Center

[Report Center Home Page](#)

CA XOssoft Replication

DIFFERENCE REPORT

Comparison mode	BlockSynchronization (ignore files with the same size and modification time)
Scenario	File Server
Master host	QA95-W2K3-EX1(1)
Replica host	QA95-W2K3-EX2(2)
Scenario start time	22-Mar-09 12:53:36
Report start time	22-Mar-09 12:54:18
Report finish time	22-Mar-09 12:54:24

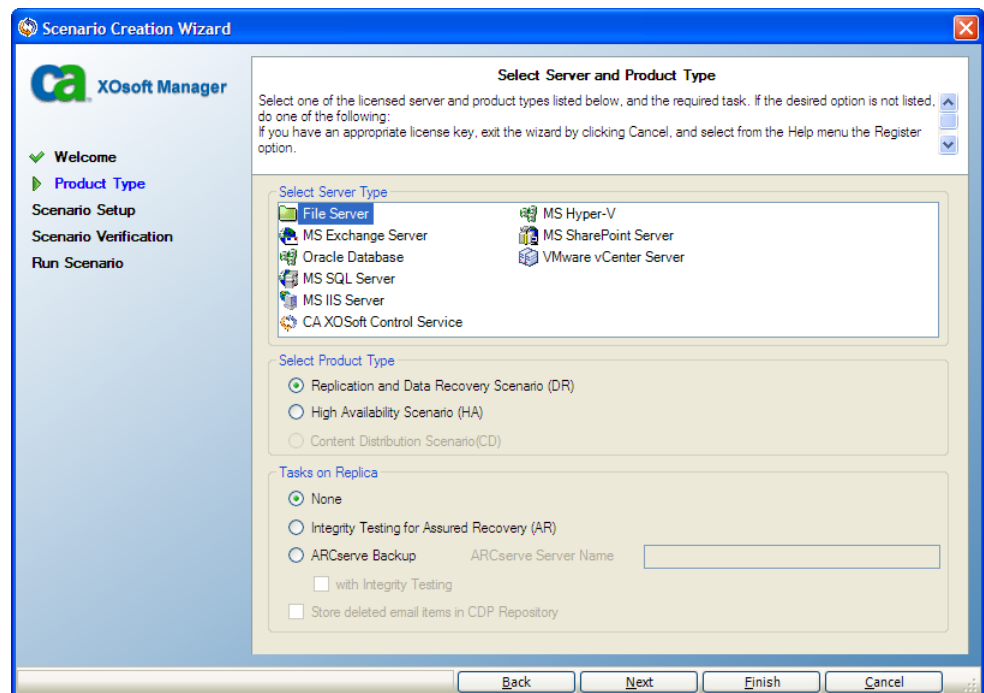
Summary:

Total number of files/directories removed	3
Total number of files modified	169
Total number of bytes changed	43.77MB
Total number of directories created	45

Chapter 6: Editing and Managing Scenarios and Hosts

This section 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.



This section contains the following topics:

[Define the Master and Replica Servers](#) (see page 104)

[Add Additional Replica Servers](#) (see page 105)

[Select Master Directories and Their Contents for Replication](#) (see page 105)

[Filter Master Directory Files](#) (see page 109)

[Synchronize Registry Keys](#) (see page 117)

[Auto-discover Database Files for all Databases](#) (see page 123)

[Select Replica Root Directories](#) (see page 125)

[Propagating Master Root Directories to Multiple Replica Hosts](#) (see page 127)

[Scenario Operations](#) (see page 129)

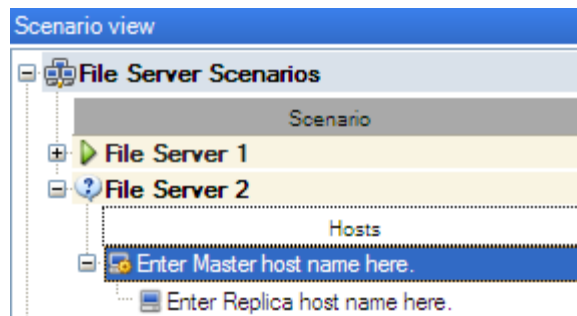
[Host Maintenance](#) (see page 132)

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.



Enter the hostname or IP address of the host.

2. Press the **Enter** key, or click anywhere outside of the text field.
3. 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, refer to [Select Master Directories and Their Contents for Replication](#) (see page 105).
- To define the Replica root directories, refer to [Select Replica Root Directories](#) (see page 125).


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 opens.



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.

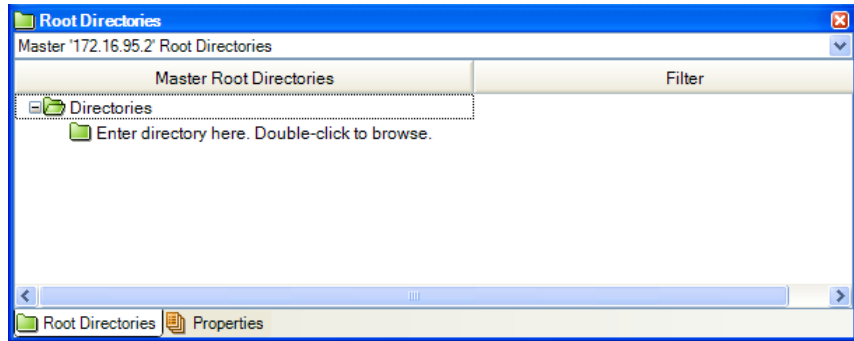
Notes:

- Working with **Master Root Directories** dialog is possible only if the CA XOsoft Engine is installed and running on the host.
- You can also select registry keys for synchronization, as described in [Synchronize Registry Keys](#) (see page 117).

Important! Special limitations apply to UNC paths (\\server\share) of remote root directories. This path type is not supported as a source (on the Master) for real-time replication. However, it can be the target for data replicated in real-time, meaning it can be used to store data on the Replica. In this case, these root directories can even support ACL replication.

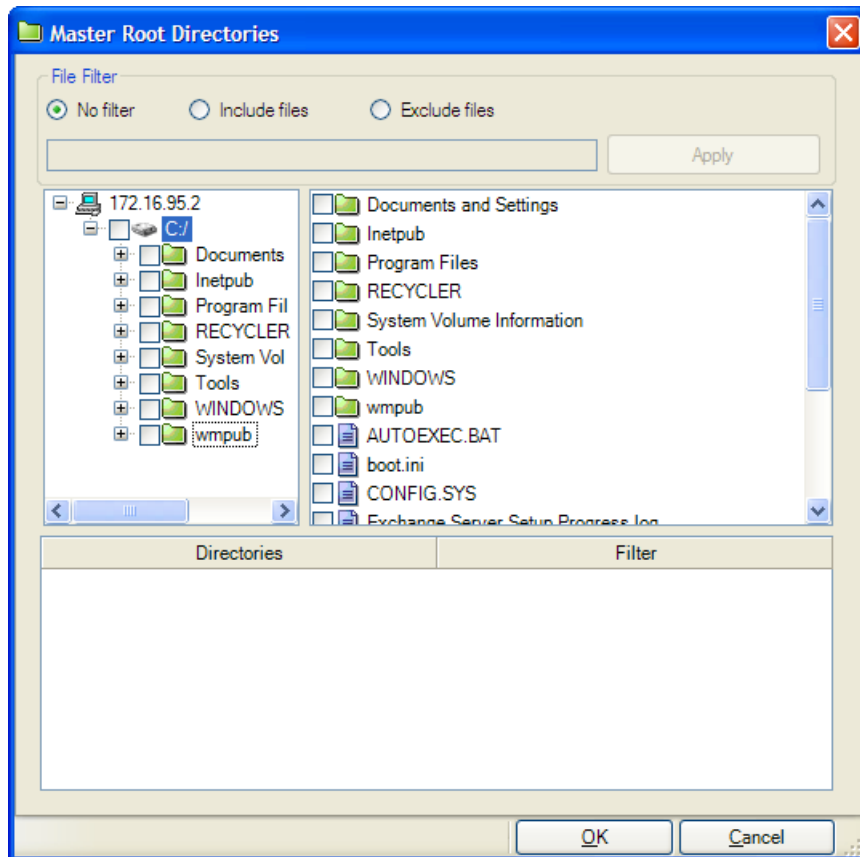
To select Master directories and their contents

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 opens 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 opens.

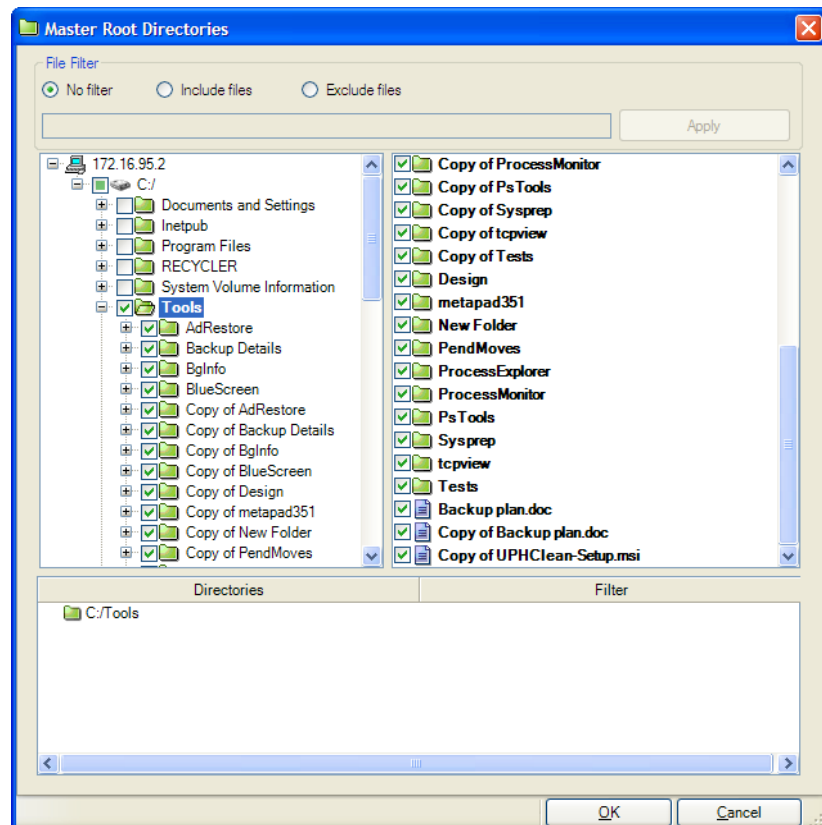


The **Master Root Directories** dialog has two areas. The left area shows only directories and sub-directories. The right area shows both directories and sub-directories, and files in those directories. The checkboxes are for you to select or clear. When selected, those directories or files will be replicated.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

- 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:

Notes:

- When you select root directories for the Master or Replica servers, the total character length of root directory plus subdirectory names should not exceed 1024 characters.
- If a root directory is a sub-directory, it remains bold and marked, and its parent directory is marked with a grayed checkmark.



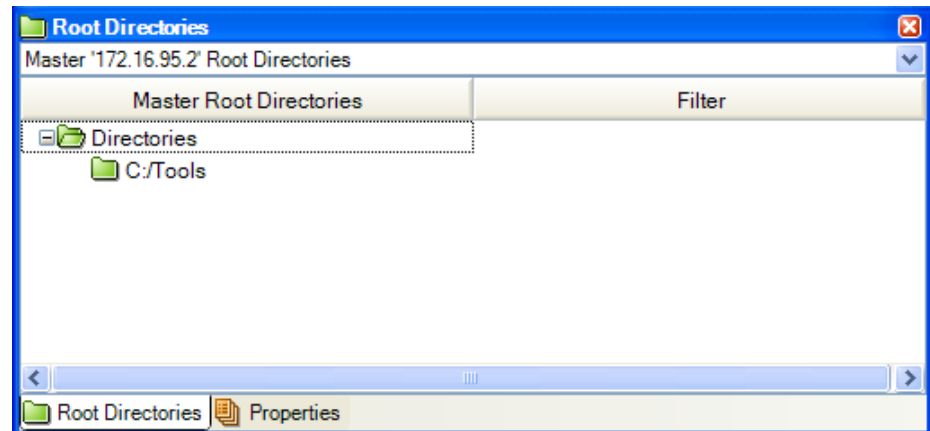
All files and sub-directories belonging to the directory that is highlighted in the left area are displayed in the right area.

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

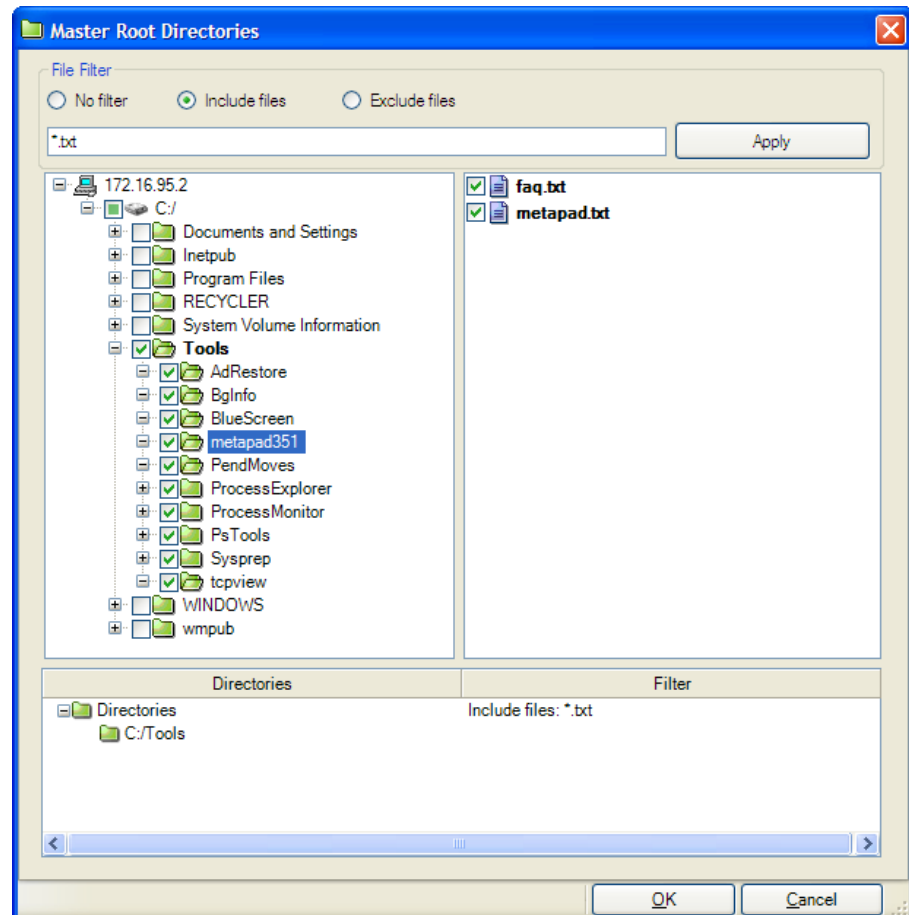
To remove a Master root directory

- 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:

Note: A "character" in this context refers only to alphabetical or numerical character.

- 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:

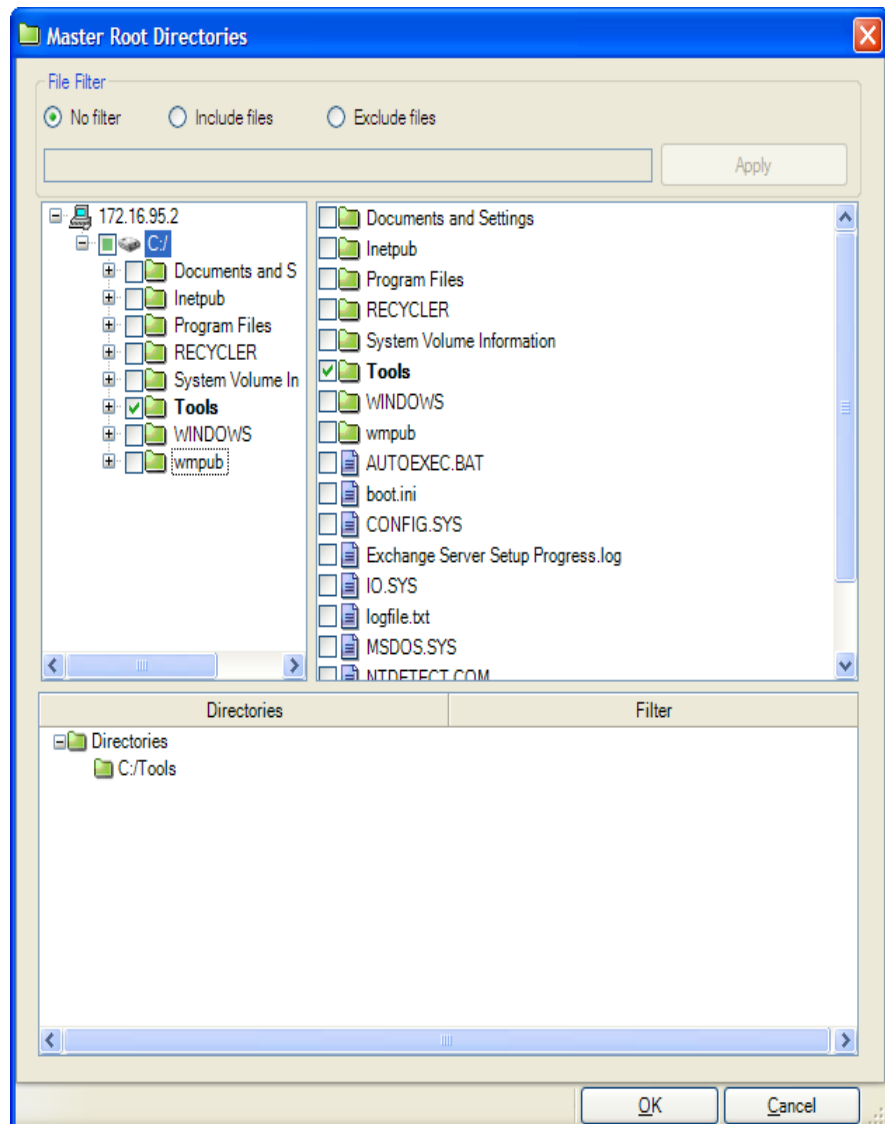
- **No filter** - all the directories and files you manually select will be replicated. This is the default option. Refer to [Select Master Directories and Their Contents for Replication](#) (see page 105).
- **Include files** - ONLY the selected files or file-types will be replicated. Refer to [Include Files](#) (see page 111).
- **Exclude files** - ONLY the selected files or file-types will be excluded from replication, and all others will be included. Refer to [Exclude Files](#) (see page 114).

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.

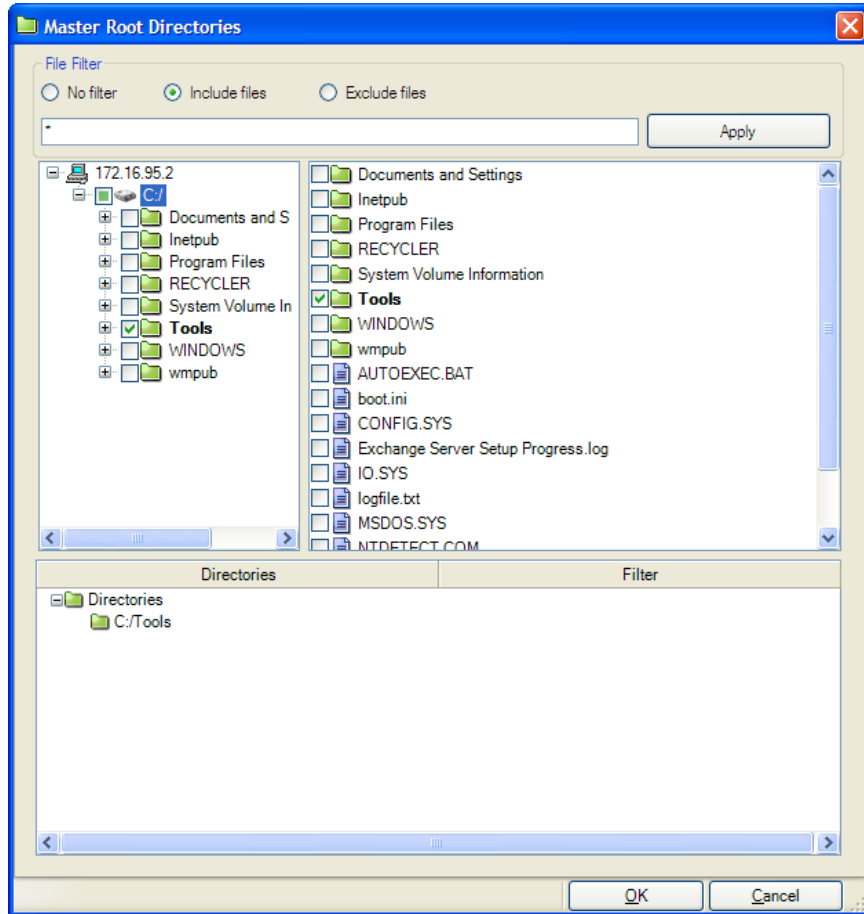
To include files

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.

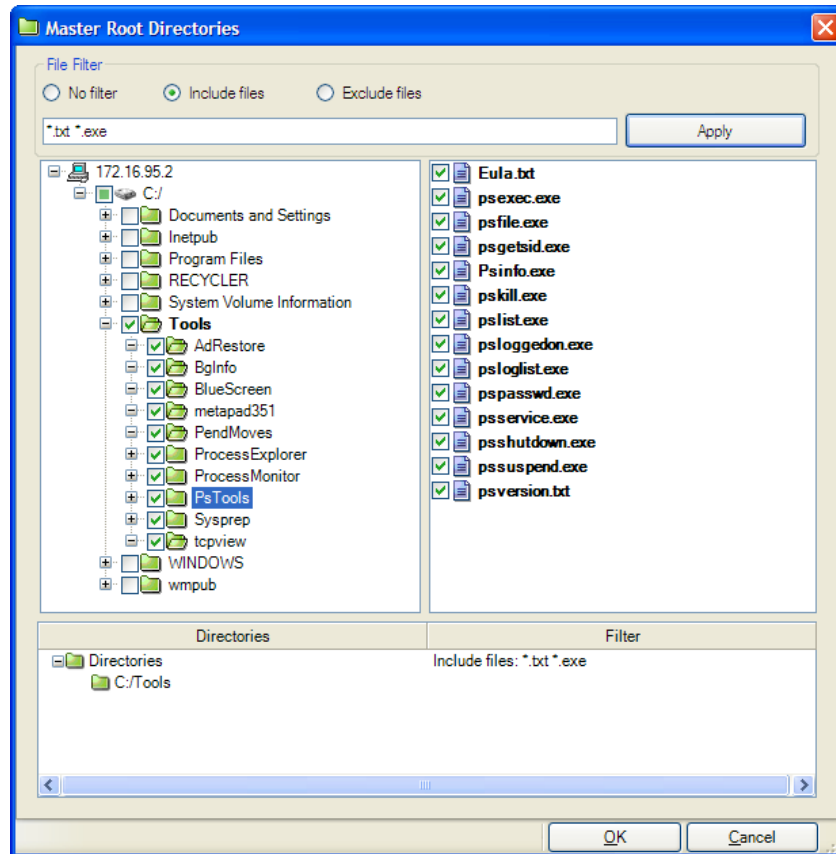
2. Click the **Include files** option button at the top of the **Master Root Directories** dialog. The Filter box is enabled with an asterisk (*) wildcard.



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 ("").

- 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.

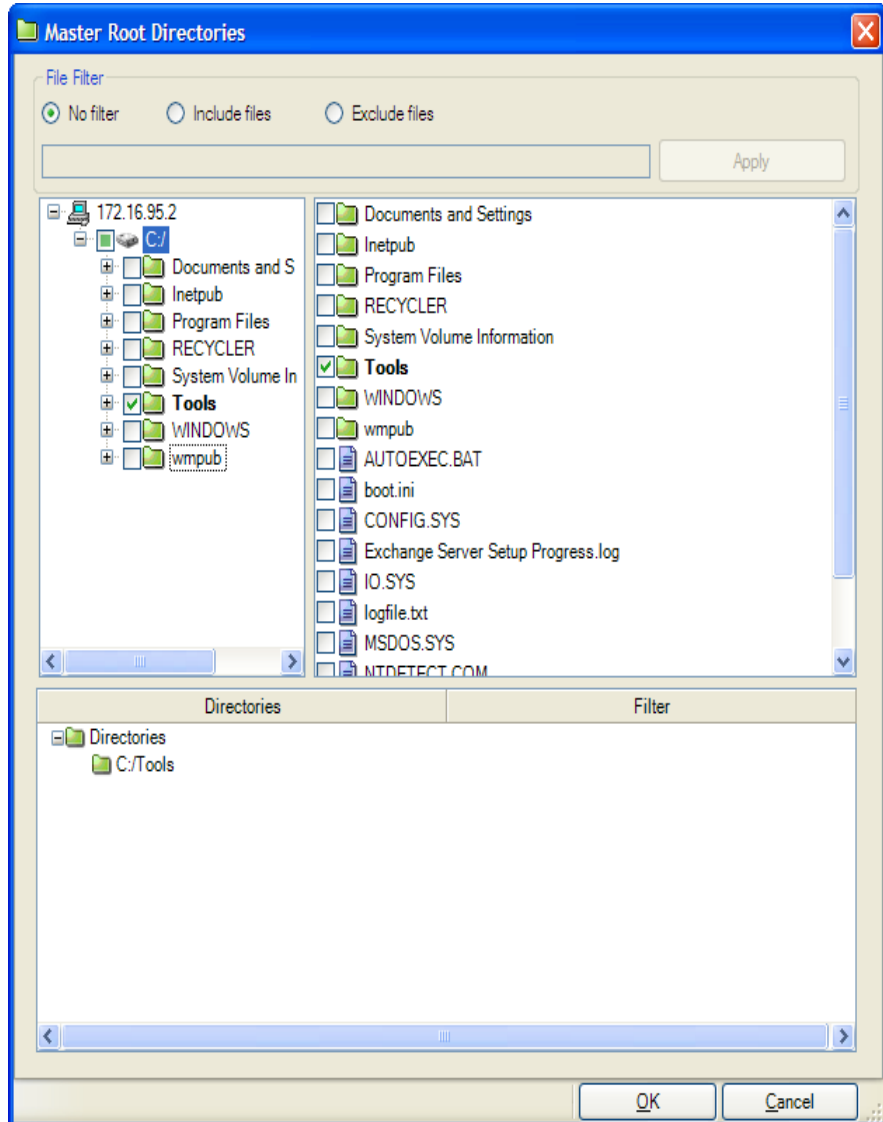
- [Optional] You can manually select or clear individual directories and files. This action overrides the **Include files** option regarding the individual directory or file.
- Click **OK** to save your root directory selection and close the **Master Root Directories** dialog.

Exclude Files

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

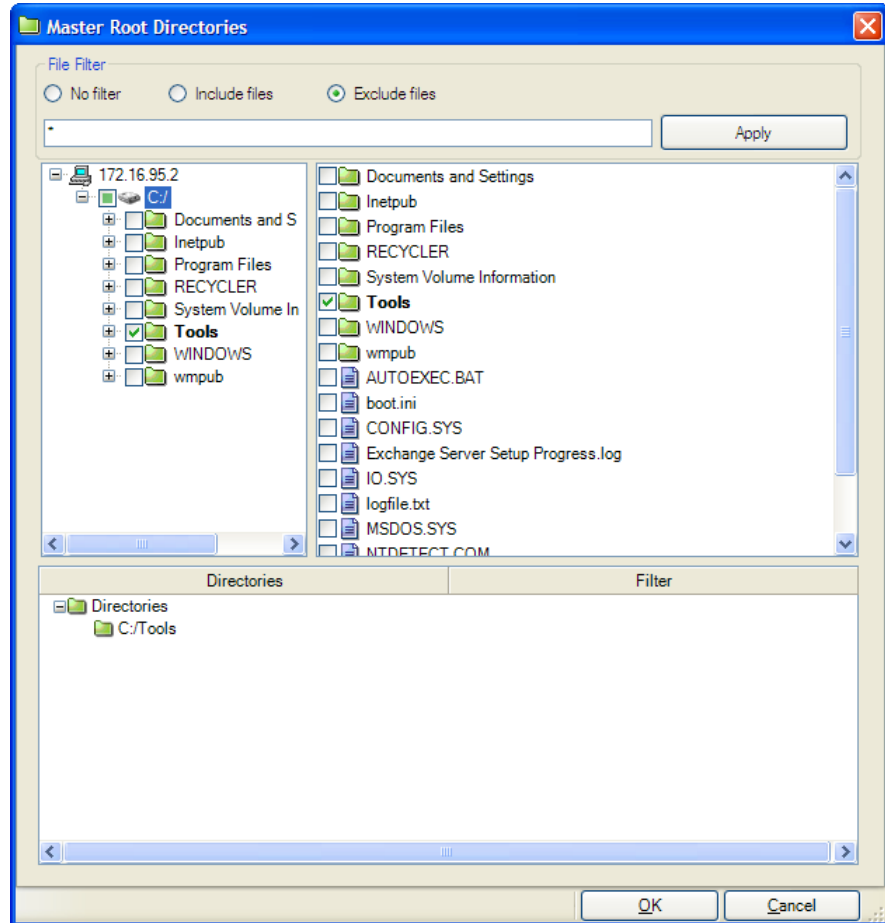
To exclude files

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.

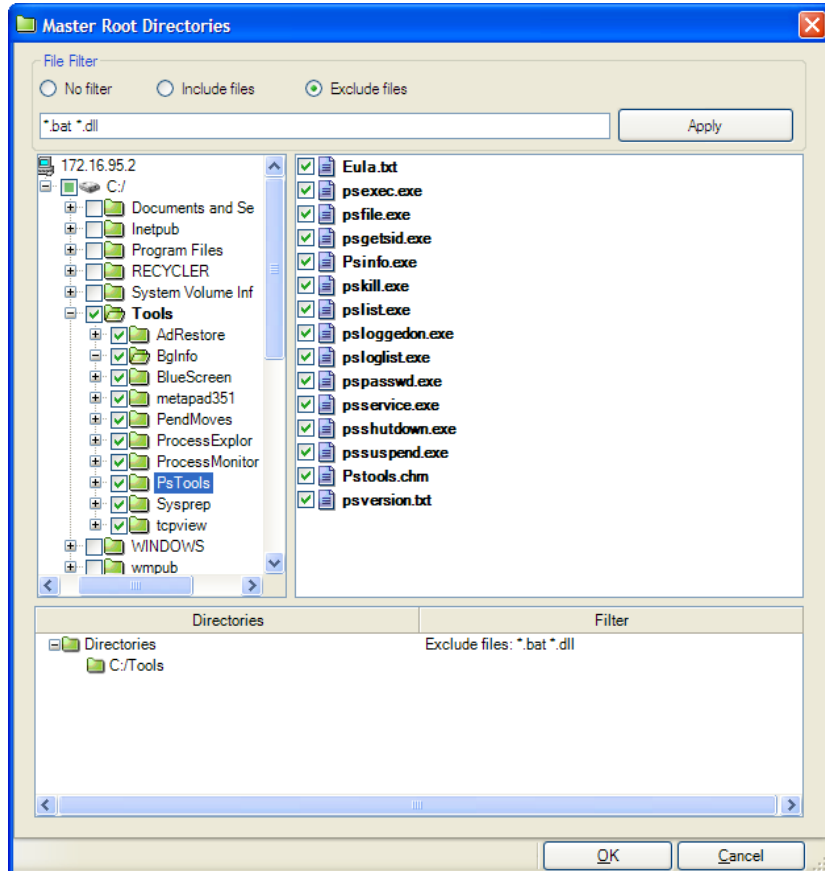
2. Click the **Exclude files** option button at the top of the **Master Root Directories** dialog. The Filter box is enabled with an asterisk (*) wildcard.



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.

Synchronize Registry Keys

In addition to synchronizing and replicating application data, CA XOssoft enables you to synchronize the Master and the Replica registry keys. Using the Registry Synchronization option, you can select which registry keys on the Master will be copied to the Replica, and define the synchronization frequency. You can copy the Master registry keys to the same location on the Replica, or you can change the name and storage path of the synchronized keys. If there are multiple Replica hosts in the replication tree, the registry synchronization process will be applied to all of them. The registry keys are not replicated in real time. They are copied from the Master to the Replica on a scheduled basis, according to the frequency you defined.

Important! Use this feature with caution. Changing registry keys may result in system failure.

Notes:

- This feature does not apply to applications that block the access to their registry keys, or to applications whose registry keys cannot be altered.
- By default, the Registry Synchronization option is disabled.

There are several steps in configuring and running the Registry Synchronization option:

1. [Activating the Registry Synchronization property](#) (see page 118).
2. [On the Master host, selecting the registry keys to be synchronized](#) (see page 119).
3. [Optional] [On the Replica host, selecting the name and storage location for the synchronized registry keys](#) (see page 121).
4. [Running the scenario to start the registry keys synchronization](#) (see page 67).

Activate the Registry Synchronization Option

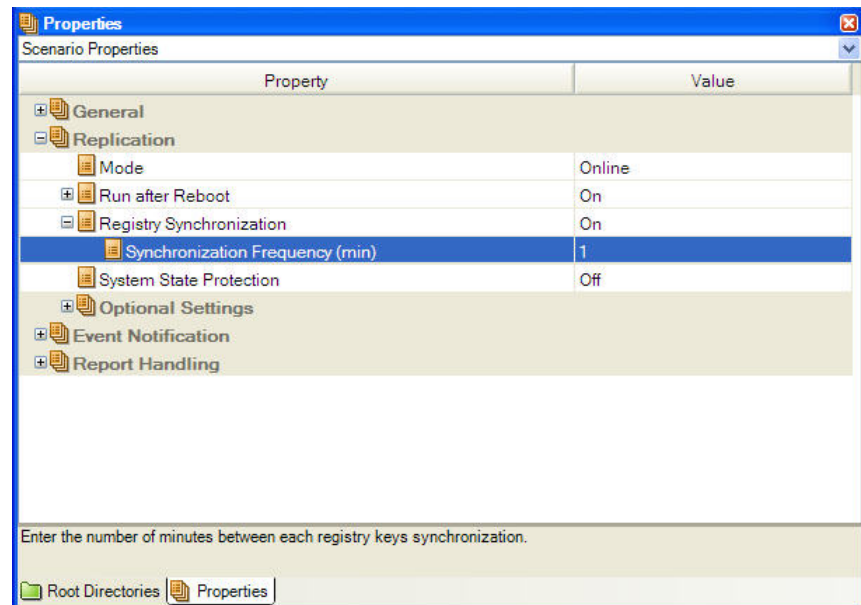
The first step in configuring and running the Registry Synchronization option is activating this option and defining its frequency.

Note: To configure Registry Synchronization properties, the scenario must be stopped.

To activate the Registry Synchronization property

1. In the Scenario pane, select the name of the scenario for which you want to activate the **Registry Synchronization** property.
2. In the Properties pane, click the **Properties** tab at the bottom.
The **Scenario Properties** list appears in the pane.
3. Open the **Replication** group, select the **Registry Synchronization** property, and set its value to On.

The **Synchronization Frequency** property appears under the **Registry Synchronization** property.



4. In the **Synchronization Frequency** value box, enter the number of minutes that will pass between each registry keys synchronization.
5. Save your configuration by clicking the **Save** button on the Standard toolbar.

Now, you need to [select on the Master host the registry keys that will be synchronized](#) (see page 119).

Select Registry Keys for Synchronization

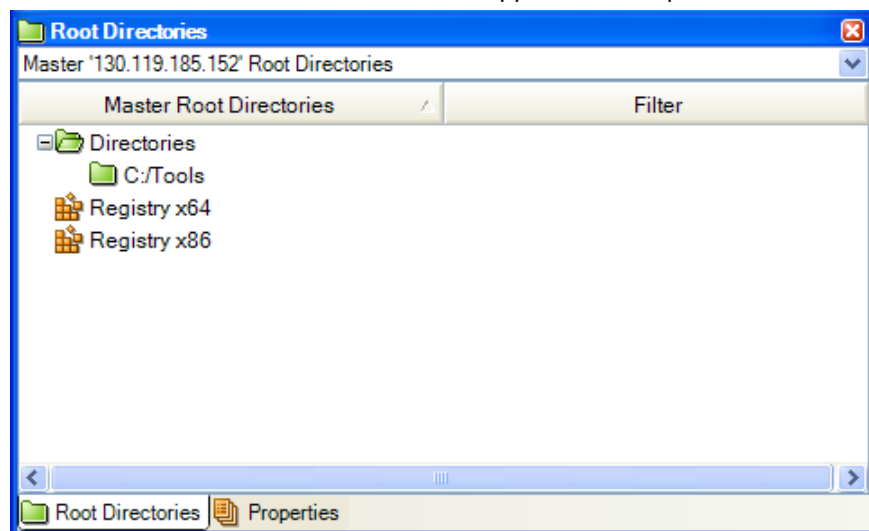
The second step in configuring and running the Registry Synchronization option is selecting on the Master host the registry keys that you want to synchronize.

Notes:

- To select registry keys for synchronization, the scenario must be stopped.
- You cannot select registry keys for synchronization through the Scenario Creation wizard, during the creation of a new scenario.
- Only keys are displayed for selection. You cannot select specific values for synchronization.

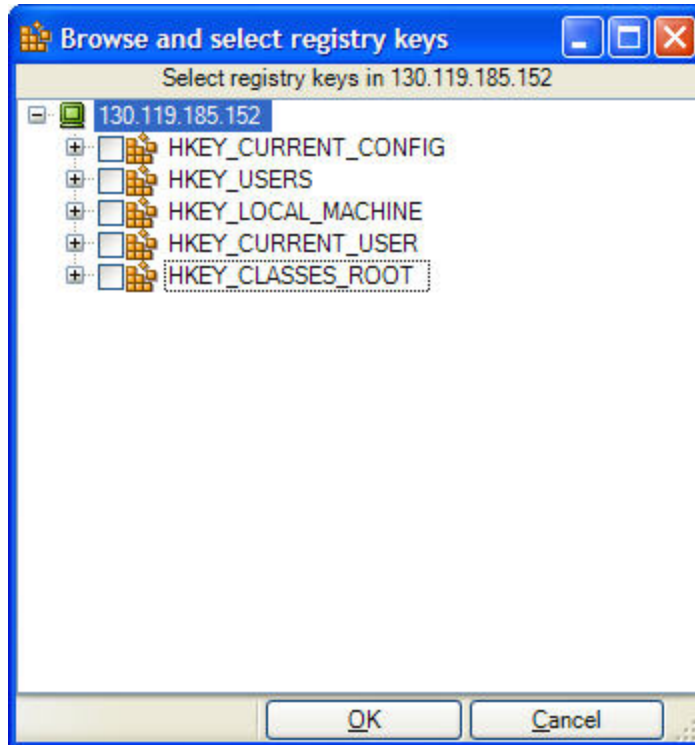
To select registry keys for synchronization on the Master

1. In the Scenario pane, select the name of the Master host whose registry keys you want to synchronize.
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 Registry Keys**. Alternatively, double-click the **Registry** folder that corresponds to your operation system - either **x86** or **x64**.

The **Browse and select registry keys** dialog opens.



The **Browse and select registry keys** dialog displays the Master host's registry keys list. The checkboxes are for you to select or clear. When selected, those registry keys will be synchronized. Those not selected are ignored.

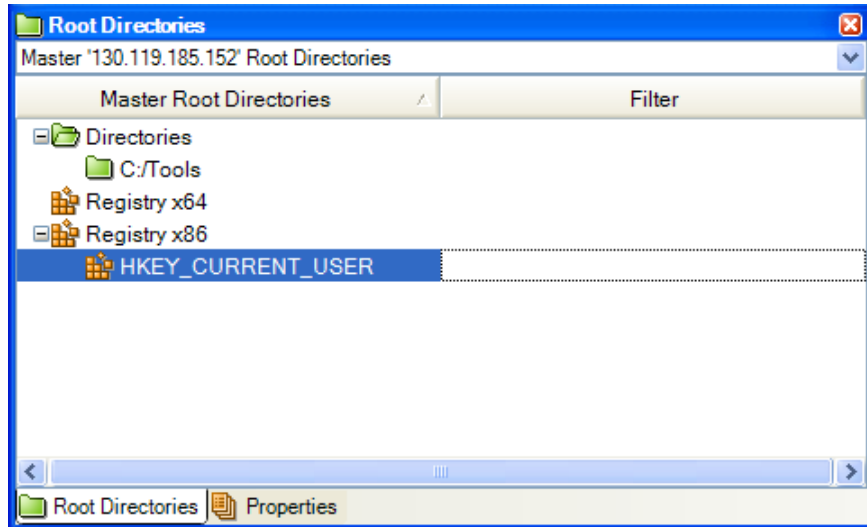
4. Select the checkboxes of the registry keys you want to synchronize, by clicking on the relevant checkboxes.

Notes:

- You cannot filter registry keys selection.
- If you are manually entering a name and a path of a registry key that does not exist on the Master, the scenario verification may be successful, but the scenario will stop running and an error message will be issued. You should only enter the details of existing registry keys for replication.

- When you have finished selecting all the registry keys you want to synchronize, click **OK**.

The selected registry keys now appear in the Root Directories pane under the **Master Root Directories** column.



- Save your configuration by clicking the **Save** button on the Standard toolbar.

By default, the system automatically configures the Replica registry keys to be the same as the selected Master registry keys. If you want to change the name and storage location of the synchronized Replica registry keys, follow the instructions described in [the next section](#) (see page 121).

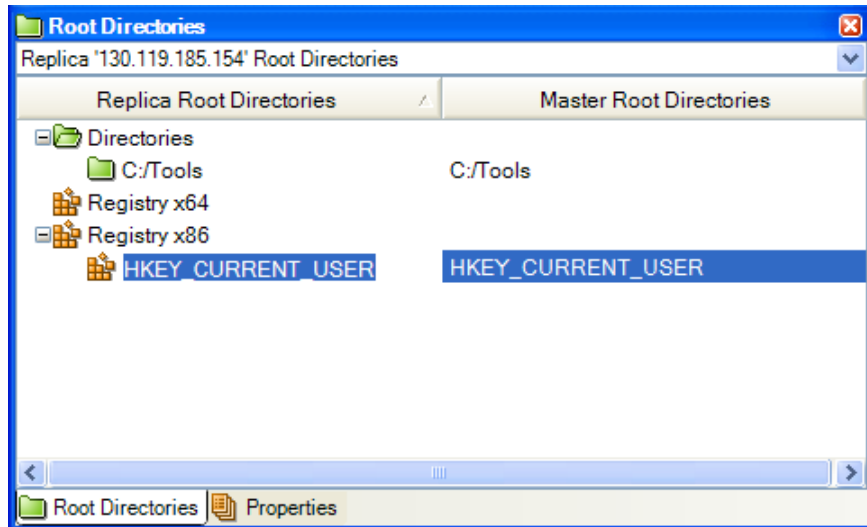
Select the Name and Storage Location of the Synchronized Registry Keys

The third step in configuring and running the Registry Synchronization option is selecting on the Replica host the name and storage location of the synchronized registry keys. Since by default the system configures the Replica registry keys to be the same as the selected Master registry keys, this step is optional.

To select the path for storing the synchronized registry keys

- In the Scenario pane, select the name of the Replica where you want to store the synchronized registry keys.

- In the Properties pane, click the **Root Directories** tab at the bottom. The Replica Root Directories information appears in the pane.



The registry keys that were selected on the Master for synchronization, appear on the Replica in the same location and under the same name.

- You can change the default path and name of the Replica registry keys in two ways:
 - Replacing the default path and name with the path and name of existing registry keys:
 - Right-click anywhere in the pane, and select **Browse and Select Registry Keys**. Alternatively, double-click the name of the specified Replica register key.
The **Browse and select registry keys** dialog appears.
 - Select the checkboxes of the registry keys you want use, and click **OK** to save your selection.
 - Manually entering new path and name for the default values: double-click the registry key name on the Root Directories pane, and manually enter a new path and name.
- Save your configuration by clicking the **Save** button on the Standard toolbar.

To start the registry keys synchronization, you need to [run the scenario](#) (see page 67).

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.

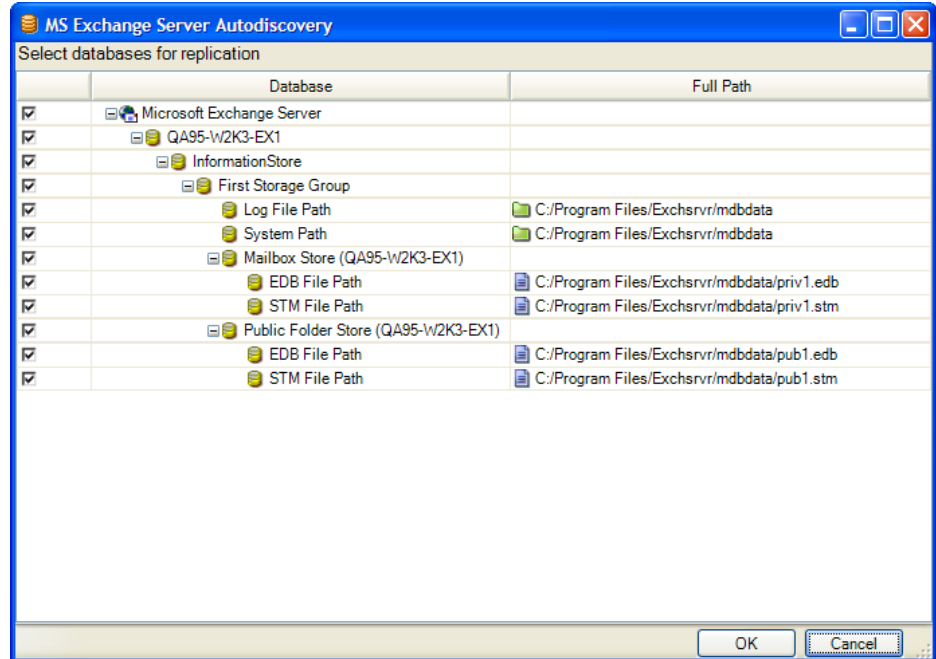
The **Auto-discovered** database files icon appears in the Master Root Directories tab.



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 opens.



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

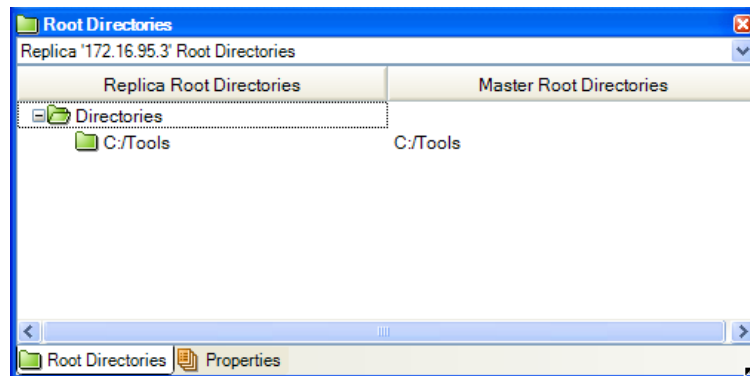
You must select the Master directories before this function becomes available. For each Master root directory, you must define a Replica root directory on each related Replica.

Important! Special limitations apply to UNC paths (\\server\share) of remote root directories. This path type is not supported as a source (on the Master) for real-time replication. However, it can be the target for data replicated in real-time, meaning it can be used to store data on the Replica. In this case, these root directories can even support ACL replication.

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.

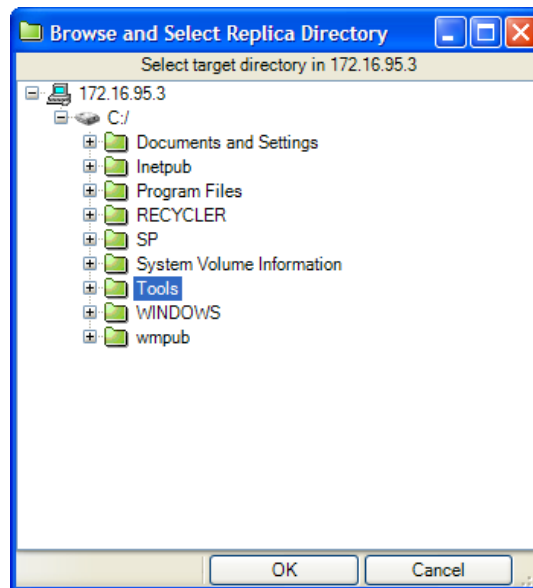


Important! The Scenario Creation Wizard automatically configures the Replica root directories to be the same as the Master root directories. If you want to keep this configuration, ensure that your Replica server has the same drive letters as the Master server, and that the selected directories on the 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 XOssoft connects to the Replica server and opens the **Browse and Select Replica Directory** dialog opens.



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 XOssoft 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 XOssoft creates it when the replication starts.

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

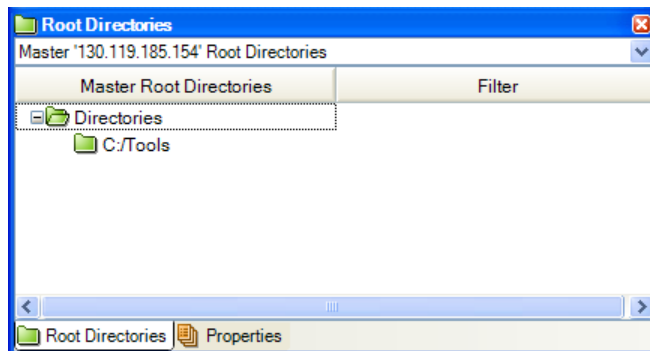
Propagating Master Root Directories to Multiple Replica Hosts

CA XOssoft enables you to propagate the root directories you set for the Master to multiple Replica hosts at once. Instead of separately configuring the root directories of each Replica host, in a click-of-a-button you can distribute the root directory of one Master to as many Replica hosts as you want. This option is particularly useful for a scenario that has many Replica hosts.

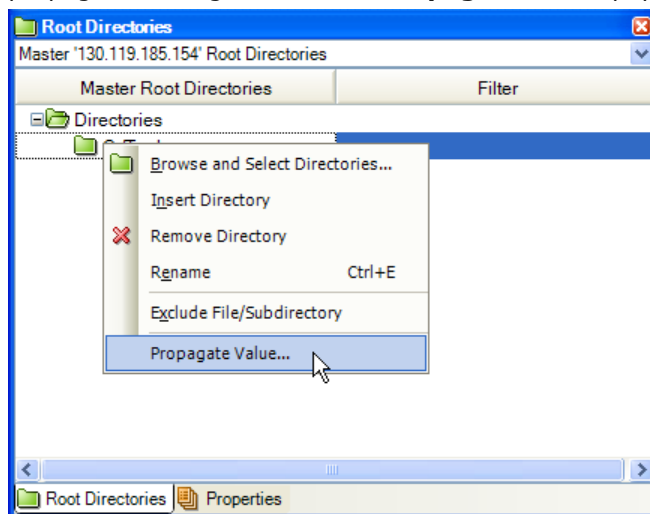
Note: The scenario must be stopped to apply root directory changes.

To propagate root directories

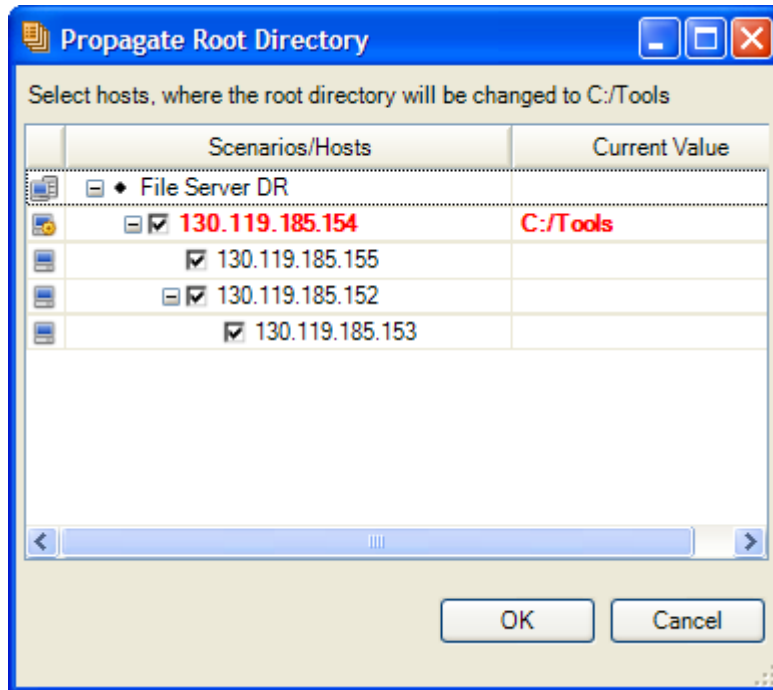
1. On the Scenario pane, select the Master whose root directories you want to propagate.
2. On the Framework pane on the left, click the **Root Directories** tab at the bottom. The Master Root Directories information appears in the pane.




3. On the Master Root Directories pane, select the root directory you want to propagate and right-click. The **Propagate Value** pop up command opens.



- Click the **Propagate Value** command. The **Propagate Value** dialog opens.



The Master and all Replica hosts in the scenario appear in the dialog, along with their selected root directories. The root directory you selected for propagation is displayed above the **Scenarios/Hosts** table, and in the **Current Value** column marked in red.

- To propagate the root directory to all Replica hosts, click **OK**.
Note: To exclude hosts from the root directory value propagation, clear their check boxes, and then click **OK**.
- After the **Propagate Value** dialog is closed, click the **Save**  button on the Standard toolbar to save and apply your changes to all hosts.

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

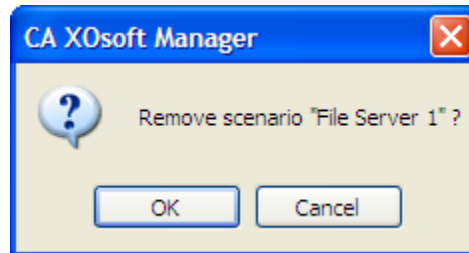
- On the Scenario pane, select the scenario and click the **Save** icon, or select the **Save** option from the **Scenario** menu.
- or -
- Click the **Save All** icon or select **Save All** from the **Scenario** menu, to save all scenarios on the Manager.

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.

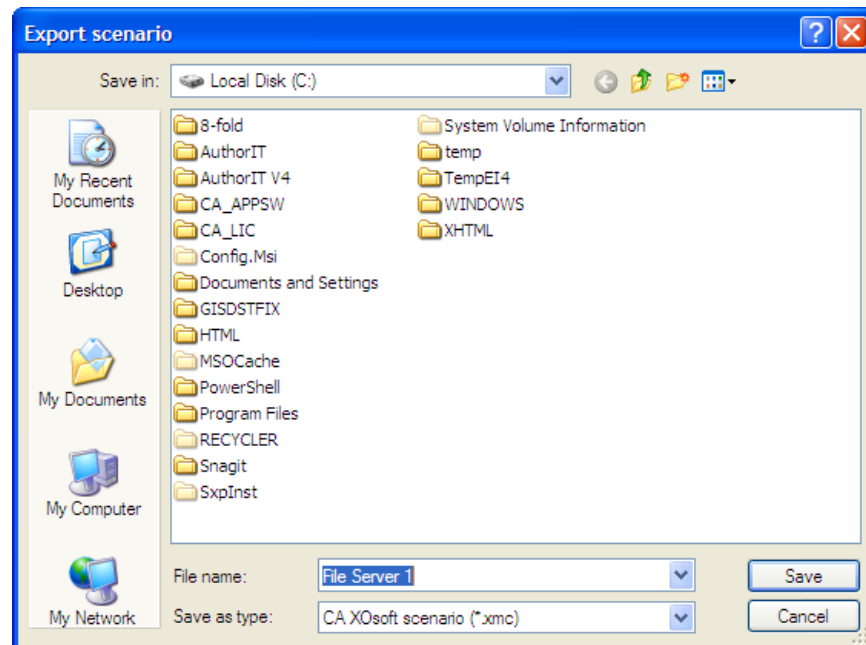
Export Scenarios

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.

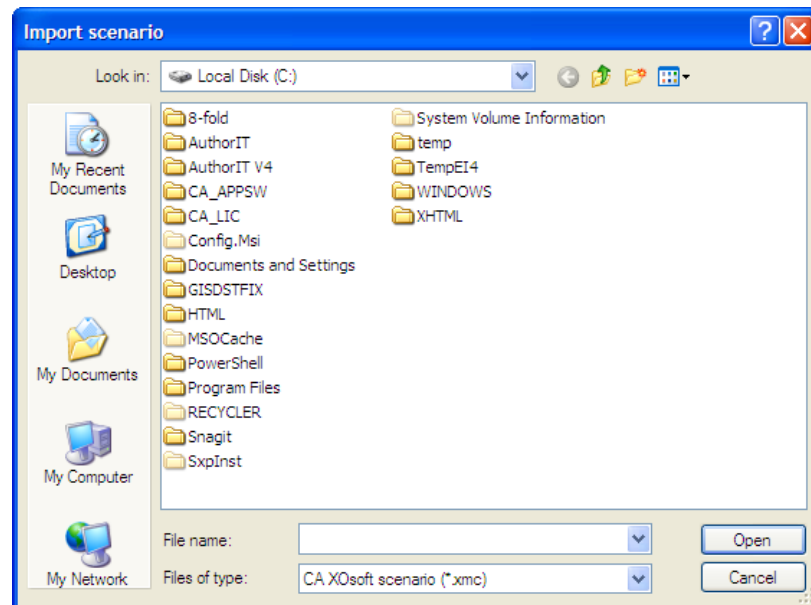
Import Scenarios

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.

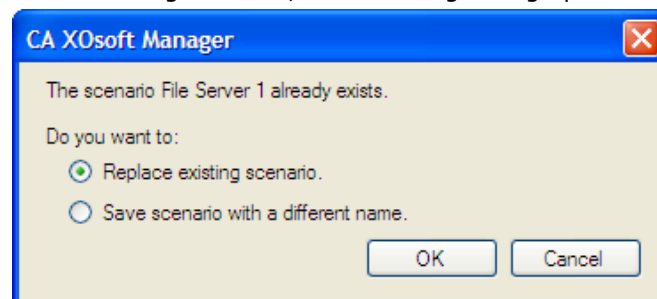
An **Import scenario** dialog opens.



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 opens.



3. 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.

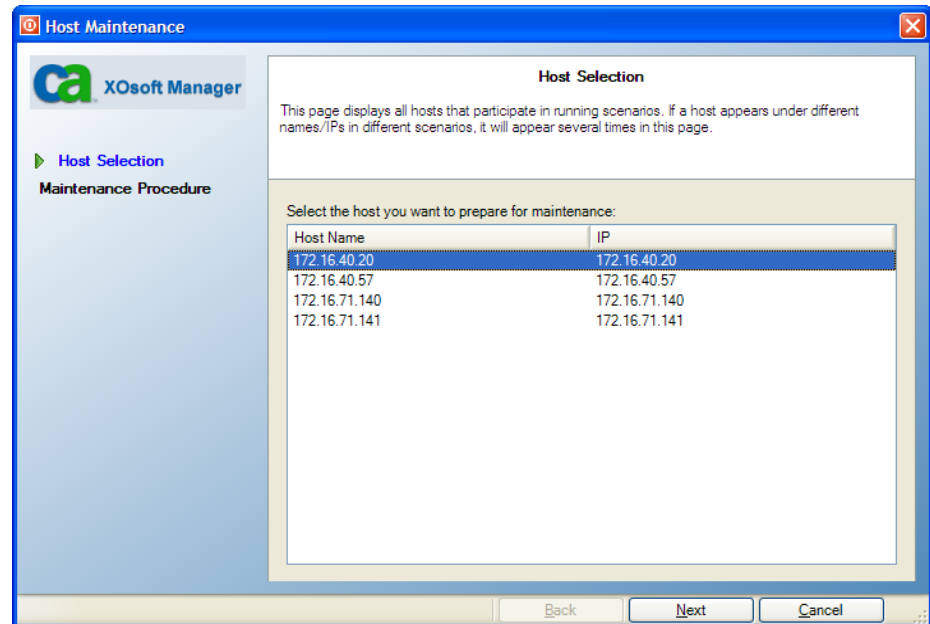
To prepare your hosts for maintenance procedures

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 opens.

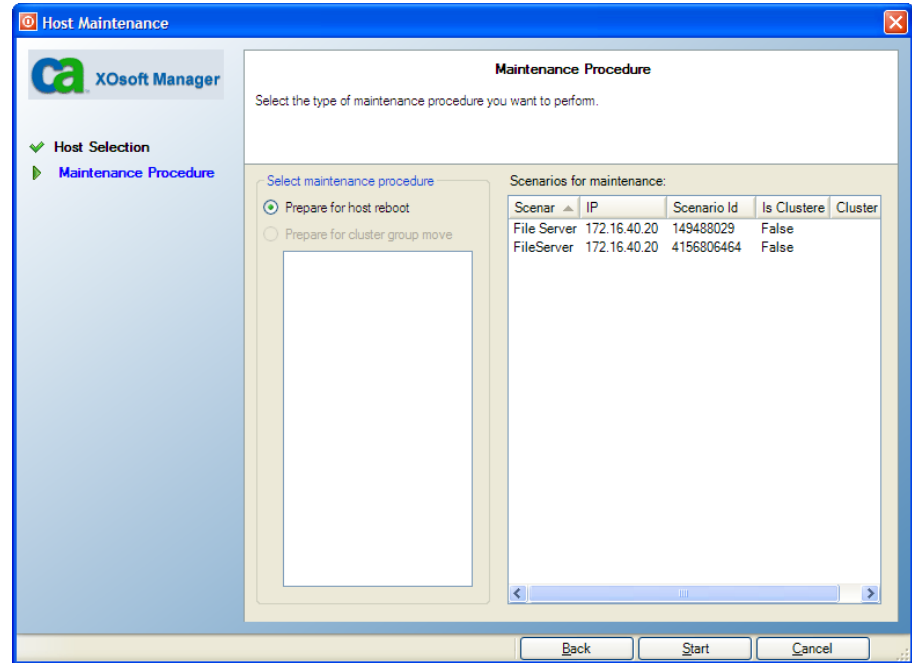


The **Host Maintenance** wizard displays all hosts that participate in the running scenarios.

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 opens.



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

- 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.**

File Server Scenarios					
Scenario	State	Connectivity	Product	Mode	
File Server	Ready for HM	DR		Online	
Hosts					
	Changed	Synchronized	Files	In spool	
172.16.40.20	213.20 MB	390.41 MB	6810	0 Bytes	
172.16.40.57	213.20 MB	390.41 MB	6810	0 Bytes	

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 opens.



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

This section contains the following topics:

[Configuring Scenario Properties](#) (see page 137)

[Understanding Scenario Properties](#) (see page 138)

[Schedule Synchronization](#) (see page 145)

[Protecting Your System State](#) (see page 149)

Configuring Scenario Properties

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

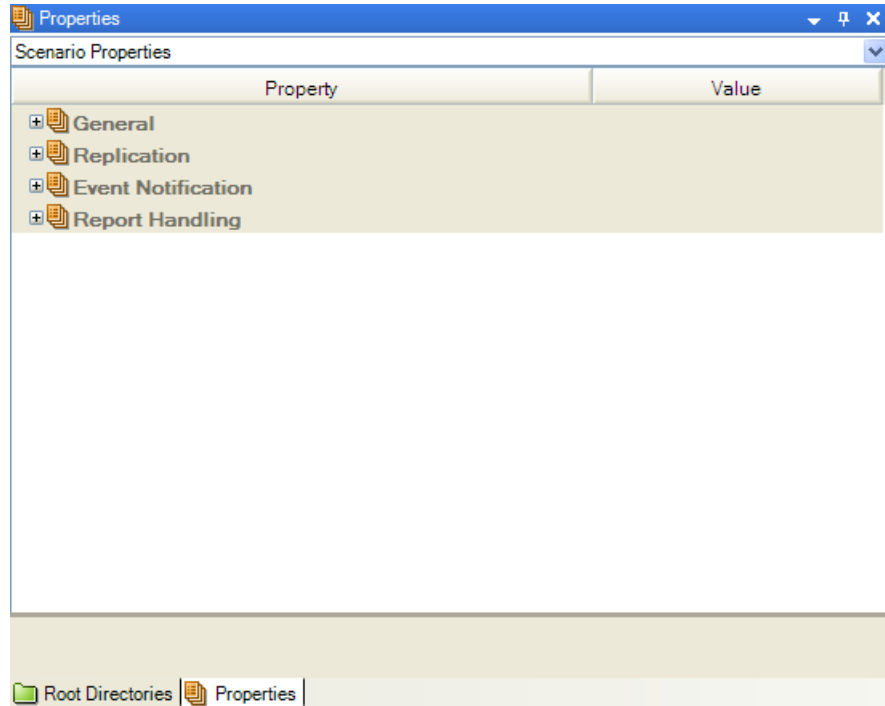
Notes:

- 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 opens.



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 XOssoft Manager, or CA XOssoft 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, refer to [Schedule Synchronization](#) (see page 145).

Run after Reboot

If the Master is rebooted, CA XOssoft 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.

Registry Synchronization

When set to On, this option enables you to synchronize the Master and the Replica registry keys on a scheduled basis.

For a detailed description of registry synchronization, refer to [Synchronize Registry Keys](#) (see page 117).

System State Protection

When set to On, this option enables you to save snapshots of the system and boot files of the Master on the Replica. To activate this option, you need to set the System State Protection schedule, and to define which Replica host will store the snapshots. For a detailed description, refer to [Protecting Your System State](#) (see page 149).

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, send an email notification, or write it to Windows event log.

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 XOssoft 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 XOssoft, which include the event details written in a notification file. Arguments entered here are static values.

Note: On Windows x64, you cannot run scripts that activate applications with UI.

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.

Note: On Windows x64, you cannot run scripts that activate applications with UI.

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:

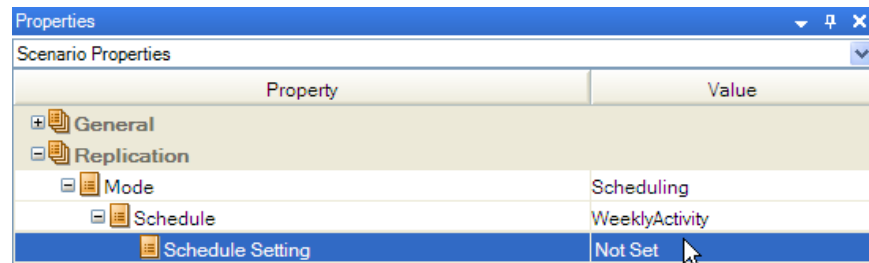
- Synchronization on selected days of the week and for specific hours in a 24 hour cycle.
- Synchronization over selected periods (e.g., once every 36 hours) in a 7 day cycle.
- Exclusion of specific dates.

To open the schedule

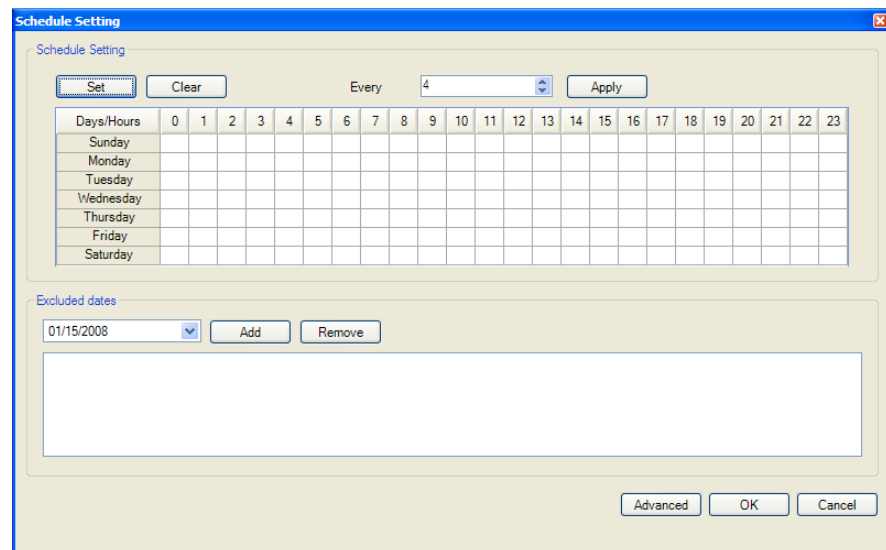
1. On the Scenario Properties list, open the **Replication** group. On the **Mode** property, select the **Scheduling** value.

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.

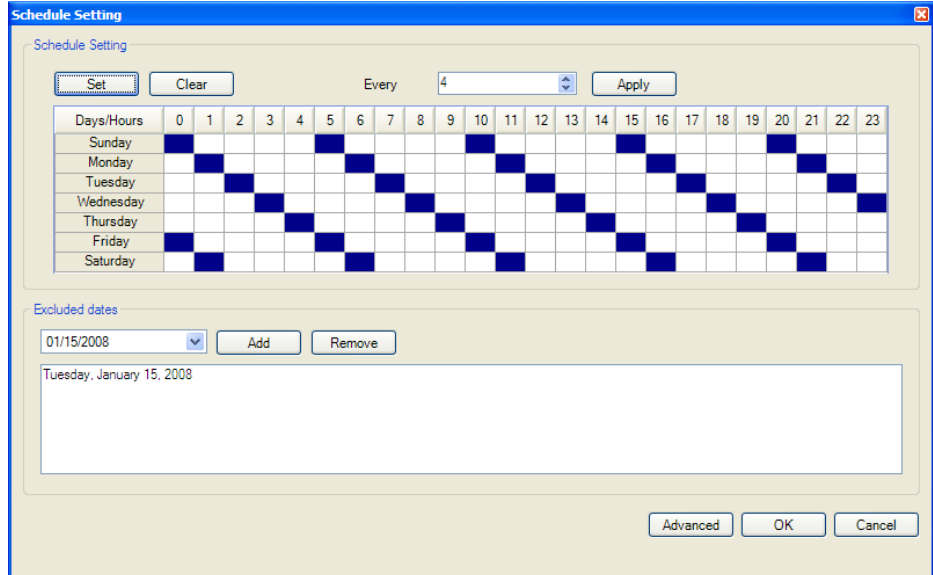


The **Schedule Setting** dialog appears:



- 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.

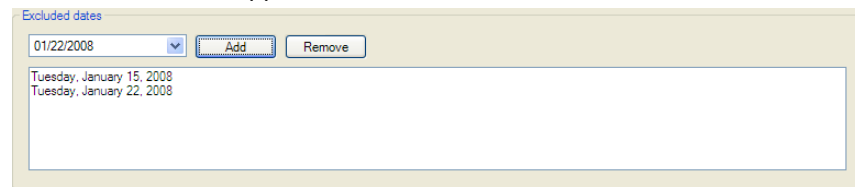
Exclude Dates from Scheduled Synchronization

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.



The screenshot shows a dialog box titled "Excluded dates". At the top, there is a text input field containing "01/22/2008", followed by a dropdown arrow, an "Add" button, and a "Remove" button. Below this is a list box containing two entries: "Tuesday, January 15, 2008" and "Tuesday, January 22, 2008".

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.

The screenshot shows the 'Schedule Setting' dialog box with the 'Advanced' mode selected. The 'Schedule Setting' section includes a dropdown for 'Tuesday' and a time input field set to '19:15', with 'Set' and 'Clear' buttons. Below this is an 'Every' section with a time input field set to '01:30' and an 'Apply' button. The main area contains a grid of schedule entries for various days and times, with a scrollbar at the bottom. The 'Excluded dates' section has a date input field set to '01/15/2008', 'Add', and 'Remove' buttons, and a list box containing 'Tuesday, January 15, 2008'. At the bottom right are 'Standard', 'OK', and 'Cancel' buttons.

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.

Protecting Your System State

System State Protection allows you to include critical system-related components in your replication and high availability scenarios, so that you can recover these components after a failure. System State Protection is server-type independent, which means you can enable this feature in any CA XOssoft scenario. When System State Protection is enabled, information about the snapshot schedule and file management is stored directly in the scenario configuration file. Snapshot files are then transferred to all Replica servers configured under the Master, including any "child" Replicas, if its Retain System Snapshot property (Transferring Path) is enabled.

Note: Even if a child Replica's Retain System Snapshot property is enabled, a child Replica will not receive system state snapshots if this property on its parent is disabled.

System State Protection lets you:

- Configure System State Protection in the Scenario Creation Wizard
- Set periodic backup intervals
- Send the snapshot file to more than one replica
- Restore from a System State snapshot

The following components are always included when System State Protection is enabled:

- boot files
- COM+ Class Registration database
- Registry service

The following components are also included in System State Protection depending upon the operating system:

- Windows XP Professional and Windows 2000 Operating Systems
 - files under Windows File Protection (WFP)
- Windows Server 2003 Operating Systems
 - All files protected by WFP
 - performance counter configurations
 - Active Directory (ADSI) on systems that are domain controllers
 - SYSVOL directory replicated by File Replication Service (FRS) on systems that are domain controllers
 - Certificate server on systems that provide Certificate Authority
 - Cluster database on systems that are a node of a Windows cluster

- Windows Vista and Windows Server 2008 Operating Systems
 - Certificate Services database
 - Active Directory Domain Services (NTDS)
 - SYSVOL directory (FRS Writer)
 - Cluster Service information
 - Microsoft Internet Information Services (IIS) meta-directory (IIS Metabase Writer/IIS Configuration Writer)
 - System files under WFP (System Writer)

Note: For Windows Server 2003 and higher, system state snapshots are taken using System Writer. Refer to the Microsoft website for more information on Backing Up and Restoring System State Under VSS.

How to Configure System State Protection

By default, System State Protection is set to Off. You can enable System State Protection during scenario creation using the wizard, or you can enable it for existing scenarios using the Scenario Properties pane of the CA XOssoft Manager.

For either method, you need to perform the following tasks:

- Enable the System State Protection property for the scenario
- Set the snapshot schedule, as desired.
- Enable the Store System State Protection property on one or all Replicas involved in the scenario.

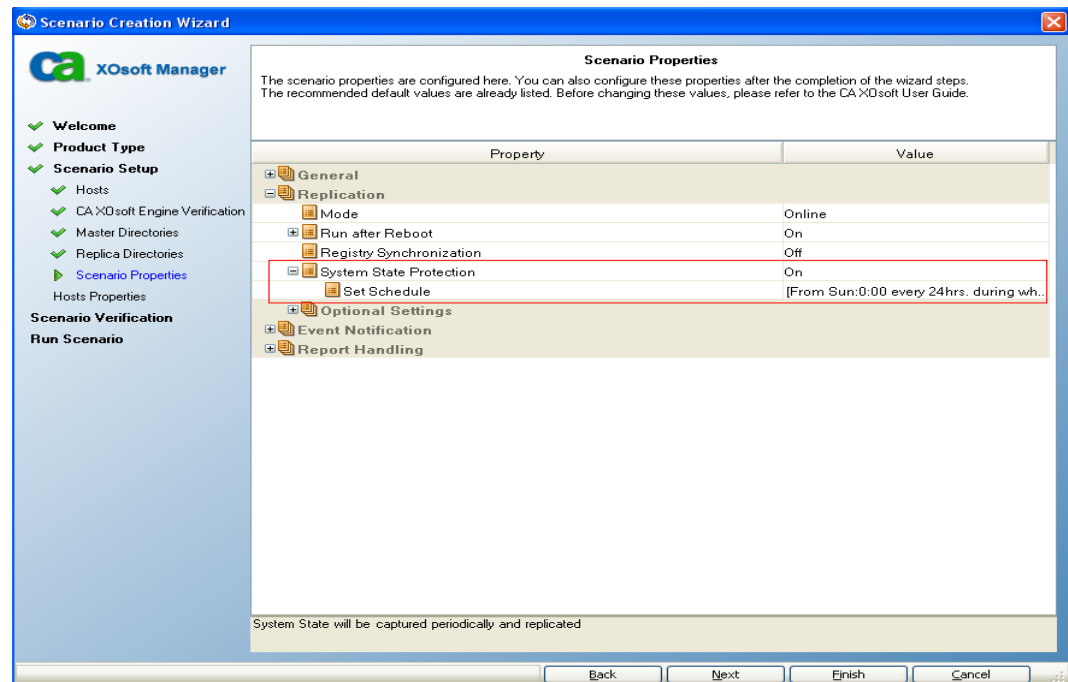
Note: After System State is restored on a Replica, that machine must be rebooted so the system state can take effect. You can set the Reboot After Recovery property to On, if desired.

Configure System State Protection in the Scenario Creation Wizard

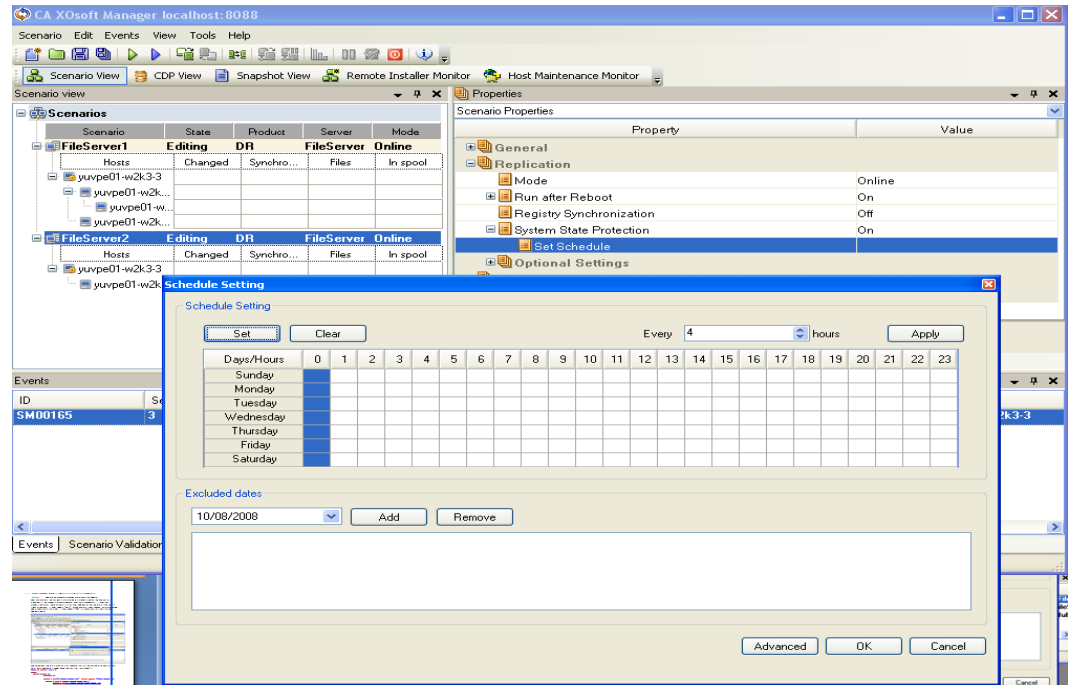
System State Protection can be enabled for any CA XOssoft supported server type, from directly inside the Scenario Creation Wizard.

To enable System State Protection in the Wizard

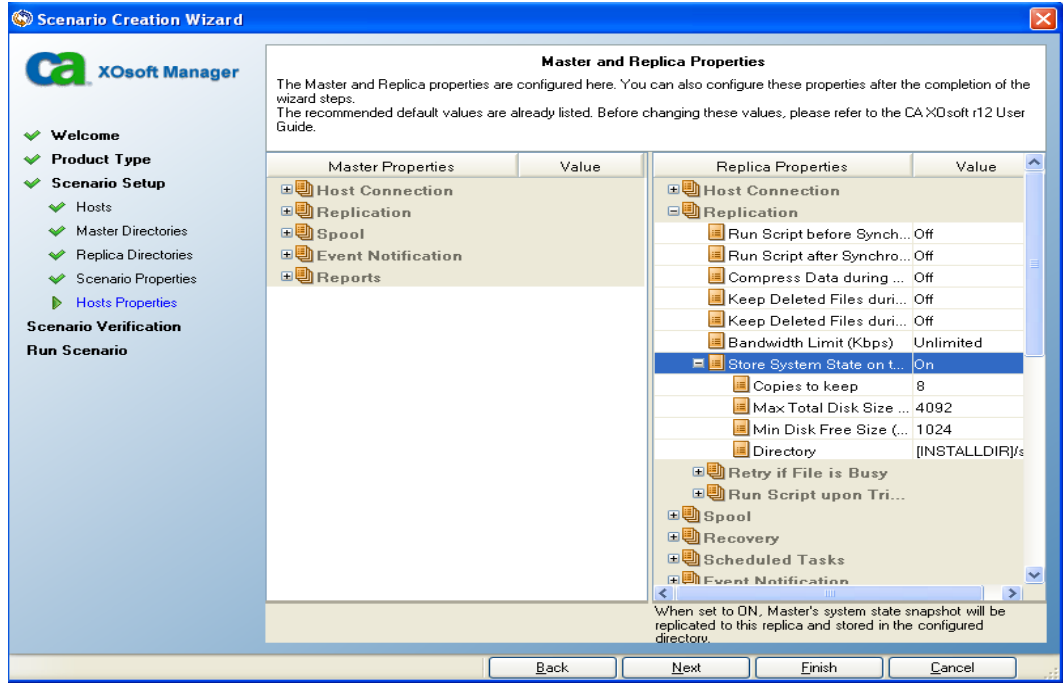
1. From the CA XOssoft Manager, start the scenario creation wizard using the toolbar button or the menu command: New, Scenario.
2. Complete the steps in the usual manner for the server type you selected until you reach the Scenario Properties dialog:



3. Under Replication properties, set System State Protection to On.
4. Under System State Protection, click the Value for Set Schedule to access the Schedule Setting dialog:



5. Set the schedule for taking system state snapshots.
6. Complete scenario creation steps as usual until you reach the Master and Replica Properties dialog:



7. Expand Replication properties on the Replica and set the Store System State on this Replica property to On. You may also set additional storage properties at this time. For more information, refer to the topic, Store System State Properties.
8. Save and run the scenario. System State Protection starts.

Configure System State Protection for Existing Scenarios

If you did not enable System State Protection when creating the Scenario, you can configure the System State Protection property outside the Scenario Creation Wizard using the following procedure.

Note: By default System State Protection is set to off.

Before performing this procedure, stop the scenario. From the CA XOsoft Manager, select the scenario and click the Stop button from the toolbar.

To configure system state protection for existing scenarios

1. From the CA XOsoft Manager select a Scenario to change its System State Protection property.
2. On the Scenario Properties panel, expand the Replication node and set the Enable System State Protection node property to On.

If the Replica is not configured, a message dialog opens.

3. Click OK.
4. Expand the System State Protection node and set a schedule, if desired. For more information, refer to the topic, [Set the System State Protection Schedule](#) (see page 155).
5. On the Replica, enable the store system state protection property. For more information, refer to the topic, [Configure System State Protection on the Replica](#) (see page 155).
6. Save the Scenario.

More information:

[Set the System State Protection Schedule](#) (see page 155)

[Configure System State Protection on the Replica](#) (see page 155)

Set the System State Protection Schedule

If you did not enable set a System State Protection schedule when creating the Scenario, you can configure the System State Protection schedule property outside the Scenario Creation Wizard using the following procedure.

Note: By default System State Protection schedule is set to off.

To set the System State Protection Schedule

1. From the CA XOssoft Manager, select a Scenario to set the System State Protection Schedule property.
2. On the Scenario Properties panel, expand the Replication node and expand the System State Protection node.
3. Click the Value column of the Set Schedule node.
The Set Schedule dialog opens.
4. Set the day, start time, frequency, and exclude dates as needed.
5. Click OK.

Configure System State Protection on the Replica

If you did not enable System State Protection on the Replica when creating the Scenario, you can configure the Store System State on this replica property outside the Scenario Creation Wizard using the following procedure.

Notes:

- You can enable System State Protection on multiple Replicas.
- By default System State Protection is set to off.

To configure system state protection of the Replica

1. From the CA XOssoft Manager, select a Replica to enable its System State Protection property.
2. On the Properties panel, expand the Replication node and set the Store System State on this replica node property to On.
3. Modify the Replica storage property values as required. For more information, refer to the topic, [Store System State Protection Properties](#) (see page 156).
4. Save the Scenario.

Store System State Protection Properties

You may set the following properties on the Replica server to manage system state snapshot storage:

Copies to keep

Specify the number of System State snapshots to keep on the Replica. The default value is eight. For an unlimited number of snapshots, enter a zero. If the number of snapshots exceeds the set value, the oldest snapshots are deleted to make space for subsequent snapshots.

Max Total Disk Size (MB)

Specify the total amount of disk space to allocate for System State snapshots.

Default Values by Operating System:

- Windows XP: 4096
- Windows 2000: 4096
- Windows 2003: 8192
- Windows 2008: 16,384

For an unlimited amount of space, enter a zero. If snapshots fill the space allocated, the oldest snapshots are deleted to make space for subsequent snapshots.

Min Disk Free Size (MB)

Specify the minimum amount of free disk space for System State snapshots. The default value is 1024. If the free disk space is less than the set value, the oldest snapshots are deleted to ensure that the minimum free disk space value is maintained.

Directory

Specify the directory where the System State snapshot is stored.

Modify Scenario System State Protection

When a scenario is stopped, you may modify its properties, including System State Protection. The changes you make take effect when the scenario is restarted.

System State Protection Property

If you disable the System State Protection property on an existing scenario, you are prompted to delete existing snapshots. If you select Yes, all snapshots from the Replica are deleted when the scenario is restarted. If you select No, all snapshots are retained.


Store System State on this Replica Property

If you disable the Store System State on this Replica property on an existing scenario, you are prompted to delete existing snapshots. Choose Yes to delete all snapshots or choose No to retain them.

You can modify all properties under the Store System State on this Replica group:


Copies to Keep

You can adjust the number of snapshots to keep. If you reduce the count to a value lower than what is already stored on the Replica, message WR00591 is displayed.



Max Total Disk Size (MB)

You can adjust the maximum total disk size setting. If you reduce this setting to a value lower than the currently occupied disk space, message WR00589 is displayed.



Min Disk Free Size (MB)

You can adjust the minimum disk free size to which snapshots are saved. If you set this value to 0, size is unlimited. If this setting is less than the size when the scenario is running, the oldest snapshot is deleted to make space for a new one.

Modify Directory to Store Snapshots

You can change the directory to which stored snapshots are saved. If you specify an invalid path, you are notified that snapshots cannot be saved. If a valid path is set, you are prompted to move old snapshots to a new location. If you select No, old snapshots are deleted.


For more information, refer to the topic, [Store System State Protection Properties](#) (see page 156).

Restore System State Data

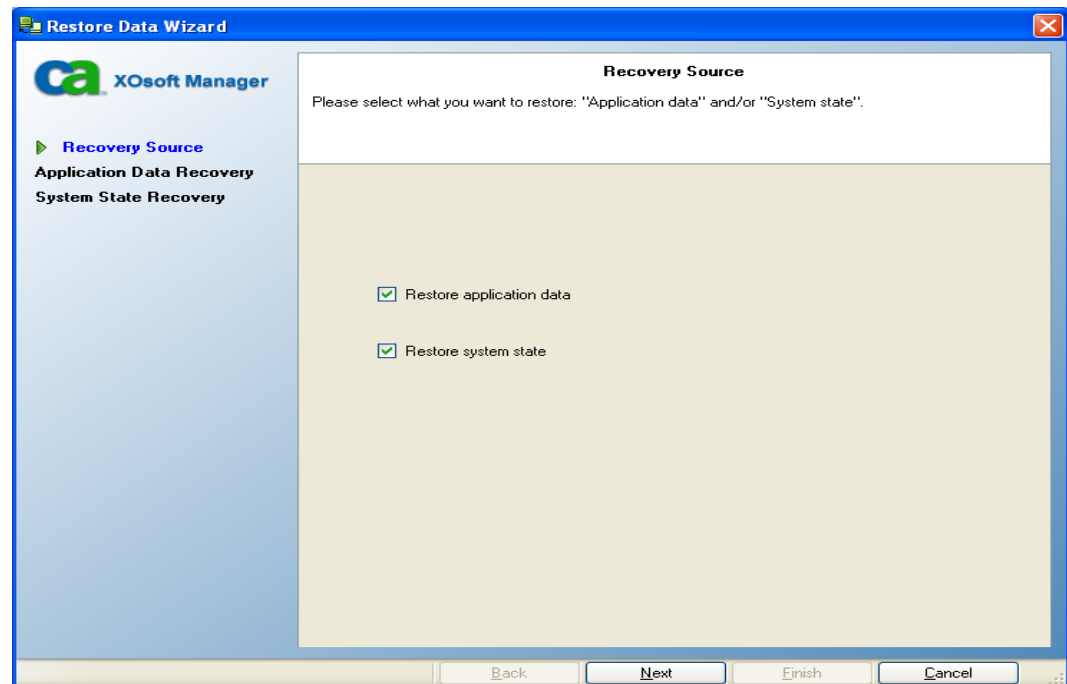
The recovery process for restoring System State data is similar to the usual recovery process, with one distinction. If System State Protection is enabled, you are prompted at recovery time to select a Recovery Source (see the following procedure).

CA XOssoft waits for application data recovery to complete first, including synchronization, before recovering the system state. You must select a Replica to act as the restore source machine. During the restore process, a new scenario is created in which the Master and Replica servers reverse roles. The scenario then transfers the System State snapshot to the original Master.

To restore system state data

1. From the CA XOssoft Manager select a Scenario to restore its System State data.
2. Select the Replica host.
3. From the toolbar, click the Restore Data  button.

The Recovery Source dialog opens if the scenario has System State Protection enabled.

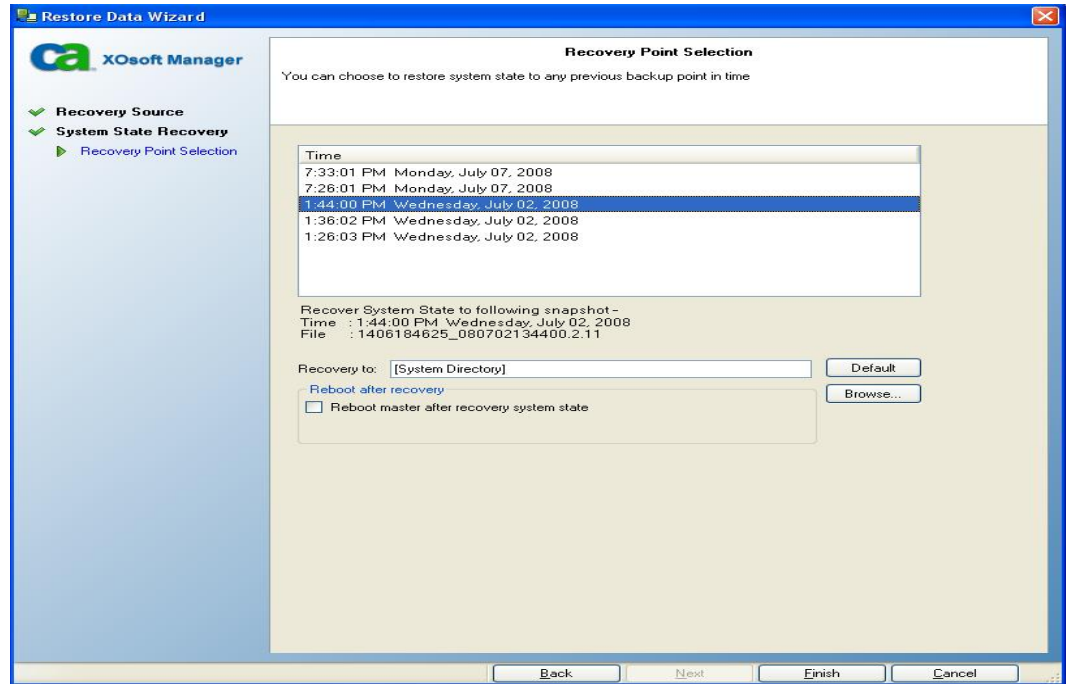


- Click Restore Application Data and Restore System State, and then click Next.

Note: If you select Restore System State only, you will not be able to choose an application recovery point. If you disable System State Protection, you will not see the Recovery Source screen.

The Rewind Point Selection dialog opens.

- From the Recovery Point Selection dialog, click Select Rewind Point to open the Recovery Point Selection dialog.



- Set the following criteria:

Time

Select any previous backup point from the list.

Recovery to

Recovery data to the default location or browse to a specific location.

- Click Finish.
- Reboot the Replica now.

Command Line Enhancements for System State Protection

The following commands have been added to PowerShell to support System State Protection:

set-properties scenario_name index value

Use the set-properties command to set System State Protection for a scenario.

To obtain index values, use the get-properties command.

set-hostproperty scenario_name replica_name index value

Use the set-hostproperty command to enable the Store System State property on a Replica.

To obtain index values, use the get-hostproperties command.

RecoveryMode [A|S|B]

Use A to recover application data only. (Default setting)

Use S to recover system state only.

Use B to recover both.

RebootAfterRecovery [0|1]

Use 0 to skip reboot (default)

Use 1 to enable Master reboot after recovery

Additional System State Information

System State Protection in High Availability Scenarios

After switchover, system state snapshots are not retained on the original master.

Chapter 8: Setting Master and Replica Properties

This section 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.

This section contains the following topics:

[Configuring Master or Replica Server Properties](#) (see page 164)

[Understanding Master Properties](#) (see page 165)

[Understanding Replica Properties](#) (see page 170)

[Schedule the Bandwidth Limit](#) (see page 176)

[Propagating Property Values](#) (see page 178)

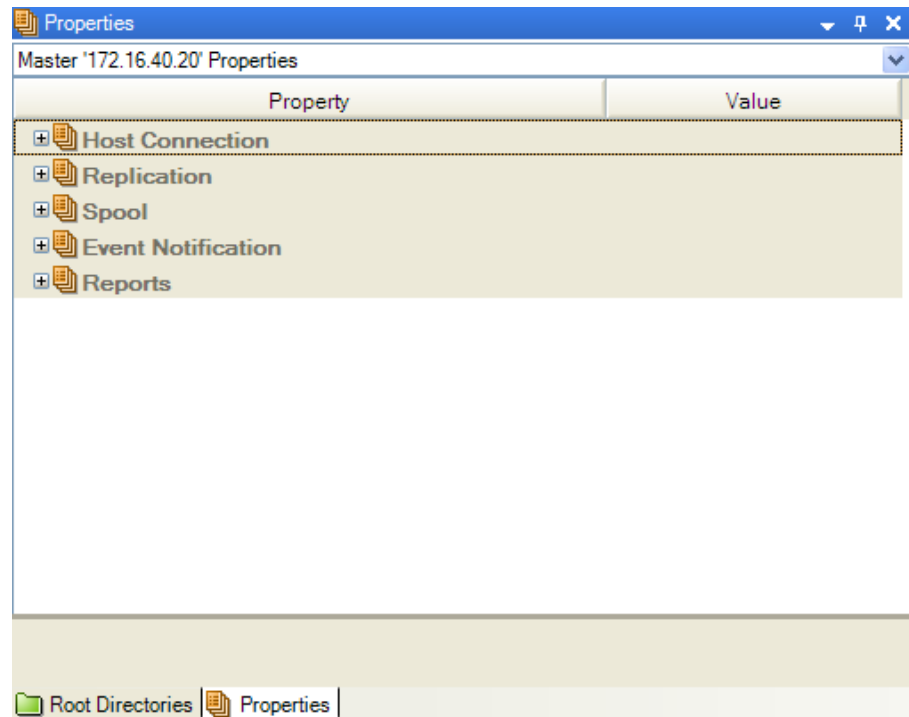
Configuring Master or Replica Server Properties

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 opens.



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.

Note: On Windows x64, you cannot run scripts that activate applications with UI.

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 XOsft 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, send an email notification, or write it to Windows event log.

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 XOssoft 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 XOssoft, 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.

Note: On Windows x64, you cannot run scripts that activate applications with UI.

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.

Bandwidth Limit (Kbps)

Controls the size of the allowed incoming bandwidth on the Replica host. You can either define one limit size that will apply to all hours of the day, or you can specify different values for different hours. The default value is **Unlimited**.

For a detailed description of bandwidth scheduling, refer to [Schedule the Bandwidth Limit](#) (see page 176).

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 XOssoft stops the database services before running the scenario. [Does not apply to HA scenarios]

Store System State on this Replica

This option can be enabled only when the **System State Protection** property in the Scenario Properties list is set to On. For more information, refer to [Protecting Your System State](#) (see page 149).

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.

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

Refer to [Schedule Replication Suspension](#) (see page 77).

Replica Integrity Testing for Assured Recovery

Refer to [Setting Assured Recovery Properties](#) (see page 284).

Event Notification

Notification

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

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 XOssoft 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 XOssoft, 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.

Schedule the Bandwidth Limit

CA XOssoft enables you to control the size of the allowed incoming bandwidth on the Replica host. You can either define one limit size that will apply to all hours of the day, or you can specify different values for different hours. By using the Bandwidth Scheduler, you can decrease the bandwidth size on busy hours and increase it on off-peak hours in order to optimize your bandwidth resources.

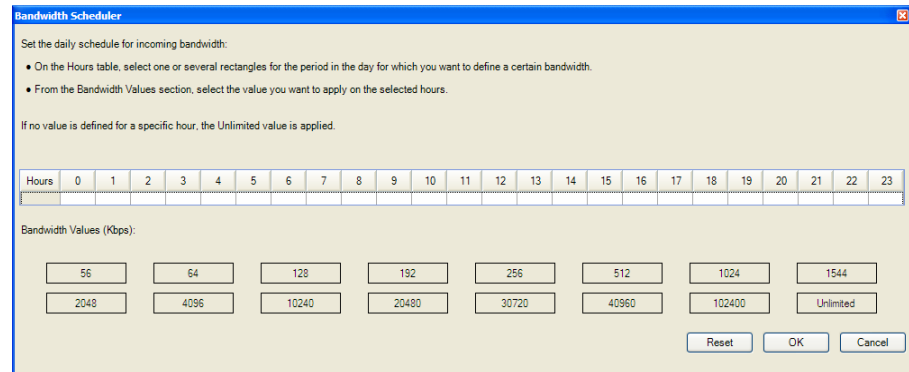
Notes:

- The bandwidth limit that you set for one Replica host does not apply to other Replica hosts that reside in the same replication tree. You need to change each Replica host definition separately.
- The default value for the Bandwidth Limit option is "**Unlimited**". It means that no restriction is imposed on the bandwidth between the Master and the Replica.

To schedule the bandwidth limit:

1. On the Replica Properties list, open the **Replication** group. On the **Bandwidth Limit** property, click the Value box that contains the **Unlimited** default value.

The **Bandwidth Scheduler** dialog opens.



2. Set the daily schedule for incoming bandwidth size according to the following guidelines:
 - On the **Hours** table, select one or several rectangles for the period in the day for which you want to define a certain bandwidth size.

Bandwidth Scheduler

Set the daily schedule for incoming bandwidth:

- On the Hours table, select one or several rectangles for the period in the day for which you want to define a certain bandwidth.
- From the Bandwidth Values section, select the value you want to apply on the selected hours.

If no value is defined for a specific hour, the Unlimited value is applied.

Hours	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23

Bandwidth Values (Kbps):

56	64	128	192	256	512	1024	1544
2048	4096	10240	20480	30720	40960	102400	Unlimited

Reset OK Cancel

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

- Once the rectangles are marked, from the **Bandwidth Values** section click the values (in Kbps) you want to apply on the selected hours.

Bandwidth Scheduler

Set the daily schedule for incoming bandwidth:

- On the Hours table, select one or several rectangles for the period in the day for which you want to define a certain bandwidth.
- From the Bandwidth Values section, select the value you want to apply on the selected hours.

If no value is defined for a specific hour, the Unlimited value is applied.

Hours	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
											102400	102400	102400	102400	102400	102400	102400	102400	102400					

Bandwidth Values (Kbps):

56	64	128	192	256	512	1024	1544
2048	4096	10240	20480	30720	40960	102400	Unlimited

Reset OK Cancel

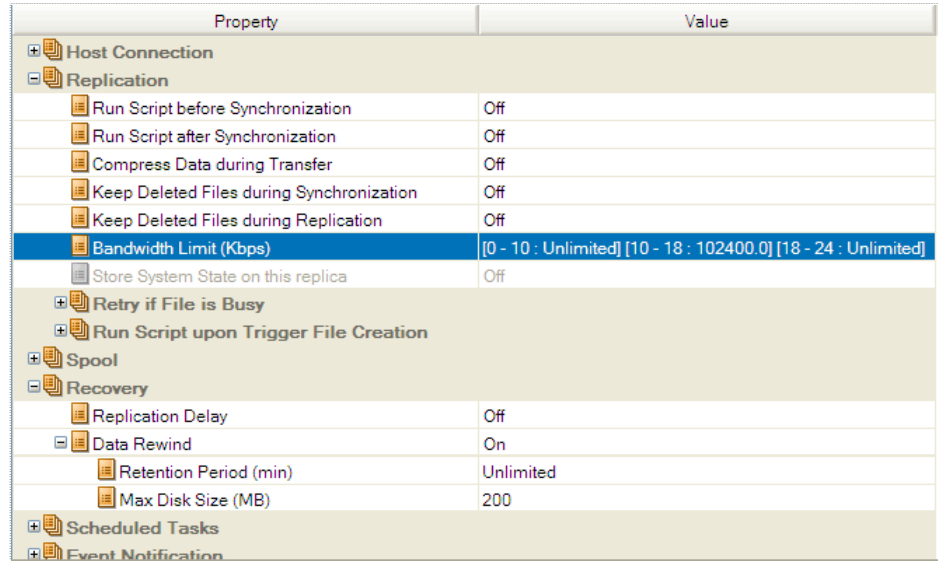
The rectangles of the selected hours now contain the value you selected.

3. You can repeat this process for all hours. If no size is defined for a specific hour, the **Unlimited** default value is applied to it.

Note: To clear the setting, click the **Reset** button.

4. Once you finished defining the Bandwidth Scheduler, click **OK** to save your setting and close the dialog.

The setting you defined now appears in the **Bandwidth Limit** value box in the Properties list.



Property	Value
Host Connection	
Replication	
Run Script before Synchronization	Off
Run Script after Synchronization	Off
Compress Data during Transfer	Off
Keep Deleted Files during Synchronization	Off
Keep Deleted Files during Replication	Off
Bandwidth Limit (Kbps)	[0 - 10 : Unlimited] [10 - 18 : 102400.0] [18 - 24 : Unlimited]
Store System State on this replica	Off
Retry if File is Busy	
Run Script upon Trigger File Creation	
Spool	
Recovery	
Replication Delay	Off
Data Rewind	On
Retention Period (min)	Unlimited
Max Disk Size (MB)	200
Scheduled Tasks	
Event Notification	

5. To save your setting, click the **Save** button on the Standard toolbar.

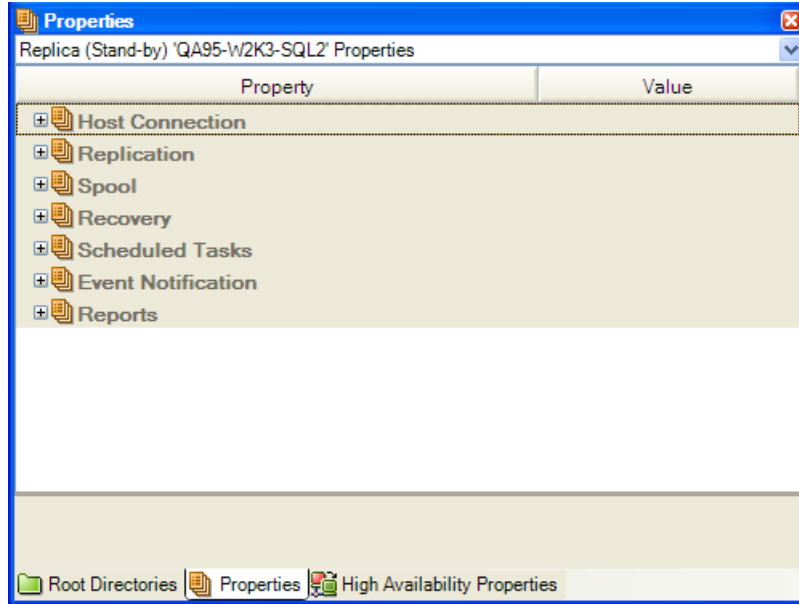
Propagating Property Values

CA XOssoft enables you to apply the values you set for one scenario to multiple scenarios at once. Instead of separately configuring the properties of each scenario, in a click-of-a-button you can propagate the value of one scenario to as many scenarios as you want. A good example would be to use this option to simultaneously change the e-mail notification address for multiple scenarios. You can propagate the values of scenarios, Master hosts and Replica hosts.

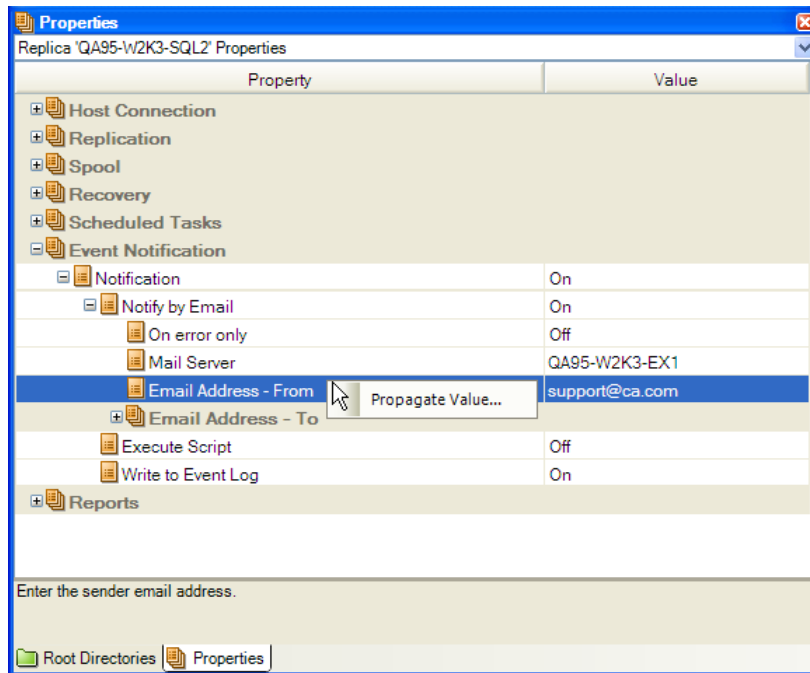
Note: The scenarios must be stopped to apply property changes.

To propagate values of properties

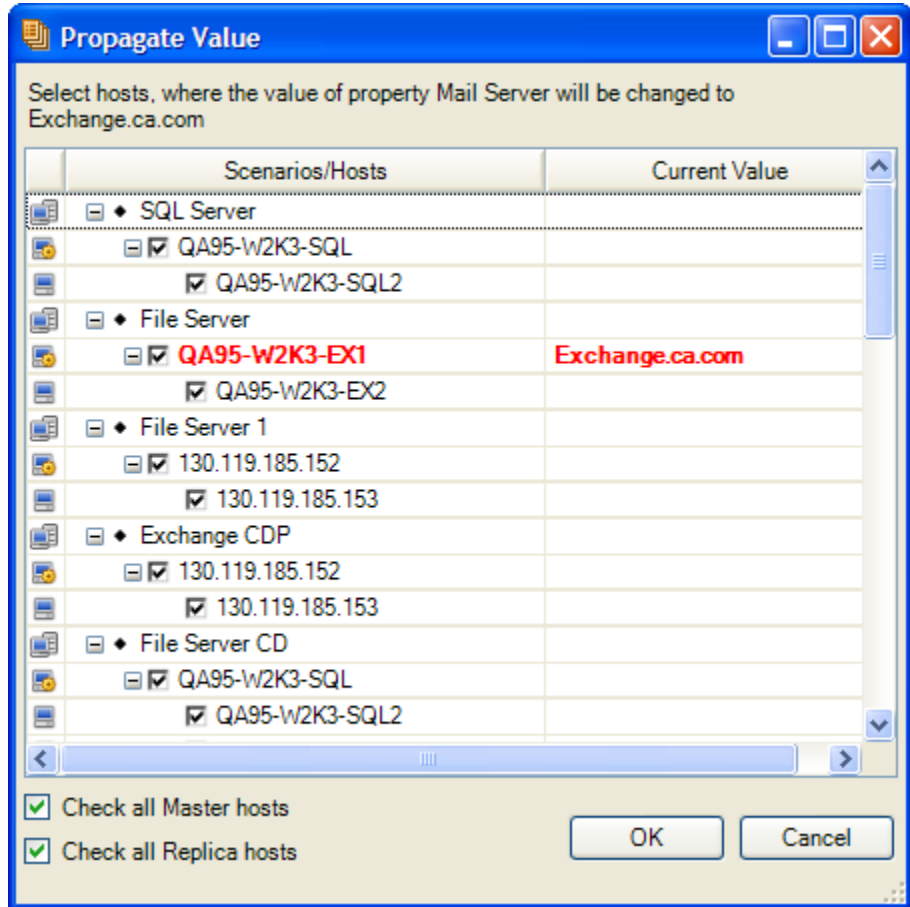
1. On the Scenario pane, select either the scenario, the Master or the Replica whose properties you want to propagate. On the Framework pane on the left, the Properties list opens.




2. On the Properties list, open the desired group, select the property its value you want to propagate and right-click. The **Propagate Value** pop up command opens.



3. Click the **Propagate Value** command. The **Propagate Value** dialog opens.



All scenarios in your Manager appear in the dialog, while the scenario whose property value you want to propagate is marked in red. The property and the value you can propagate are displayed above the **Scenarios** table, and in the **Current Value** column.

4. To propagate the value to all scenarios, click **OK**.
Note: To exclude scenarios or hosts from the value propagation, clear their check boxes, and then click **OK**.
5. After the **Propagate Value** dialog is closed, click the **Save All**  button on the Standard toolbar to save and apply your changes to all scenarios.

Chapter 9: Data Recovery

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

This section contains the following topics:

[The Data Recovery Process](#) (see page 181)

[Recover Lost Data from Replica](#) (see page 182)

[Setting Bookmarks](#) (see page 186)

[Data Rewind](#) (see page 187)

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 recover data in two ways:

- Recover lost data from the Replica to the Master -- this option is a synchronization process in the reverse direction and requires you to stop the scenario. (This option is not recommended for Oracle, SQL or Exchange scenarios.)
- Recover lost data from a certain event or point in time (Data Rewind) -- This option uses a process of stamped checkpoints and user-defined bookmarks to roll corrupt data on the Master back to a time before corruption occurred.


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

Recover Lost Data from Replica

To 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.

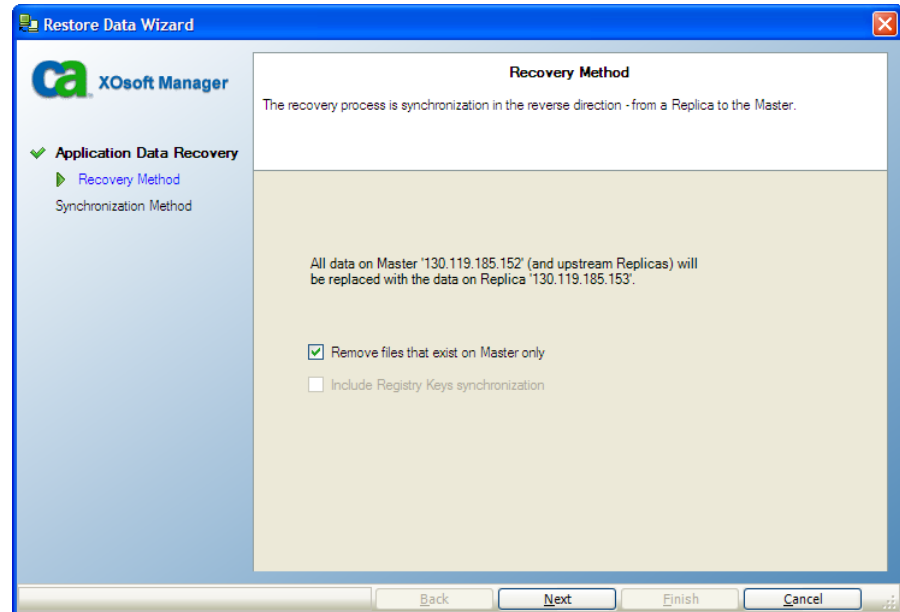


The **Restore Data** option is enabled.

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



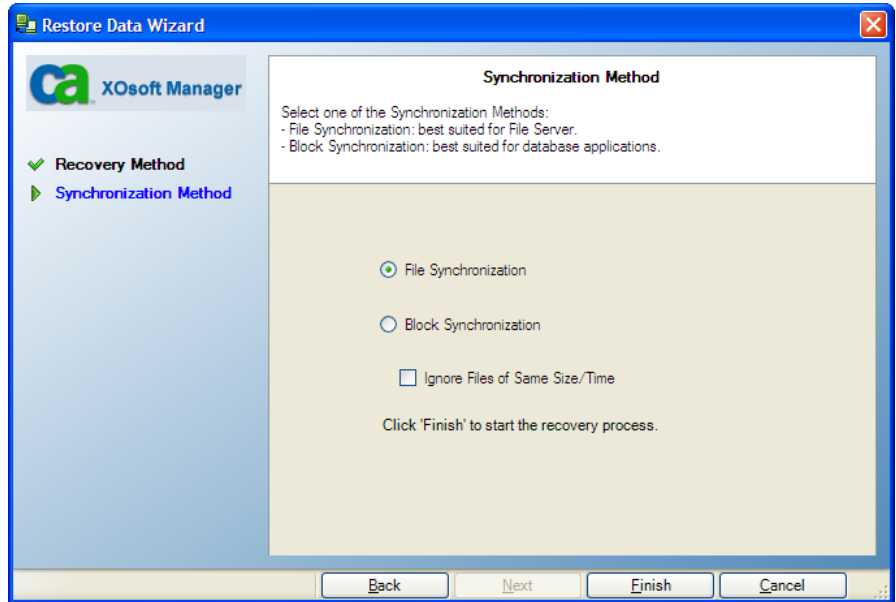
The **Recovery Method** page of the Restore Data wizard opens.



Notes:

- If the **Data Rewind** property is set to On, another **Restore Data** dialog will appear. In this case, select the first option - Replace all data on Master with the data on Replica.
- The **Include Registry Keys synchronization** checkbox is enabled, only if you activated the [Registry Synchronization property](#) (see page 117) before starting the scenario. If the checkbox is enabled, you can select it to include the synchronized Registry Keys in the recovery process.

5. Click **Next**. The **Synchronization Method** page opens.

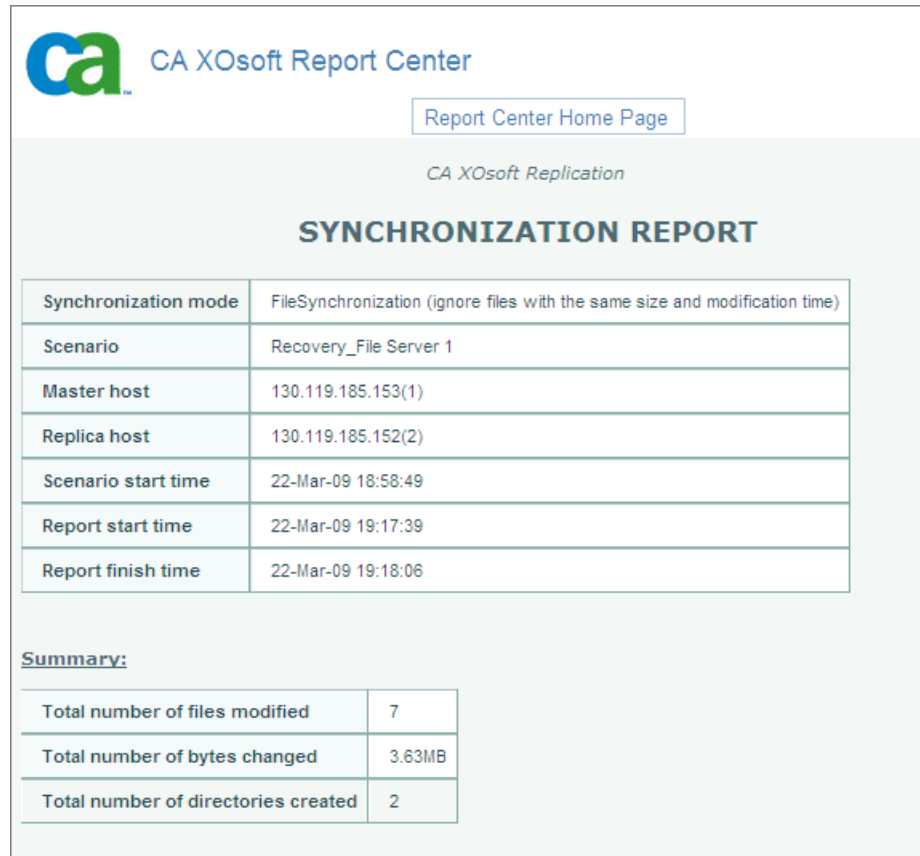


6. Make sure that the **File Synchronization** method is selected, and click **Finish**.

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 XOssoft 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.

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



The screenshot displays the CA XOssoft Report Center interface. At the top left is the CA logo and the text "CA XOssoft Report Center". A button labeled "Report Center Home Page" is located to the right. Below this, the text "CA XOssoft Replication" is centered. The main heading is "SYNCHRONIZATION REPORT".

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	Recovery_File Server 1
Master host	130.119.185.153(1)
Replica host	130.119.185.152(2)
Scenario start time	22-Mar-09 18:58:49
Report start time	22-Mar-09 19:17:39
Report finish time	22-Mar-09 19:18:06

Summary:

Total number of files modified	7
Total number of bytes changed	3.63MB
Total number of directories created	2

Now, the Replication process can restart on 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 the **Recovery - Data Rewind** option to **On** (default setting is Off).
- You cannot set bookmarks during the synchronization process.

To set a bookmark

1. When the required scenario is running, on the Scenario pane select the Replica host from which you want to rewind data.
2. From the **Tools** menu, select the **Set Rewind Bookmark** option.

The **Rewind bookmark** dialog opens.



The text that appears in the **Rewind bookmark** dialog will appear in the **Rewind Points Selection** dialog as the bookmark's name. The default name includes date and time.

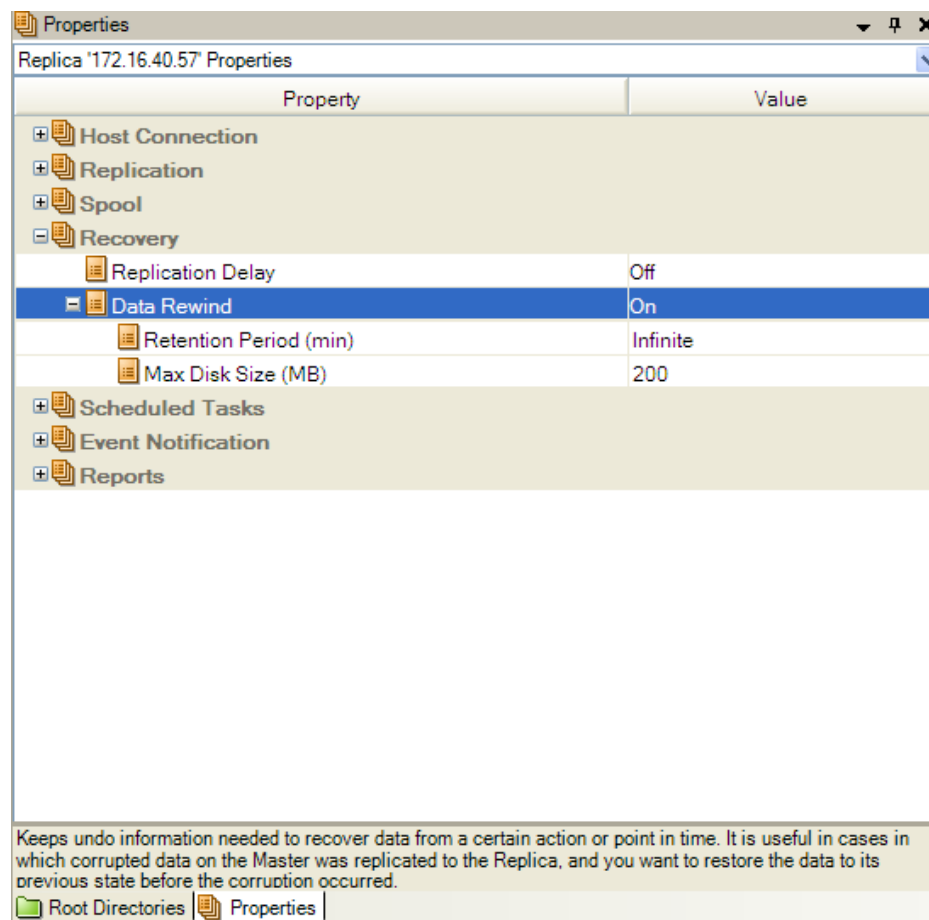
3. Accept the default name, or enter a new name for the bookmark. It is recommended to give a meaningful name that will later help you recognize the required bookmark. Then, click **OK**.

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**.



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 cannot manually set bookmarks during synchronization. In the following example, a File Server scenario is used, but the steps are the same for all scenario types.

To recover lost data using rewind points

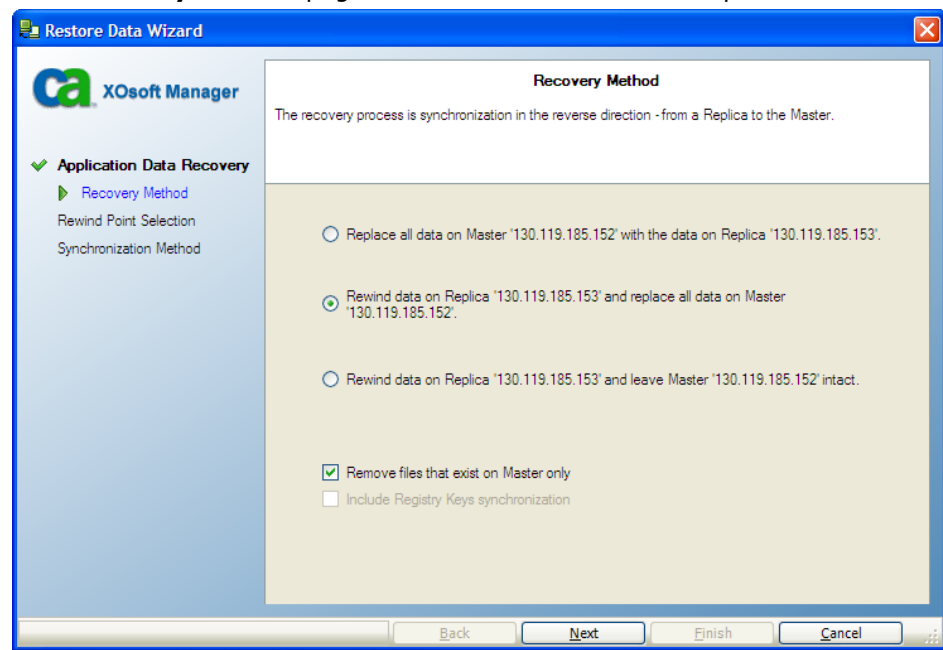
1. On the Manager, from the Scenario pane select the desired scenario and stop it.
2. [For database applications only] stop the database services on the Master host.
3. On the Manager, from the scenario folder select the Replica host:

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



4. From the **Tools** menu, select **Restore Data**, or click the **Restore Data**  button. If you are prompted for user credentials, enter the appropriate information and click OK.

The **Recovery Method** page of the Restore Data Wizard opens.



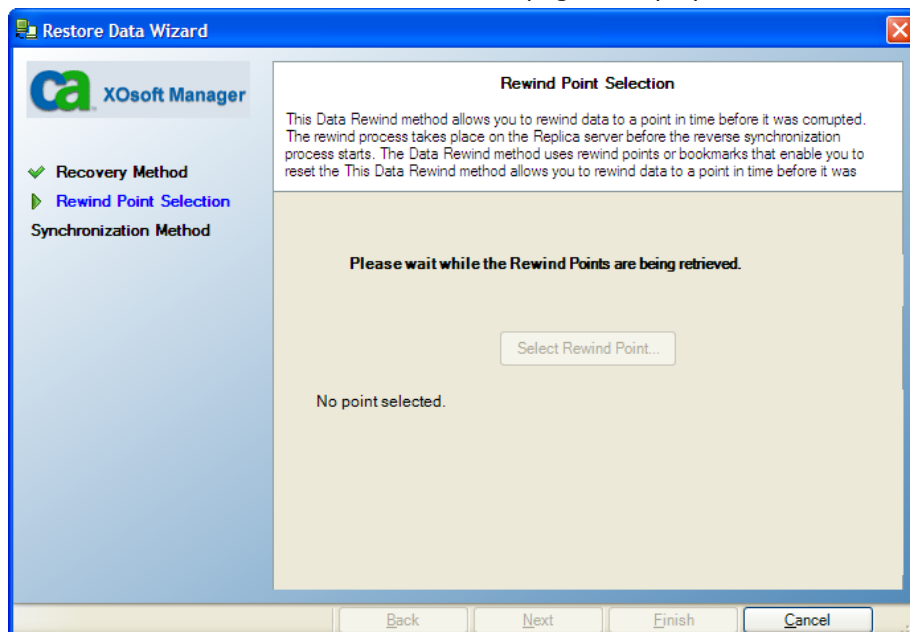
5. Select one of the Rewind data options, depending on whether you want the rewind data synchronized back to the Master (option 2) or left on the Replica only (option 3).

Notes:

- 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 **Include Registry Keys synchronization** checkbox is enabled, only if you activated the [Registry Synchronization property](#) (see page 117) before starting the scenario. If the checkbox is enabled, you can select it to include the synchronized Registry Keys in the recovery process.


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 opens.



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

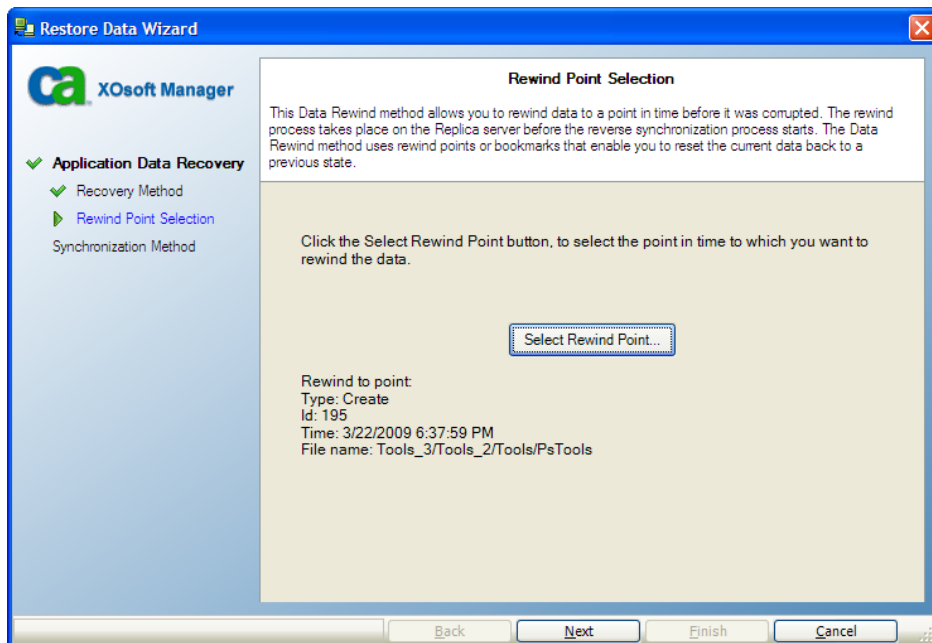
The list can be filtered according to the rewind point type or other criteria, using the **Filter Rewind Points** pane on the left.

Notes:

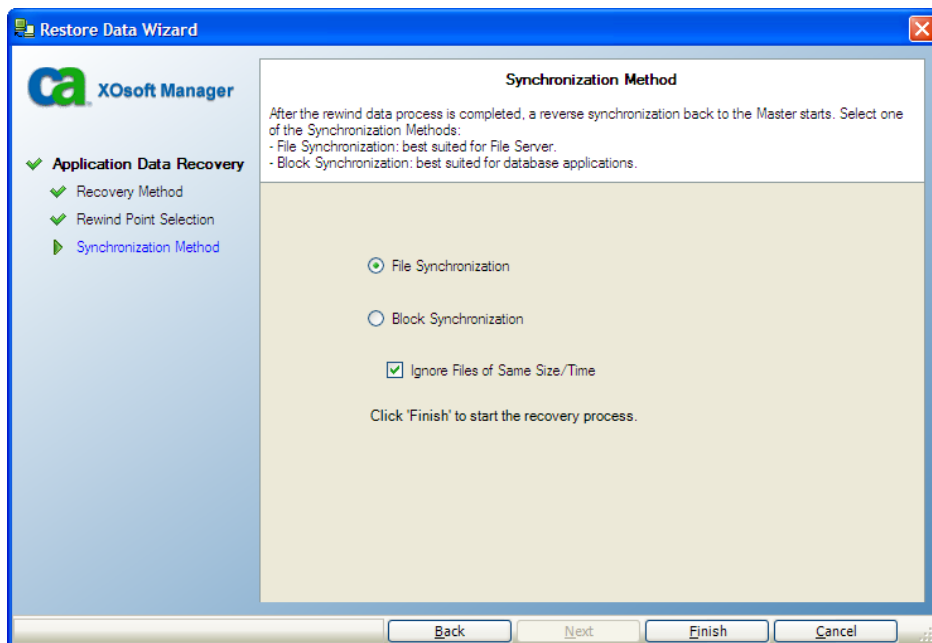
- If the **Select Rewind Points** dialog is empty, make sure that the Data Rewind [property](#) (see page 173) is enabled.
 - The entire list can be exported to an Excel file by clicking the **Export to Excel** button on the bottom-left corner.
8. Select the required rewind point, and click **OK**.

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.



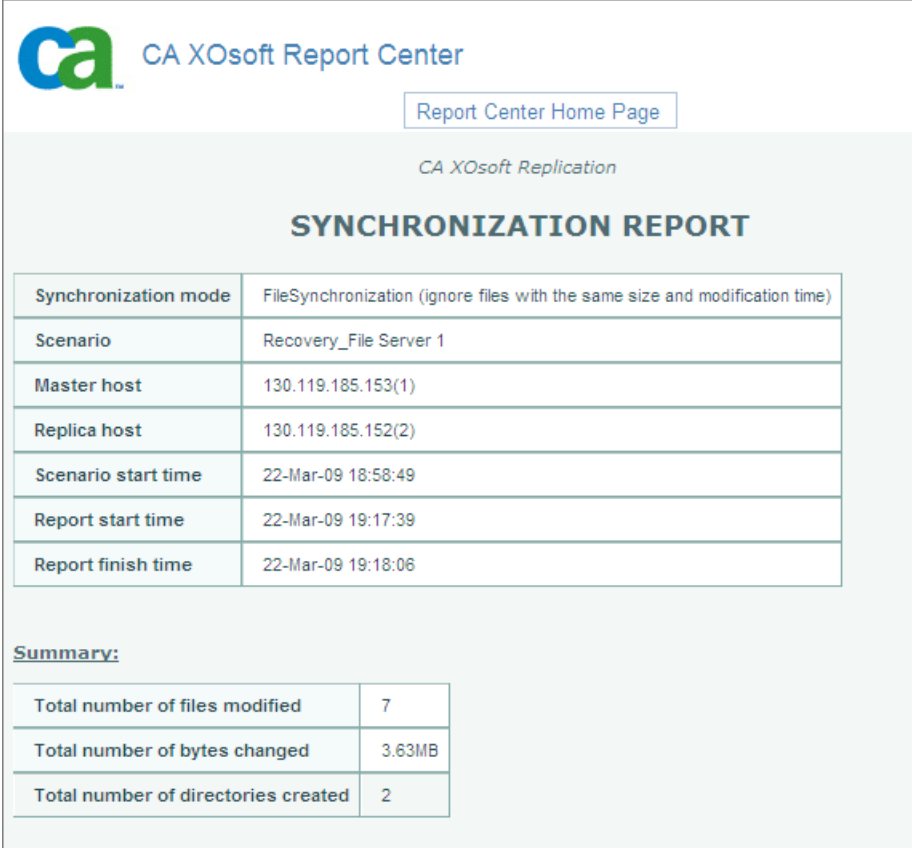
10. Select the **File Synchronization** method and click **Finish**.

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 Report Center

[Report Center Home Page](#)

CA XOsoft Replication

SYNCHRONIZATION REPORT

Synchronization mode	FileSynchronization (ignore files with the same size and modification time)
Scenario	Recovery_File Server 1
Master host	130.119.185.153(1)
Replica host	130.119.185.152(2)
Scenario start time	22-Mar-09 18:58:49
Report start time	22-Mar-09 19:17:39
Report finish time	22-Mar-09 19:18:06

Summary:

Total number of files modified	7
Total number of bytes changed	3.63MB
Total number of directories created	2

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

Chapter 10: Switching Over and Switching Back

This section 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.

This section contains the following topics:

[Understanding the High Availability System and the Switchover and Switchback Procedures](#) (see page 193)

[Creating High Availability Scenario](#) (see page 195)

[Switchover](#) (see page 208)

[Switchback](#) (see page 211)

[Recover Active Server Using the Manager](#) (see page 215)

[Recover Active Server from Outside the Manager](#) (see page 216)

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 XOsoft HA 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 XOsoft HA 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.

Note: After a connection loss, during the attempt to reconnect, a node (either Master or Replica) tries to determine its role. If the two nodes establish themselves as Masters, upon reconnection the newest active Master will continue to act as the Master, while the older one will turn into the standby Replica.

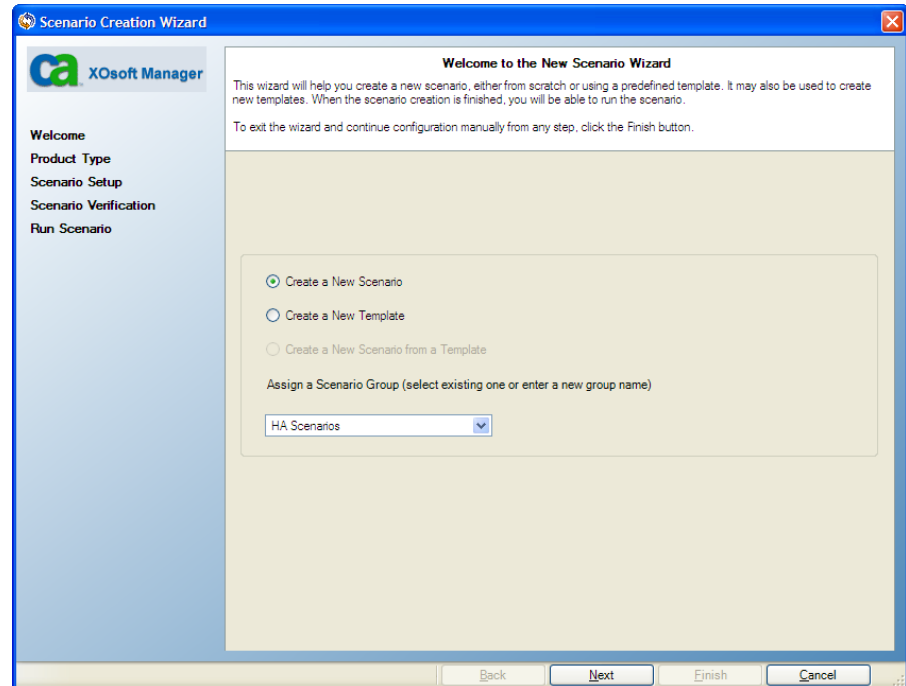
Creating High Availability Scenario

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 XOssoft HA performs automatically, and you need to configure several High Availability properties.

To create a High Availability scenario

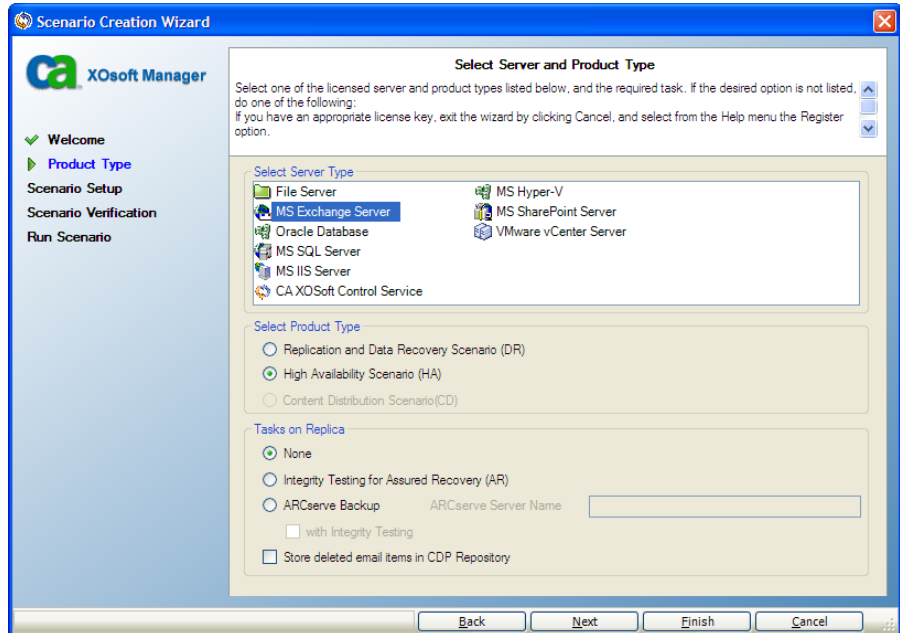
1. Open the CA XOssoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New**  button on the Standard toolbar.

The **Scenario Creation Wizard** opens.



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 opens.



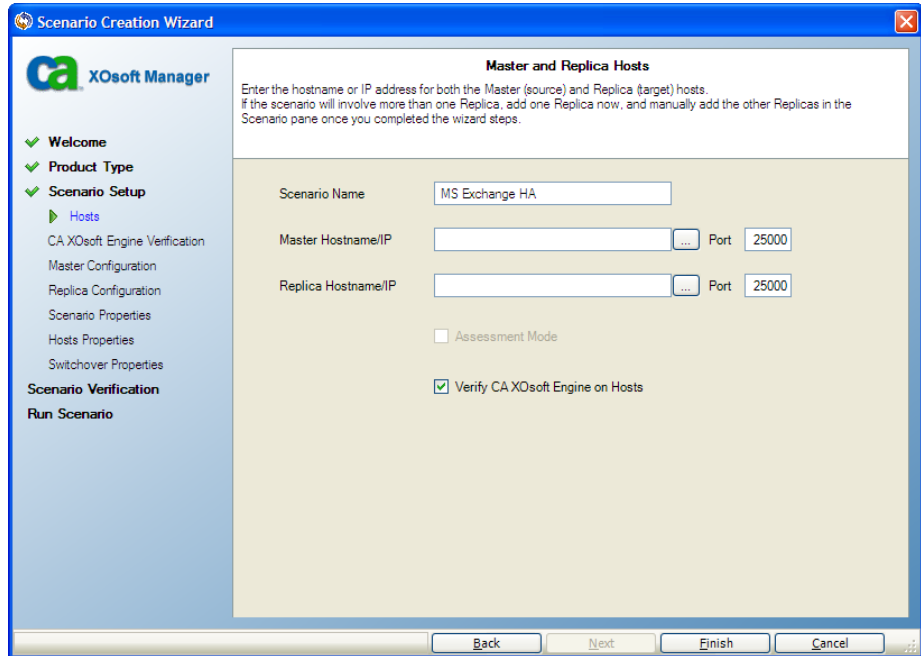
A list of available applications and scenario types is presented.

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 (HA)**.
- [Optional - a license is needed] From the **Tasks on Replica** options, select the tasks you want to implement in this scenario.

- Click **Next**. The **Master and Replica Hosts** page opens.



The screenshot shows the 'Master and Replica Hosts' step of the Scenario Creation Wizard. The window title is 'Scenario Creation Wizard' and the CA XOssoft Manager logo is visible. The left sidebar shows the progress: Welcome, Product Type, Scenario Setup (with 'Hosts' selected), CA XOssoft Engine Verification, Master Configuration, Replica Configuration, Scenario Properties, Hosts Properties, Switchover Properties, Scenario Verification, and Run Scenario. The main area contains the following fields and options:

- Scenario Name: MS Exchange HA
- Master Hostname/IP: [Empty] ... Port: 25000
- Replica Hostname/IP: [Empty] ... Port: 25000
- Assessment Mode
- Verify CA XOssoft Engine on Hosts

Buttons at the bottom: Back, Next, Finish, Cancel.

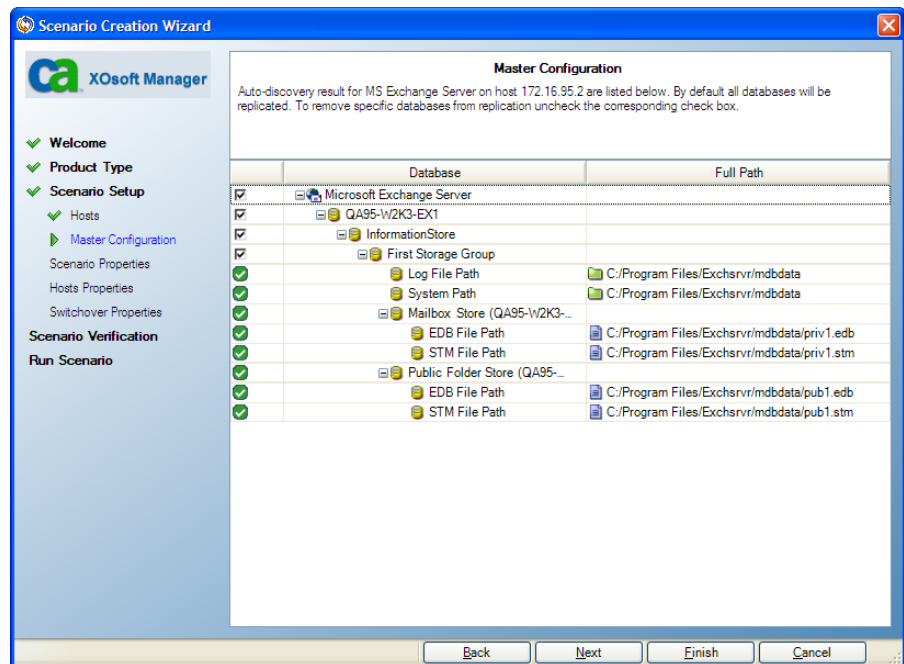
- Enter the following information:

- 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 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](#) (see page 105).
 - 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.
6. After you selected the desired options, click **Next**. The **Master Configuration** page opens.

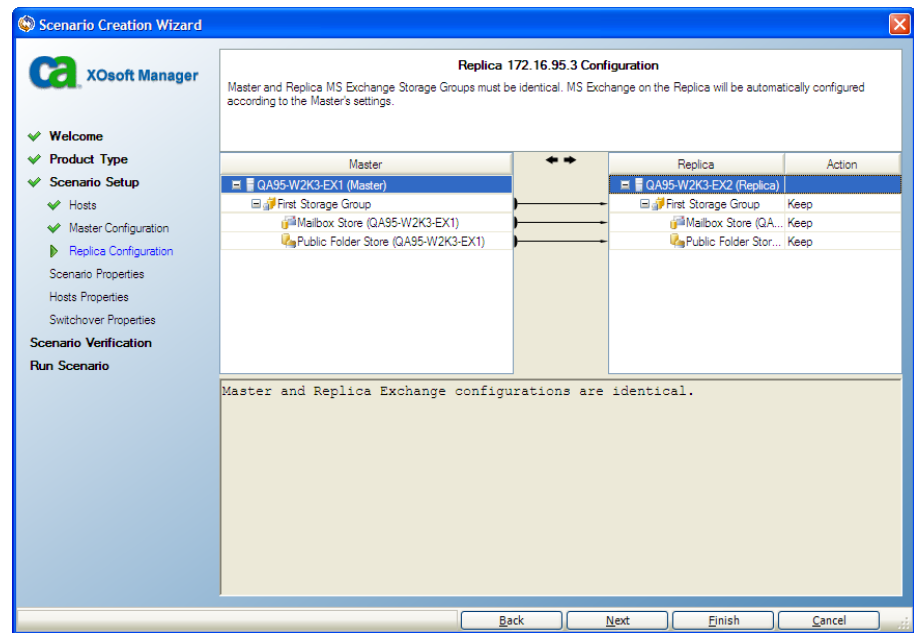


CA XOssoft auto-discovery component automatically displays the Exchange databases that are on the Master server. These are the databases that can be replicated and protected.

- 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.

[For Exchange 2007 with Outlook 2003 support, Exchange 2003 and earlier] **Important!** When selecting Exchange Storage Groups for replication, you must include at least one Public Folder for replication.

- Click **Next**. The **Replica Configuration** page opens.



CA XOssoft 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 XOssoft 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.

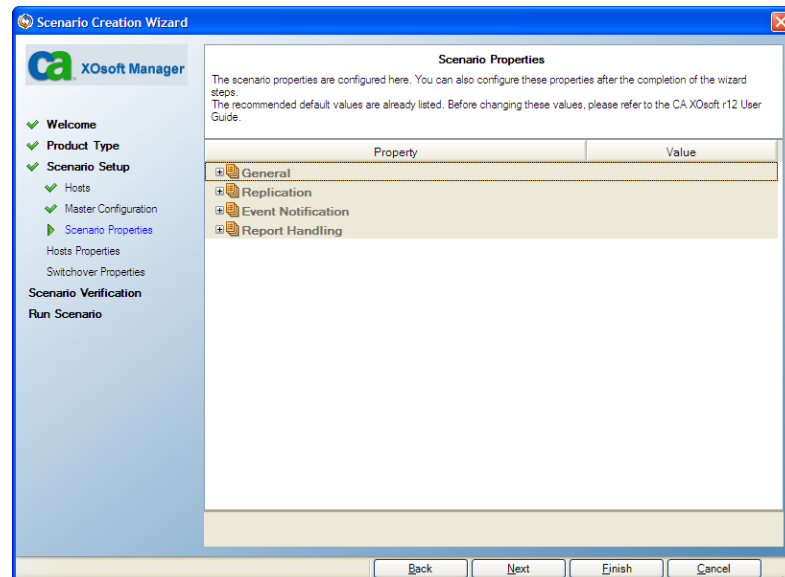
Note: The actions that can be performed in the auto-configuration process are as follows:

- **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.

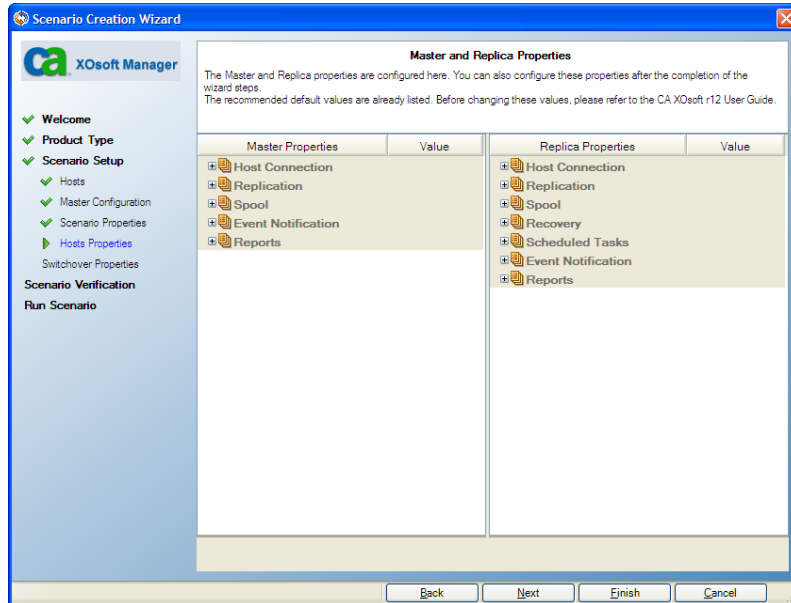
10. Click **Next** to start the Replica configuration process. The **Scenario Properties** page opens.



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](#) (see page 138). To configure the scenario properties at a later stage, refer to [Configuring Scenario Properties](#) (see page 137).

11. Click **Next**. The **Master and Replica Properties** page opens.

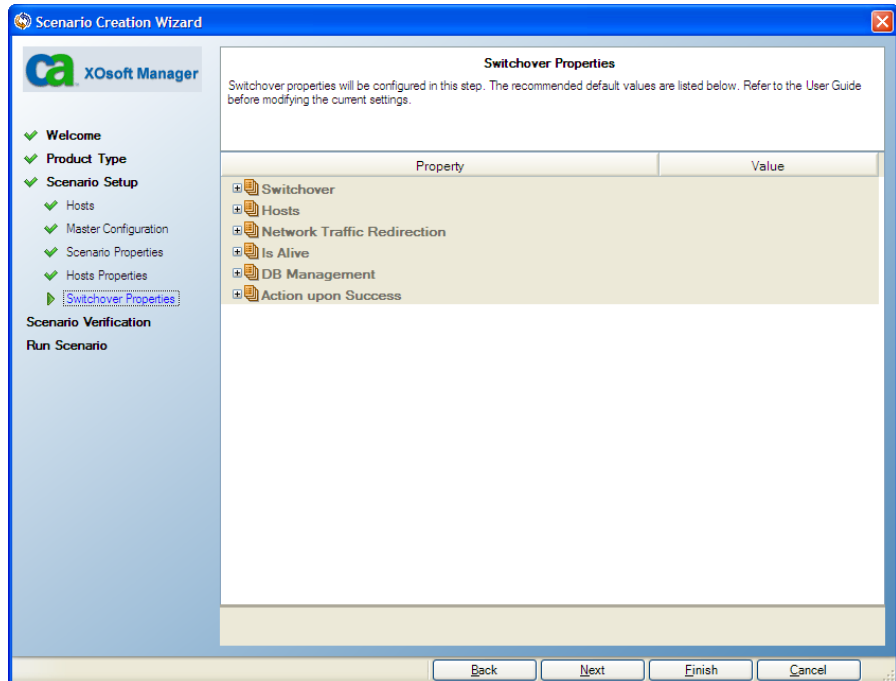


The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

If you want to configure the Master and Replica properties at this stage, refer to [Setting Master and Replica Properties](#) (see page 163). To configure the Master and Replica properties at a later stage, refer to [Configuring Master or Replica Server Properties](#) (see page 164).

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 information](#) (see page 167) for configuration details.

12. Once you are satisfied with the Master and Replica properties, click **Next**.
The **Switchover Properties** page opens.

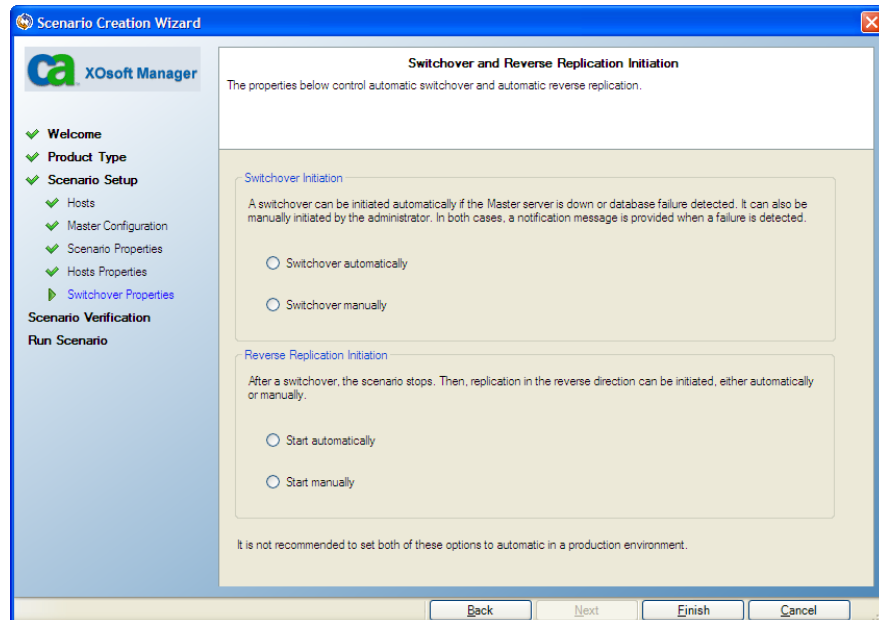


The **Switchover Properties** page allows you to modify switchover parameters. As with the prior steps, no changes are required.

If you want to configure the switchover properties at this stage, refer to [Understanding High Availability Properties](#) (see page 223). To configure the switchover properties at a later stage, refer to [Configuring High Availability Properties](#) (see page 221).

Note: If you plan to use automatic switchover, note that the **Is Alive Timeout (sec)** property controls how long to wait after a failure is detected before triggering a switchover. The default is 300 seconds. For more information see [Is Alive](#) (see page 229).

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



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 XOssoft HA, 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](#) (see page 229).

When **Reverse Replication Initiation** is set to **Start automatically**, CA XOssoft HA 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.

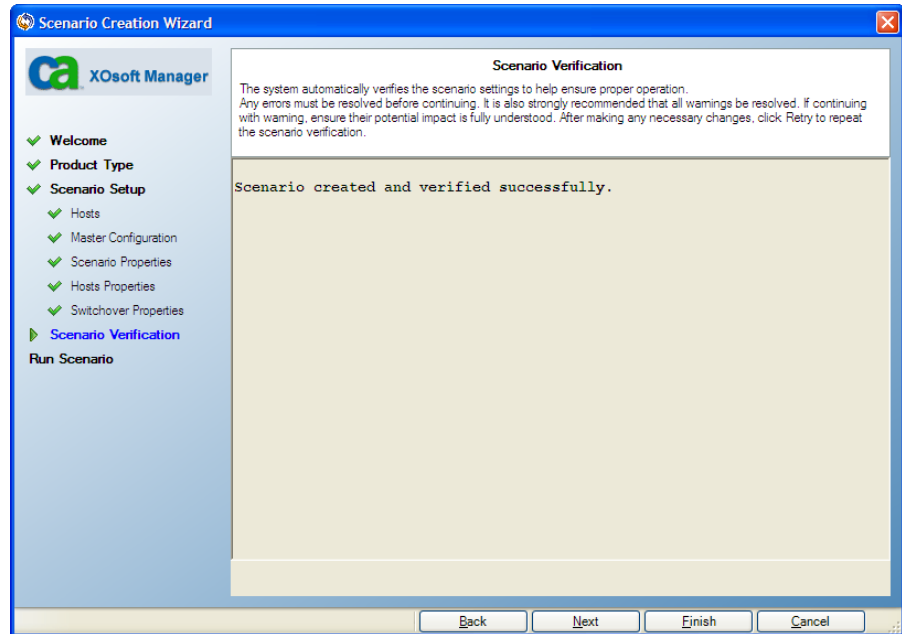
Important! We advise you against setting both of these options to automatic in a production environment. While these options are individually beneficial, it is best practice to set only one of these options to automatic.

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

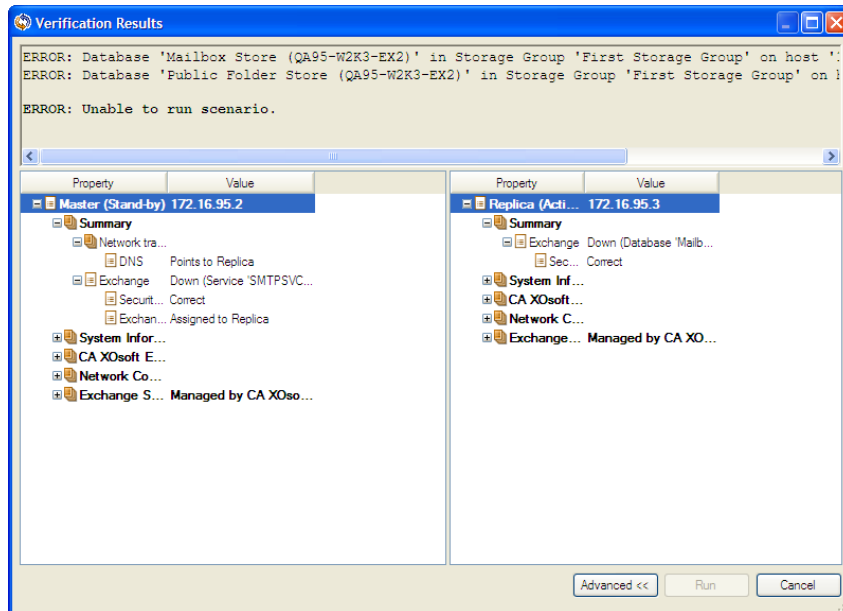
Note: Numerous checks are performed at this stage. Broadly, they fall into three categories:

- **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 opens.



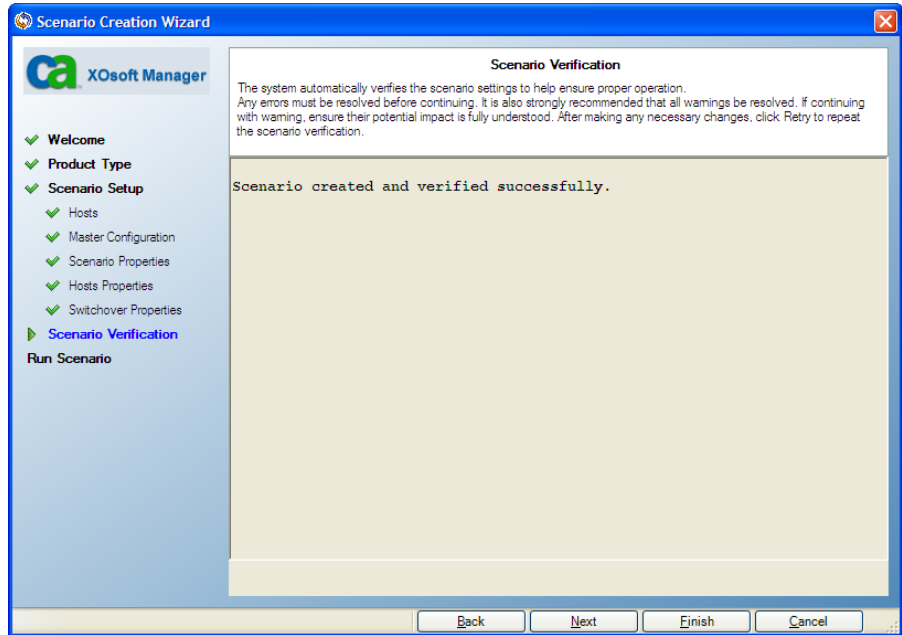
17. If the scenario was not set up correctly, or problems occurred in the participating hosts or the connection between the CA XOsoft components, the **Verification Results** dialog opens, listing all the errors and warnings detected:



- 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.

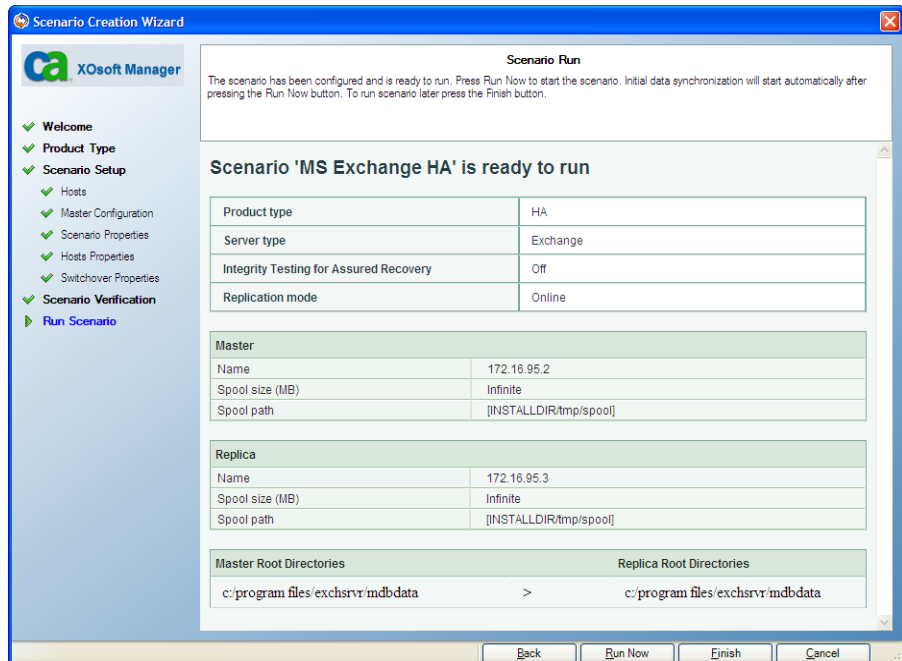
Note: The **Advanced** button opens an additional pane with detailed information about the checks performed to help diagnose problems. It is intended to help you resolve any issues encountered in running the software. You can also contact Technical Support for further assistance.

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



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

The **Scenario Run** page opens.

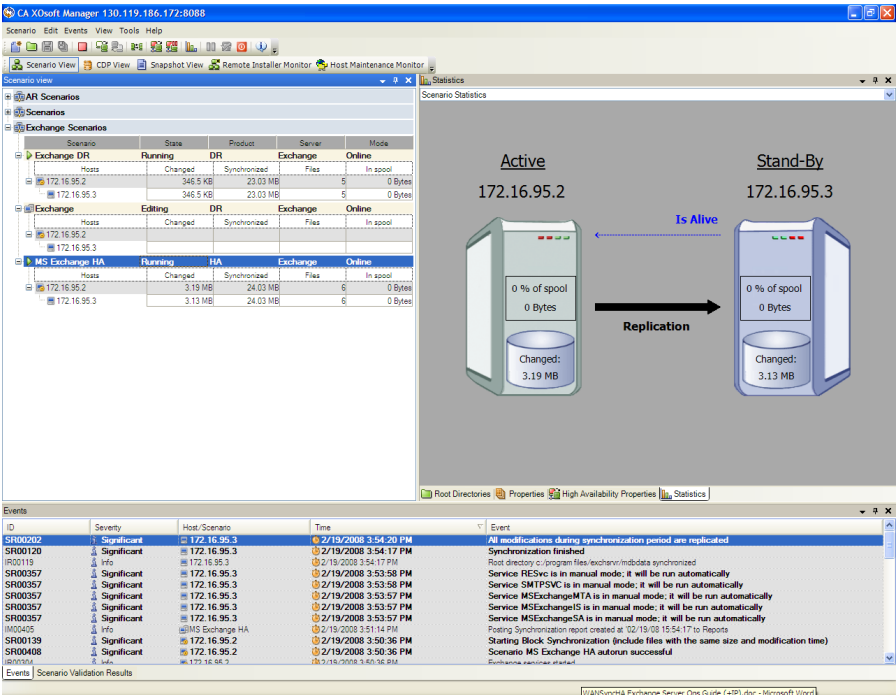


20. The scenario configuration is now completed and you are prompted to run it. Running the scenario starts the data synchronization process, following by replication and is alive checks.

- 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 period 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 XOssoft HA when it detects that the Master is unavailable. Alternatively, CA XOssoft HA can simply alert you to the problem, and then you can manually initiate switchover from the CA XOssoft 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, 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.

If you choose to initiate switchover automatically, after the Master is considered to be down, CA XOssoft HA automatically tries to restore the services and databases on it to their active state. First, CA XOssoft HA tries to restart the Exchange services that it previously checked. If the services are running, it then tries to mount the databases. If all attempts failed, CA XOssoft HA initiates a switchover. These attempts to restore the services and databases are not performed if the switchover is initiated manually.

Note: The Exchange services that are checked and managed by CA XOsoft are as follows -

- Exchange 2003:
 - Microsoft Exchange System Attendant
 - Microsoft Exchange Information Store
 - Microsoft Exchange MTA Stacks
 - Simple Mail Transport Protocol (SMTP)
 - Microsoft Exchange Routing Engine
- The following services are managed only if they are in automatic startup mode before the scenario is initiated:
 - Microsoft Exchange POP3
 - Microsoft Exchange IMAP4
 - Microsoft Exchange Management
 - Microsoft Exchange Site Replication Service
- Exchange 2007:
 - Microsoft Exchange Information Store
 - Microsoft Exchange Search Indexer


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

To initiate switchover

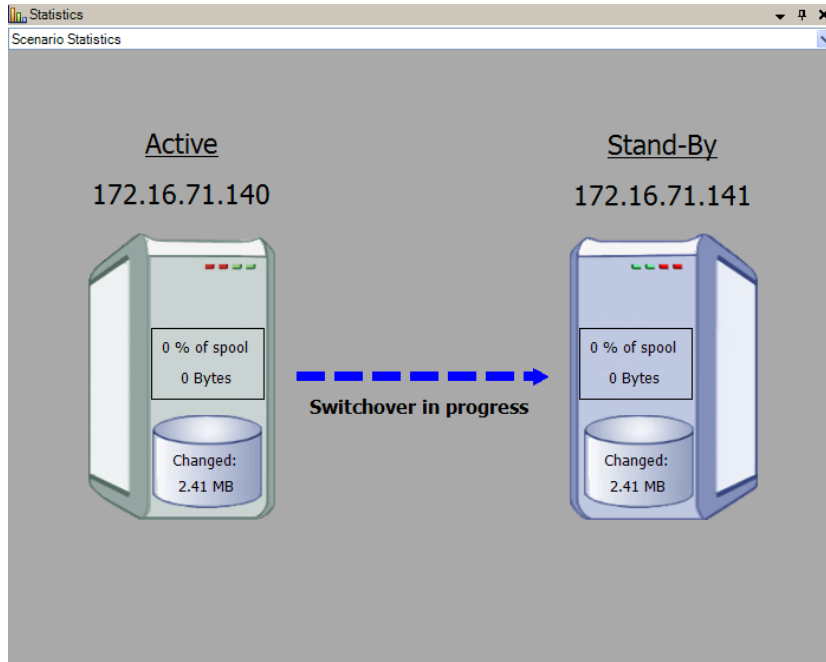
1. Open the Manager and select the desired scenario from the Scenario pane. Verify it is running.
2. Click on the **Perform Switchover** button, or select from the **Tools** menu the **Perform Switchover** option.



A confirmation message appears:



- 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.

- After the switchover is completed the scenario stops.

HA Scenarios					
Scenario	State	Product	Server	Mode	
MS Exchange HA	Stopped on Aut...	HA	Exchange	Online	
Hosts					
172.16.71.140	Changed	Synchronized	Files	In spool	
172.16.71.141					

Note: The only case in which the scenario may continue to run after switchover is when **automatic reverse replication** is defined as **Start automatically**.

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").

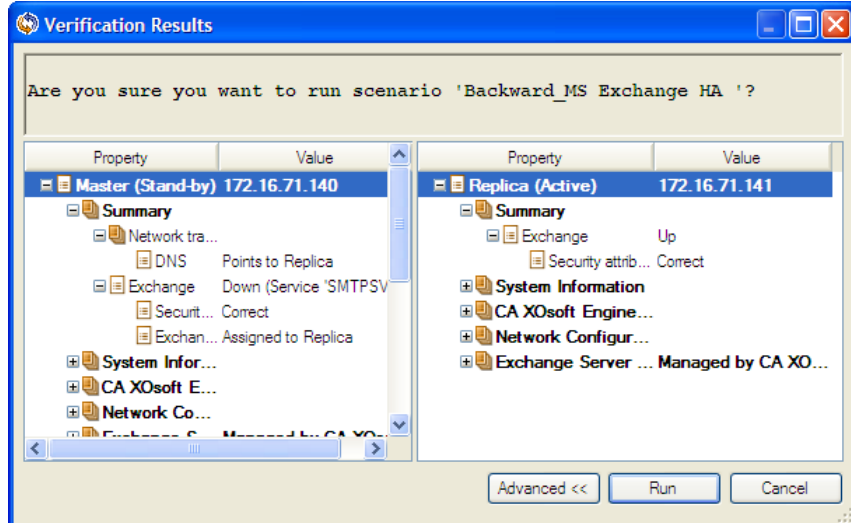
During the creation of the HA scenario, you defined how you want the reverse scenario to be initiated. If you selected in the **Switchover and Reverse Replication Initiation** page the **Initiate Reverse Replication automatically** option, replication in the reverse direction (from Replica to Master) 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** (see page 215) option.

To initiate switchback

1. Ensure that both Master and Replica servers are available on the network and that the CA XOssoft 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.]

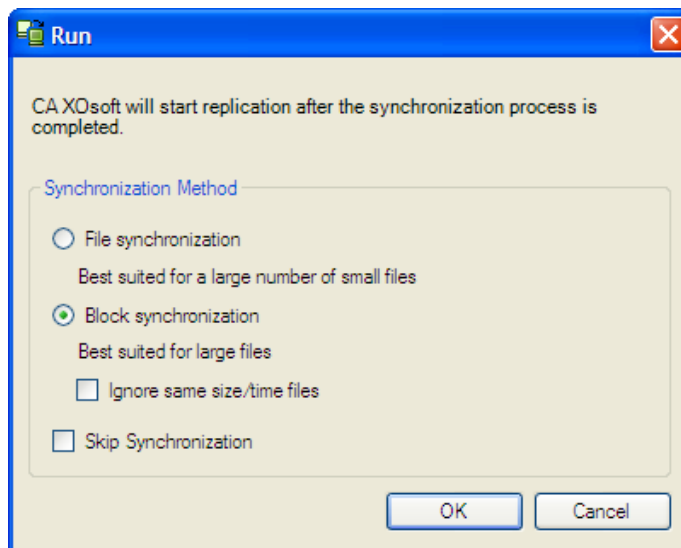
If the backward scenario is not running, select the **Run** button to start the scenario. CA XOsoft HA detects that a switchover has occurred, verify its state and configuration, and prompts you to approve the running of the backward scenario.



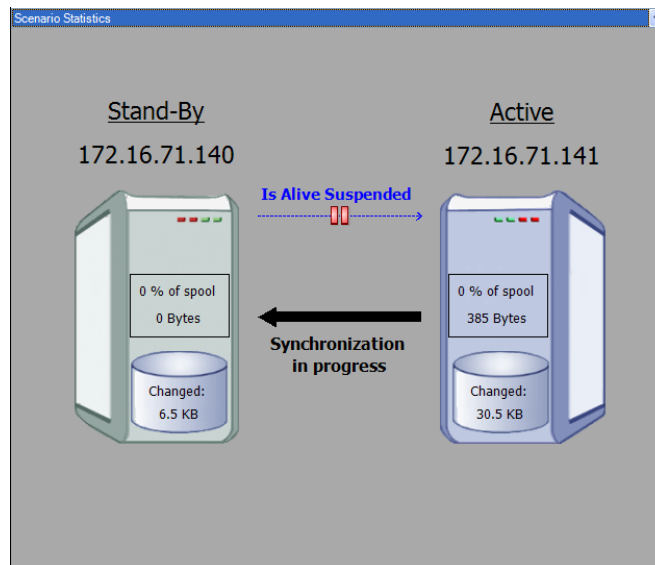
Note: The **Advanced** button opens an additional pane with detailed information about the hosts that participate in the scenario.

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

The **Run** dialog opens.

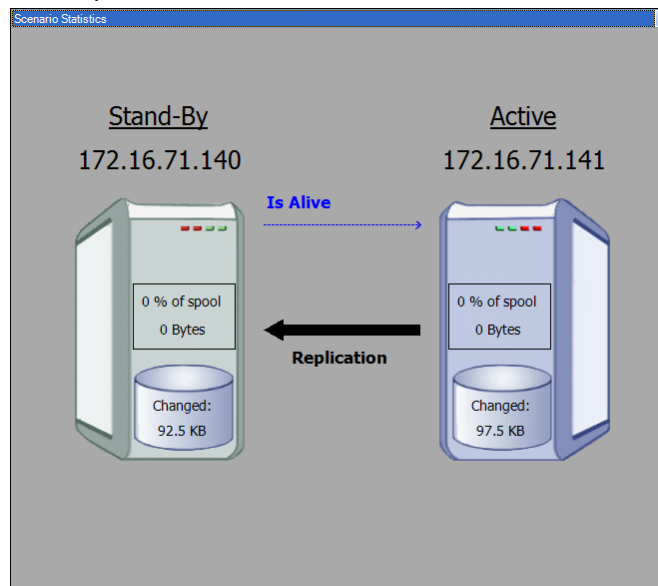



- For MS Exchange, select **Block Synchronization**, and click **OK**. The resynchronization starts.




Wait until the resynchronization is completed.

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

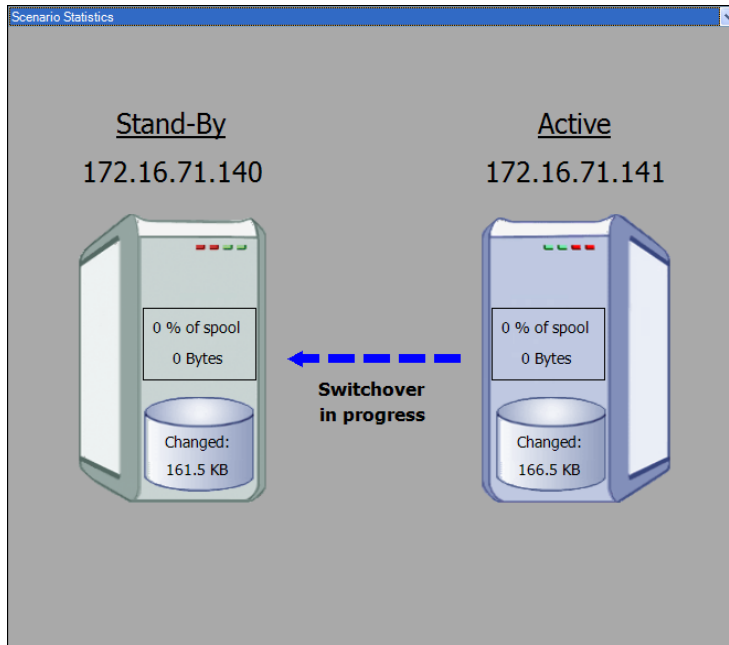


- Now, you can reverse back the roles between the Master and Replica servers. To reverse the roles, while the backward scenario is running, click the **Perform Switchover**  button, or select the **Perform Switchover** option from the **Tools** menu.

A confirmation message opens.



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



- After the switchover is completed, and the server roles are reversed back, the scenario automatically stops.

HA Scenarios				
Scenario	State	Product	Server	Mode
MS Exchange HA	Stopped on Aut...	HA	Exchange	Online
Hosts				
172.16.71.140	Changed	Synchronized	Files	In spool
172.16.71.141				

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

Recover Active Server Using the Manager

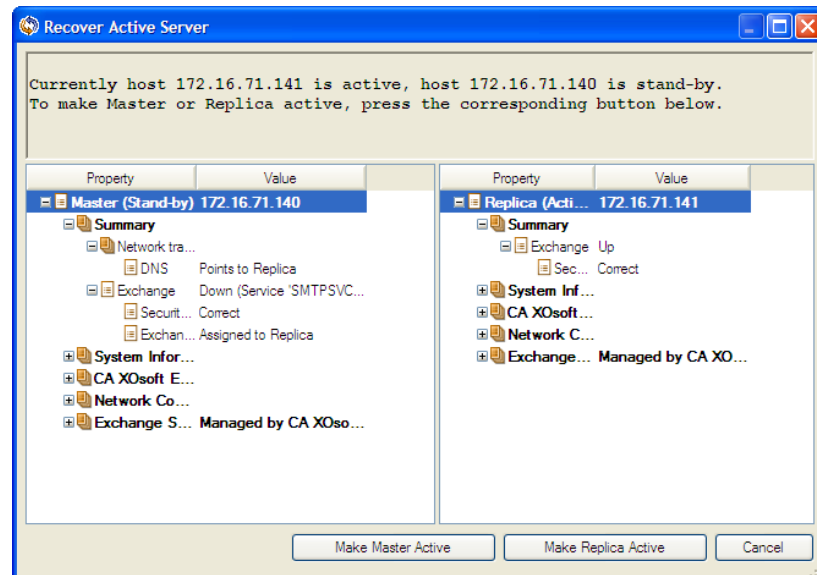
In certain circumstances, it may be necessary to intentionally make the Master or Replica server the active server without completing the synchronization process. This may happen when, for example, a switchover occurred but no data was changed on the Replica server, and you might even have newer data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

Important! While this option is the right choice in certain situations, use it with caution. If used improperly, data loss can occur. Normally, CA XOssoft 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 the **Recover Active Server** option, CA XOssoft 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 XOssoft 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.

Recover Active Server from Outside the Manager

If the switchover process does not complete correctly for some reason, and using the Recover Active Server option from the Manager does not resolve the issue

- First, perform the Recover Active Server procedure. For more information, refer to the topic, Recover Active Server.
- If the Recover Active Server procedure does not resolve the issue, try one or more of the following manual tasks appropriate to the redirection method you use:
 - If IP Redirection is used, manually remove the IP. You cannot use this method for scenarios that do not support Move IP redirection (Hyper-V HA, CS HA). For more information, refer to the topic, Manually Recover a Failed Server when IP Redirection is used.
 - If Switch Computer Name Redirection is used, manually switch the names. You cannot use this method for scenarios that do not support Switch Computer Name Redirection (Hyper-V HA, Exchange HA, vCenter HA if Oracle is used). For more information, refer to the topic, Manually Recover a Failed Server when Switch Computer Name Redirection is used.

If both IP and Switch Computer Name Redirection methods are used, manually remove the IP and switch the computer names. You cannot use this method for scenarios that do not support Move IP and Switch Computer Name redirection (Exchange, CS HA). For more information, refer to the topic, Manually Recover a Failed Server - IP and Switch Computer Name is used.

Manually Recover a Failed Server - Move IP Address

To recover a failed server when Move IP redirection is used

1. Boot the Master server without a network connection, to avoid IP conflicts.
2. From the **TCP/IP Properties** dialog, remove the additional IP address.
3. Reboot the Master server, and reconnect to the network.
4. If the HA scenario is not already running, start the scenario from the CA XOssoft Manager by clicking the **Run** button.

If the **Run Reverse Replication Scenario after Switchover** property was set to On, the scenario runs in backward mode so that the Replica server is now active and the Master server is on standby.

5. Wait for synchronization to complete.
6. To return the active role to the Master server, perform a manual switchover by clicking the **Perform Switchover** button on the Standard toolbar.

Note: We recommend that you perform the manual switchover outside of normal business hours.

Manually Recover a Failed Server-Switch Computer Name

To manually recover a failed server when the Switch Computer Name redirection method is used

1. Boot the Master server without a network connection, to avoid duplicate network names.
2. Rename the server to <NewServerName>-XO, and move it to a temporary workgroup.

For example, if the server is called "Server1", rename it to "Server1-XO".

3. You are required to reboot this machine.

After the reboot completes, the following error appears: "At least one Service could not be started." Ignore this, it is normal under these circumstances because the CA XOssoft Engine usually runs in a domain account.

4. Connect to the network.
5. Rejoin the domain, ensuring that you use the -XO name assigned in step 2.
6. Reboot the computer.
7. If the HA scenario is not already running, start the scenario from the CA XOssoft Manager by clicking the **Run** button on the Standard toolbar.

If the **Run Reverse Replication Scenario after Switchover** property was set to On, the scenario runs in backward mode so that the Replica server is now active and the Master server is on standby.

8. Wait for synchronization to complete.
9. To return the active role to the Master server, perform a manual switchover by clicking the **Perform Switchover** button on the Standard toolbar.

Note: We recommend that you perform the manual switchover outside of normal business hours.

Manually Recover Failed Server-IP and Switch Name

To manually recover a failed server when both IP and Switch Computer Name Redirection are used

1. Repair any hardware problems that could have caused the switchover, if any.
2. Reboot the server without a network connection to prevent IP conflicts.
3. From the **TCP/IP Properties** dialog, remove the additional IP address.
4. From the System Properties, Computer Name dialog, change the Computer Name to <ServerName>-XO. For example, if your server is called Server 3, rename it to Server 3-XO.
5. Assign the server to a temporary workgroup.
6. Restart the computer to enable your changes to take effect. When rebooting completes, reconnect to the network now. Ignore the message, "At least one service failed during system startup." This is normal because the CA XOsoft Engine runs in a domain, which is not currently available.
7. Rejoin the domain, making sure you use the -XO name, and reboot again.
8. The reverse scenario begins and the Replica server assumes the active role. Wait while synchronization completes.
9. Perform a manual switchover by clicking the **Perform Switchover** button from the toolbar, to return the active role to the Master server.

Chapter 11: Setting High Availability Properties

This section 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.

This section contains the following topics:

[Configuring High Availability Properties](#) (see page 221)

[Understanding High Availability Properties](#) (see page 223)

[Move IP Redirection](#) (see page 233)

Configuring High Availability Properties

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

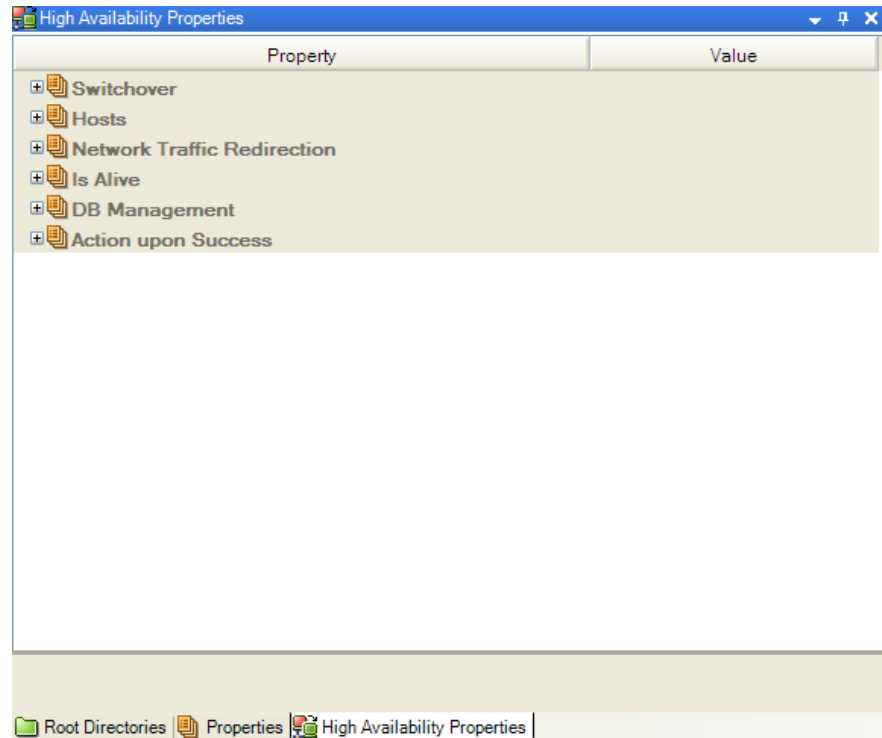
Notes:

- 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 opens.



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 (as defined in the [Is Alive properties](#) (see page 229)), 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 XOssoft 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 unsuspending 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**. Refer to [Is Alive](#) (see page 229).

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 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 XOssoft 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 XOssoft 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 XOssoft 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**.

When you use the **Switch Computer Name** redirection method on Windows 2008 systems, a reboot is required after switchover and switchback. We recommend setting the **Reboot After Switchover and Switchback** property to On when using this method.

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.

Note: The **Active to Standby Redirection Script** is also automatically executed when the HA scenario starts running. At this stage, the script runs on the standby Replica.

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.

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. For more information, refer to [Redirection properties](#) (see page 225).

- 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 XOssoft 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/Application/Shares Management

Automatic

If you want CA XOssoft to manage services on your DB Server or shares on your File Server, set this option to On. Then:

1. During the scenario creation, the list of application (DB) services or shares that must be managed are auto-discovered.
2. Once the scenario is running:
 - [DB] DB services on the active host are initiated (if they are not running), and they are stopped on the standby host (if they are running).
 - [FS] File shares are enabled on the active host (if they are disabled), and they are disabled on the standby host (if they are enabled).
3. During the switchover:
 - [DB] DB services on the active host are stopped, and they are started on the standby host.
 - [FS] File shares are disabled on the active host, and they are enabled on the standby host. The service that manages shares is stopped on the active host for a short period of time, and the restarted.

User-Defined Scripts

Start DB/Start Application/Enable Shares Script

If set to On, runs a user-defined script to augment or replace the start of DB services/the start of applications/the enabling of folder sharing. This action occurs during a scenario run on the active host, or during a switchover on the standby 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/Stop Application/Disable Shares Script

If set to On, runs a user-defined script to augment or replace the stop of DB services/the stop of applications/the disabling of folder sharing. This action occurs during a scenario run on the standby host, or during a 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 steps required for adding Move IP redirection to the High Availability scenario.

Important! Use this method only when both servers are on the same IP subnet.

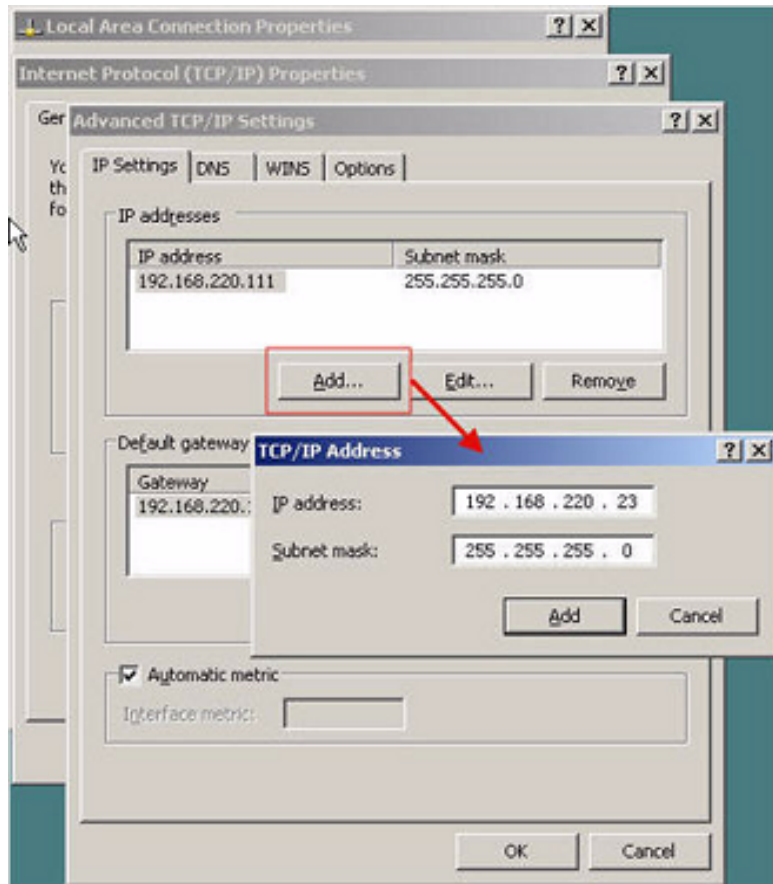
Adding IP Address to the Master Host

You need to add an additional IP address to the Master host, to use Move IP redirection in your HA scenarios. (This additional IP address is denoted as **XO-IP** in the following steps). This new IP address is used for CA XOsoft internal communication and replication. This is necessary because once switchover occurs, the original Master IP address is no longer available on the Master - it switches to the Replica host.

To add IP address to the Master host

1. Select **Start, Settings, Control Panel, Network Connections, Local Area Connection**.
The **Local Area Connection Status** dialog appears.
2. On the **General** tab, click the **Properties** button.
The **Local Area Connection Properties** dialog appears.
3. On the **General** tab, select **Internet Protocol (TCP/IP)**, and then click the **Properties** button.
The **Internet Protocol (TCP/IP) Properties** dialog appears.
4. On the **General** tab, click the **Advanced** button.
The **Advanced TCP/IP Settings** dialog appears.
5. On the **IP Settings** tab, click the **Add** button.

6. The **TCP/IP Address** dialog appears.



Note: In the screenshot, the XO-IP IP address is 192.168.220.23 and the current Master IP address is 192.168.220.111.

7. In the **TCP/IP Address** dialog, enter the additional IP address (XO-IP). Then, click the Add button.

The additional IP address is saved, and the **TCP/IP Address** dialog is closed.

8. Click **OK** on all the open dialogs, until you close all dialogs and exit the Lan settings.

Configuring the Move IP Method through the Manager

After you add the additional IP address to the Master host, you must add the XO-IP to your HA scenarios. There are two ways to add the XO-IP address to an HA scenario:

- For new scenarios, directly from the Scenario Creation Wizard.
- For existing scenarios, by modifying the Master host name.

The procedures for both ways follow.

Adding XO-IP to New Scenarios

To add XO-IP to a new scenario for Move IP redirection method

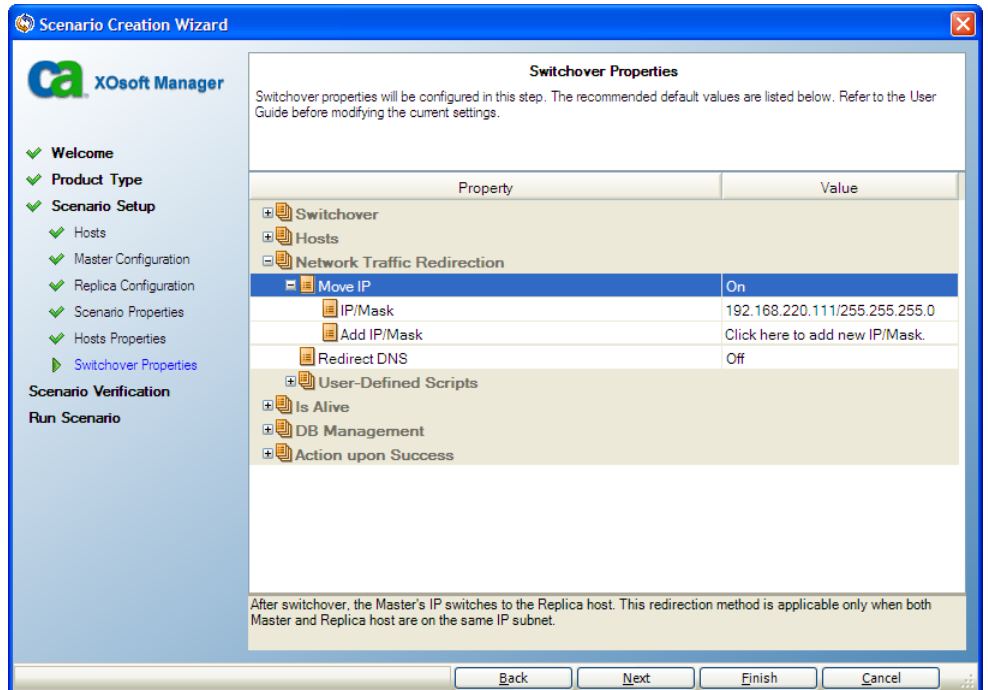
1. On the Scenario Creation Wizard, in the **Master and Replica Hosts** page, enter the following:
 - On the **Master Hostname/IP** box- enter the additional IP address (XO-IP).
 - On the **Replica Hostname/IP** box - enter the IP address of the Replica host, NOT its hostname.

The screenshot shows the 'Scenario Creation Wizard' window in XOsoft Manager. The main pane is titled 'Master and Replica Hosts' and contains the following fields and options:

- Scenario Name:** MS Exchange Server
- Master Hostname/IP:** 192.168.220.23 (with a red arrow pointing to it from the label 'XO-IP')
- Port:** 25000
- Replica Hostname/IP:** 192.168.220.24 (with a red arrow pointing to it from the label 'Replica Server IP Address')
- Port:** 25000
- Assessment Mode
- Verify CA XOsoft Engine on Hosts

The left sidebar shows the progress of the wizard, with 'Scenario Setup' and 'Hosts' selected. The bottom of the window has 'Back', 'Next', 'Finish', and 'Cancel' buttons.

- Click **Next**, and continue defining the scenario as usual until the **Switchover Properties** page appears.




- On the **Switchover Properties** page, open the **Network Traffic Redirection** group, select the **Move IP** property, and set its value to On. By default, the second IP address of the Master host appears here in the **IP/Mask** box.
Note: If the Master host has only one IP address, the **IP/Mask** box would be empty.
- If you have end users who connect to the Master host using its hostname, use the **Redirect DNS** or **Switch Computer Name** methods along with the **Move IP**. If you do not need to use the Master hostname, disable the **Redirect DNS** option by setting its value to Off.
- After setting the redirection method, click **Next**, and continue defining the HA scenario as usual.

Adding XO-IP to Existing Scenarios

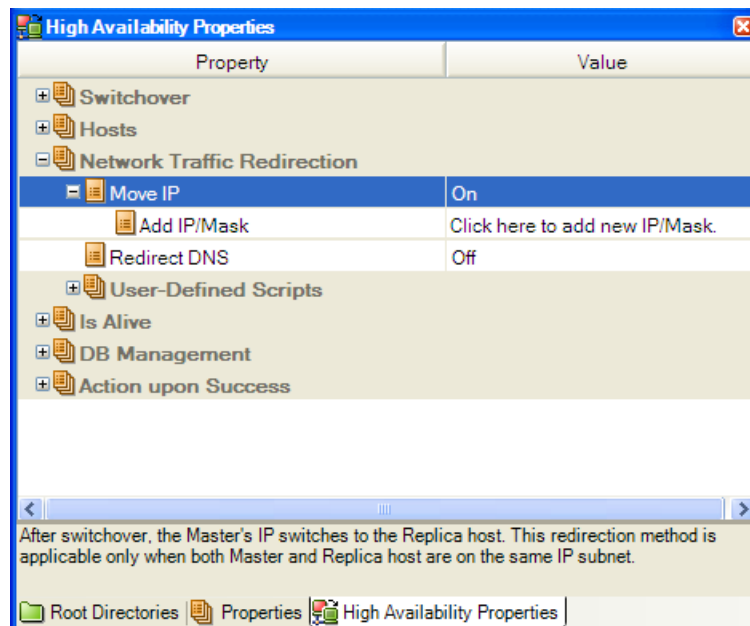
To add XO-IP to an existing scenario for Move IP redirection method

1. On the Scenario pane, select the required Master host.

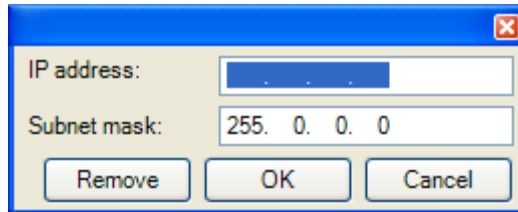


2. Right-click the Master and select **Rename** from the pop-up menu. Then, enter the **XO-IP** address.
3. Make sure that the Replica host is defined by its IP address and NOT by its hostname. If necessary, enter the Replica IP address instead of its hostname.
4. On the Framework pane, select the **High Availability Properties** tab..
5. Open the **Network Traffic Redirection** group, select the **Move IP** option, and set its value to On.

The **IP/Mask** property appears.



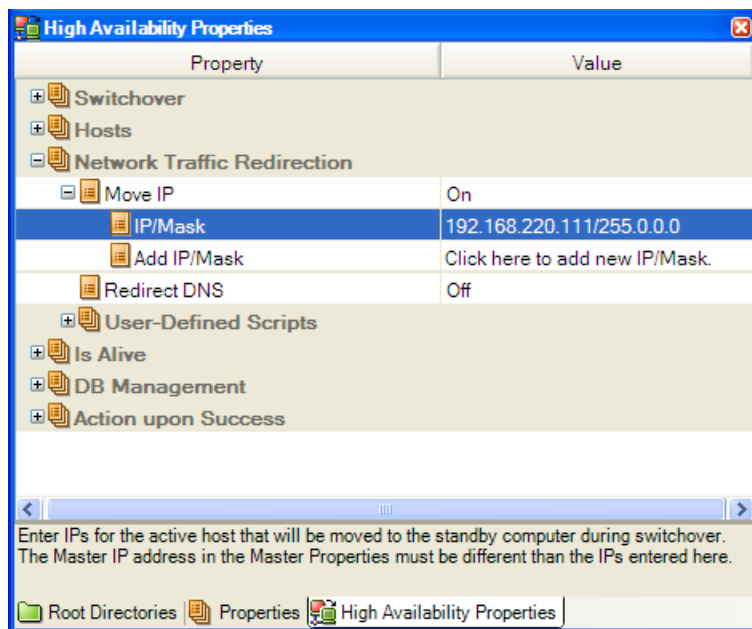
- Click the **IP/Mask** value box. The **IP Address** dialog appears.



The IP Address dialog box is a small window with a blue title bar and a close button (X) in the top right corner. It contains two text input fields: 'IP address:' and 'Subnet mask:'. The 'Subnet mask:' field contains the text '255. 0. 0. 0'. Below the input fields are three buttons: 'Remove', 'OK', and 'Cancel'.

- Enter the original IP address of the Master host. This IP address will be moved to the standby computer during switchover. Then, click **OK**.

Note: 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**.



The High Availability Properties dialog box is a larger window with a blue title bar and a close button (X) in the top right corner. It features a tree view on the left and a table on the right. The tree view includes 'Switchover', 'Hosts', 'Network Traffic Redirection', 'Move IP', 'IP/Mask', 'Add IP/Mask', 'Redirect DNS', 'User-Defined Scripts', 'Is Alive', 'DB Management', and 'Action upon Success'. The table on the right has two columns: 'Property' and 'Value'. The 'IP/Mask' row is selected and highlighted in blue, showing the value '192.168.220.111/255.0.0.0'. The 'Add IP/Mask' row contains the text 'Click here to add new IP/Mask.'. Below the table is a scrollable text area with the following text: 'Enter IPs for the active host that will be moved to the standby computer during switchover. The Master IP address in the Master Properties must be different than the IPs entered here.' At the bottom of the dialog, there are three buttons: 'Root Directories', 'Properties', and 'High Availability Properties'.

- If you have end users who connect to the Master host using its hostname, use the **Redirect DNS** or **Switch Computer Name** methods along with the **Move IP**. If you do not need to use the Master hostname, disable the **Redirect DNS** option by setting its value to Off.
- Click the **Save** button on the Standard toolbar to save your setting.

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

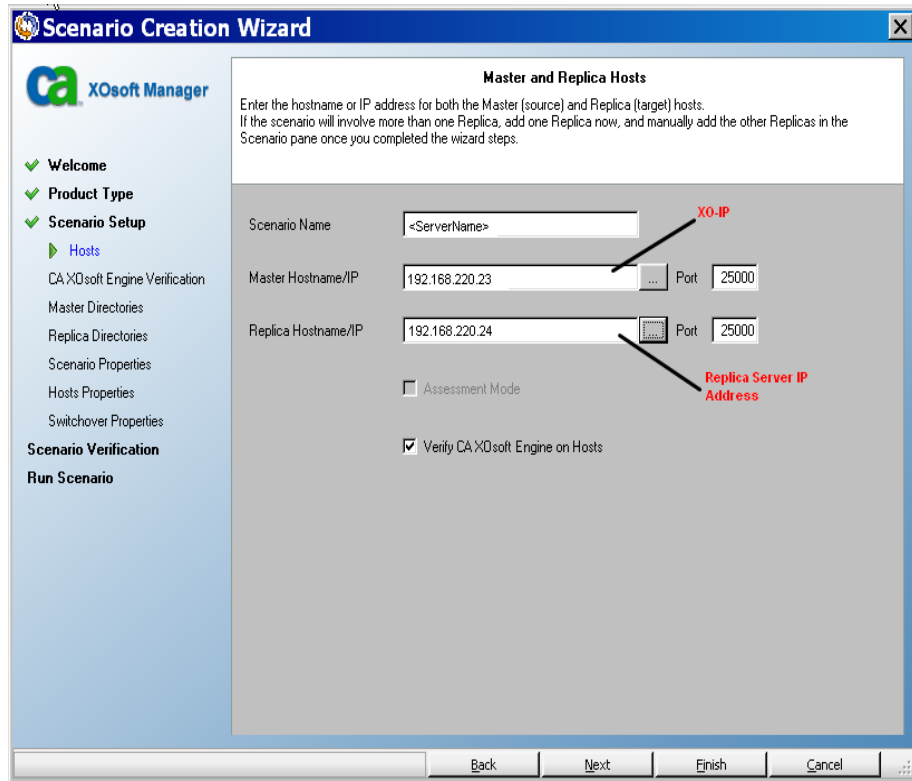
1. Open the Cluster Administrator.
2. In the Master Cluster Resource Group, create a new IP resource and name it **XO-IP**.
3. Bring this resource online, and verify it is visible from the Replica via the ping command. This new IP address is used for CA XOssoft HA internal communication and replication. 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. The following screen shows the XO-IP entered in the Master Hostname/IP field and the Replica Server IP address entered in the Replica Hostname/IP field.



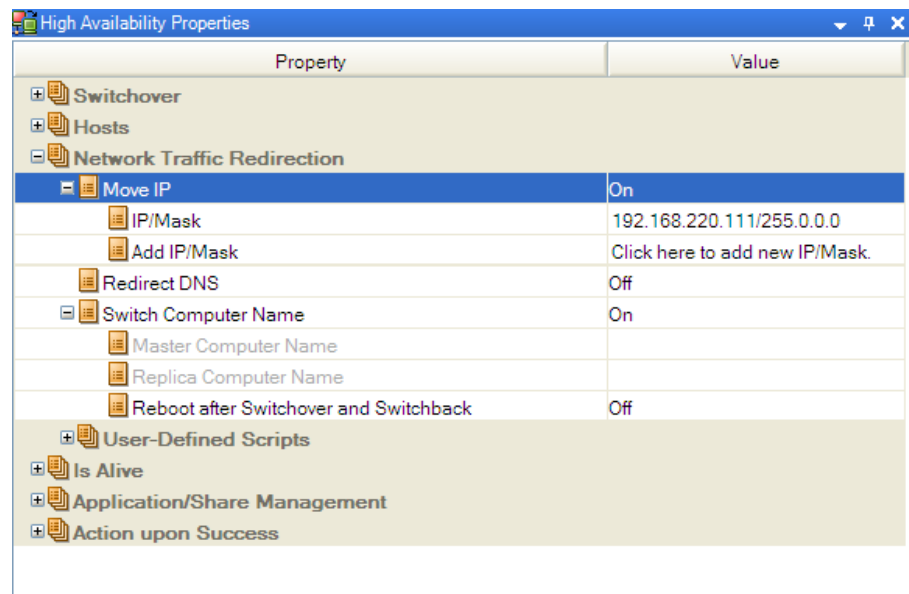
For Existing Scenarios

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: Data Replication and Role Switching of Control Services

This section explains how to replicate the Control Service data, and how to switch the roles of two Control Services when the active Control Service is down. The section describes in detail the following operations: creating and using HA Control Service scenario, and performing Control Service switchover and switchback.

This section contains the following topics:

[Understanding the Control Service Scenario](#) (see page 243)

[Opening the Manager for Using the HA Control Service Scenario](#) (see page 245)

[Creating High Availability Scenario for the Control Service](#) (see page 247)

[Switching the Roles of the Active and Standby Control Services](#) (see page 258)

Understanding the Control Service Scenario

The CA XOssoft Control Service functions as the single-point-of-control of the CA XOssoft operation, and it contains the entire data of the existing scenarios. In one CA XOssoft scenario-system, one Control Service manages all scenario-related-tasks, and the Managers that are connected to it enable you to monitor the CA XOssoft activities. If the Control Service is down, the scenario functioning is not affected. However, you are not able to control, manage and monitor the state and operation of the scenarios during this time. To overcome the danger of losing the Control Service data or losing the ability to manage and monitor your scenarios, CA XOssoft offers you the DR and HA Control Service scenarios. These scenarios enable you to protect the Control Service data and functionality, in the same way you protect other supported applications.

CA XOssoft enables you to replicate the Control Service data, and to save the replicated data on a Replica host. In order to perform this, you need to create a DR Control Service scenario. The DR scenario also enables you to activate the Rewind option, and to recover lost Control Service data if necessary.

In addition, CA XOssoft enables you to apply the HA solution to the Control Service. This means that if the active Control Service is down, you can switch the roles between the active Control Service and a standby Control Service, and make the standby Control Service the active one. For switching over and switching back the roles of two Control Services, you need to create an HA Control Service scenario.

Important! If you are running a DR Control Service scenario, you cannot use a second Control Service to manage your scenario-related-tasks. To use a second Control Service when the first one is down, you need to initially install two Control Services, one as the active Control Service and the second as the standby Control Service. You also need to install two Engines, one on each Control Service machine, and to verify that they are running. Only then you can create and run HA Control Service scenario.

The creation of DR and HA scenarios for CA XOsoft Control Service is similar to the creation of DR and HA scenarios for application and database servers. In both you are using the same step-by-step Scenario Creation wizard. However, there are some differences in the creation of DR and HA scenarios for the CA XOsoft Control Service, as follows:

- [DR and HA scenarios] Running only one scenario per Control Service - you can run only one Control Service scenario at a time for a specific Control Service.
- [DR and HA scenarios] No special license is needed - you do not need a special license for creating a Control Service scenario, either DR or HA. However, you do need [to register the CA XOsoft product](#) (see page 43) before creating a scenario for the Control Service.
- [DR and HA scenarios] Master details cannot be changed - In the **Master and Replica Hosts** page in the Scenario Creation Wizard, where you enter the IP address/hostname of the Master and Replica hosts, the Master host details are entered automatically by the system and cannot be changed. The Master Control Service details that appear in the wizard are the ones you entered in the Web browser for connecting the Control Service to the Overview Page.
- [HA scenario] Control Service items cannot be excluded from replication - in the **Master Configuration** page in the Scenario Creation Wizard, the auto-discovery results are read-only. You cannot exclude Control Service items from the replication process.
- [HA scenario] Move IP redirection method cannot be used - there are only two network traffic redirection methods you can use: Redirect DNS and Switch Computer Name. You cannot use the Move IP redirection method.
- [HA scenario] Automatic switchover and backward scenario cannot be disabled - you cannot disable the automatic initiation of a switchover when the Master is down, and the automatic initiation of a backward scenario. Therefore, the **Switchover and Reverse Replication Initiation** page in the Wizard and the corresponding properties are either not displayed or disabled. However, you can manually initiate a switchover and a switchback by using the **Perform Switchover** button on the Standard toolbar.
- [HA scenario] To create HA scenario for a Control Service, you need to install two Control Services: one should function as the active Control Service, and the other should function as the standby Control Service. For more information, refer to *CA XOsoft Installation Guide*.

To learn how:

- To create DR Control Service scenario, use the instructions for [Creating a Basic DR Scenario for File Server](#) (see page 46), along with the qualifications specified above.
- To recover Control Service data, refer to the [Data Recovery chapter](#) (see page 181).
- To create HA Control Service scenario, refer to [Creating High Availability Scenario for the Control Service](#) (see page 247).
- To manually initiate a switchover, refer to [Manually Initiating a Control Service Switchover](#) (see page 259).
- To handle a loss of connection and the switchover process, refer to [The Switchover and Backward Scenario Processes](#) (see page 260).
- To reverse back the Control Services to their original states, refer to [Switching Back the Control Services Roles](#) (see page 263).

Opening the Manager for Using the HA Control Service Scenario

To properly work with the HA Control Service scenario, it is important that you open the Overview Page, and from it the Manager, by using the Control Service hostname, instead of its IP address. If you will use the Control Service IP address, after a switchover CA XOssoft HA will not be able to automatically reconnect the Overview Page and the Manager to the new active Control Service.

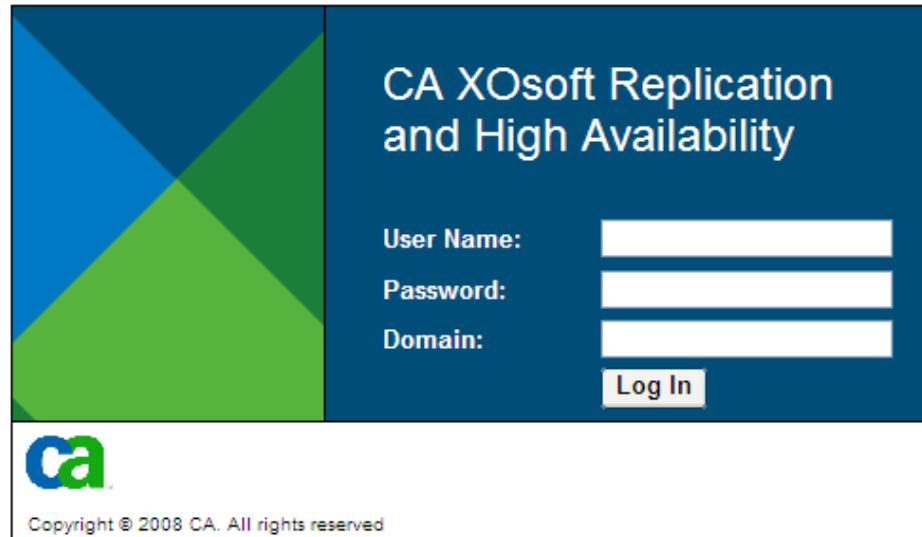
In addition, if you intend to work with HA Control Service scenario, you should NOT open the Manager from the machine where you installed a Control Service. Open the Manager from a third machine, which does not act as either the active or standby Control Service.

To open CA XOssoft Manager for working with HA Control Service scenario

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

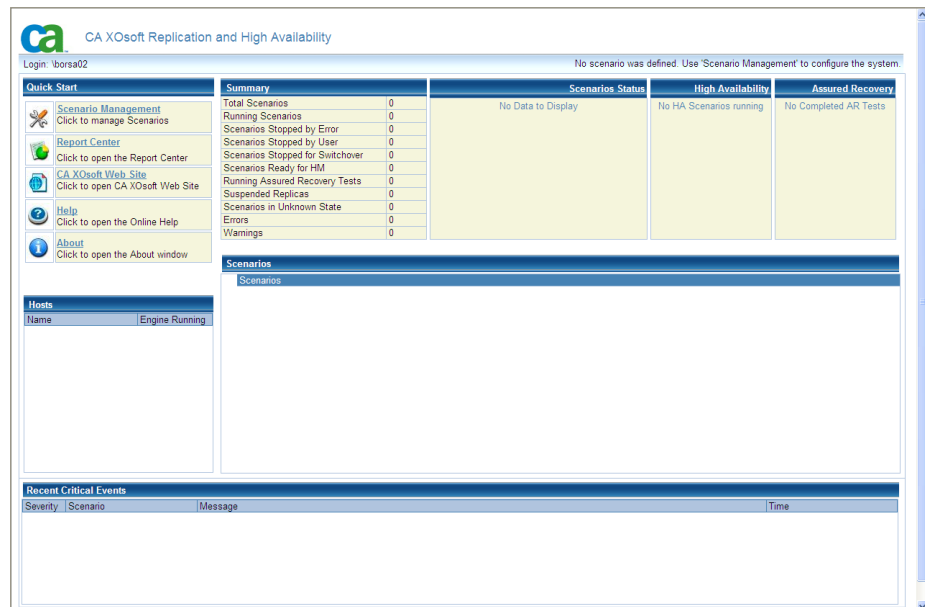
Note: If you selected the **SSL Configuration** option during the installation of the Control Service, enter the Control Service hostname and Port Number as follows: `https://host_name:port_no/start_page.aspx`

The **Login** dialog opens.



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

The **Overview** page opens.



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 opens.

Now, you can now start [creating the HA Control Service scenario](#) (see page 247).

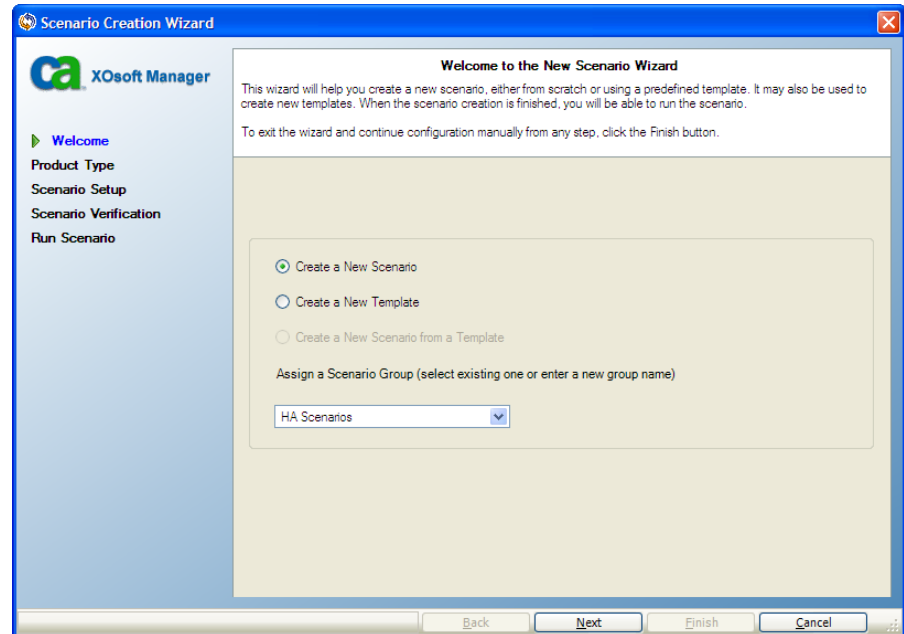
Creating High Availability Scenario for the Control Service

Important! Before you create HA scenario for the Control Service, you should verify that you have two Control Services installed, one as the (active) Master Control Service and one as the (standby) Replica Control Service. The Replica Control Service should be down. In addition, an Engine should be installed and running on both the Master and Replica hosts.

To create a High Availability scenario for the Control Service

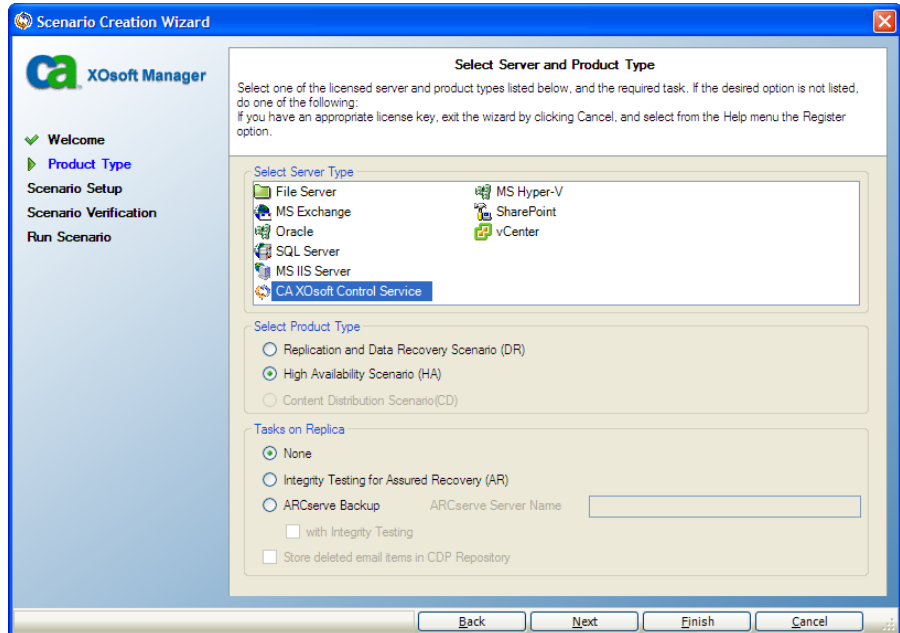
1. Open the CA XOssoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New**  button on the Standard toolbar.

The **Scenario Creation Wizard** opens.



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 opens.



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 **CA XOssoft Control Service**.
- From the **Select Product Type** options, select **High Availability Scenario (HA)**.
- [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 opens.

The screenshot shows the 'Master and Replica Hosts' configuration window. On the left is a sidebar with a tree view containing: Welcome, Product Type, Scenario Setup (with 'Hosts' selected), CA XOssoft Engine Verification, Master Configuration, Scenario Properties, Hosts Properties, Switchover Properties, Scenario Verification, and Run Scenario. The main panel has a title bar 'Master and Replica Hosts' and a subtitle 'Enter the hostname or IP address for both the Master (source) and Replica (target) hosts. If the scenario will involve more than one Replica, add one Replica now, and manually add the other Replicas in the Scenario pane once you completed the wizard steps.' Below this are several input fields: 'Scenario Name' with the value 'Control Service HA'; 'Master Hostname/IP' with the value 'qa95-w2k3-sql' and a 'Port' field with '25000'; and 'Replica Hostname/IP' with an empty field and a 'Port' field with '25000'. There are two checkboxes: 'Assessment Mode' (unchecked) and 'Verify CA XOssoft Engine on Hosts' (checked). At the bottom are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

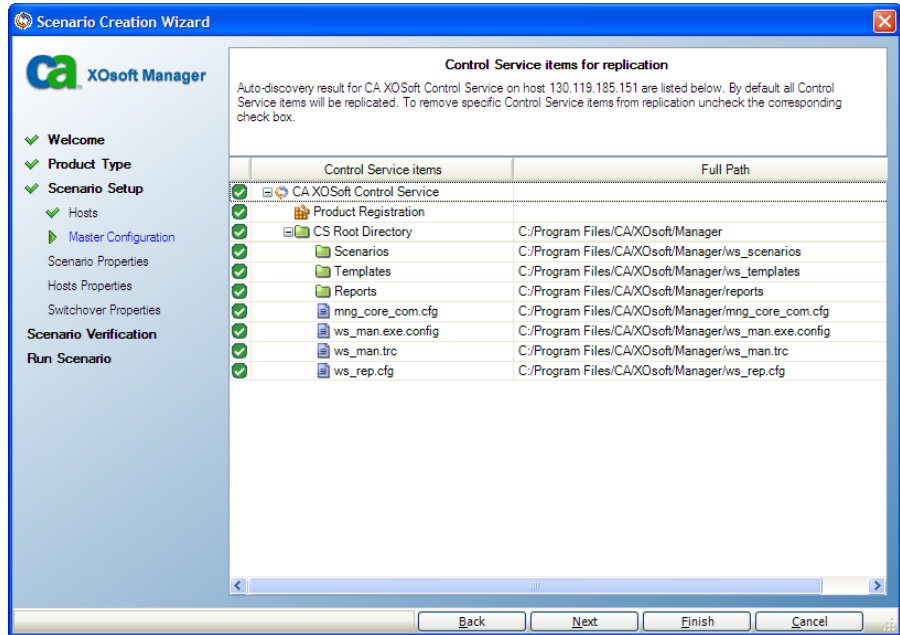
5. Enter the following information:

- 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 Hostname/IP** box - the system automatically enters the hostname or IP address of the (active) Master Control Service, based on the Control Service details you entered for opening the Overview Page. These details cannot be changed here. To use a different Control Service, you need to connect it to the Overview Page, and then reopen the Manager.
- In the **Replica Hostname/IP** box - enter the hostname or IP address of the Replica (standby) Control Service, or use the **Browse** buttons to find it.
- In the **Port** boxes - accept the default port no. (25000) or enter a new port numbers for the Master and Replica.

Notes:

- The **Assessment Mode** check box - verify that it is NOT selected.
- The **Verify CA XOssoft 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 [Creating a Scenario](#) (see page 46), step 8.

- After you selected the desired options, click **Next**. The **Master Configuration** page opens.



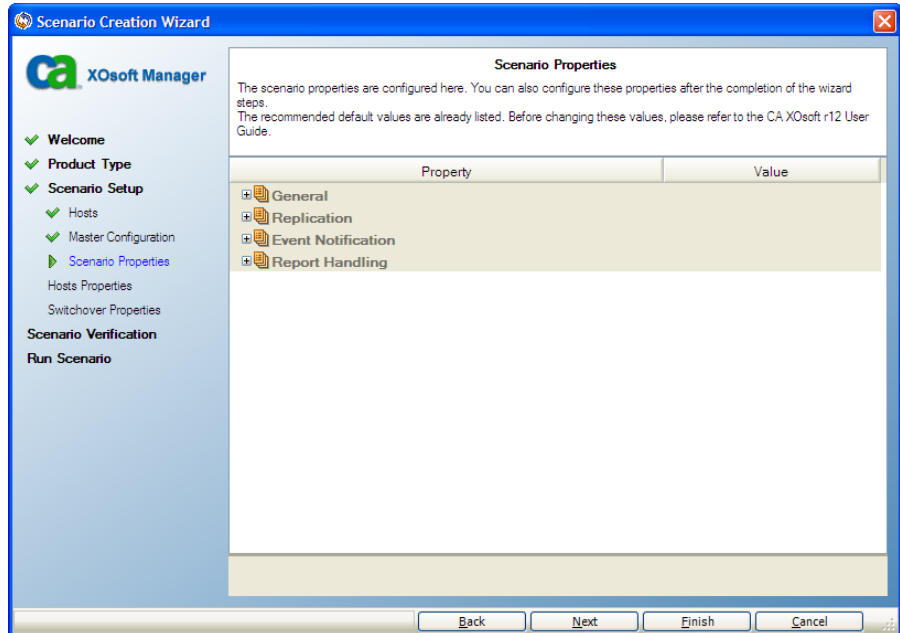
CA XOssoft auto-discovery component automatically displays the directories and files that are on the active Control Service. These directories and files are the data that will be replicated and protected.

The replicated Control Service items include:

- Product Registration - product registry keys
- Scenarios - xmc files of scenario definitions
- Templates - xmc files of user-defined templates
- Reports - files of scenario reports
- Configuration and Management files

Note: In HA Control Service scenario, you cannot exclude Control Service items from replication. In DR Control Service scenario, you can exclude the items you do not want to replicate, by clearing their check boxes.

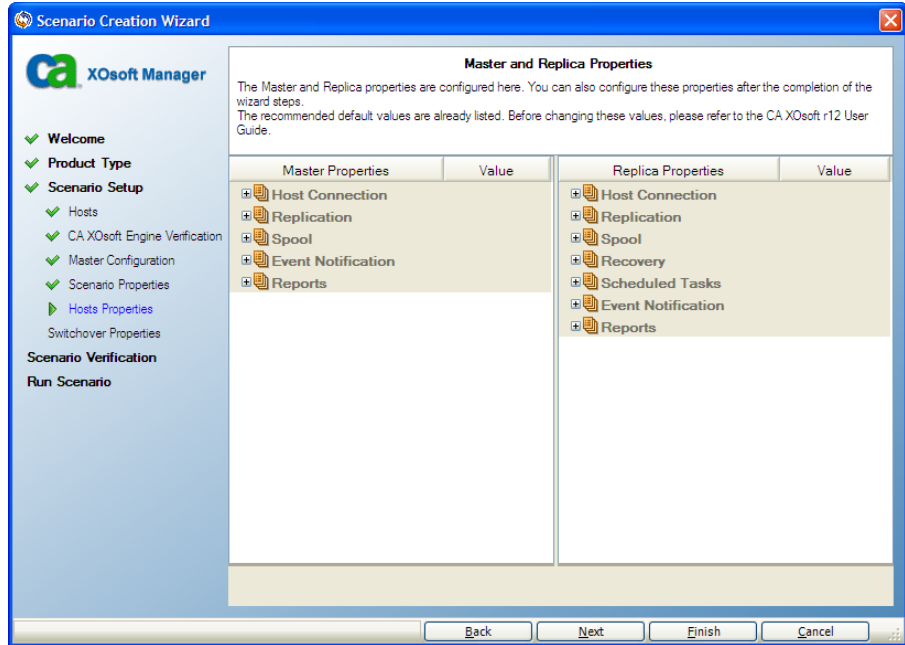
- Click **Next**. The **Scenario Properties** page opens.



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](#) (see page 138). To configure the scenario properties at a later stage, refer to [Configuring Scenario Properties](#) (see page 137).

- Click **Next**. The **Master and Replica Properties** page opens.

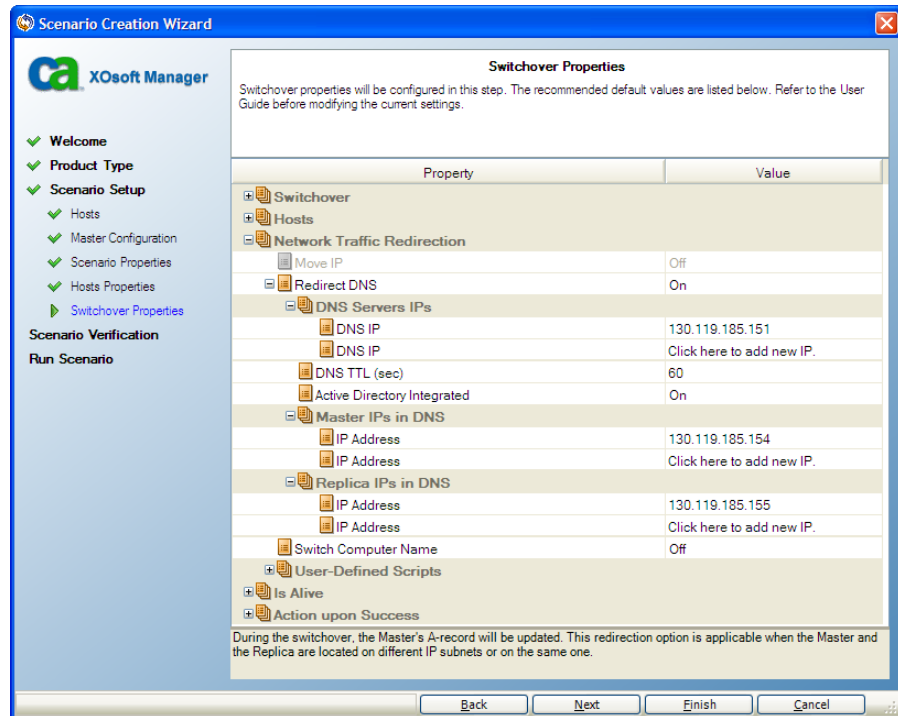


The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

If you want to configure the Master and Replica properties at this stage, refer to [Setting Master and Replica Properties](#) (see page 163). To configure the Master and Replica properties at a later stage, refer to [Configuring Master or Replica Server Properties](#) (see page 164).

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 information](#) (see page 167) for configuration details.

9. Once you are satisfied with the Master and Replica properties, click **Next**. The **Switchover Properties** page opens.



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](#) (see page 223). To configure the switchover properties at a later stage, refer to [Configuring High Availability Properties](#) (see page 221).

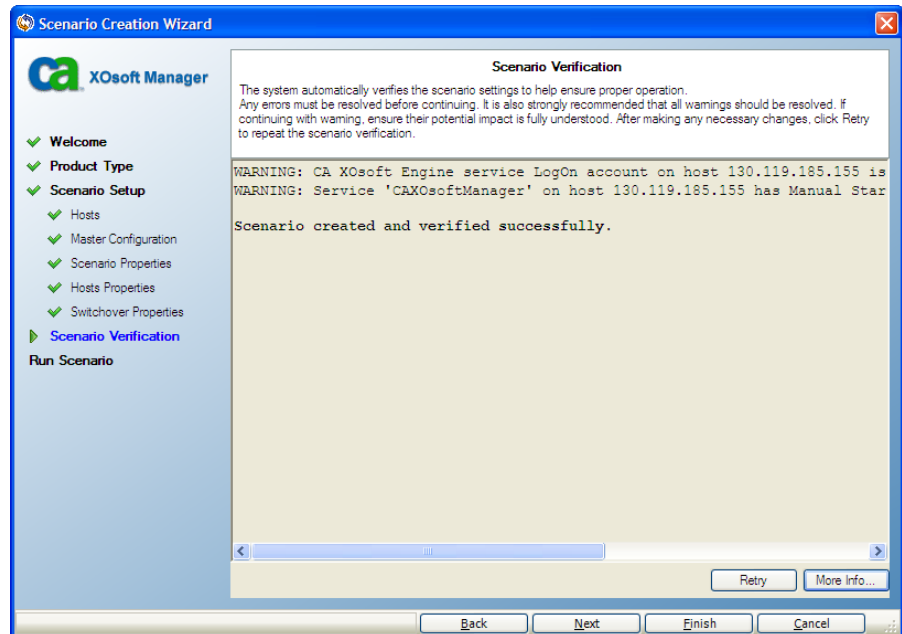
Notes:

- When selecting the **Network Traffic Redirection** method, there are only two methods you can use for this scenario: **Redirect DNS** and **Switch Computer Name**. You cannot use the **Move IP** redirection method.
- The **Is Alive Timeout (sec)** property controls how long to wait after a failure is detected before triggering an automatic switchover. The default is 300 seconds. For more details, review the [Is Alive information](#) (see page 229).

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

Note: In HA Control Service scenario, once a Master failure is detected, a switchover and a backward scenario are always initiated automatically. You cannot disable this automatic initiation. However, you can also manually initiate a switchover, by clicking the **Perform Switchover** button.

11. Once the verification is completed the **Scenario Verification** page opens.

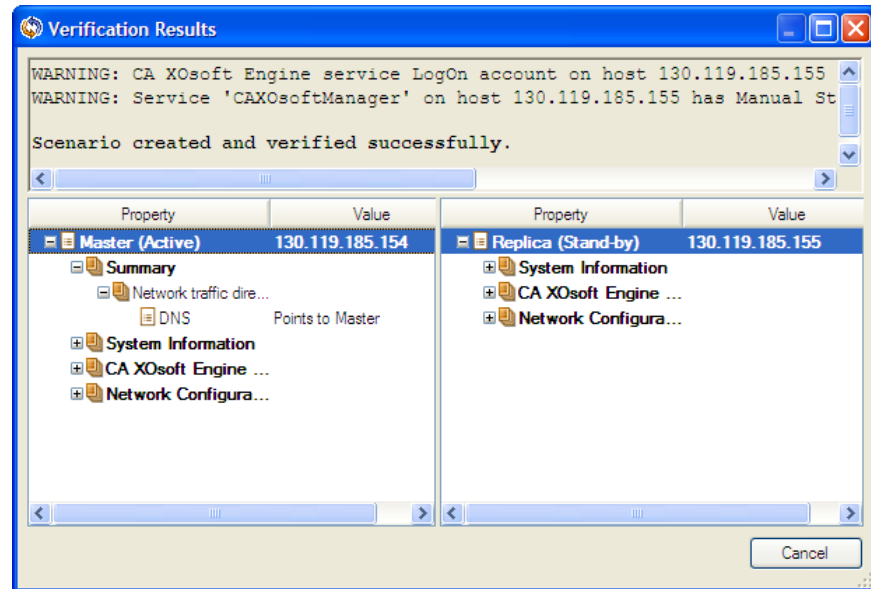


12. If the scenario was not set up correctly, or problems occurred in the participating hosts or the connection between the CA XOssoft components, the errors and warnings detected are displayed, and two additional buttons appear: **Retry** and **More Info**.

13. To repeat the verification process, click the **Retry** button.

14. To view additional information about the errors and warnings, click the **More Info** button.

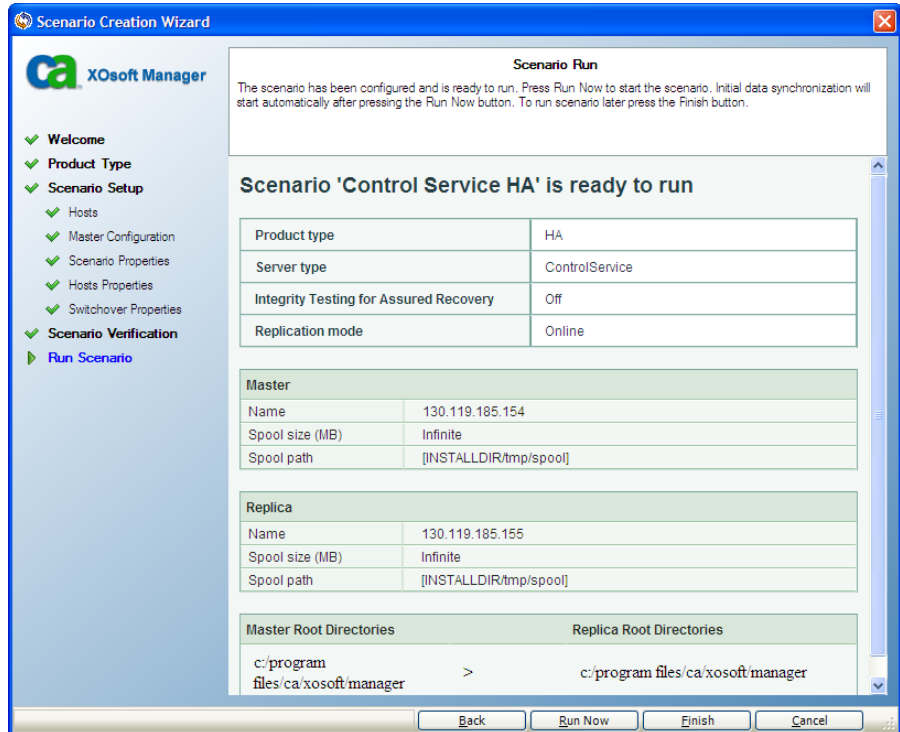
The **Verification Results** dialog opens, listing all the errors and warnings detected.



The **Verification Results** dialog provides you 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.

- 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. 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. To close the dialog and return to the **Scenario Verification** page, click the **Cancel** button.
- When the scenario is verified successfully, on the **Scenario Verification** page click **Next** to continue.

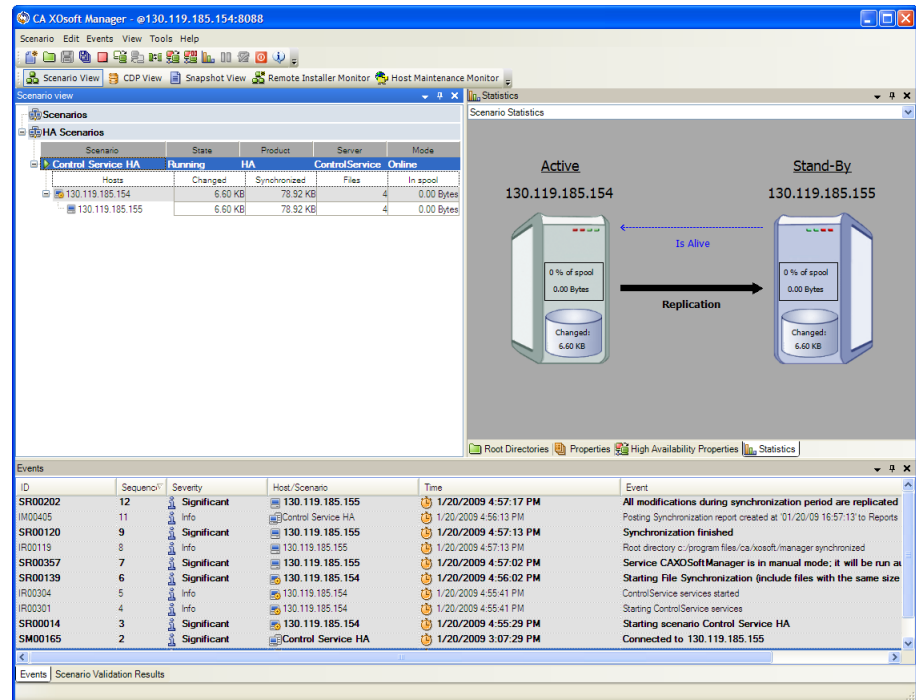
The **Scenario Run** page opens.



15. 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**.

- 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 period are replicated.**

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



Switching the Roles of the Active and Standby Control Services

The shutting down of a Control Service, either intentionally or due to a crash, presents a unique problem. When the Control Service is down, the Overview Page and the Manager are disconnected. Consequently, they can no longer receive updated information, and they cannot display a visible indication to the type of event that occurred. Although the Overview Page and the Manager are up, you cannot see that the Control Service is down, and you cannot manually initiate a switchover at this stage.

CA XOssoft handles the disconnection problem by automatically trying to restore the Control Service service ("CAXOssoftManager") to its active state. If the attempt fails, and the active Control Service is still detected as down, CA XOssoft automatically initiates a switchover. During the switchover, the standby Control Service becomes the active Control Service. Following that, the Overview Page and the Manager are automatically reconnected to the new active Control Service, and once again they display the updated state of your system. During the reconnection, you may be prompted to log in again.

When the original active Control Service is up again, CA XOssoft automatically initiates a backward scenario. The backward scenario is a replication in the reverse direction: from the new active Control Server to the new standby Control Server. At this stage, you can reverse back the roles of the Control Services. All you need to do is to manually initiate a switchback, meaning, a switchover in the opposite direction.

In addition to the default automatic switchover, you can also manually initiate a switchover between the active and standby Control Services. Once triggered, whether manually or automatically, the switchover process itself is fully automated.

There are several stages in the role switching of the active and standby Control Services:


1. [Initiating a switchover](#) (see page 259) - this can be done either automatically by the system, when it detects that the active Control Service is down, or manually by you.
2. [The switchover process and the automatic initiation of a backward scenario](#) (see page 260) - these processes are performed automatically and you cannot disabled them.
3. [Initiating a switchback](#) (see page 263) - this can only be done manually by you, when you decide that the original active Control Service can become the active server again.

Manually Initiating a Control Service Switchover

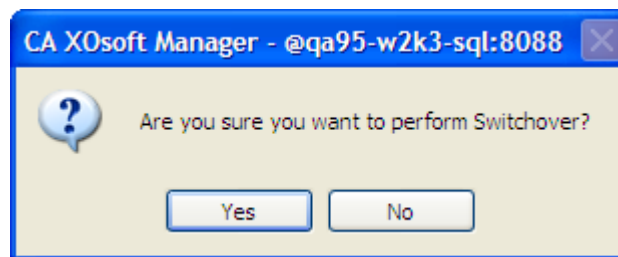
Once CA XOsft HA detects that the active Control Service is down, it automatically tries to restart the Control Service service, and once this attempt fails, it initiates a switchover. However, you can also manually initiate a switchover, when the active Control Service is still up.

To manually initiate a switchover

1. Open the Manager and select the desired Control Service 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 opens.




3. Click **Yes** on the **Perform Switchover** confirmation message. This procedure initiates a switchover from the active Control Service to the standby Control Service.

From this stage on, [the switchover process](#) (see page 260) is the same for both manual and automatic initiation.

The Switchover and Backward Scenario Processes

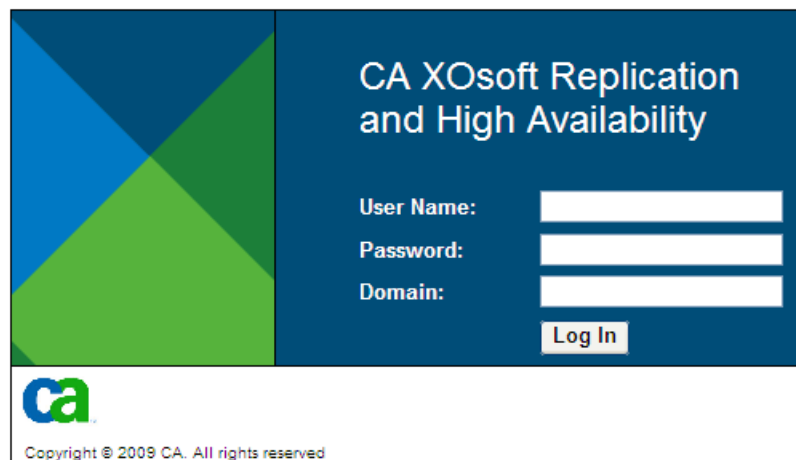
Understanding the switchover and backward scenario process

1. Since the original active Control Service is down, the Overview Page and Manager are no longer connected to it. Therefore, they no longer receive and display updated information, and the changes that occur following the switchover initiation are not shown in them, as they are shown in a regular switchover.
2. When the Overview Page loses its connection to the original active Control Service, the following message opens.

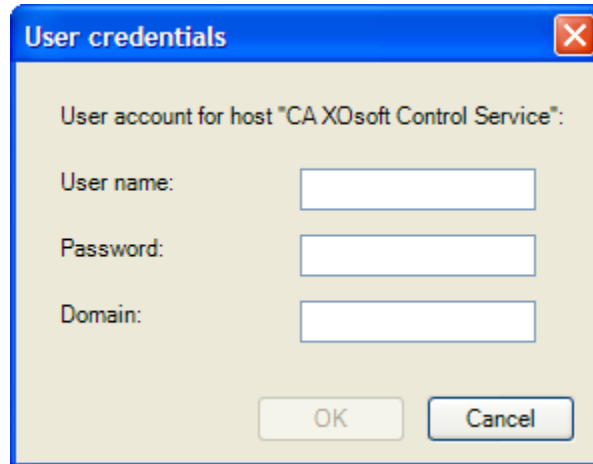


This message indicates that the original active Control Service is down, and therefore it is no longer connected to the Overview Page.

3. Click **OK** to close the message. This message may appear several times until the original standby Control Service becomes active, and a connection to it is established.
4. When the original standby Control Service is up and functioning as the new active Control Service, the Overview Page is automatically reconnected to it, and the **Login** dialog appears, prompting you to login to the new active Control Service.



5. Enter your User Name, Password and Domain and click the **Log In** button.
The **Overview page** re-appears, and it is now connected to the new active Control Service.
6. On the Manager, the **User Credentials** dialog may open.



The **User credentials** dialog prompts you to login to the new active Control Service. If this dialog appears, enter the necessary details and click **OK**.

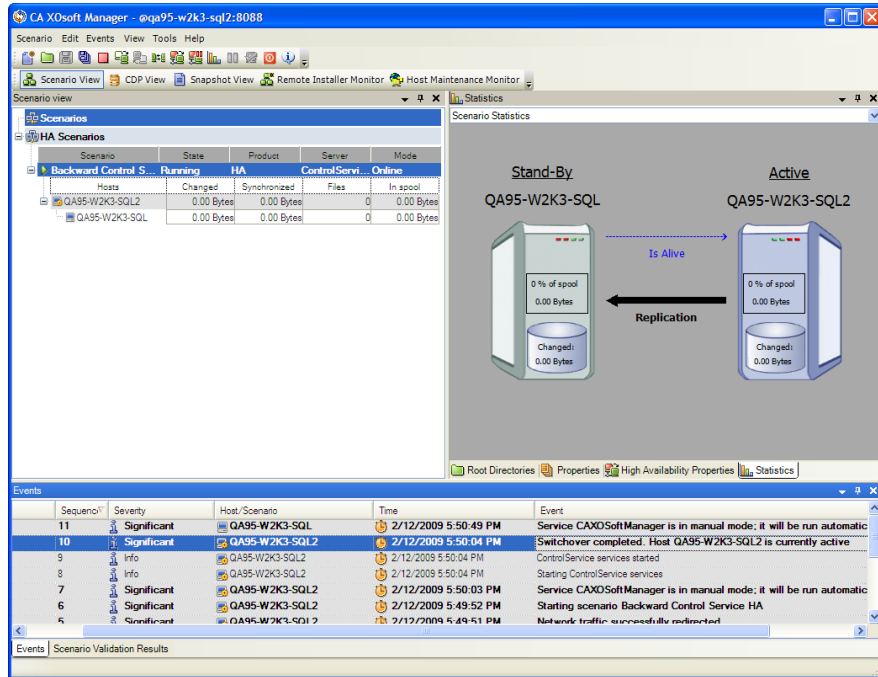
Note: The appearance of the **User credentials** dialog is related to internal caching settings, and it is not necessarily an indication to the progress of the switchover process. The switchover may take place even if the **User credentials** dialog does not appear.

7. The two Control Services have switched roles. Now, the Manager is no longer connected to the original active Control Service but to the standby Control Service, which became active following the switchover. The switchover related-events are displayed in the Event pane.

Note: The "Split Brain" problem and solution

After a connection loss and a switchover, the original standby Control Service is functioning as the active Control Service. However, the original active Control Service may still be up. Upon reconnection, both Control Services may try to act as the active Control Service. In order to solve this potential problem, CA XOsoft HA keeps built-in numerical parameter in each Control Service, and the switchover process increases the number of the newly active Control Service. All connection requests are sent with this parameter, and when a Control Service receives a connection request, it checks whether it contains a lower or higher number than the one it carries. The Control Service that carries the lower number, shuts itself down, and becomes the standby Control Service.

8. After the switchover, a backward scenario is automatically initiated by the system.




9. The backward scenario starts running once the original active Control Service is up. It replicates data from the new active Control Service to the new standby Control Service, while overwriting the data on the current standby Control Service.
10. Now, you can [switch back the roles of the active and standby Control Service](#) (see page 263), and make the original Master the active server again and the Replica the standby server.


Switching Back the Control Service Roles

When the original active Control Service is up again, and the backward scenario is running, you can switch back the roles of the standby and active Control Service and reverse them back to their original states.

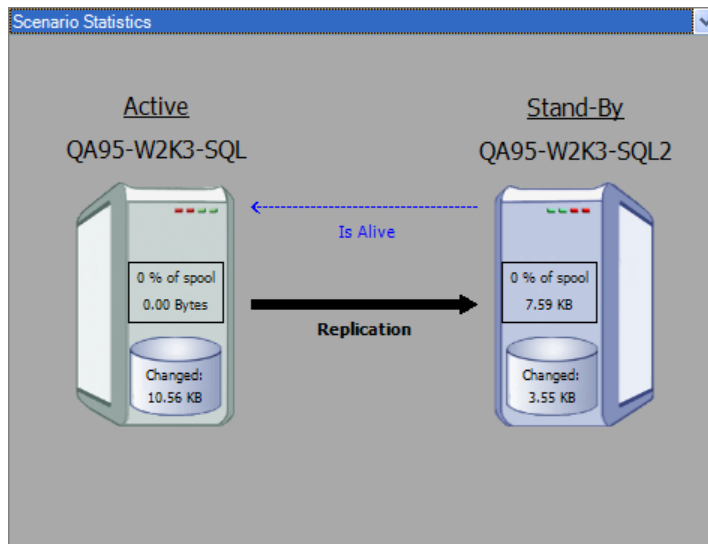
To initiate a switchback

1. To reverse back the roles of the Control Services, while the backward scenario is running, click the **Perform Switchover**  button or select the **Perform Switchover** option from the **Tools** menu.

A confirmation message opens.



2. Click **Yes** on the **Perform Switchover** confirmation dialog. This procedure initiates a switchback from the original Replica server to the Master server. Again, you are not able to see the process of the switchback on the Overview Page and the Manager, since they are disconnected from the active Control Service. But once the Overview Page and Manager are reconnected, you can see that the Control Services have switched their roles and returned to their original states.



Now, the Control Service HA scenario is running in its original direction.

Chapter 13: Testing Assured Recovery and Managing VSS Snapshots

This section 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 section describes how to set up VSS snapshot creation and how to manage snapshots.

This section contains the following topics:

[About Assured Recovery](#) (see page 265)

[Creating Assured Recovery Testing Scenario](#) (see page 267)

[Performing Assured Recovery Test](#) (see page 275)

[Setting Assured Recovery Properties](#) (see page 284)

[Using VSS Snapshots](#) (see page 289)

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.

Note: The Assured Recovery option is not available for backward (reverse replication) scenarios.


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

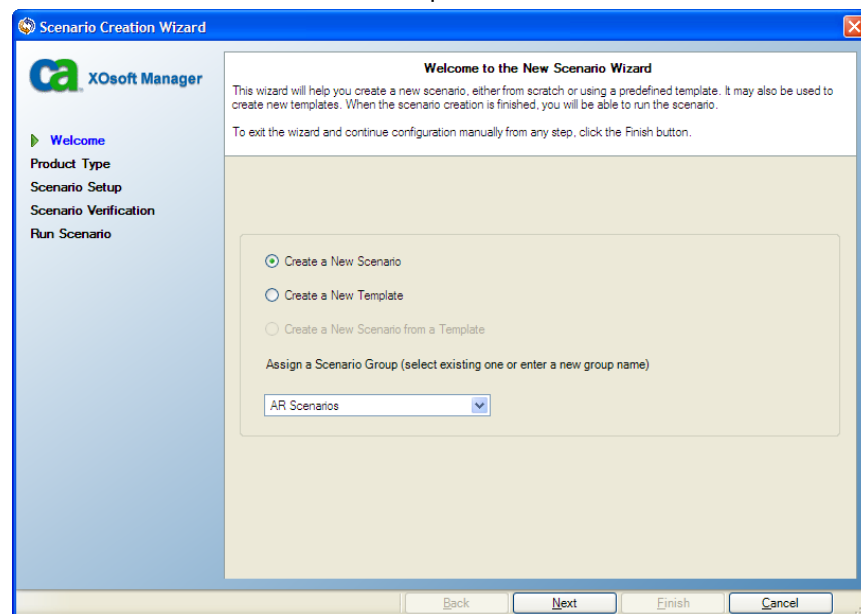
The Assured Recovery testing feature must be enabled during the creation of the scenario that will later use it. For this reason, you can not perform Replica testing within a DR or HA scenario that is already running, and was not configured to use the Assured Recovery option. To use Assured Recovery, 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

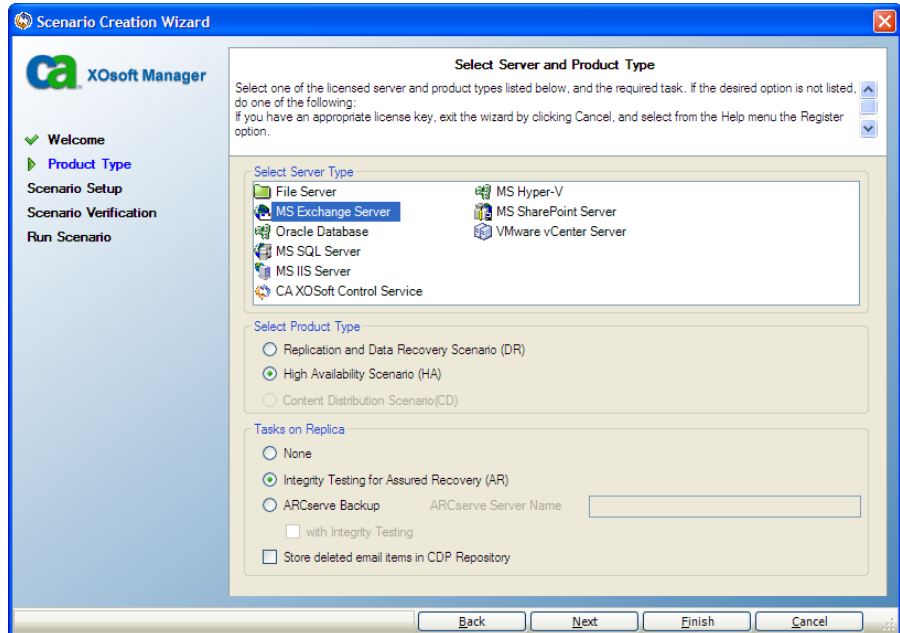
1. Open the CA XOssoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New**  button on the Standard toolbar.

The **Scenario Creation Wizard** opens.



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.

- Click **Next**. The **Select Server and Product Type** page opens.



- 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 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 opens.

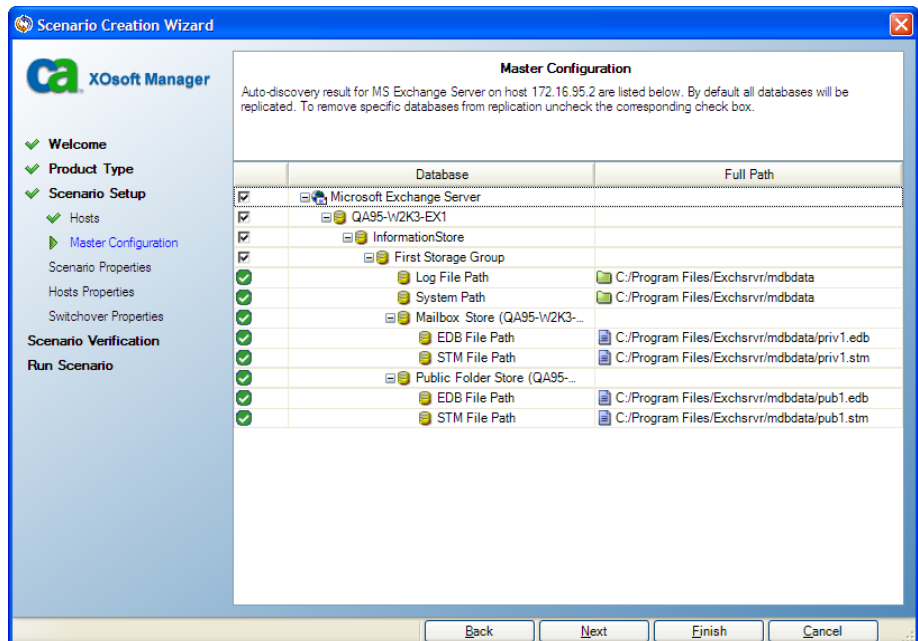
6. Enter the following information:

- 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 Virtual Server Name or IP address as the Master and/or Replica name (instead of the physical node's name/IP).

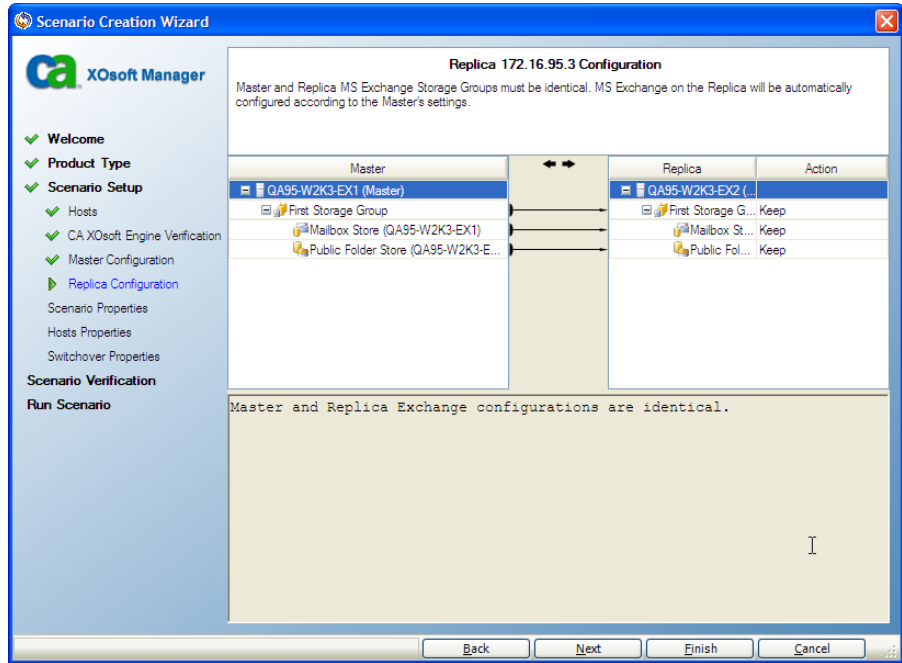
- In the **Port** boxes: accept the default port no. (25000), or enter a new port numbers for the Master and Replica.
 - The **Verify CA XOssoft Engine on Hosts** check box - select this check box if you want the system to verify whether Engines are installed and running on the Master and Replica hosts 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, refer to [Creating a Scenario](#), (see page 46) step 8.
7. After you entered or selected the desired options, click **Next**. The **Master Configuration** page opens.



The CA XOssoft 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.

- Click **Next**. The **Replica Configuration** page opens.



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.

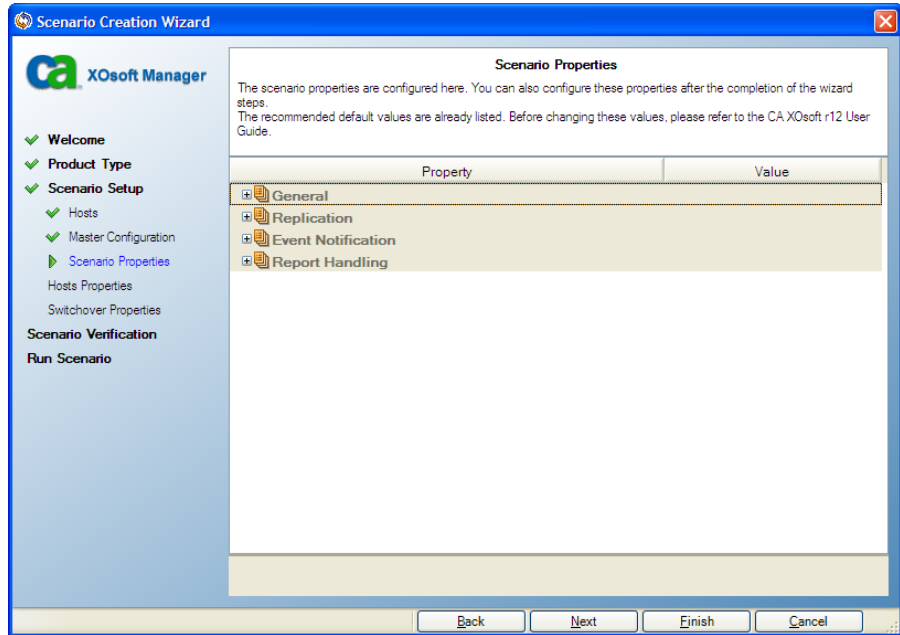
Note: For more information about the actions that can be performed in the Replica auto-configuration process, refer to [Creating High Availability Scenario](#) (see page 195), step 8.

- 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.

Important! You cannot use UNC paths as root directories on the Replica host for Assured Recovery scenario.

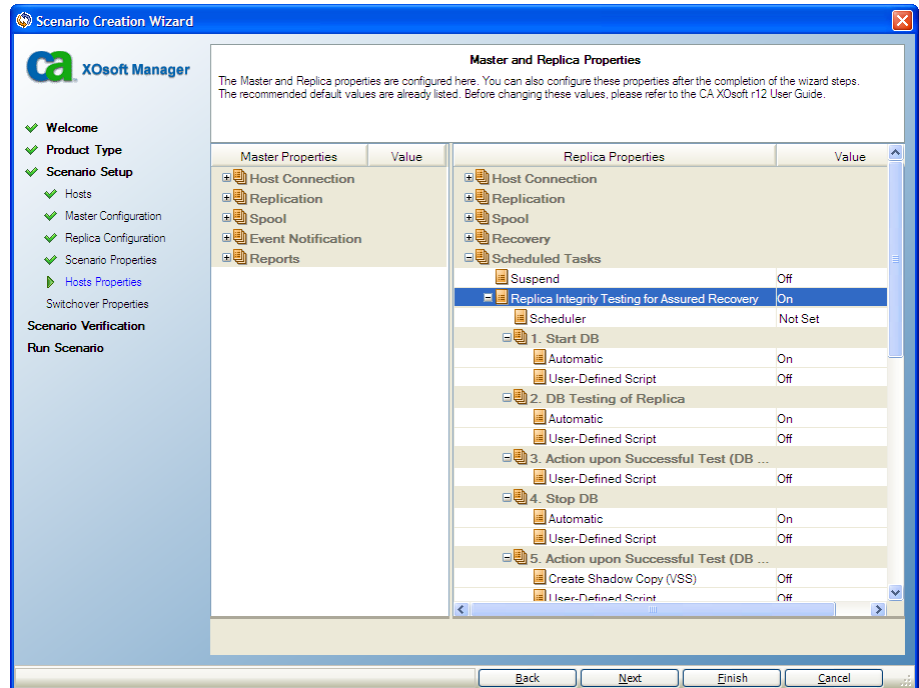
11. Click **Next** to start the Replica configuration process. The **Scenario Properties** page opens.



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](#) (see page 138). To configure the scenario properties at a later stage, refer to [Configuring Scenario Properties](#). (see page 137)

12. Click **Next**. The **Master and Replica Properties** page opens.



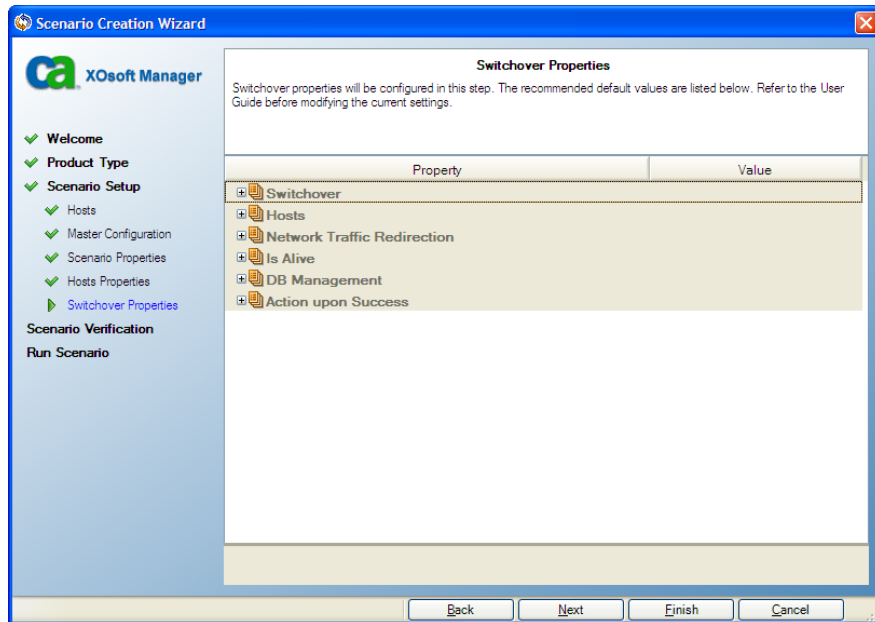
The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

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](#) (see page 286).

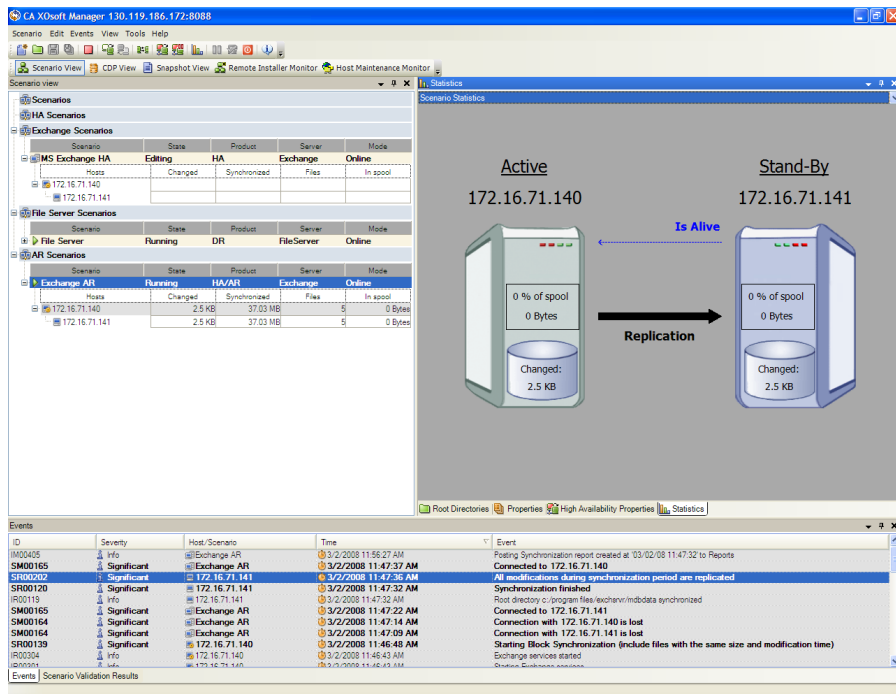
If you want to configure the Master and Replica properties at this stage, refer to [Setting Master and Replica Properties](#) (see page 163). To configure the Master and Replica properties at a later stage, refer to [Configuring Master or Replica Server Properties](#) (see page 164).

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 information](#) (see page 167) for configuration details.

- Click **Next**. If you selected HA solution the **Switchover Properties** page opens.



- From this stage, set up the scenario as you would normally following the instructions in [Creating DR Scenario](#) (see page 46), [Creating HA Scenario](#) (see page 195), 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 XOssoft 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 XOssoft at the beginning of each AR test.
3. Initiating a rewind component on the tested Replica - this step is performed automatically by CA XOssoft. 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 XOssoft. However, it can be switched off, 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.

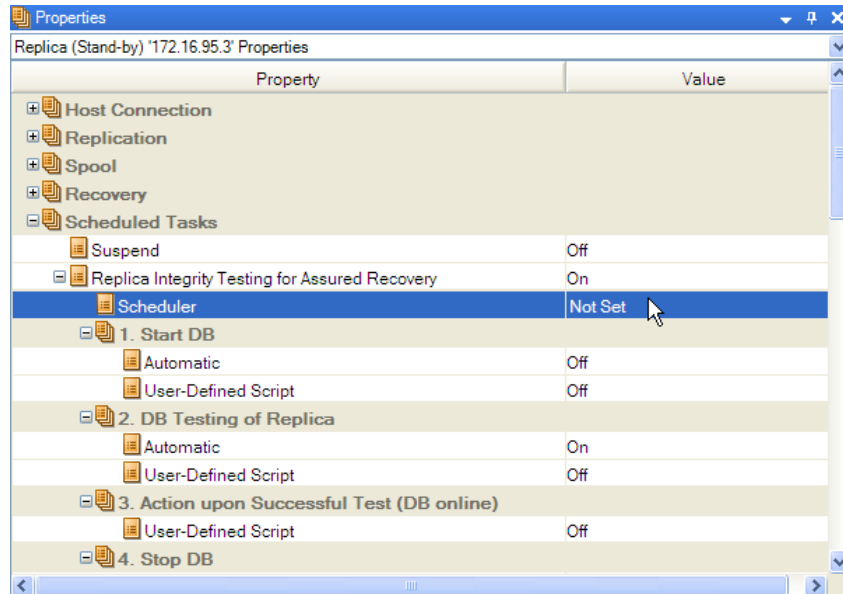
The AR test schedule can be set when the scenario is created or at a later stage.

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**.

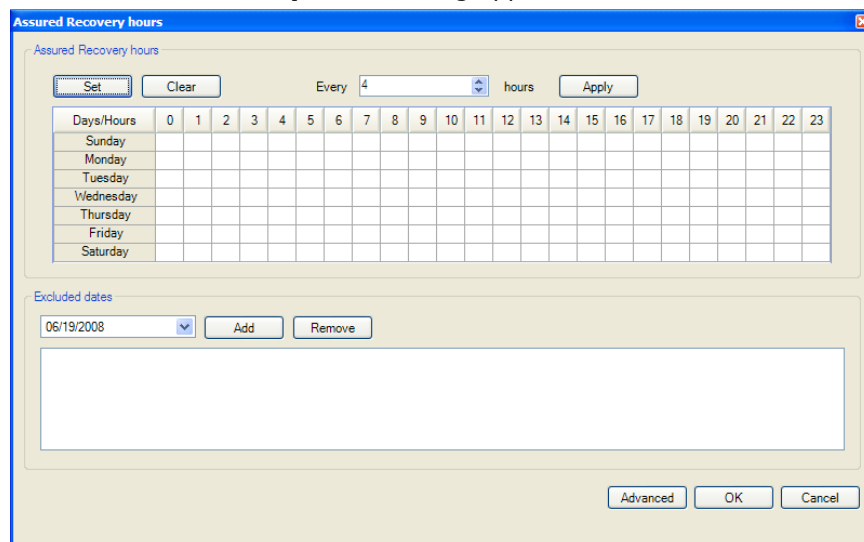
To schedule the AR test

1. On the Scenario pane, select the Replica you want to test.
On the Framework pane on the left, select the properties tab.
The Replica Properties list appears.
2. If the scenario is running, click the **Stop** button on the Standard toolbar.
The scenario is stopped.

- 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](#) (see page 145).

4. Set the schedule for automatic AR testing in the **Assured Recovery hours** dialog, and click **OK** to save your schedule and close the dialog.
5. 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 XOssoft 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 XOssoft 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 XOssoft 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

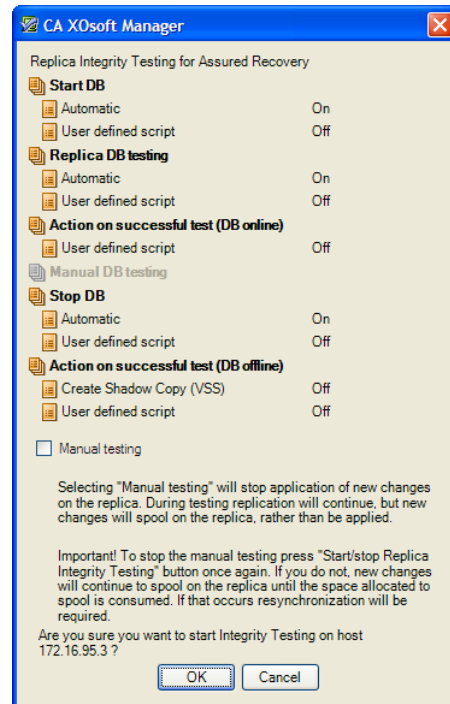
1. On the Manager, verify that the AR scenario is running.

Exchange AR	Running	HA/AR	Exchange	Online
Hosts	Changed	Synchronized	Files	In pool
172.16.95.2	74.5 KB	47.01 MB	6	0 Bytes
172.16.95.3	74.5 KB	47.01 MB	6	0 Bytes

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 opens.

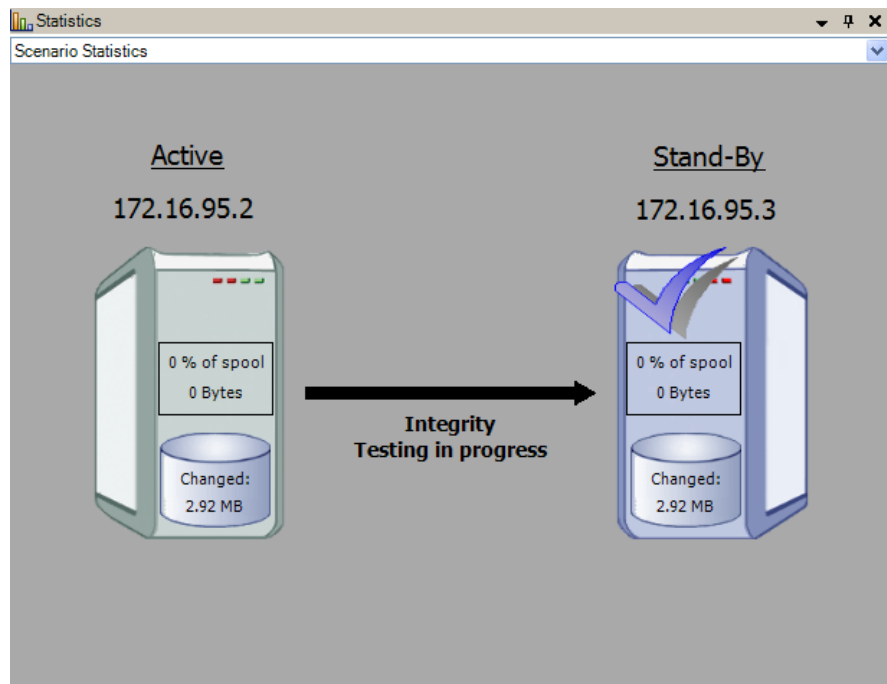


In this dialog, the configuration you set for the AR test is displayed.

- 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](#) (see page 284).
 - To manually perform the AR test, select the **Manual testing** check box, click **OK**, and refer to [Perform Assured Recovery Test Manually](#) (see page 282).
- After you initiate the AR testing, the **Replica Integrity Testing for Assured Recovery** dialog is closed. Then, before the test begins to run, CA XOsft verifies that no synchronization, AR test or replication suspension are currently in progress on any of the hosts that participate in the current scenario.
 - Once the verification stage ends, the AR test begins.



The steps of the test are displayed as messages in the Event pane.

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

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 Report Center

[Report Center Home Page](#)

CA XOsoft replication r12

Assured Recovery Report

Scenario name:	Exchange AR
Scenario id:	3236985457
Scenario type:	Exchange
Master host:	130.119.185.152
Replica host:	130.119.185.153
Integrity test start time:	23/03 19:14:07
Integrity test finish time:	23/03 19:16:47

✔ Start Exchange 23/03 19:14:07 Success
✔ Start Services 23/03 19:16:04 Success
✔ Testing Exchange 23/03 19:16:04 Success
✔ Stop Exchange 23/03 19:16:05 Success
✔ Stop Services 23/03 19:16:46 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.
- To view the report, refer to [Viewing a Report](#) (see page 94).

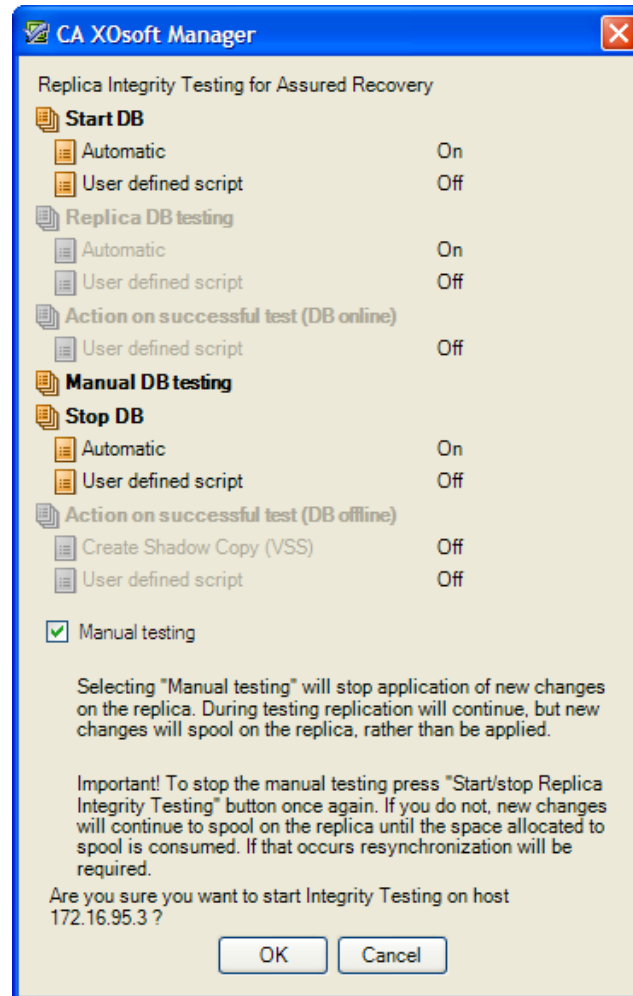
All the tasks that were performed during the test are listed in the AR Report, along with their activation time and status.

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 opens.




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](#) (see page 284).
 - To automatically perform the AR test, clear the **Manual testing** check box, click **OK** and refer to [Perform Assured Recovery Test Automatically](#) (see page 279).
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 pool 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 opens.



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.

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

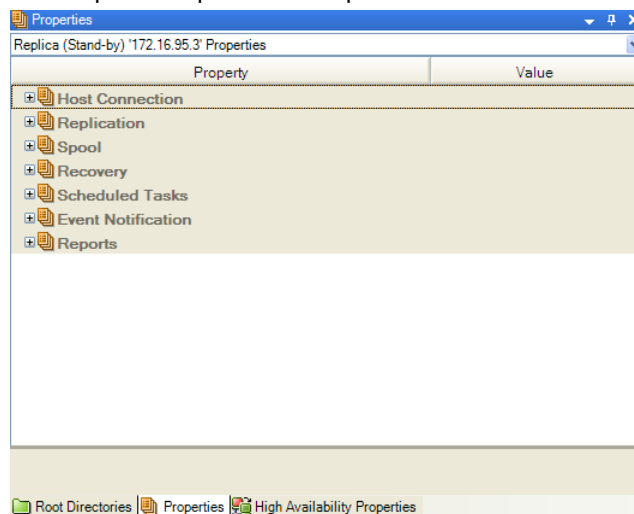
To configure AR properties, the scenario must be stopped.

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

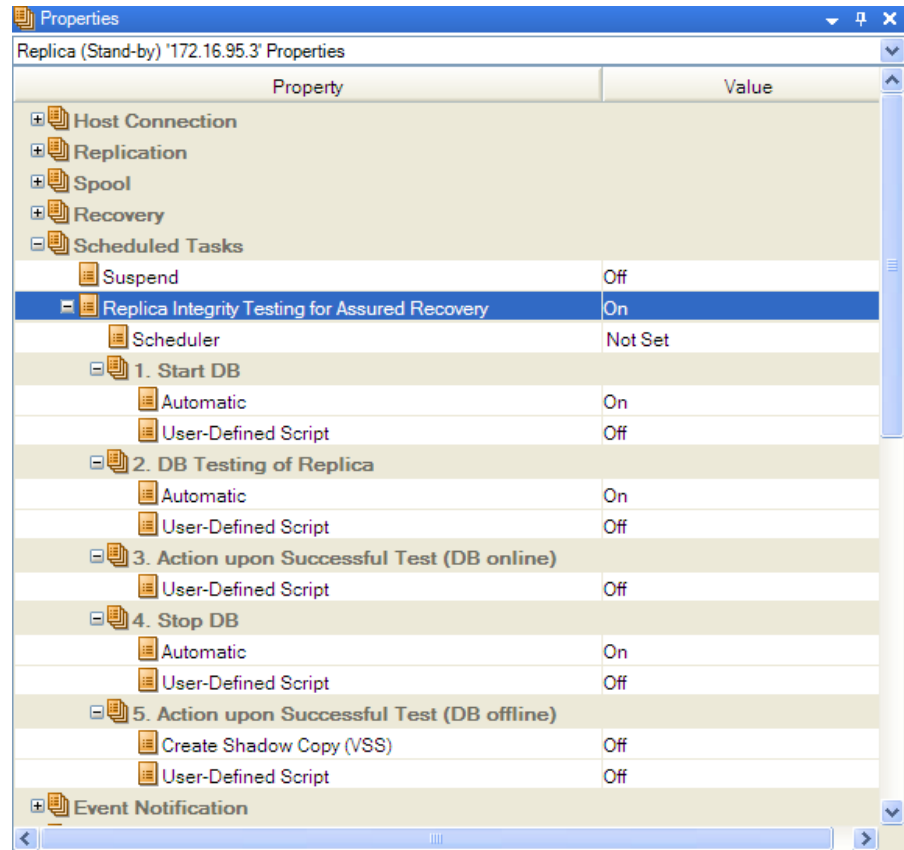
1. On the Scenario pane, select the Replica that you want to test and whose properties you want to configure.
2. On the Framework pane on the left, select the Properties tab.


The Replica Properties list opens.



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 Standard 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.

Note: On Windows x64, you cannot run scripts that activate applications with UI.

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.

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 upon 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 are 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 upon Successful Test (DB Offline)

As noted in **Actions upon 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](#) (see page 289)).

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 XOssoft 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 XOssoft is integrated with ARCserve, a VSS snapshot is automatically created upon each ARCserve backup. All these snapshots are displayed in CA XOssoft Snapshots Management window, which allows you to monitor and manage them.

Configuring VSS Snapshot Creation

By default, CA XOssoft 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.

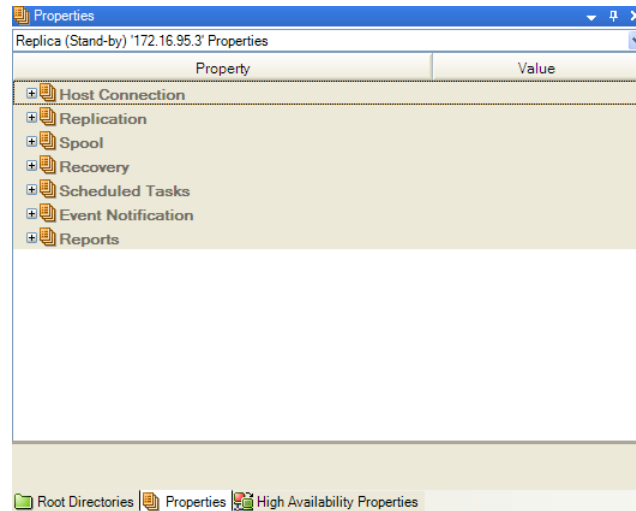
Note: Manual suspension cannot cause the creation of VSS snapshots. VSS snapshots will be created automatically only when associated with scheduled suspension.


Setting Up Snapshot Creation

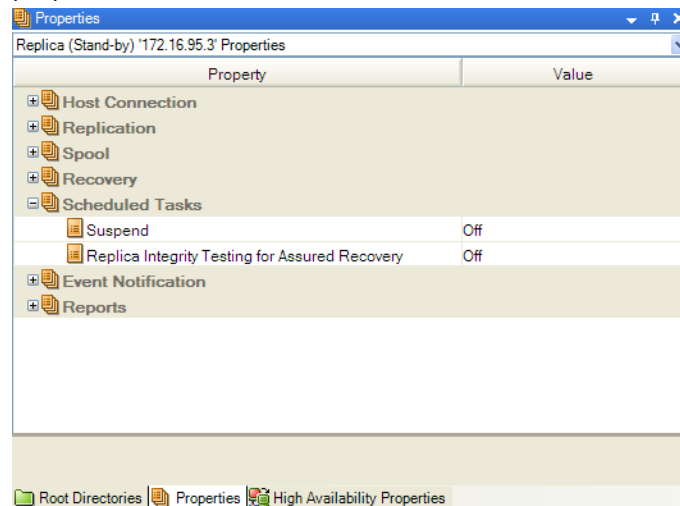
To set up a snapshot creation

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 opens.

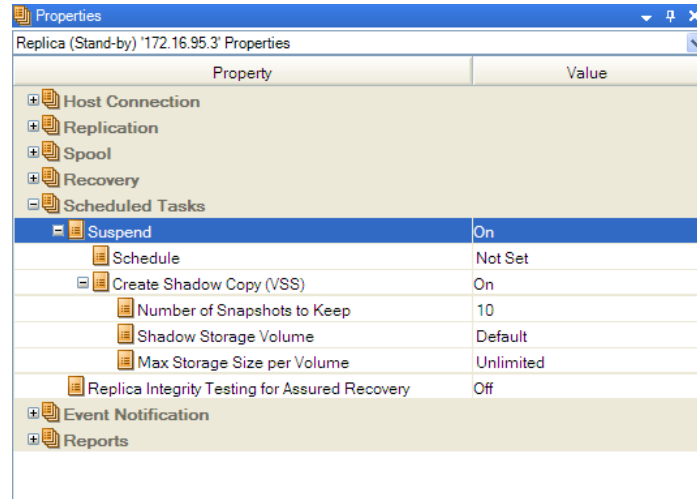


3. If the scenario is running, click the **Stop**  button on the toolbar. The scenario is stopped.
4. On the Replica Properties list, open the **Scheduled Tasks** group to display the **Suspend** and the **Replica Integrity testing for Assured Recovery** properties.




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 opens along with its related properties.



Notes:

- 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.
 - To associate VSS snapshot creation with the **Suspend** property, you need to schedule the suspension. Manual suspension will not create a VSS snapshot.
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](#) (see page 293).
 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.

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.

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. Other internal VSS reasons can cause the number of saved snapshots to be higher than you specified. 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

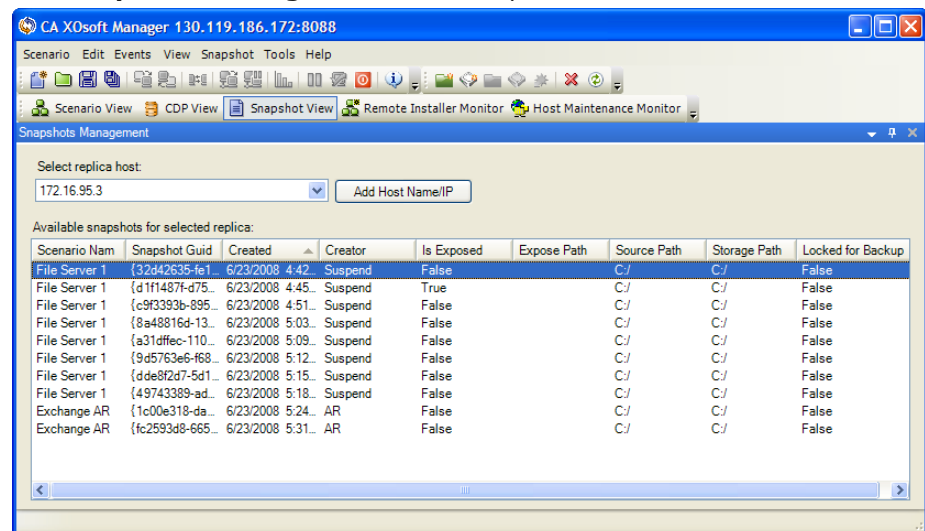
CA XOssoft provides you with a special window for managing your VSS snapshots.

Viewing Snapshots

To open the Snapshots Management window

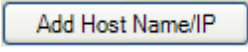
- On the Manager, click the **Snapshot View**  button on the Viewing toolbar.

The **Snapshots Management** window opens.



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:

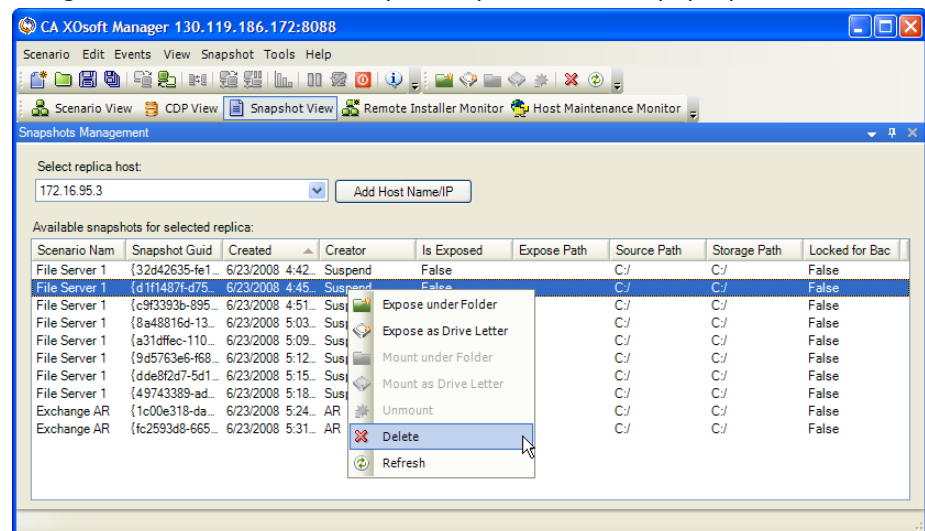
- **Scenario name** - the scenario in which the snapshot was created.
- **Snapshot Guid** - the unique ID that identifies the snapshot.
- **Created** - the date and time of the snapshot creation.
- **Creator** - the type of operation that is associated with the creation of the snapshot. Two types are 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 **Mount** 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 14: Using the CDP Repository

This section provides instructions for creating, managing and using the CDP Repository module.

This section contains the following topics:

[Understanding the CDP Repository](#) (see page 297)

[Configuring the CDP Repository](#) (see page 299)

[Setting CDP Scenario Properties](#) (see page 313)

[Understanding the CDP Repository Statistics and Reports](#) (see page 319)

[Retrieving Deleted Outlook Items Using the E-mail Retrieval](#) (see page 327)

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 XOssoft 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 XOssoft 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 extracted data. 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 log on as a DB owner to the SQL Server (not through the CA XOssoft 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 XOssoft 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.

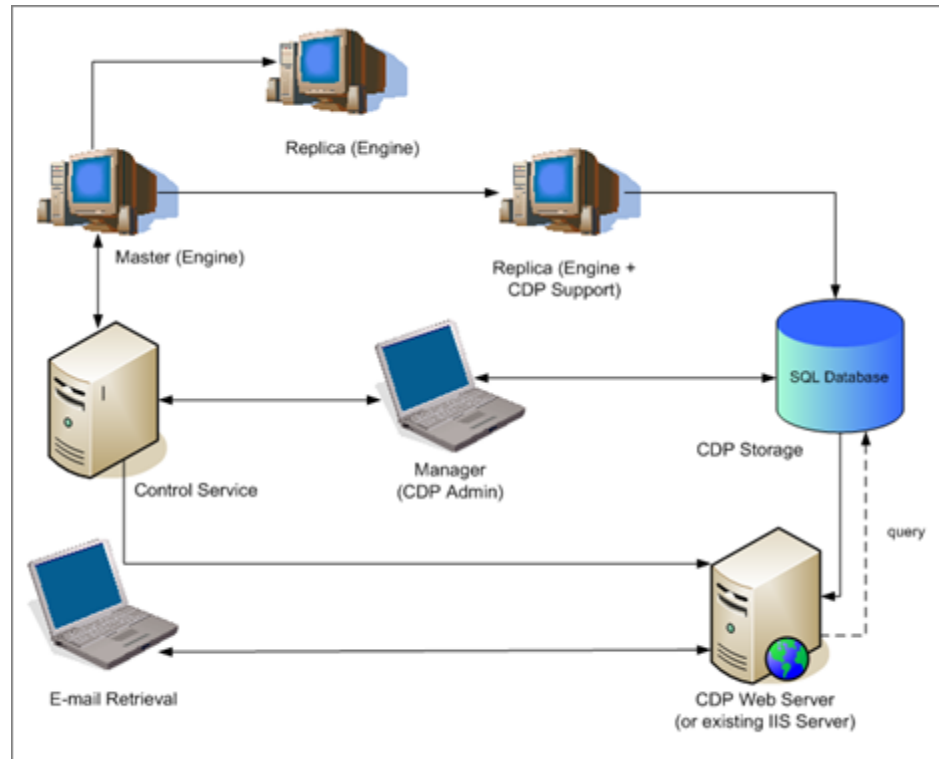
CDP Repository Components

The CA XOssoft 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.

A typical deployment of CDP Repository shows the component residing on its own web or IIS server.



Configuring the CDP Repository

There are three steps in configuring the CDP Repository

1. [Defining the CDP Database](#) (see page 300).
2. [Creating an Exchange scenario with the CDP option](#) (see page 305).
3. Running the scenario.

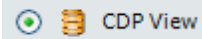
The CDP Database configuration needs to be done before the scenario creation.

Defining the CDP Database


The CDP Database configuration consists of three tasks:

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

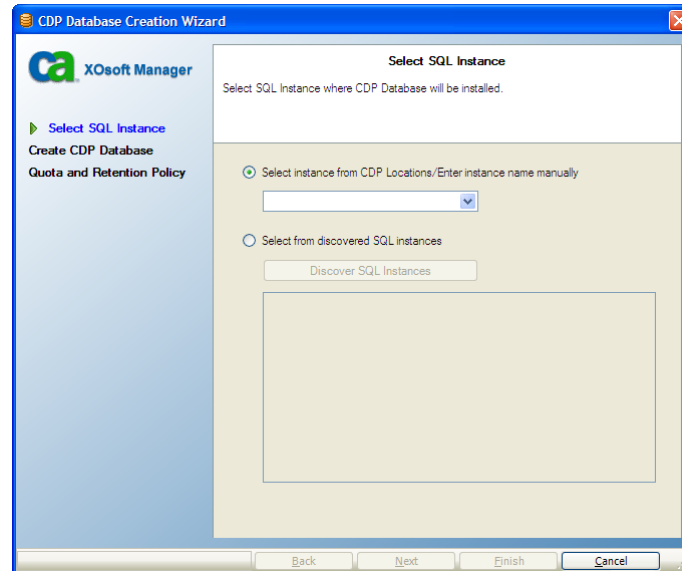
1. On the Manager, open the CDP View by selecting the **CDP View**  option button on the Viewing toolbar. Alternatively, select from the **View** menu the **Active View - CDP View** option.

The **CDP View** opens.



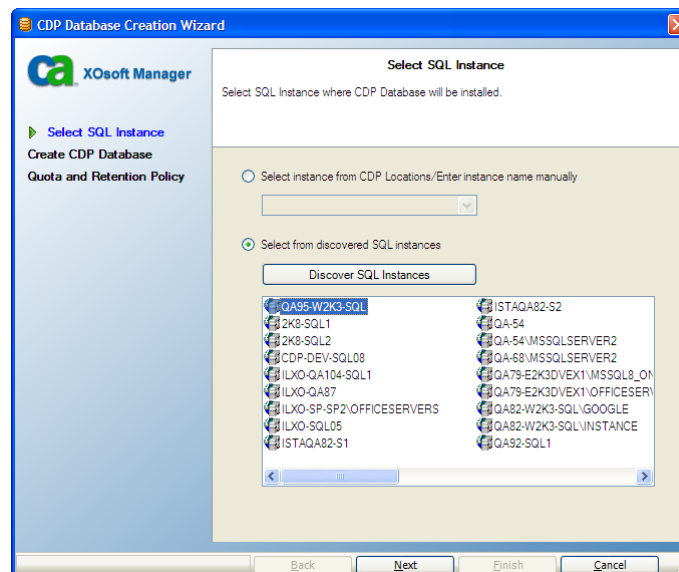
2. To create a CDP Database, select from the **CDP** menu the **Add Database** option.

The **Select SQL Instance** dialog opens.



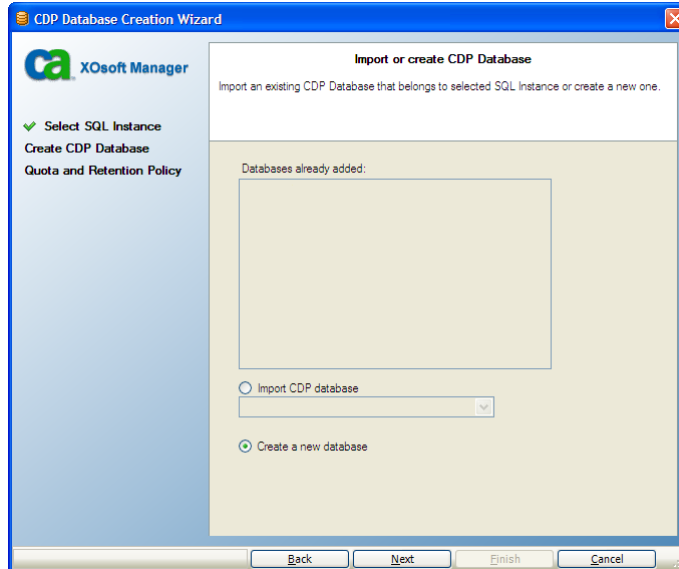
In this page, you select the SQL instance where the CDP database will be created.

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 XOssoft automatically discovers the existing SQL instances, and displays them.



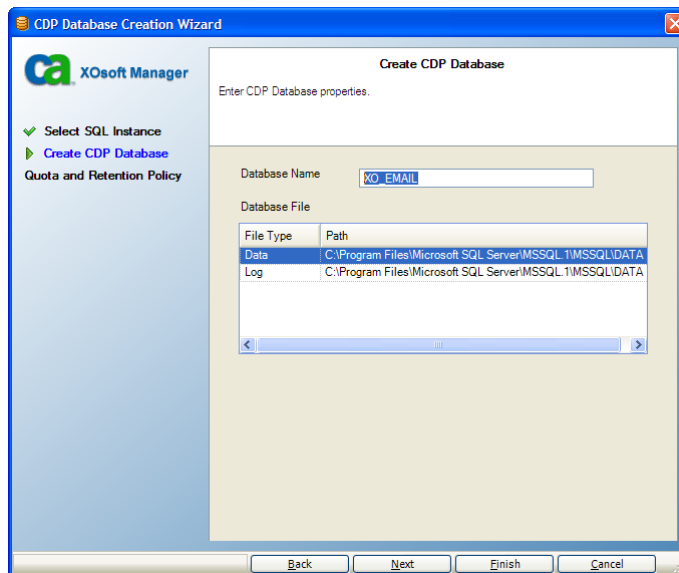
Select from the list the required SQL instance.

4. After defining the SQL instance, click **Next**. The **Import or create CDP Database** page opens.




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 opens.



In this page, you can see the default CDP Database name and its storing path.

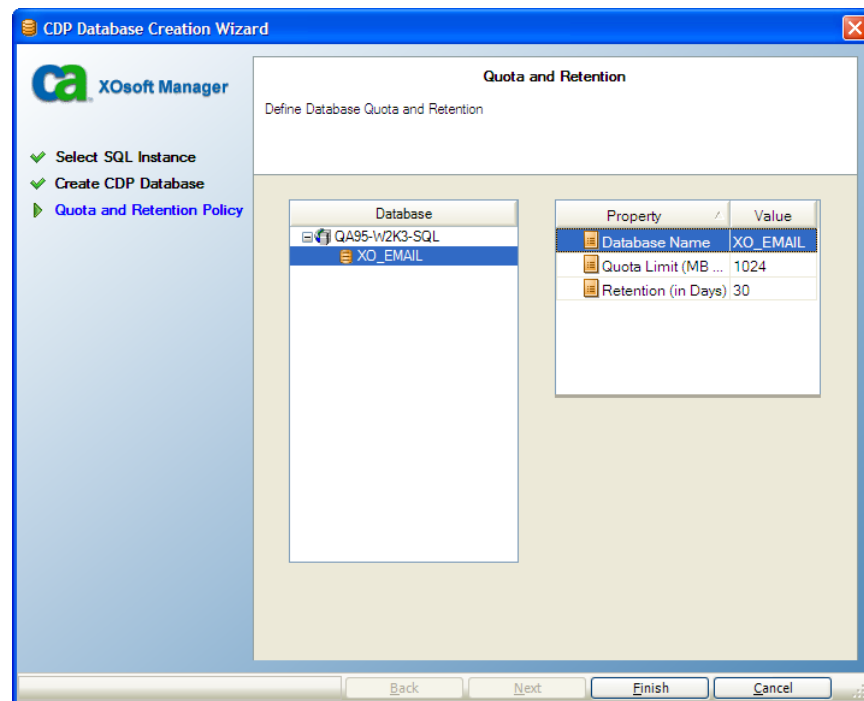
Note: If you click **Cancel** at this stage, the following message opens.



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.

- In the **Create CDP Database** page, keep the default database values or change them. Then, click **Next**.

The **Quota and Retention** page opens.

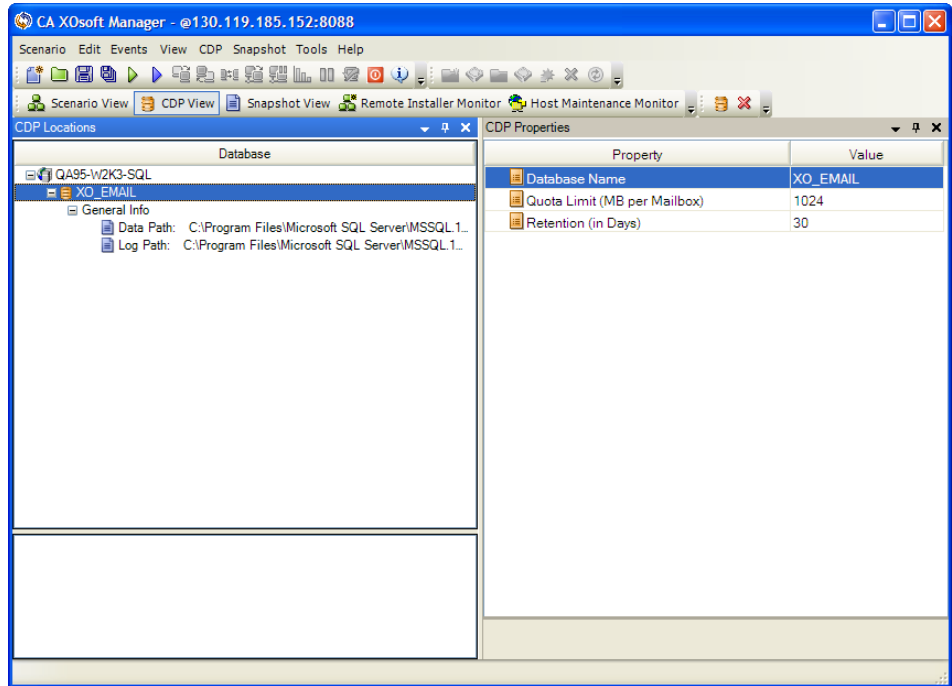


In this page, you define for the entire database its **Quota Limit**, meaning the maximum disk space allocated for the deleted messages. Once this size is reached, new messages will not be inserted into the database. You can also define a **Retention** period for the database, meaning for how many days the deleted messages will be kept in the CDP database.

- 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 **Retention** 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.

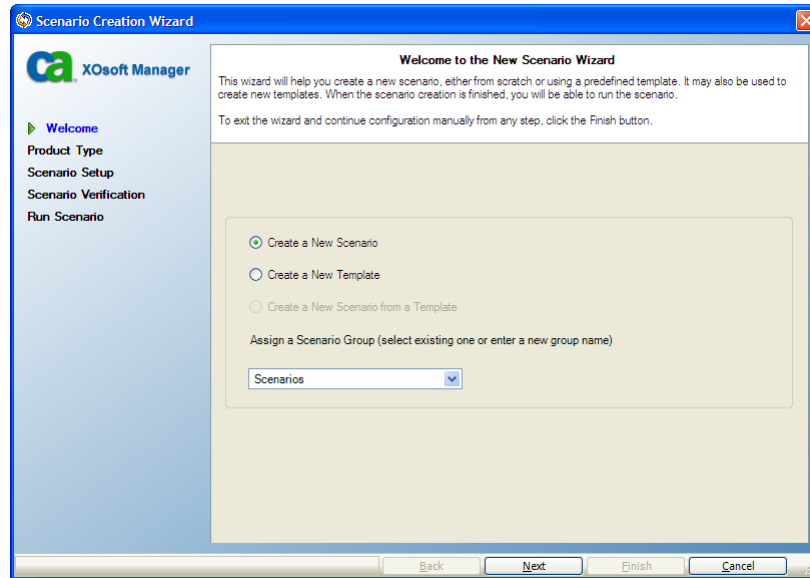
To configure the CDP database from the CDP View, select the database on the Database pane on the left, and change its property values on the Properties pane on the right.

Creating an Exchange Scenario with the CDP Option

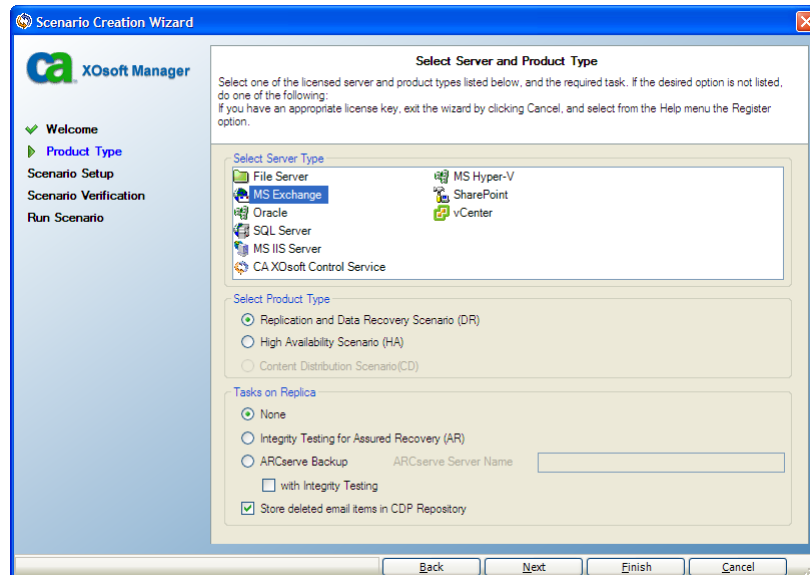
To activate the CDP Repository module, you need to create an Exchange scenario with the CDP option enabled.

To create an Exchange scenario with CDP

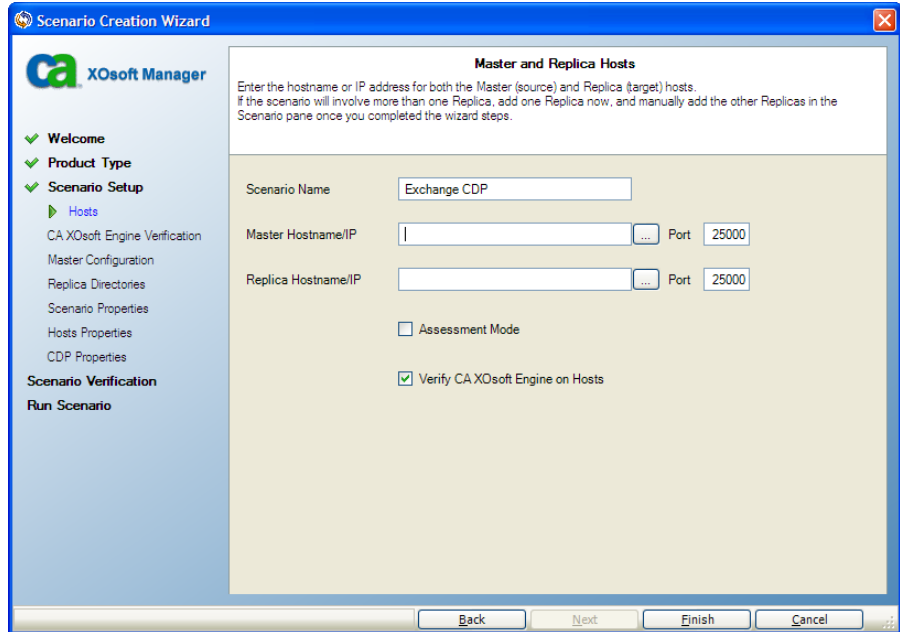
1. On the Manager, click the **New** button on the Standard toolbar.
The **Scenario Creation Wizard** opens.



2. Select the **Create a New Scenario** option button, and click **Next**.
The **Select Server and Product Type** page opens.



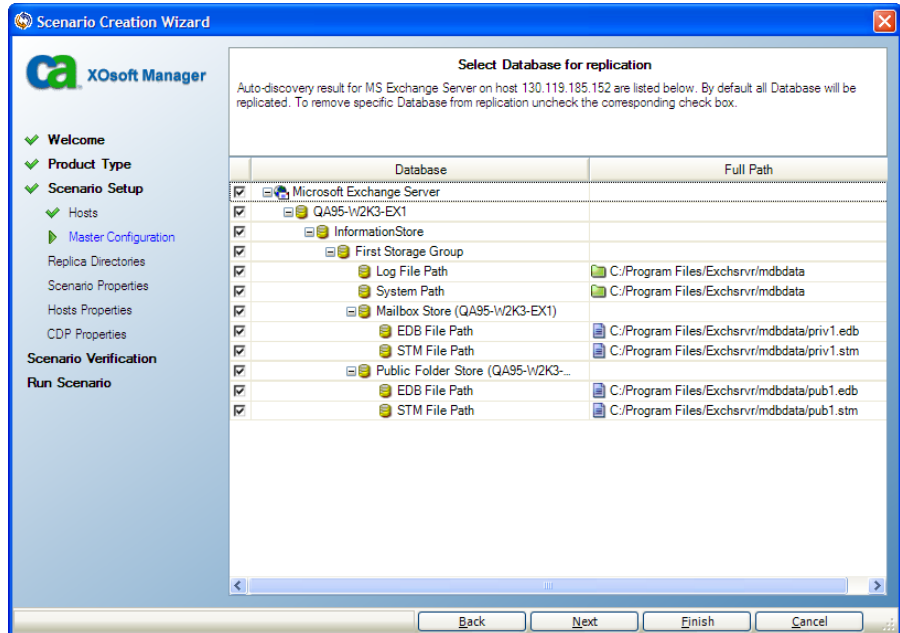
3. Select the required scenario options, as follows:
 - From the **Select Server Type** list, select **MS Exchange**.
 - From the **Select Product Type** options, select either **Replication and Data Recovery** or **High Availability Scenario**.
 - From the **Tasks on Replica** options, select Stored deleted email items in **CDP Repository**.
4. Click **Next**. The **Master and Replica Hosts** page opens.



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 XOssoft 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**.

6. Click **Next**. The **Master Configuration** page opens.

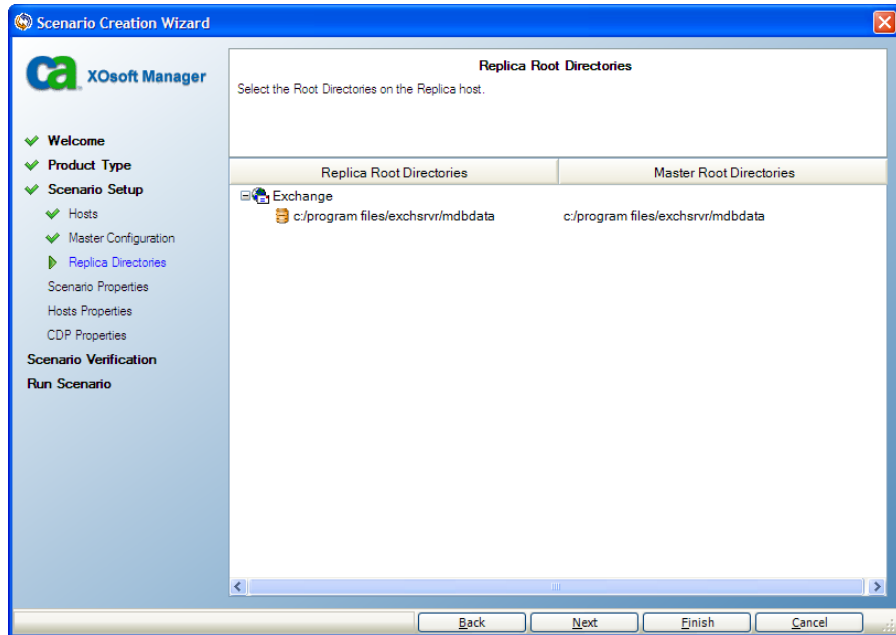


CA XOssoft auto-discovery component automatically discovers all Exchange databases on your Master server. These are the databases that can be replicated and protected.

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.

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


The **Replica Root Directories** page opens.



CA XOssoft auto-configuration component verifies that the Exchange Server configuration on the Master and Replica servers will be identical during the replication procedure.

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

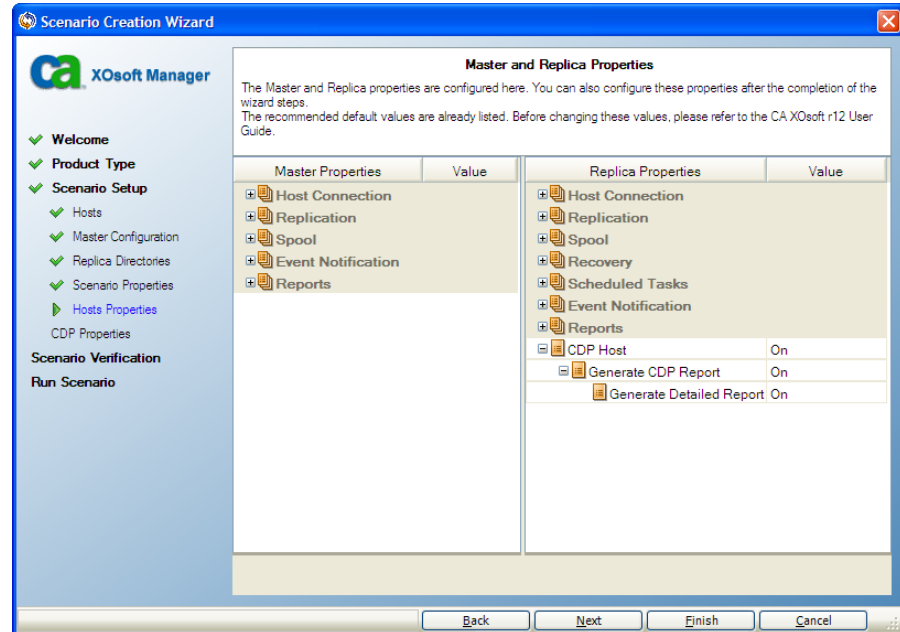
The **Scenario Properties** page opens.



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](#) (see page 138). To configure the scenario properties at a later stage, refer to [Configuring Scenario Properties](#) (see page 137).

10. Click **Next**. The **Master and Replica Properties** page opens.



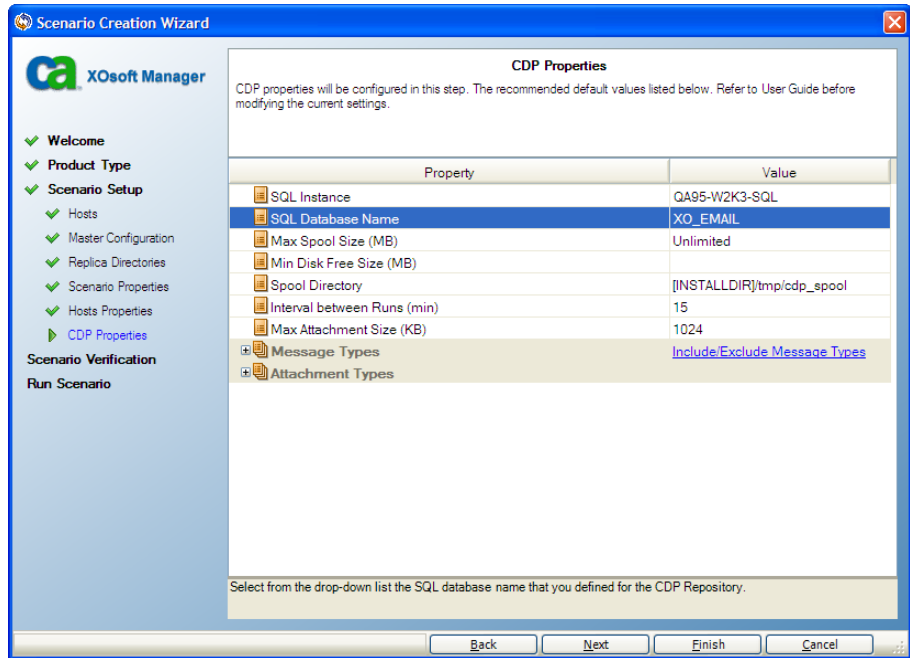
The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient. All you need to verify at this stage is that on the **Replica Properties** list, the **CDP Host** property is **On**.

Note: in a CDP Exchange scenario, one Replica, and only one, should function as the CDP Replica, meaning the Replica that participates in the CDP Repository process.

If you want to configure the Master and Replica properties at this stage, refer to [Setting Master and Replica Properties](#) (see page 163). To configure the scenario properties later, refer to [Configuring Master or Replica Server Properties](#) (see page 164).

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 information](#) (see page 167) for configuration details.

11. Once you are satisfied with the Master and Replica properties, click **Next**.
The **CDP Properties** page opens.



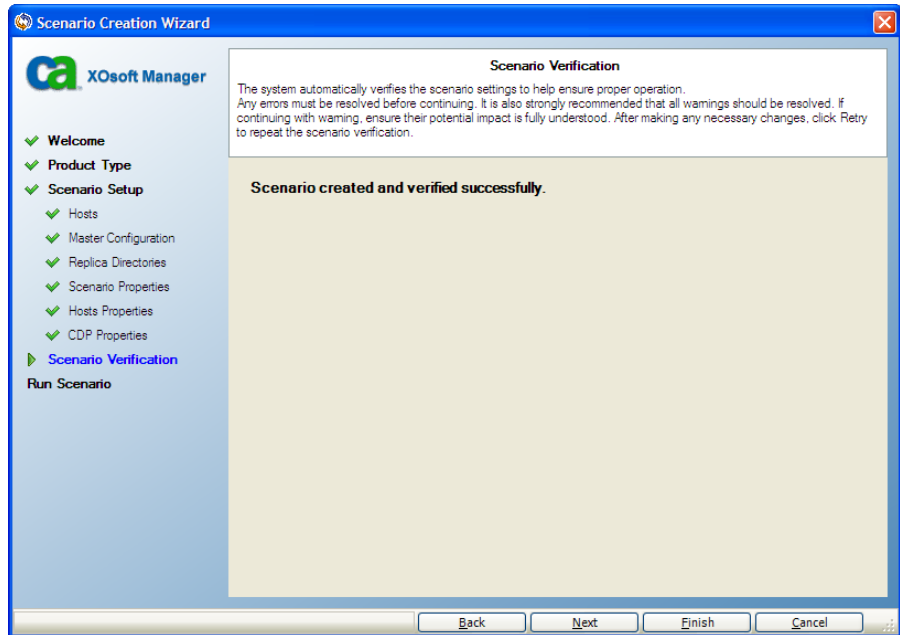
The **CDP Properties** page enables you to configure the properties that are related to the CDP Repository.

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.

Typically, the other default values are sufficient. If you want to configure the other CDP properties at this stage, refer to [Understanding CDP Scenario Properties](#) (see page 315). To configure the CDP scenario properties at a later stage, refer to [Configuring CDP Scenario Properties](#) (see page 314).

13. Once you set the CDP Repository properties, click **Next**.

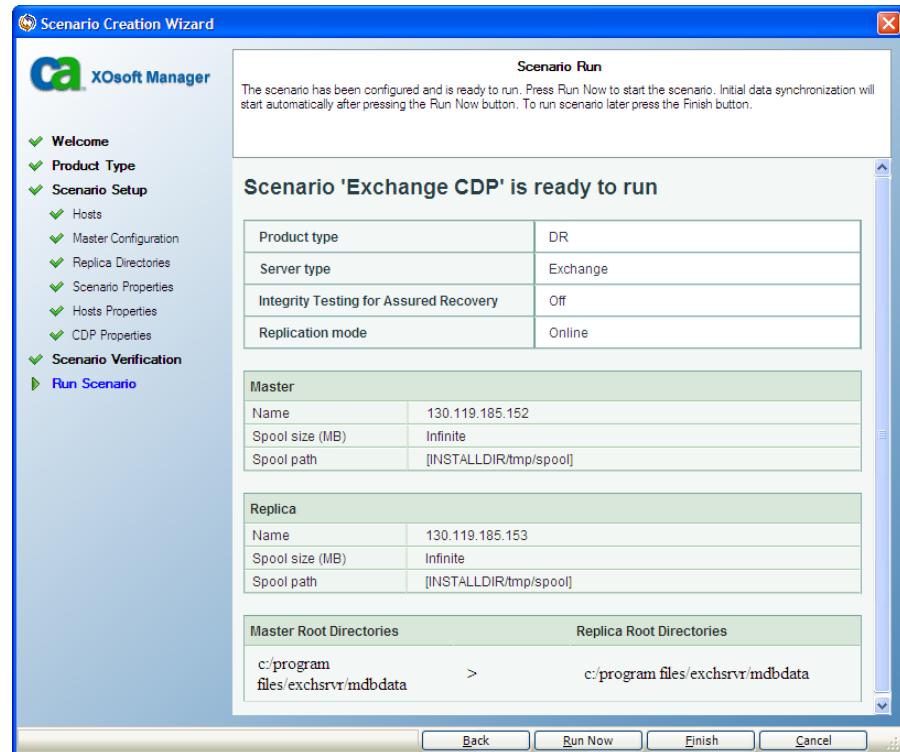
CA XOsoft verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful replication and data recovery process. Once the verification is completed, the **Scenario Verification** page opens.



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 opens.



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.

The screenshot shows the CA XOssoft Report Center interface. At the top left is the CA logo and the text "CA XOssoft Report Center". To the right is a button labeled "Report Center Home Page". Below this is the text "CA XOssoft Replication". The main heading is "SYNCHRONIZATION REPORT".

Synchronization mode	BlockSynchronization (include files with the same size and modification time)
Scenario	Exchange CDP
Master host	130.119.185.152(1)
Replica host	130.119.185.153(2)
Scenario start time	11-Mar-09 15:33:11
Report start time	11-Mar-09 15:33:19
Report finish time	11-Mar-09 15:35:51

Summary:

Total number of files modified	2
Total number of bytes changed	5.01MB

Note: For more information about opening a report, see [Viewing a Report](#) (see page 94).

Setting CDP Scenario Properties

This section describes how to configure the CDP Repository properties, and provides the list of properties, corresponding values, and an explanation of each property. Some of the CDP properties are configured in the CDP Scenario Properties list, while others are configured in the CDP Replica Properties List.


Configuring CDP Scenario Properties

After you defined an Exchange scenario with the CDP option, you can set or modify its properties.

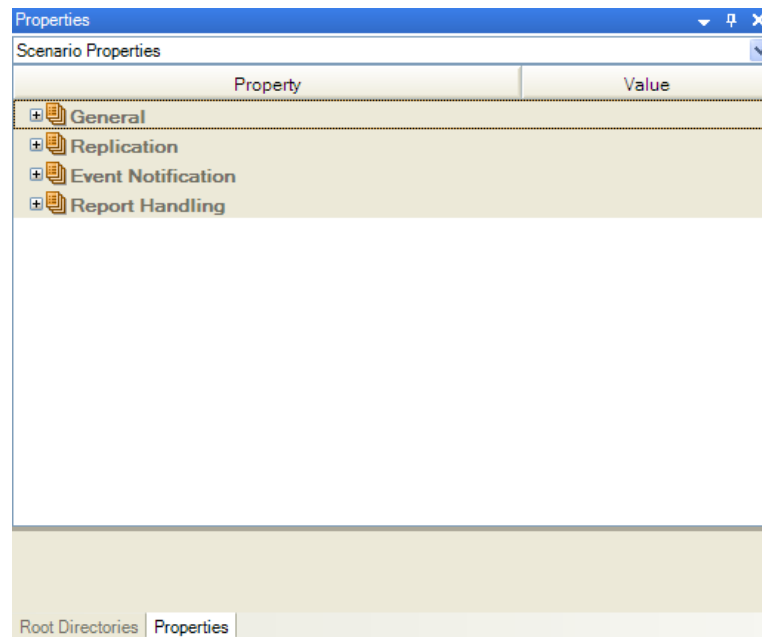
Note: To configure the CDP scenario properties, the scenario must be stopped.

To set CDP scenario properties

1. If you are not in the Scenario View, open it by selecting the **Scenario**

View  Scenario View option button on the Viewing toolbar.

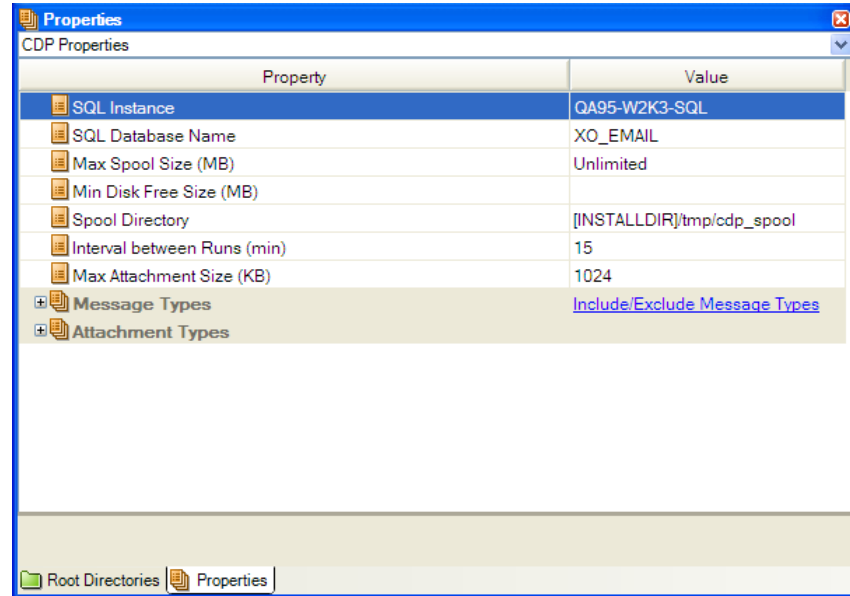
2. On the Scenario pane, select the Exchange scenario its CDP properties you want to configure. On the Framework pane on the left, the Scenario Properties list opens.




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

- From the **Properties** drop-down list on the Framework pane, select **CDP Properties**.

The **CDP Properties** list opens.



- On the CDP Properties list, 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 drop-down list.
- After you set the required properties, click the **Save**  button on the Standard toolbar to save and apply your changes.

Understanding CDP Scenario Properties

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

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.

Max Spool Size (MB)

Enter the maximum spool size allowed in MB. 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. When this threshold is exceeded, the system issues an error message, and stops the CDP checkup session.

Min Disk Free Size (MB)

Enter the free disk space threshold in MB. When reaching this threshold, the system issues an error message, and stops the CDP checkup session.

Spool Directory

Enter the directory to be used to store the spool for the CDP Repository on the Replica.

Interval between Runs (min)

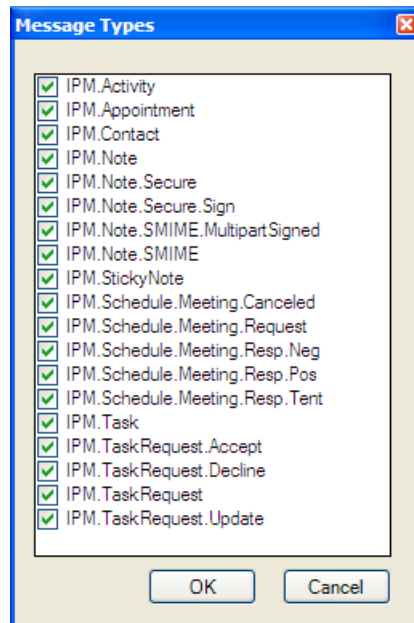
Enter the number of minutes between each check the system will perform in the Exchange database, in search for new deleted messages. The minimum allowed value is 15 minutes.

Max Attachment Size

Enter the attachment size allowed for storage. Deleted attachments that exceed this size will not be stored in the CDP SQL database.

Message Types

Specify whether deleted message of this type will be stored in the CDP SQL 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 opens.

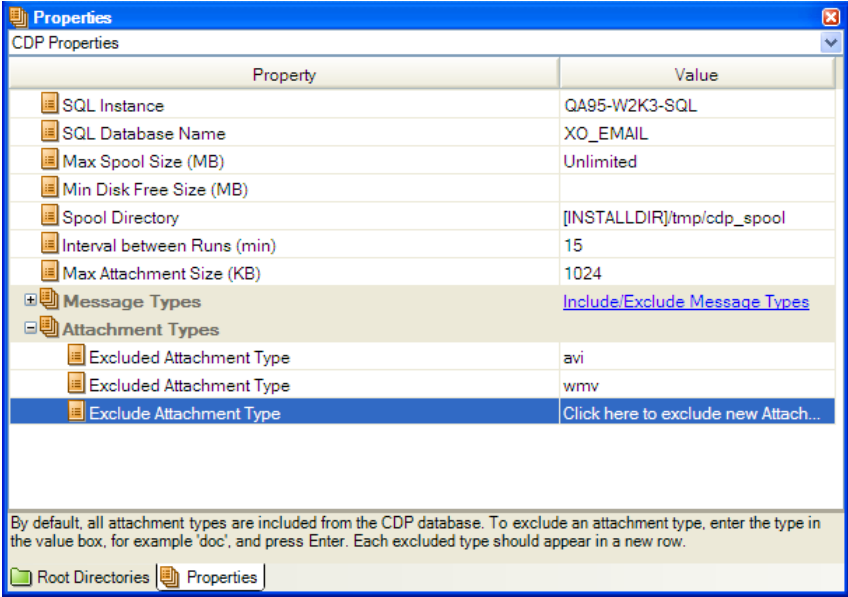


To exclude message types from storage, clear their check boxes.

Attachment Types

Enter the type of attachments to be stored. By default, all attachment types are included in the CDP SQL database. To exclude an attachment type, enter the type in the value box, for example **avi**, and press the **Enter** key.

Each included type should appear in a new row.



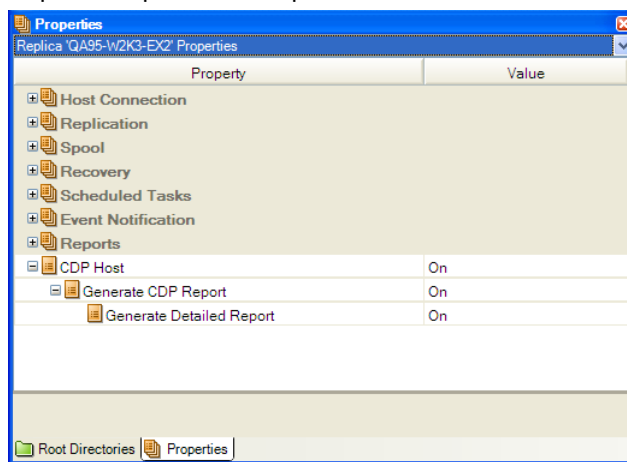
Configuring CDP Replica and Report Properties

After you defined an Exchange scenario with the CDP option, and you defined the scenario properties, you can set or modify its Replica and Report properties.


Note: To configure the CDP Replica and Report properties, the scenario must be stopped.

To set CDP Replica and Report properties

1. On the Scenario pane, select the Replica of the CDP Exchange scenario its properties you want to configure. On the Framework pane on the left, the Replica Properties list opens.



The CDP Replica and Report properties appear at the bottom of the list.

2. Set the CDP Replica and Report properties as follows:
 - **CDP Host** – in a CDP Exchange scenario, one Replica, and only one, should function as the CDP Replica, meaning the Replica that participates in the CDP Repository process. Set the value to On to select this Replica as the CDP Replica in the CDP Exchange scenario.
 - **Generate CDP Report** – set to On to automatically generate a [Summary CDP Repository Report](#) (see page 324), when deleted items are found during a CDP backup session.
 - **Generate Detailed Report** – set to On to automatically generate a [Detailed CDP Repository Report](#) (see page 326), when deleted items are found during a CDP backup session.
3. After you set the required properties, click the **Save**  button on the Standard toolbar to save and apply your changes.

Understanding the CDP Repository Statistics and Reports

CA XOsoft checks at regular intervals whether the Exchange database contains new deleted items. On each checkup session, CA XOsoft updates the CDP Statistics on the Manager, and by default generates two reports: a **Summary Email CDP Repository Report** and a **Detailed Email CDP Repository Report**. The CDP Statistics is updated only during a new checkup session, and the CDP Reports are generated only when new deleted items are found in the checkup session.

Notes:

- You can change the checkup interval through the **Interval between Runs** property in the [CDP Properties list](#) (see page 315). The minimum allowed value is 15 minutes.
- You can cancel the CDP Report generation by setting to Off the **Generating CDP Report** and **Generate Detailed Report** [properties](#) (see page 318).

CDP Repository Statistics

The Statistics tab in the Framework pane displays live statistics. Different statistics information is displayed for a scenario, a Master and each Replica host. The CDP Statistics is displayed for a CDP Replica, in addition to the replication statistics that is provided for this Replica.

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

To view the CDP Repository Statistics

1. On the Scenario pane, select the CDP Replica whose statistics you want to view.
2. On the Framework pane, select the **Statistics** tab.

The CDP Statistics appears at the lower half of the pane.

The screenshot shows a window titled 'Statistics' for 'Replica '130.119.185.153''. It contains several data sections:

State	Running
Start of replication	3/11/2009 3:33:11 PM
Version	12.5.0.136

Spool space:

Size	% of threshold
0Bytes	0%

Online file changes per root directory:

Root Directory	Size	Folders Created	Changed	Removed	Renamed
c:/program files/.../mdbdata	4.91MB	0	7	0	0

Email CDP Repository Statistics:
Repository Table messages_batch_2009_3_11_16_1_44_960

Start Time	Duration (sec)	Status	Items Inserted Successfully	Total Size (MB)	Items Filtered Out	Items Failed
3/11/2009 16:00:20	75	Success	11	3.33	0	0

Email CDP Repository Statistics Summary (From Scenario Start):

Items Inserted Successfully	Total Size (MB)
11	3.33

The CDP Repository Statistics consists of two tables:

- The **Email CDP Repository Statistics** [table](#) (see page 321)
- The **Email CDP Repository Statistics Summary** [table](#) (see page 322)

Email CDP Repository Statistics Table

The **Email CDP Repository Statistics** table provides information about the deleted items that were found at the current CDP checkpoint session.

Email CDP Repository Statistics:						
Repository Table messages_batch_2009_3_11_16_1_44_960						
Start Time	Duration (sec)	Status	Items Inserted Successfully	Total Size (MB)	Items Filtered Out	Items Failed
3/11/2009 16:00:20	75	Success	11	3.33	0	0

This table contains the following information:

- **Repository Table** (heading) - the name of the SQL table that contains the deleted items.
- **Start Time** - the start date and time of the current checkpoint session.
- **Duration (sec)** - the number of seconds that the current checkpoint session lasted.
- **Status** - the status of the current checkpoint session. The status can be one of the following: **Success**, **Initiation Failed**, **Extract Failed**, **Feed Failed**, **Statistics Failed**, **Unknown**.

Note: Even if some deleted items have failed to be inserted into the database, the overall **Execution Status** can still be **Success**. Only if the entire process has failed, the status will be **Failed** or **Unknown**.

- **Items Inserted Successfully** - the number of deleted items that were inserted into the SQL database in the current checkpoint session.
- **Total Size (MB)** - the total size in MB of the inserted items, including attachments.
- **Items Filtered Out** - the number of deleted items that were not inserted into the database, either because the quota limit was exceeded, or because these items were excluded by the definitions set in the **Message Types** property.

Note: Attachments that were not inserted into the database due to restrictions set in the **Attachment Types** property, will not be counted here, since their messages might have been inserted.

- **Items Failed** - the number of items that were not inserted into the database due to system error.

When no deleted items were found in the checkup session, the only columns that appear in the table are: **Start Time**, **Execution Status** and **Deleted Items Found**.

Email CDP Repository Statistics:		
Start Time	Status	Deleted Items Found
3/11/2009 15:45:14	Success	0

Note: The **Deleted Items Found** column appears only when no deleted items were found, and therefore it should always display a "0" value.

Email CDP Repository Statistics Summary Table

The **Email CDP Repository Statistics Summary** table provides information about the deleted messages that were found since the CDP scenario was initiated.

Email CDP Repository Statistics Summary (From Scenario Start):	
Items Inserted Successfully	Total Size (MB)
11	3.33

This table contains the following information:

- **Items Inserted Successfully** - the number of deleted items that were inserted into the SQL database since the initiation of the CDP scenario.
- **Total Size (MB)** - the total size in MB of the inserted items, including attachments, since the initiation of the scenario.

CDP Repository Reports

If deleted messages are found in a CDP checkup session, by default a Summary Email CDP Repository Report and a **Detailed** Email CDP Repository Report are generated.

Note: For detailed instructions on how to open and use the Report Center, refer to [Viewing a Report](#) (see page 94).

To view an Email CDP Repository Report

- On the Report Center, select from the **Available Reports per Scenario** table the Exchange CDP scenario that this report represents. Then, from the **Reports** table, click the CDP report you want to open, either **Summary** or **Detailed**.

The screenshot shows the CA XOsoft r12 Report Center interface. At the top, it says 'Updated: Monday, January 12, 2009 1:12:08 PM'. Below this is a section titled 'Available Reports per Scenario' which contains a table with the following data:

Scenario Name	Synchroniza	Difference	Replication	Assessment Mo	Assured Recovery	CDP	Total Reports
Exchange CDP	2	0	0	0	0	4	6
File Server	1	0	0	0	0	0	1

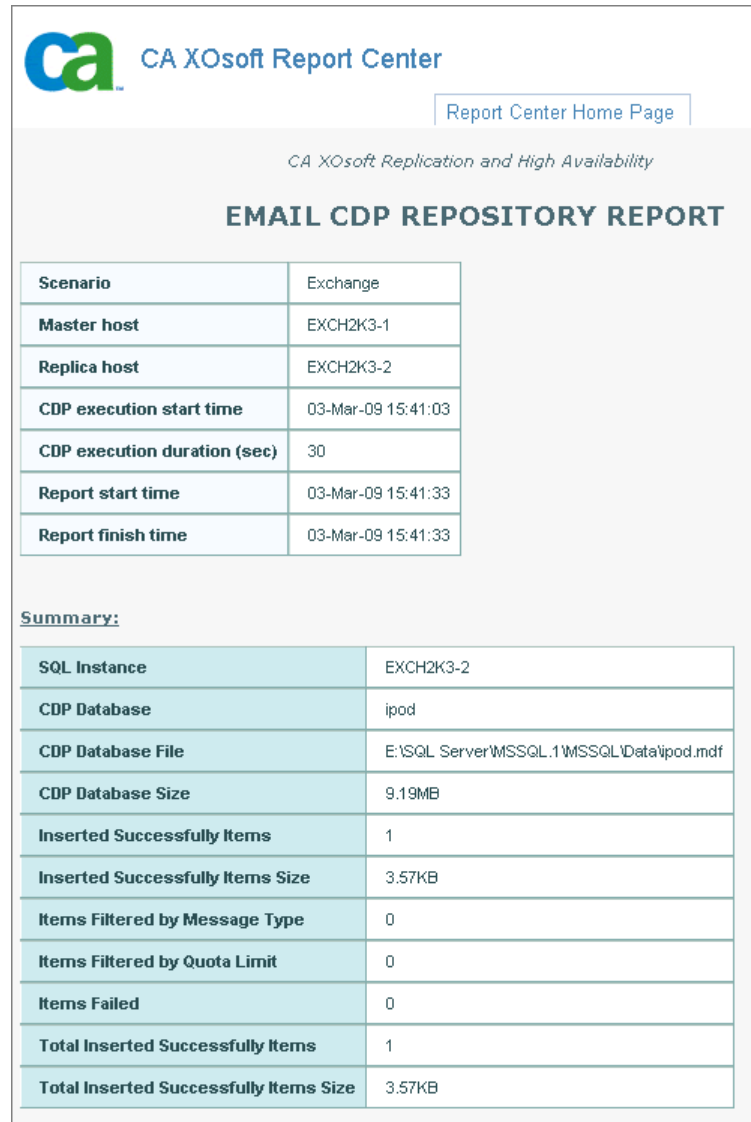
Below this is a section titled 'Reports' with a sub-header 'Drag a column header here to group by that column'. It contains a table with the following data:

Host	Changes	Date	Time	Type	Summary	Detailed	Size (bytes)
130.119.185.153	UnKnown	Today	12:01:06	CDP			1714
130.119.185.153	Changes found	Today	11:47:30	Synchronization			2382
130.119.185.153	UnKnown	01/08/09	00:00:27	CDP			1535
130.119.185.153	UnKnown	01/07/09	17:45:24	CDP			1539
130.119.185.153	UnKnown	01/07/09	17:31:40	CDP			1544
130.119.185.153	No changes made	01/07/09	17:27:59	Synchronization			1507

The report you selected opens.

Email CDP Repository Report

If deleted messages are found in a CDP checkup session, the summarized Email CDP Repository Report is generated.



Scenario	Exchange
Master host	EXCH2K3-1
Replica host	EXCH2K3-2
CDP execution start time	03-Mar-09 15:41:03
CDP execution duration (sec)	30
Report start time	03-Mar-09 15:41:33
Report finish time	03-Mar-09 15:41:33

Summary:

SQL Instance	EXCH2K3-2
CDP Database	ipod
CDP Database File	E:\SQL Server\MSSQL.1\MSSQL\Data\ipod.mdf
CDP Database Size	9.19MB
Inserted Successfully Items	1
Inserted Successfully Items Size	3.57KB
Items Filtered by Message Type	0
Items Filtered by Quota Limit	0
Items Failed	0
Total Inserted Successfully Items	1
Total Inserted Successfully Items Size	3.57KB

The first few lines (at the top) include: the scenario name, the Master and Replica hosts, the start time and duration of the CDP checkup session, and the start and finish time of the report.

The **Summary** table contains the following information:

- **SQL Instance** - the SQL instance name that contains the CDP SQL database.
- **CDP Database** - the SQL database name that you defined for the CDP Repository.
- **CDP Database File** - the path and name of the CDP database file.
- **CDP Database Size** - the size in MB of the CDP database.
- **Inserted Successfully Items** - the number of deleted items that were inserted into the SQL database in the last checkup session.
- **Inserted Successfully Items Size** - the total size in MB of the inserted items, including attachments.
- **Items Filtered by Message Type** - the number of items that were not inserted into the database, because they were excluded by the **Message Types** restrictions.
- **Items Filtered by Quota Limit** - the number of items that were not inserted into the database because the quota limit was exceeded.
- **Items Failed** - the number of items that were not inserted into the database due to system error.
- **Total Inserted Successfully Items** - the number of deleted items that were inserted into the SQL database since the initiation of the CDP scenario.
- **Total Inserted Successfully Items Size** - the total size in MB of the inserted items, including attachments, since the initiation of the scenario.

Detailed Email CDP Repository Report

The **Detailed Email CDP Repository Report** presents the same information that appears in the **Summary Email CDP Repository Report** (see page 324), along with additional information about the Mailboxes whose deleted items were found in the last checkup sessions. For each Mailbox, the following information is provided:

- **Mailbox Name** - the Mailbox name
- **Inserted Successfully Items** - the number of deleted items that were inserted into the SQL database in the last checkup session from this Mailbox.
- **Size** - the total size in MB of the inserted items of the Mailbox, including attachments.


CA XOsoft Report Center

[Report Center Home Page](#)

CA XOsoft Replication and High Availability

EMAIL CDP REPOSITORY REPORT

Scenario	Exchange
Master host	EXCH2K3-1
Replica host	EXCH2K3-2
CDP execution start time	03/16/2009 17:34:04
CDP execution duration (sec)	24
Report start time	03/16/2009 17:34:29
Report finish time	03/16/2009 17:34:29

MAILBOX NAME	INSERTED SUCCESSFULLY ITEMS	SIZE
user1-1	1	3.14KB

Summary:

SQL Instance	EXCH2K3-2
CDP Database	nano
CDP Database File	E:\SQL Server\MSSQL.1\MSSQL\Data\nano.mdf
CDP Database Size	9.94MB
Inserted Successfully Items	1
Inserted Successfully Items Size	8.25KB
Items Filtered by Message Type	0
Items Filtered by Quota Limit	0
Items Failed	0
Total Inserted Successfully Items	1
Total Inserted Successfully Items Size	8.25KB

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.

Logging Into the Email Retrieval

CA XOssoft Email Retrieval does not require any component or application installed in advance. It can be opened 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 XOssoft Email Retrieval

1. Select **Start, Programs, CA, XOssoft, Email Retrieval**.

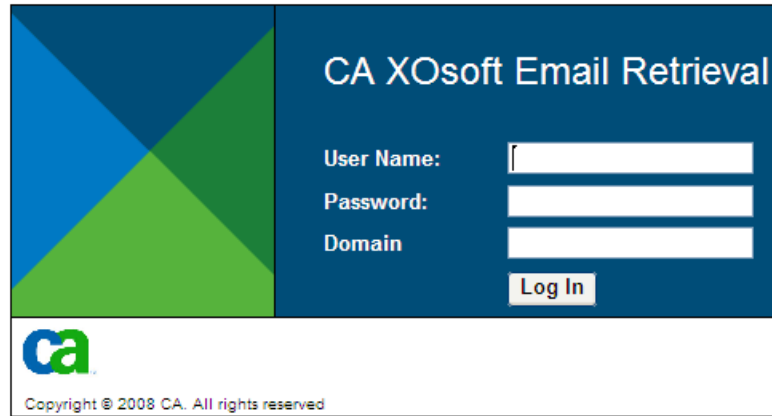
- or -

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/pages/entry_point.aspx`

Notes:

- If you are opening the Email Retrieval from the machine where the Control Service is installed, you can use the default parameters:
`http://localhost:8086/pages/entry_point.aspx`
- If you selected the **SSL Configuration** option during the installation of the Control Service, when you open the Email Retrieval 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/pages/entry_point.aspx`

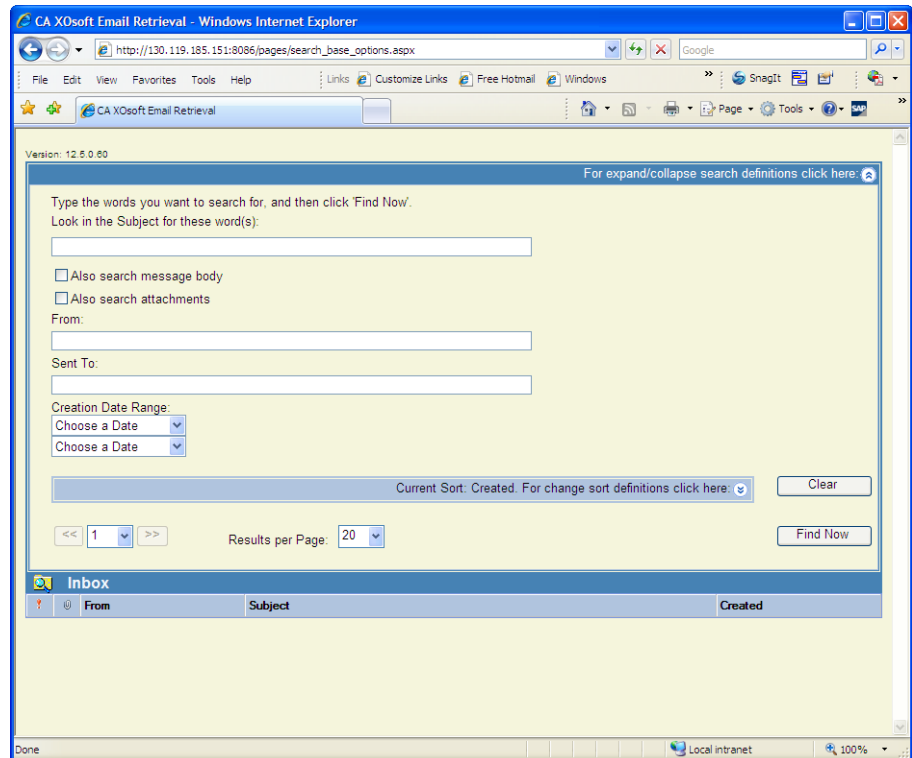
The **Login** dialog opens.



The login dialog features a dark blue header with the text "CA XOssoft Email Retrieval" in white. On the left, there is a graphic of four triangles (blue, green, yellow, red) meeting at a central point. Below the header is a white box containing the CA logo and the text "Copyright © 2008 CA. All rights reserved". The main area is dark blue and contains three white input fields labeled "User Name:", "Password:", and "Domain:". A white "Log In" button is positioned below the "Domain" field.

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

The **Email Retrieval** window opens.



Search for Deleted Items and Retrieve Them

To search and retrieve a deleted item:

1. On the Email Retrieval window, 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.

CA XOssoft E-Mail Retrieval - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Refresh Home Search Favorites

Address http://localhost:8086/Pages/search_base_options.aspx Go Links

For expand/collapse search definitions click here [Close](#)

Enter search criteria, and then click 'Find Now'.

Look in the Subject for these word(s):

Also search message body

Also search attachments

From:

Sent To:

Creation Date Range:

Choose a Date

Choose a Date

Current Sort: Created. For change sort definitions click here:

Clear

Find Now

Page Size: 20

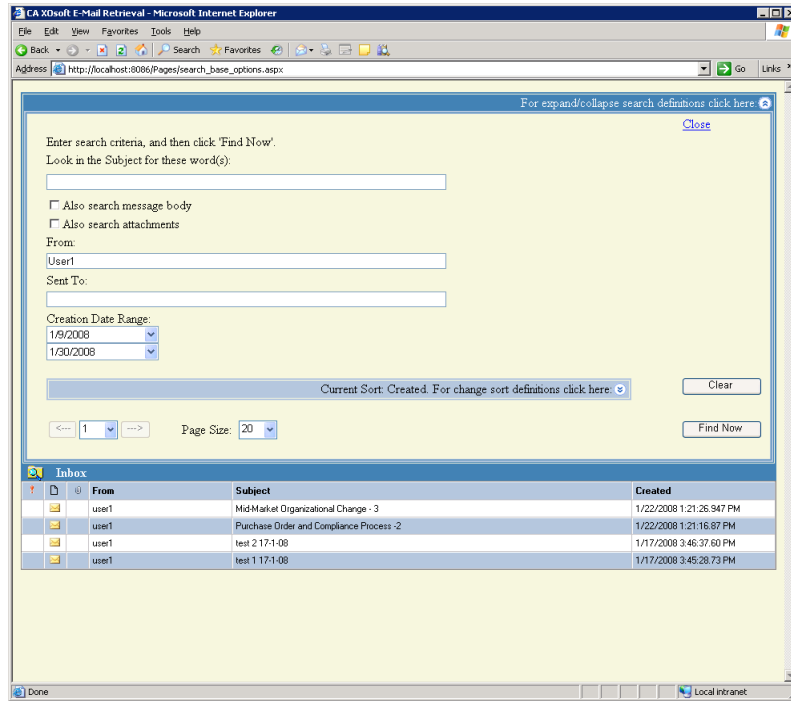
Inbox

From	Subject	Received

Done Local intranet

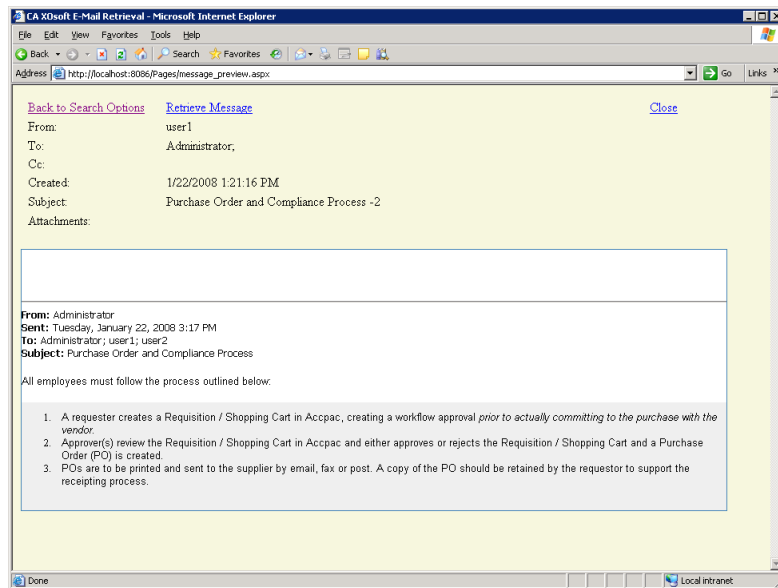
2. [Optional] Define the date range in which you want to conduct the search by using the **Creation Date Range** drop-down menus.
3. After you entered the criteria, click the **Find Now** button.

The search results open.



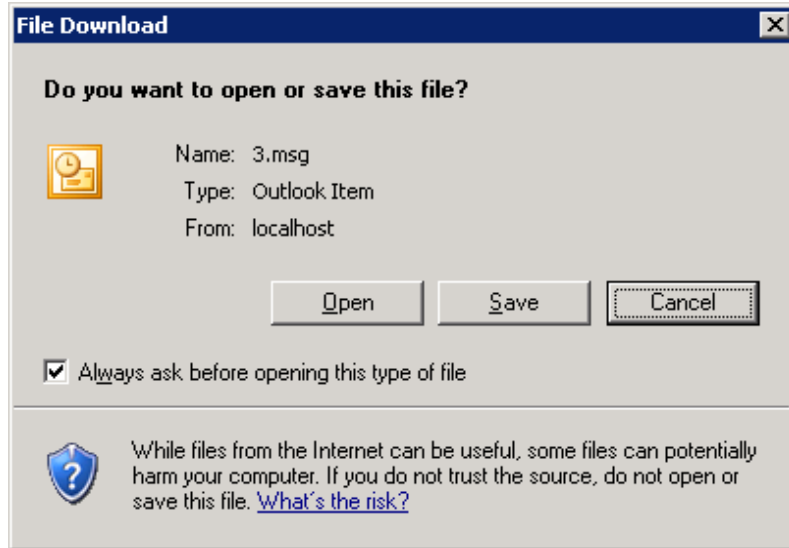
4. To view a deleted message, double-click it.

The message opens.

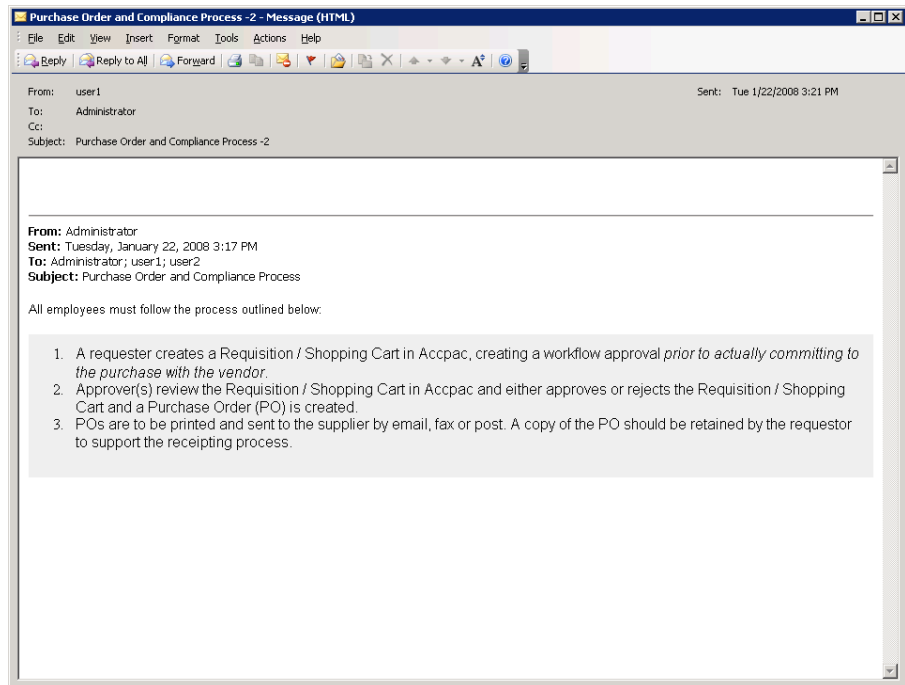


- To retrieve the message, click the **Retrieve Message** link at the top of the page.

A standard **File Download** dialog opens.



- To open the message as an Outlook item, click **Open**.



Chapter 15: Using the Content Distribution Solution

This section provides instructions for creating, managing and using the Content Distribution solution.

This section contains the following topics:

[Understanding the Content Distribution Solution](#) (see page 333)

[Creating a Content Distribution Scenario](#) (see page 336)

Understanding the Content Distribution Solution

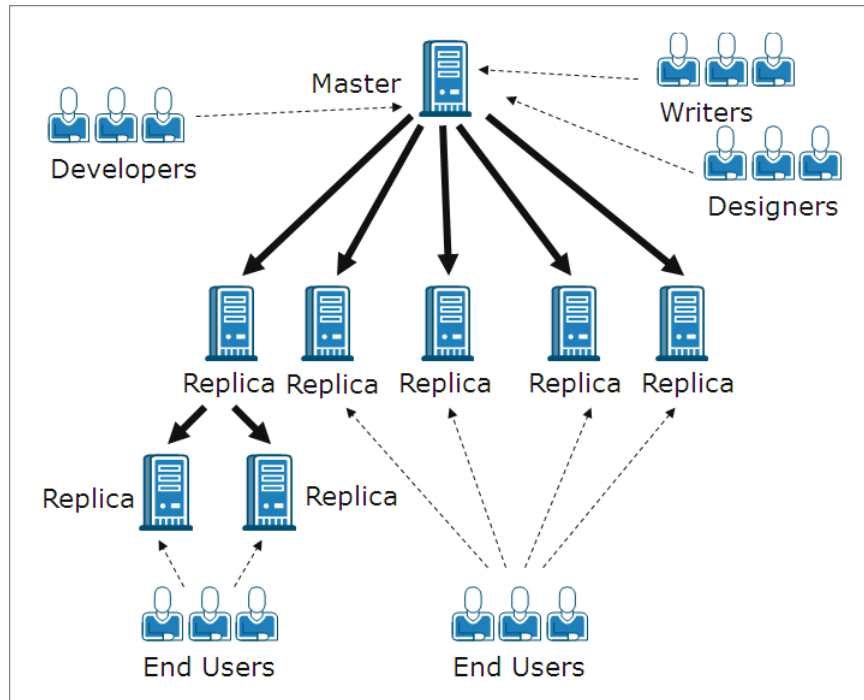
Important! The Content Distribution solution requires a special license.

The Content Distribution solution is aimed at addressing the need of reliably spreading and managing information across a highly distributed environment. In a highly distributed IT environment, many servers contain the same or similar content that they receive from a single repository, and they simultaneously serve many end users. An example of such a distribute environment can be a large organization, which needs to deliver, synchronize and consolidate corporate information among internal users who reside in multiple locations and branch offices. This information can include price lists, policies, sales materials, manuals, and news. With the Content Distribution solution, employees and representatives in the field always have the right information at the right time.

The Content Distribution solution is also a powerful content delivery and web publishing solution that can serve your external customers. Through portals and web sites, you can deliver to your customers any information that is stored in files, from music to movies to documents to news. A good example is a service provider, who distributes content to dozens, hundreds or thousands of e-shops across the globe.

In a regular DR or HA scenario, the Master is usually the active or production server, while the Replica hosts are mainly a storage place for replicated data or standby servers. Unlike this role structure, in a CD scenario the Replica hosts are usually the active hosts, which directly provide information to end users, while the Master host only acts as the initial provider of updated data. The content is maintained in a single repository on the Master, and changes to the Replica hosts are delivered immediately or on a scheduled basis. When applying the CD solution to a large organization, multiple CD scenarios can use the same or overlapping root directories, apply different filtering options, and replicate data to different set of Replica hosts.

The CD solution is designed for one-to-many scenarios, meaning, scenarios that have one Master host and a large number of Replica hosts. These scenarios may replicate many files, or work with a small number of very large files. In this type of scenario, many of the Replica hosts are organized horizontally, as siblings on the same level, and not in hierarchal order as in parent-child relations.



In a regular scenario that contains several Replica hosts on the same level, if more than one Replica host need re-synchronization following a reboot or some connection failure, all other Replica hosts will be re-synchronized as well. However, such a procedure might cause a performance problem when there are hundreds and thousands of Replica hosts. Therefore, in a CD scenario, if more than one Replica host needs re-synchronization, only the hosts that actually need synchronization will be re-synchronized.

Another feature of regular scenarios that might cause problems in a highly distributed environment, is the online replication mode. In a regular online replication mode, changes that occur on the Master are immediately transferred to the Replica, and overwrite the data that exist there. This process is useful for keeping the most up-to-date data on the Replica, but when users are directly using the data that is stored on the Replica, their work might be interrupted by the ongoing and continuous updates. To overcome this problem, a CD scenario can run in a special **On File Close** replication mode, a replication mode that is available only for CD scenarios.


In the **On File Close** mode, all data that is accumulated on the Master is transferred to the Replica, but it does not immediately overwrite the existing Replica data. Instead, data that was changed and transferred to the Replica is saved there as a temporary copy of the original file, and stored in a hidden directory. Once the original file is closed on the Master, the temp copy on the Replica is renamed. When the copy on the Replica receives the original file name, it replaces the older file that is stored on the Replica, and keeps the data on the Replica updated. This method allows for an update process that does not interrupt the user's work. However, if the **On File Close** mode does not suit your environment needs, you can also use either the online or scheduled replication mode for your CD solution.

Creating a Content Distribution Scenario

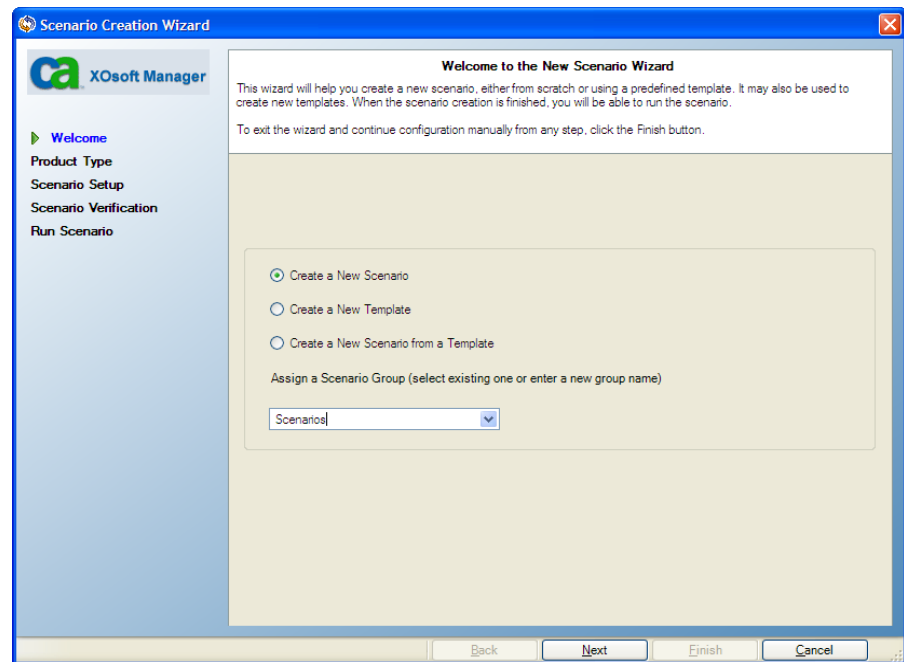
The creation of a CD scenario is similar to the creation of a DR scenario for application and database servers. In both you are using the same step-by-step Scenario Creation wizard. The only major difference between them is that when you select the replication mode of a CD scenario, you have an additional replication mode. This replication mode, **On File Close**, is available only for Content Distribution purposes.

Note: This section demonstrates the configuration of a generic File Server Content Distribution scenario. For more detailed instructions involving scenarios tailored to specific applications, see the appropriate Operation Guide.

To create a Content Distribution scenario

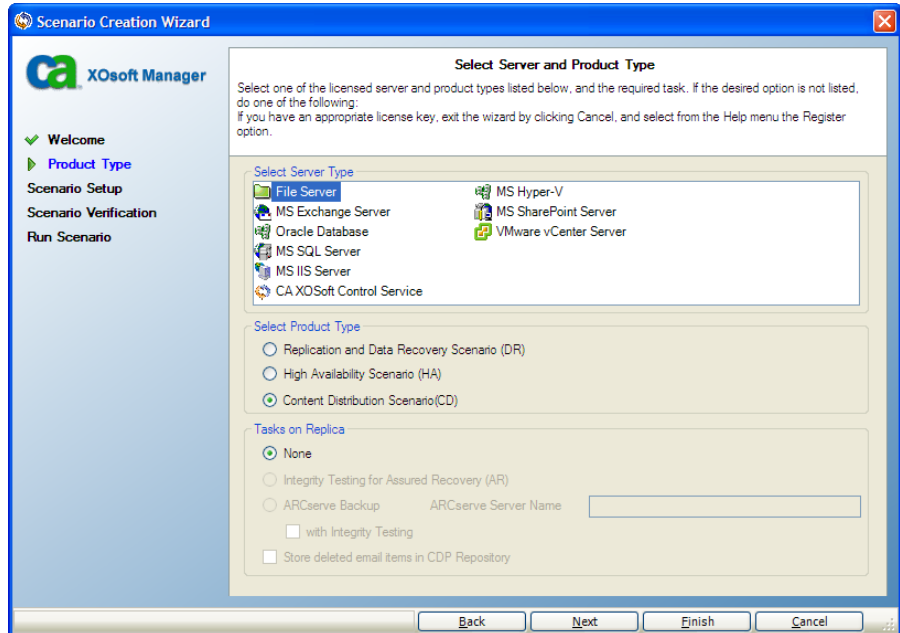
1. Open the CA XOssoft Manager. Then, select from the **Scenario** menu the **New** option, or click the **New**  button on the Standard toolbar.

The **Scenario Creation Wizard** opens.



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.



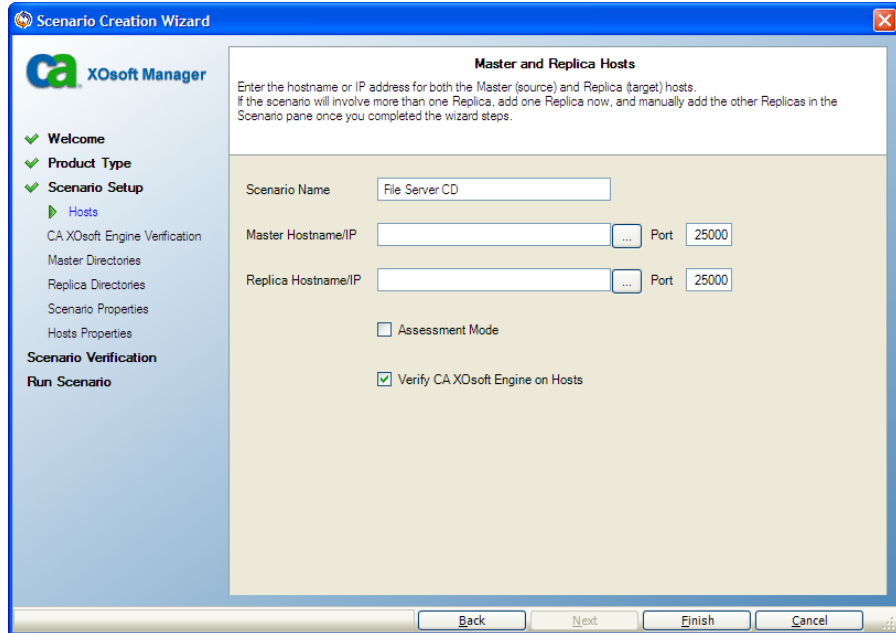
A list of available applications and scenario types is presented.

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

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 scenario.
 - From the **Select Product Type** options, select **Content Distribution Scenario**.

Note: The **Tasks on Replica** options are not available for the CD solution.

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

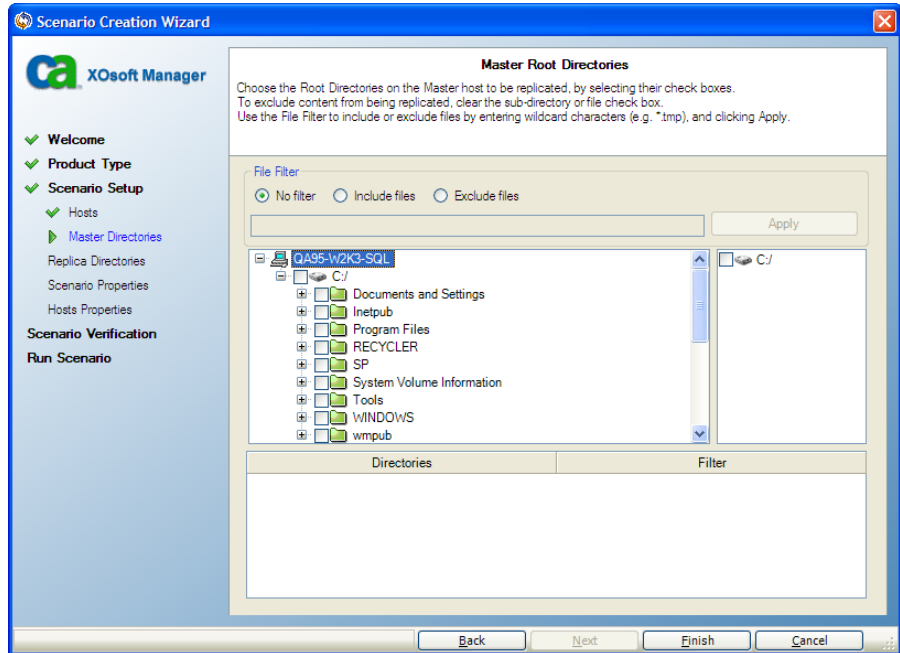


- Enter the following information:
 - In the **Scenario Name** box - accept the default name or enter a new name for the scenario. When entering a name, choose a unique name, since you cannot use the same name for more than one scenario.
 - In the **Master** and **Replica Hostname/IP** boxes - enter 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 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 in [Add Additional Replica Servers](#) (see page 105).

- [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.

8. [Optional] Select the **Verify CA XOssoft 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. For more information about the **Host Verification** page, see [Creating a Scenario](#) (see page 46), step 8.
9. After you selected the desired options, click **Next**. The **Master Root Directories** page opens.



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

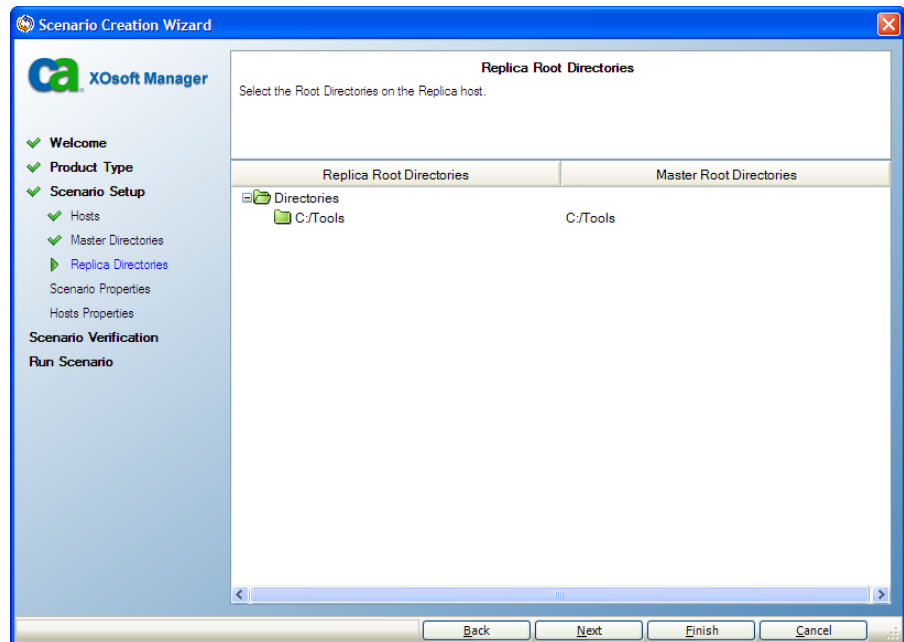
10. Choose the directories and files you want to replicate from the Master to the Replica by selecting their check boxes. You can exclude folders and files from replication by clearing their check boxes.

Notes:

- For more information about selecting and filtering root directories, refer to [Creating a DR Scenario](#) (see page 46), step 12.
- After you finish creating the scenario through the wizard, you can also select registry keys for synchronization, as described in [Synchronize Registry Keys](#) (see page 117).

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

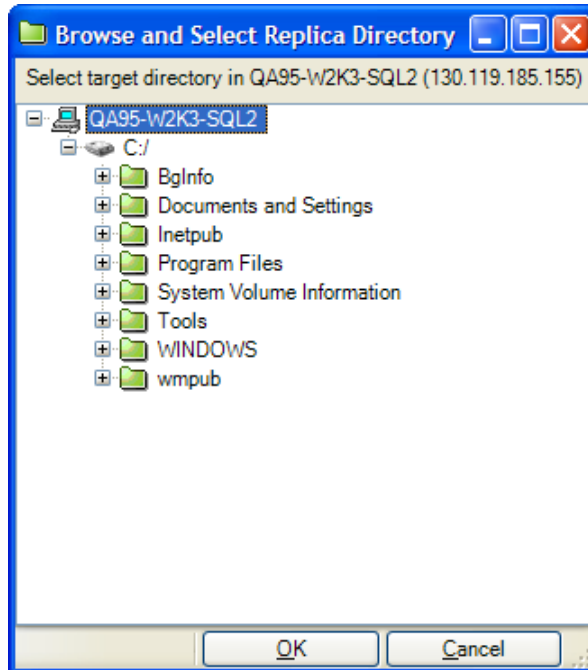
The **Replica Root Directories** page is displayed.



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](#). (see page 125)

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



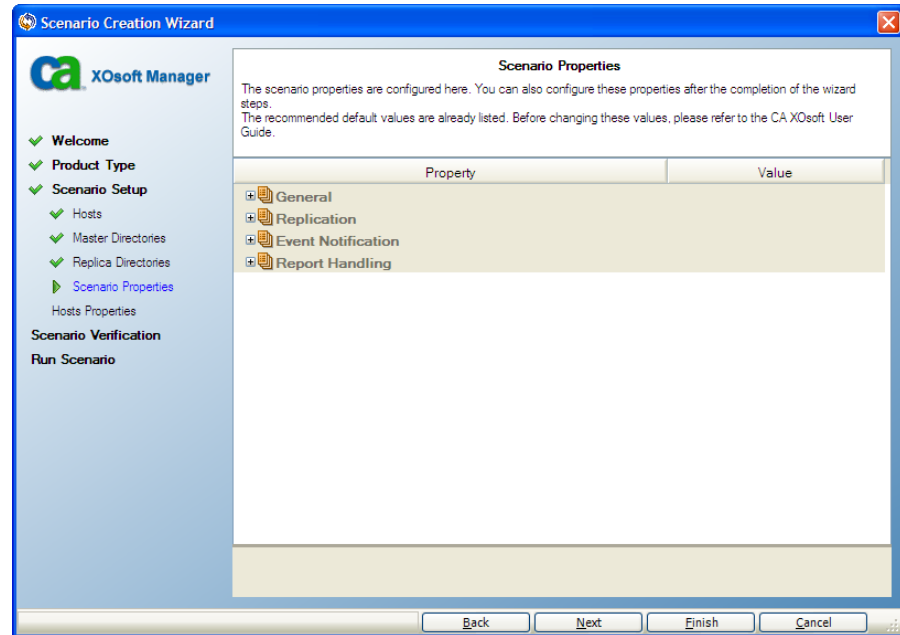
- Select the directory on the Replica in which the replicated data will be stored, and click **OK**.

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 XOssoft creates it automatically.

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

The **Scenario Properties** page opens.



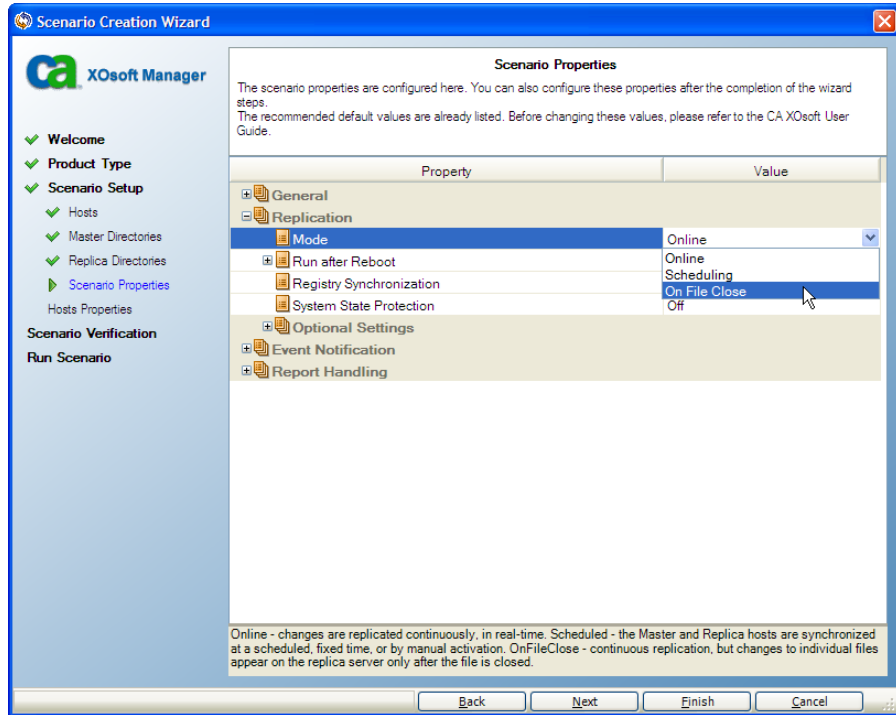
The **Scenario Properties** page enables you to configure the scenario properties that affect the entire scenario. Typically, the default values are sufficient.

If you want to configure the scenario properties at this stage, refer to [Understanding Scenario Properties](#) (see page 138). To configure the scenario properties at a later stage, refer to [Configuring Scenario Properties](#) (see page 137).

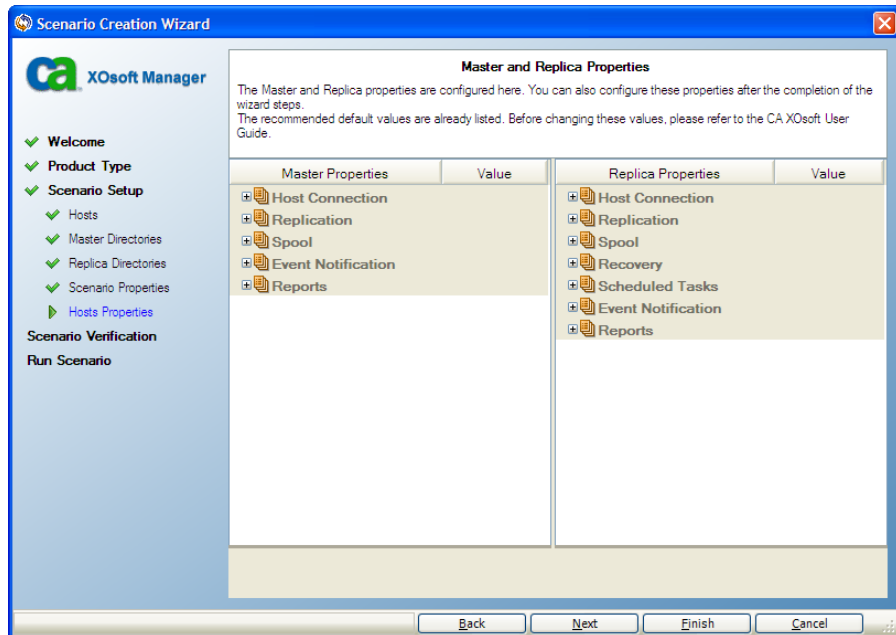
15. In the **Scenario Properties** page you can set the replication mode of the scenario. In addition to the two standard replication modes, **Online** and **Scheduling**, CA XOssoft provides you with another replication mode that is especially designed for the CD scenario, the **On File Close** mode. The **On File Close** mode is similar to the **Online** replication mode with one difference: while in the **Online** mode data changes that are transferred from the Master to the Replica immediately overwrite the existing Replica data, in the **On File Close** mode changes to individual files will appear on the Replica only after the original file on the Master is closed. This way, if users are working directly with data that is stored on the Replica, their work will not be interrupted by constant updates.

Note: The default replication mode is **Online**.

To set the replication mode, open the **Replication** group and select the **Mode** property. Then, select the required replication mode from the drop-down list:



16. Once you set the scenario properties, click **Next**. The **Master and Replica Properties** page opens.



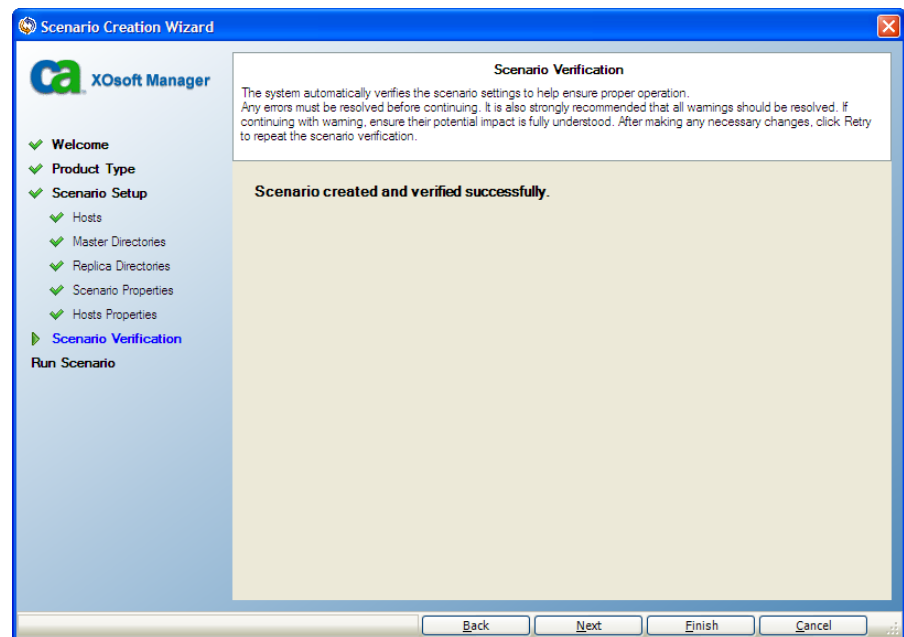
The **Master and Replica Properties** page enables you to configure the properties that are related to either the Master or Replica host. Typically, the default values are sufficient.

If you want to configure the Master and Replica properties at this stage, refer to [Setting Master and Replica Properties](#) (see page 163). To configure the Master and Replica properties at a later stage, refer to [Configuring Master or Replica Server Properties](#) (see page 164).

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 information](#) (see page 167) for configuration details.

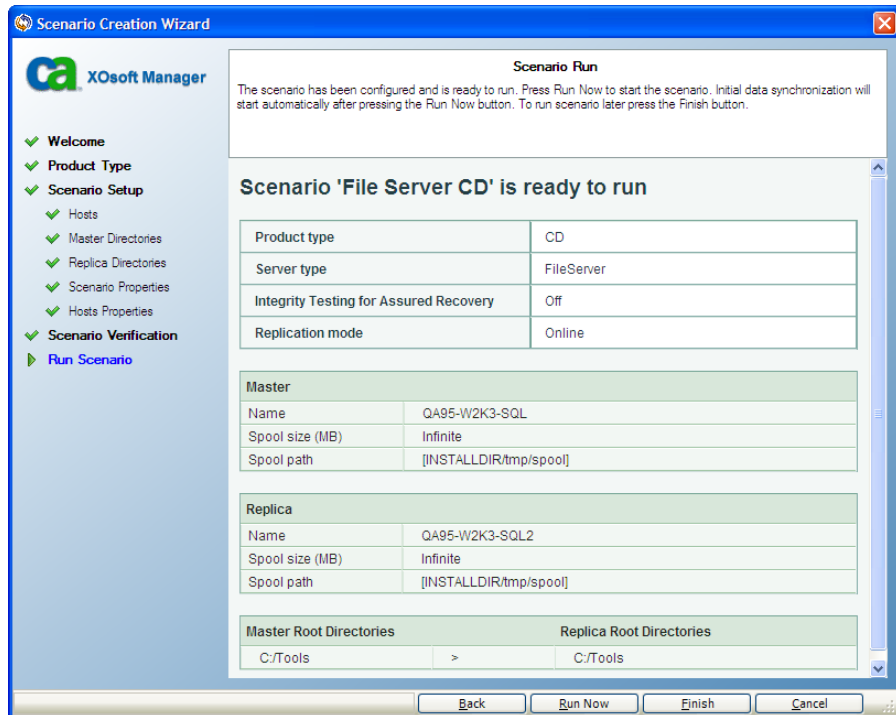
17. After you set the Master and Replica properties, click **Next**.

CA XOssoft verifies the validity of the new scenario and checks many different parameters between the Master and Replica servers to ensure a successful replication, distribution, and data recovery processes. Once the verification is completed, the **Scenario Verification** page opens.



Note: Although CA XOssoft 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.

18. If the scenario is verified successfully, click **Next**.
The **Scenario Run** page opens.



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

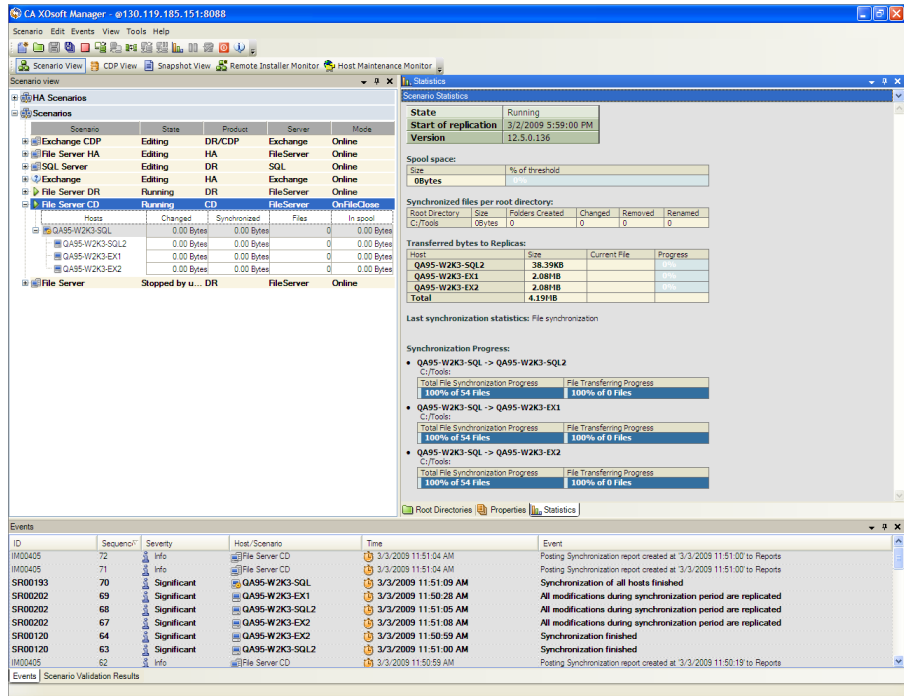
- To add more Replica hosts to the scenario and to run it later, select **Finish**.

Note: CA XOssoft offers you two methods for adding Replica hosts to the scenario:

- Through CA XOssoft Manager, by manually adding each host to the scenario, as described in [Add Additional Replica Servers](#) (see page 105).
- Through CA XOssoft PowerShell, by using the **Add-Replica** and **Add-Replicas** commands. For more information about using CA XOssoft PowerShell commands, refer to *CA XOssoft PowerShell Guide*.

- To run the scenario now, click **Run Now**.
The synchronization process starts.

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



Note: When the scenario has more than one Replica host, the **Scenario Statistics** tab does not display a graphic overview of the scenario state but scenario statistics organized in tables.

21. By default, once a synchronization occurs, a synchronization report is generated. For each Replica host that participates in the scenario, a separate Synchronization report is generated. For more information about opening reports, see [Viewing a Report](#) (see page 94).

Chapter 16: Managing Users

CA XOssoft lets you manage a user's access rights through setting the scenario file ACL properties. ACL is Access Control List, a list of security protections that applies to the scenario file.

A special license is needed.

This section contains the following topics:

[How Delegated Security Works](#) (see page 347)

[Prerequisite Tasks](#) (see page 348)

[How to Manage Users](#) (see page 352)

How Delegated Security Works

Delegated Security lets you control each user's access rights by setting the scenario file access control list (ACL) properties.

The ACL-based authentication model is a role based authentication model in CA XOssoft. There are four pre-defined roles. Each role has pre-defined permissions defining what can be done to a scenario. The roles are:

- Super User
- Admin
- Control
- View-only

A Super User has full control rights for a scenario, while Admin, Control, and View-only have limited rights to the scenario. Only the Super User has the right to create a new scenario.

A user must be assigned one of the four roles to gain access to a scenario. A Super User or Admin can assign users or groups to any scenario and delegate rights to the users or groups. When a user tries to access a scenario through CA XOssoft Manager or CA XOssoft PowerShell, the current role is checked and, based on the permission of the role; the operation is allowed or denied.

Access Rights Considerations

Before assigning user permissions, you should consider the following general conditions.

- All users are Windows domain or local users.
- A Super User has the right to create a new scenario.
- A Super User or Admin can assign users or groups of users to any scenario and delegate rights to the users or groups through CA XOssoft Manager.
- The user or group list with the respective rights is stored in the standard NTFS ACL, applied to the scenario file.
- Super User can change the Super User group. However, after the change, the rights to all existing scenarios must be reassigned.
- Users are allowed to set or change a Super User group which is recorded in an encrypted registry. All Super Users belong to the group.
- The Access Control List is controlled indirectly via Control Service. Since Multiple GUI sessions may connect to one Control Service, the impersonation of each user account becomes indispensable.

Prerequisite Tasks

You must perform the following prerequisite tasks before setting user rights or modifying user groups.

Create a User Group

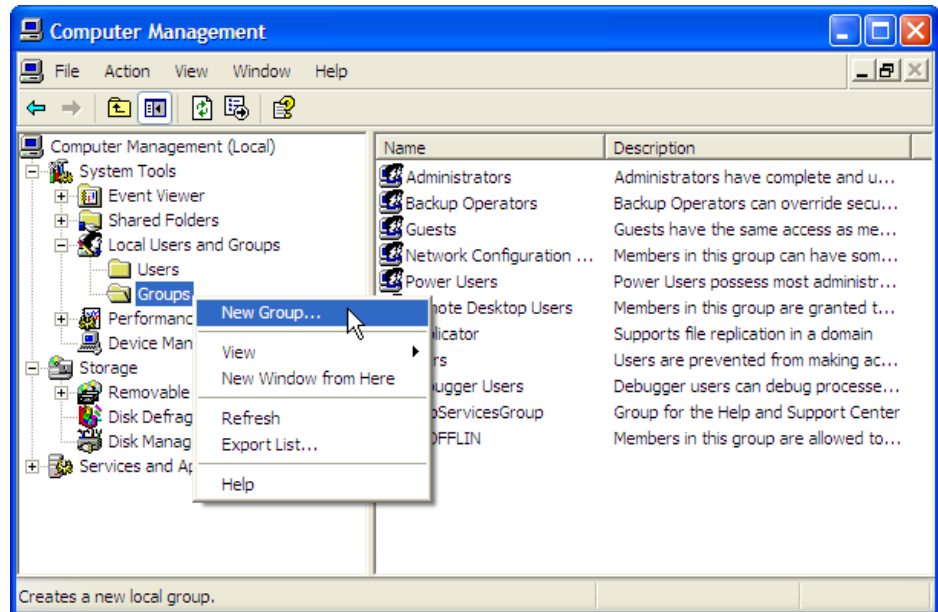
When using ACL authorizations, before you can open the Overview Page and the Manager, you need to create a Local Group. You need to define a Local Group with the name CA XOssoft Users on the Control Service machine, and on all machines that run the CA XOssoft engine, where you want users or groups of users to have permission to add and modify replica hosts or access root directory of hosts.

To create a CA XOssoft User Local Group

1. On the Control Service machine, select Start, Settings, Control Panel, Administrative Tools, Computer Management.

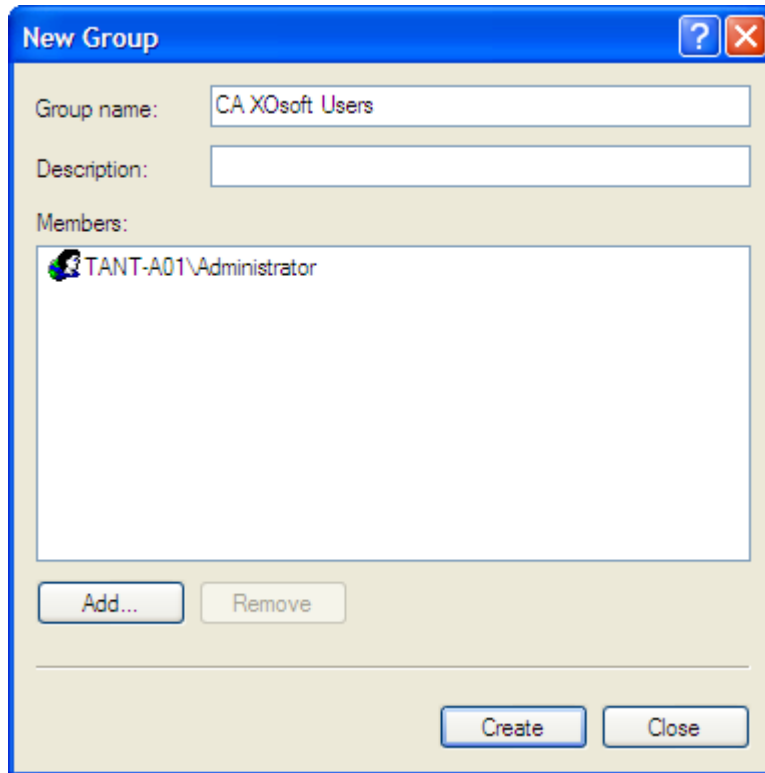
The Computer Management dialog opens.

2. Select the **Local User and Groups** folder, and then open the **Groups** sub-folder.



3. Right-click on **Groups** and select **New Group**.
The **New Group** dialog opens.
4. In the **Group name** box, enter CA XOsoft Users.

5. To add the administrator user, click the **Add** button.



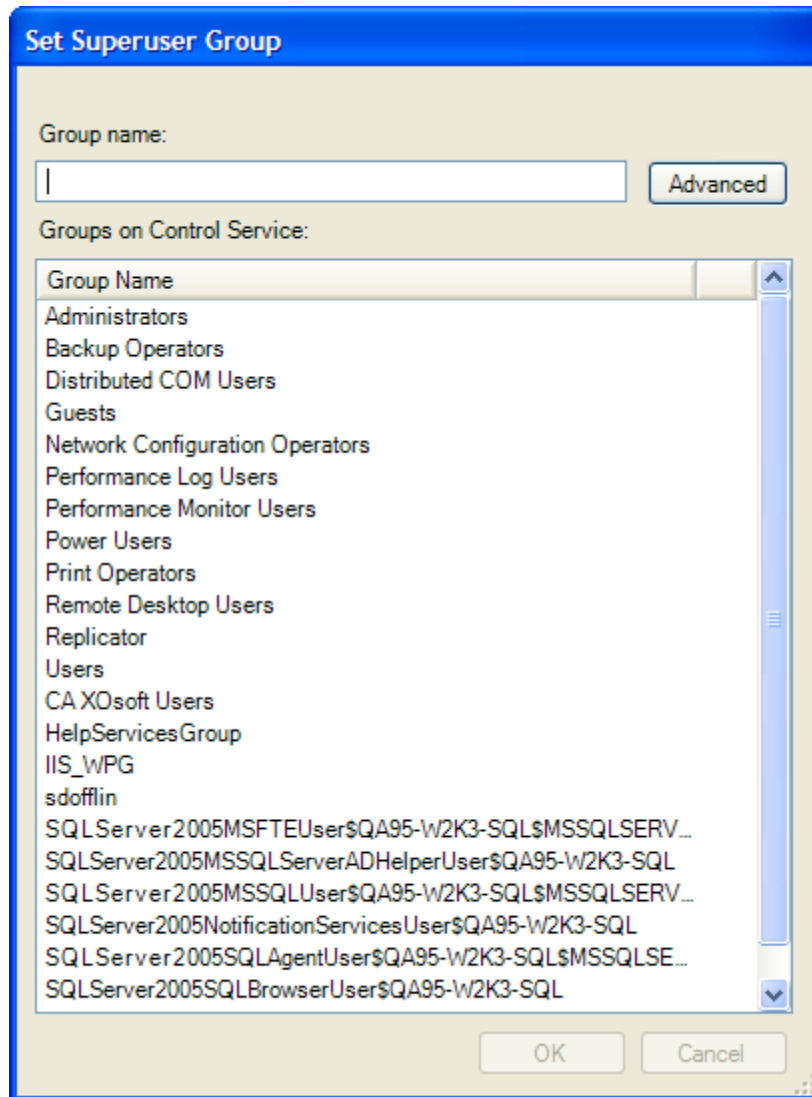
6. Click the **Create** button to create the new Local Group, and then click the **Close** button to close the dialog.

The new group is added to the Local Groups list on the Control Service machine.

Initial Group Selection

You need to define a Local Group with the name CA XOsoft Users on the Control Service machine, and on all machines that run the CA XOsoft engine, where you want users or groups of users to have permission to add and modify replica hosts or access root directory of hosts.

When you open the Manager for the first time, the Manager detects whether a Super User Group already exists. If no Super User group is defined, the **Set Superuser Group** dialog appears, asking you to select:



The Set Superuser Group dialog displays the list of Groups that exist as local groups on the Control Service machine. You need to select the group that will include the members that will be defined as Super Users. You can later change this group.

Set Up a User Group

To set up your environment to use ACL-base delegated security, CA XOssoft lets you use your existing infrastructure or create a new network and local group. There are four groups required:

- Super User
- Admin
- Control
- View

You can assign users to a specific group depending on the required privileges for the individual user. For more information on user permissions, see [Delegation of Rights](#) (see page 353).

Note: You can set up additional groups and designate them Super User, Admin, Control, View or use existing groups in the network.

On each computer participating in the scenario (Master, Replicas, and the Control Service), build a local group with the pre-defined name CA XOssoft Users. Add groups and users of the organization to the CA XOssoft Users local group as required.

When you open the user interface, if a Superuser group has not perviously been selected, you will be required to select one.

Note: Only a Super User can modify a Master server. Replica servers can be modified by a Super User, an Admin, or a Control role.

How to Manage Users

CA XOssoft lets you manage user permissions for scenarios by assigning individual users or groups delegated permissions.

The Super User or Admin manages users rights for a each individual scenario. From the user rights section of the user interface, you can set admin, control, or view permissions for a specific user or group for each scenario. This group or user then has the relevant permission for a particular scenario and can manage the scenario based on the rights they have been assigned. For example, a user or group can have user rights for one scenario and admin rights for an other scenario.

Delegation of Rights

User rights are set per user for using the Manager to make modifications to the Master host, or to the Replicas on its replication tree. User rights are assigned per scenario.

You can assign user permissions based on the following:

Operation	Super User	Admin	Control	View only
Set user rights	Yes	Yes	No	No
Edit Master host	Yes	No	No	No
Edit replication mode	Yes	Yes	No	No
Edit schedule mode	Yes	Yes	Yes	No
Edit Master spool size	Yes	No	No	No
Modify reports on Master	Yes	Yes	No	No
Edit replica host	Yes	Yes	Yes	No
Edit replica spool size	Yes	Yes	No	No
Run a scenario	Yes	Yes	Yes	No
Stop a scenario	Yes	Yes	Yes	No
Synchronize a scenario	Yes	Yes	Yes	No
Restore data	Yes	Yes	No	No
Modify Master notification	Yes	Yes	No	No
Modify Replica notification	Yes	Yes	No	No
Generate a difference report	Yes	Yes	Yes	Yes
Set a bookmark	Yes	Yes	Yes	No
Show difference report	Yes	Yes	Yes	Yes
Run high availably resources	Yes	No	No	No

Operation	Super User	Admin	Control	View only
Check a scenario state	Yes	Yes	Yes	Yes
Suspend a replica	Yes	Yes	Yes	No
Modify reports on replica	Yes	Yes	No	No
Modify Master trigger file	Yes	Yes	Yes	No
Modify Replica trigger file	Yes	Yes	Yes	No

Set User Rights

You can set or reset user rights for a specific scenario.

Note: For all scenarios with licenses other than a delegated security license, you must reset the user rights.

To set user rights

1. From the CA XOssoft Manager Scenario menu, select Scenario, User Rights.

Important! The Scenario menu contains the extra option User Rights. This option is only available to users with Super User or Admin rights.

The Security window opens displaying the security rights of each scenario.

2. Click Add.

The Select Users or Groups window is opens.

3. From the Look in field drop-down list, select a Domain.
4. Select the required user or group.

Note: Multiple user or group selections are not supported.

5. Click Add and then click OK.
6. From the Permission column, set access rights for a user or group from the drop-down list.

Important! If you remove yourself (Admin) from the list in the security window, you will no longer have any user rights in the current scenario. After restarting the CA XOssoft Manager or waiting for approximately 10 seconds, the scenario no longer appears in the list of scenarios.

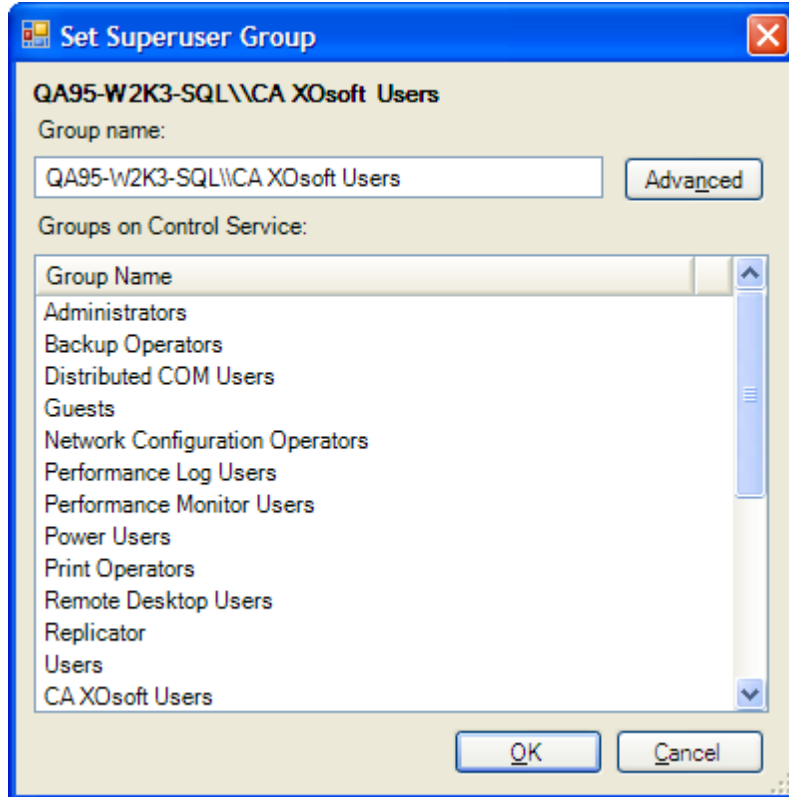
Setting the Super User Group

You can change a Super User group at any time.

To modify the super user group

1. On the Manager, open the **Scenario** menu and select the **Set Superuser Group** option.

The **Set Superuser Group** opens.



2. From the **Groups on Control Service** list, select the group to which you want to assign the Super Users group.

Index

A

- Abbreviations • 14
- About Assured Recovery • 48, 75, 265
- About this Guide • 13
- Access Rights Considerations • 348
- Actions upon Success • 232
- Actions upon successful test
 - AR, DB offline • 288
 - HA • 232
- Actions upon Successful Test (DB Offline) • 288
- Actions upon Successful Test (DB Online) • 287
- Actions, automatic configuration • 195
- Activate the Registry Synchronization Option • 117, 118
- Active and standby hosts • 193, 223
- Active and Standby Hosts • 223
- Active Directory Integrated • 225
- Active server, recovering • 215
- Active to Standby Redirection Script • 225
- Add Additional Replica Servers • 49, 105, 198, 338, 345
- Adding
 - IP/Mask for Switchover • 225
 - Replica server to scenario • 105
- Adding IP Address to the Master Host • 233
- Adding XO-IP to Existing Scenarios • 237
- Adding XO-IP to New Scenarios • 235
- Additional System State Information • 161
- Advanced
 - scheduling • 148
 - verification HA scenario results • 195, 211
- Application and database servers supported • 19
- ARCserve Backup
 - Integration with CA XOsoft • 46
 - VSS Snapshots creation • 289
- Assessment mode • 22, 40, 46
- Assured Recovery
 - defining for scenario • 267
 - defining Replica for • 267
 - manual testing • 279, 282
 - non-scheduled mode • 278
 - overview • 265
 - performing test • 275

- properties, setting • 284
- report, define generation • 175
- report, viewing • 279
- scenario creation • 267
- scheduled mode • 276
- Scheduler • 286
- Start DB • 286
- steps • 275
- supported applications • 265
- test, performing • 275

- Auto-discover Database Files for all Databases • 123
- Autodiscovering database files • 123
 - in scenario creation • 195
- Automatic Replica testing • 265
- Automatic Switchover • 195
- Automatic synchronization • 22
 - preventing upon error • 139
 - scheduling • 145
 - setting • 139
- Automatic Synchronization • 22

B

- Backward scenario • 195
 - report • 100
- Backward Scenario Report • 100
- Bandwidth limit, scheduling • 170, 176
- Block Synchronization • 20, 139
- Bookmarks
 - setting • 186
 - term • 14

C

- CA XOsoft
 - ARCserve Backup Integration • 46
 - auto-configuration • 195
 - autodiscovery • 195
 - components • 25
 - deployment • 29
 - High Availability • 18, 24
 - limitations • 18
 - logging in • 31
 - Manager • 31
 - overview • 17
 - registration • 40

-
- replication • 17
 - solutions • 19
 - CA XOssoft Components • 25
 - CA XOssoft Control Service • 25
 - CA XOssoft Deployment • 29
 - CA XOssoft Engine • 26
 - CA XOssoft High Availability • 18
 - CA XOssoft PowerShell • 28
 - CA XOssoft Replication • 17
 - CA XOssoft Replication and High Availability Overview • 17
 - CA XOssoft Solutions • 19
 - CDP Admin • 298
 - CDP Database, defining • 300
 - CDP Repository • 28
 - CDP Admin • 298
 - CDP Storage • 298
 - CDP Support • 298
 - CDP Web Server • 298
 - components • 298
 - configuring • 299
 - creating scenario for • 305
 - defining the CDP database • 300
 - E-mail Retrieval • 298
 - properties • 305, 314
 - quota and aging policy • 300
 - retrieving deleted items • 327
 - understanding • 297
 - using • 46, 297
 - CDP Repository Components • 298
 - CDP Repository Reports • 11, 323
 - CDP Repository Statistics • 11, 320
 - CDP Storage • 298
 - CDP Support • 298
 - CDP Web Server • 298
 - Change Configuration when a Scenario is Running • 72
 - Changing configuration during replication • 72
 - Check method for Is Alive • 229
 - Close and Open the CA XOssoft Manager during Replication • 74
 - Closing Manager during replication • 74
 - Cluster Move IP • 239
 - Clusters • 195, 305
 - Command Line Enhancements for System State Protection • 161
 - Components, CA XOssoft • 25
 - Compress Data during Transfer on Replica • 170
 - Configuration, changing during replication • 72
 - Configure System State Protection for Existing Scenarios • 154
 - Configure System State Protection in the Scenario Creation Wizard • 151
 - Configure System State Protection on the Replica • 154, 155
 - Configuring
 - CDP Repository • 299
 - Master server • 195
 - Replica server • 195
 - scenario properties • 137
 - Configuring Assured Recovery Properties • 284
 - Configuring CDP Replica and Report Properties • 318, 319
 - Configuring CDP Scenario Properties • 310, 314
 - Configuring High Availability Properties • 202, 221, 253
 - Configuring Master or Replica Server Properties • 57, 164, 201, 252, 273, 309, 344
 - Configuring Scenario Properties • 56, 137, 200, 251, 272, 309, 342
 - Configuring the CDP Repository • 299
 - Configuring the Move IP Method through the Manager • 235
 - Configuring VSS Snapshot Creation • 288, 289
 - Contact CA • iii
 - Control Service • 25
 - Copy Events • 91
 - Copying events • 91
 - Create a New Scenario using an Existing Template • 63
 - Create a Template • 60
 - Create a User Group • 348
 - Creating
 - new storage groups and mailbox stores, automatically • 195
 - scenario group • 45
 - scenario, using template • 63
 - scenario, using wizard • 46
 - Shadow Copy (VSS) • 293
 - template • 60
 - Creating a Content Distribution Scenario • 336
 - Creating a Scenario • 40, 45, 46, 62, 65, 245, 249, 270, 274, 339, 340
 - Creating a Scenario Group • 40, 45
 - Creating an Exchange Scenario with the CDP Option • 299, 305
 - Creating Assured Recovery Testing Scenario • 267
-

Creating High Availability Scenario • 195, 271, 274

Creating High Availability Scenario for the Control Service • 245, 246, 247

Creating Replication Scenarios • 45

Customizing

Scenario pane • 37, 82

Customizing the Scenario Pane • 37, 82

D

Data Recovery • 40, 181, 245

Data recoverysee Recovering data • 23

Data Replication and Role Switching of Control Services • 11, 243

Data rewind

activating option • 173

how it works • 24

max disk size for • 173

retention period, defining • 173

settings • 173

solution • 24

term • 14

Data Rewind • 187

Database servers supported • 19

DB Management/Application/Shares Management • 231

DB Testing of Replica • 287

Define the Master and Replica Servers • 104

Defining

Master server, manually • 104

Replica server, manually • 104

Defining the CDP Database • 42, 299, 300

Delay replication • 173

Delegation of Rights • 352, 353

Deleting

reports • 96

scenario • 129

VSS Snapshots • 295

Deleting Reports • 96

Deployment

CA XOssoft • 29

Detailed Email CDP Repository Report • 11, 318, 326

Detailed report

replication, define generation • 168

synchronization, define generation • 168

Detailed reports • 94

Detailed Synchronization Report • 98

Difference report • 101

Difference Report • 41, 101

Directory

spool in Master • 167

spool in Replica • 173

Disaster Recovery, defining for scenario • 46

DNS

Active Directory Integrated • 225

Key Filename • 225

Master/Replica IPs in DNS • 225

redirect for Switchover • 225

TTL • 225

DNS redirection • 24

Docking

panes • 38

tool • 38

Documentation Changes • 11

Documentation, related • 16

E

Edit Directory Names • 108

Editing

root directory names • 108

Editing and Managing Scenarios and Hosts • 103

Email CDP Repository Report • 11, 318, 324, 326

Email CDP Repository Statistics Summary Table • 11, 320, 322

Email CDP Repository Statistics Table • 11, 320, 321

E-mail Retrieval

Overview • 298

using • 327

Engine

verification • 46

Errors, before running • 67, 195

Event Notification • 143, 167, 174

Master • 167

Replica • 174

scenario • 143

Events • 90

copy • 91

filter • 92

pane • 90

viewing • 90

viewing in separate window • 90

viewing incoming, using pop-up window • 91

Exchange Server

- configuration verification • 195, 305
- Exchange Virtual Server Name • 195, 305
- Exclude Dates from Scheduled Synchronization
 - 147
- Exclude Files • 110, 114
- Excluding
 - dates from scheduled synchronization • 147
 - files from replication • 114
 - files from replication in HA scenario • 195
- Exploring CA XOssoft Manager • 31
- Exploring the CA XOssoft Manager Screen • 34, 42
- Export Scenarios • 130
- Exporting scenarios • 130
- Exposing VSS Snapshots • 295

F

- File Synchronization • 139
- Files to replicate
 - exclude • 114
 - include • 109, 111
- Filter Events • 11, 92
- Filter Master Directory Files • 53, 109
- Filtering • 109
 - exclude files • 114
 - include files • 111
 - Master directories • 109
- For Existing Scenarios • 240
- For New Scenarios • 240

G

- General • 139
- Generating reports
 - replication, defining for Master • 168
 - replication, defining for Replica • 175
 - synchronization, defining for Master • 168
- Graphical view of replication • 67
- Group, scenario • 45
- GUISee Manager • 46

H

- Heartbeat Frequency for Is Alive checks • 229
- Hiding panes • 38
- High Availability
 - defining for scenario • 195
 - Is Alive check • 229
 - process • 24
 - properties • 221
 - properties, setting • 221

- properties, understanding • 195, 223
- recover active server • 215
- solution • 18
- High Availabilitysee also Switchover • 225
- Host Connection • 165, 170
- Host Connection, setting for Replica • 170
- Host Maintenance • 41, 42, 132
- Hosts • 224
- Hosts, setting for switchover • 224
- How Data Rewind Works • 24
- How Delegated Security Works • 347
- How High Availability Works • 24
- How Recovery Works • 23
- How Replication Suspension Works • 24
- How Replication Works • 22, 40
- How Synchronization Works • 19, 73, 101
- How to Configure System State Protection • 150
- How to Manage Users • 352

I

- ID, scenario • 139
- Identify Network Traffic Direction Script • 225
- Ignore Files of Same Size/Time • 21, 139
- Import Scenarios • 131
- Importing scenarios • 131
- Include Files • 110, 111
- Including files in Master root directories • 111
- Initial Group Selection • 351
- Initiate Replication • 40, 67, 117, 122
- Initiating
 - replication • 67
 - scenario • 67
 - synchronization • 73
- Installing Engine remotely Engine
 - install remotely • 46
- Integrity Testing for Assured Recovery,
 - defining for scenario • 267
- Introduction • 13
- Is Alive • 41, 202, 203, 224, 225, 229, 253
 - check method • 229
 - Heartbeat Frequency • 229
 - properties • 229
 - timeout • 229

K

- Keep Deleted Files during Replication, on Replica • 170

Keep Deleted Files during Synchronization, on
Replica • 170

L

License Registration • 43, 244

Limitations • 18

Live statistics • 83

Live Statistics • 83

Locked for Backup • 294

Logging in to the Management Center • 31

Logging Into the Email Retrieval • 327

Logging Into the Management Center • 31

M

Maintenance, Host • 132

Management Center • 27

logging in • 31

Manager • 27

Overview Page • 27

Report Center • 27

Manager

arranging panes • 38

close/open • 74

creating scenario using wizard • 46

exploring • 31

opening • 31

overview • 27

panessee Panes • 38

setting bookmarks • 186

viewing options • 37

Manager Screen Areas • 35

Managing Snapshots • 295

Managing Users • 11, 347

Manually

refreshing statistics display • 89

testing Assured Recovery • 279, 282

Manually Initiating a Control Service

Switchover • 245, 258, 259

Manually Recover a Failed Server - Move IP

Address • 217

Manually Recover a Failed Server-Switch

Computer Name • 218

Manually Recover Failed Server-IP and Switch

Name • 219

Mask, add for Switchover • 225

Master server

compress data during transfer • 165

configuration • 195

defining in Scenario Creation Wizard • 46

defining, manually • 104

event notification • 167

generating replication report for, defining •
168

max spool size • 167

min disk free size • 167

properties • 165

properties, setting • 164

report handling • 168

report settings • 168

root directories, selecting manually • 104,
105

spool directory • 167

spool setting • 167

synchronization report for, defining • 168

term • 14

Master/Replica IPs in DNS • 225

Max Disk Size for Rewind journal • 173

Max Spool Size

Replica • 173

Methods of synchronization • 20

Min Disk Free Size of Spool

Master • 167

Replica • 173

Mode, Replication • 139

Modify Scenario System State Protection • 157

Monitor Multiple Scenarios • 82

Monitoring

state information • 83

statistics • 83

Monitoring Replication • 81

Mount points replication • 67

Mounting VSS Snapshots • 295

Move IP

setting for Switchover • 225

Move IP Redirection • 233

N

NetBIOS name, using for Switchover • 225

Node • 14

Notificationsee Event Notification • 143

O

Online replication mode • 22, 139

Opening

Management Center • 31

Manager • 31

Overview page • 31

Report Center • 94

Opening the Manager for Using the HA Control Service Scenario • 245

Overview page
opening • 31
overview • 27

P

Panes

docking • 38
docking tool • 38
Events • 90
hiding • 38
rearranging • 38
Scenario, customizing • 37, 82
stacking • 38
Statistics • 84

Perform Assured Recovery Test Automatically • 279, 283

Perform Assured Recovery Test Manually • 280, 282

Perform Switchover Automatically • 224

Performing Assured Recovery Test • 275

Performing Assured Recovery Test in a Non-Scheduled Mode • 41, 278

Performing Assured Recovery Test in a Scheduled Mode • 276

Pop-up window on Incoming Events, viewing • 91

PowerShell • 28

Preferred Number of Snapshots to Keep • 293

Prerequisite Tasks • 348

Prevent Automatic Re-sync upon Error • 139

Propagating Master Root Directories to Multiple Replica Hosts • 11, 127

Propagating Property Values • 11, 178

Properties

Assured Recovery • 284
CDP Repository • 305, 314
High Availability • 221
Master • 165
Replica • 170
scenario • 137
Switchover • 195, 221
VSS Snapshots • 293

Protecting Your System State • 11, 141, 149, 172

Q

Quota and Aging policy, of CDP Repository • 300

R

Rearranging Manager panes • 38

Rearranging Panes • 38

Reboot after Switchover and Switchback • 225

Recover Active Server • 215

Recover Active Server from Outside the Manager • 216

Recover Active Server Using the Manager • 211, 215

Recover Lost Data from Replica • 182

Recovering data

how it works • 23

process • 181

solution • 23

term • 14

Using • 181

Recovery • 173, 190

Redirect DNS, setting for Switchover • 225

Redirection • 225, 229

Redirection methods

Active to Standby Redirection Script • 225

Identify Network Traffic Direction Script • 225

Move IP • 225

redirect DNS • 225

settings • 225

Standby to Active Redirection Script • 225

Switch Computer Name • 225

user-defined scripts • 225

Refresh rate, statistics • 89

Refresh Statistics Display Automatically • 89

Refresh Statistics Display Manually • 41, 89

Refresh, Snapshots Management window • 295

Registration, license • 40

Registry Keys • 117

option • 118

registry select • 121

registry selection • 119

Related documentation • 16

Related Documentation • 13, 16

Remove Master Root Directories • 109

Remove Scenarios • 129

Removing

Master root directories • 109

reports • 96

-
- scenario • 129
 - Removing storage groups and mailbox stores • 195
 - Replica Integrity TestingSee Assured Recovery • 265
 - Replica server
 - adding to scenario • 105
 - compress data during transfer • 170
 - configuration • 195
 - defining in Scenario Creation Wizard • 46
 - defining, manually • 104
 - event notification • 174
 - generating replication report for, defining • 175
 - host connection • 170
 - keep deleted files during replication • 170
 - keep deleted files during synchronization • 170
 - max spool size • 173
 - min disk free size • 173
 - properties • 170
 - properties, setting • 164
 - report handling • 175
 - report settings • 175
 - retry if file is busy • 170
 - root directories, selecting manually • 125
 - spool • 173
 - spool directory • 173
 - stop database on run • 170
 - suspending • 75
 - term • 14
 - testing Assured RecoverySee Assured Recovery • 46, 265
 - Replicate NTFS ACL • 139
 - Replicate NTFS Compress Attribute • 139
 - Replication • 73, 139, 165, 170
 - assessment modeSee Assessment mode • 22
 - changing configuration during • 72
 - close/open Manager during • 74
 - delay • 173
 - excluding files from, in HA scenario • 195
 - graphical view • 67
 - how it works • 22
 - initiating • 67
 - keep deleted files on Replica during • 170
 - limitations • 18
 - mode, setting • 139
 - monitoring • 81
 - mount points • 67
 - online mode • 139
 - overview • 17
 - reports • 100
 - reportsee Replication report • 168
 - reportssee Replication reports • 99
 - resume after suspension • 77
 - running • 67
 - scheduled mode • 139
 - solution • 22
 - stopping • 71
 - suspending • 75
 - term • 14
 - Replication report
 - generating, define for Master • 168
 - generating, define for Replica • 175
 - viewing • 99
 - Replication Report • 69, 99
 - Replication tree, term • 14
 - Report Center
 - deleting reports • 96
 - Overview • 27
 - using • 94
 - Report Centersee also Reports • 94
 - Report handling
 - Master • 168
 - Replica • 175
 - scenario • 144
 - Report Handling • 93, 96, 144
 - Reporting synchronization differences • 22
 - Reporting Synchronization Differences • 22
 - Reports • 93, 99, 168, 175
 - Assured Recovery, define generation • 175
 - Assured Recovery, viewing • 279
 - backward scenario • 100
 - deleting • 96
 - detailed and summary • 94
 - Difference report • 101
 - notify by email • 144
 - replication • 99
 - replication, define generation for Master • 168
 - replication, define generation for Replica • 175
 - setting for Master • 168
 - setting for Replica • 175
 - setting for scenario • 144
 - synchronization report • 97
 - synchronization, define generating for Master • 168
 - using • 93
-

- viewing • 94
- Restore System State Data • 159
- Resume replication • 14
- Resume Replication after Manual Suspension • 77
- Resume replication after suspension • 77
- Resynchronization
 - automatic, prevent upon error • 139
- Retrieving deleted messages using E-mail Retrieval • 327
- Retrieving Deleted Outlook Items Using the E-mail Retrieval • 327
- Retry if File is Busy • 170
- Reverse Replication • 195
- Rewind bookmarkSee Bookmarks • 186
- Rewind points
 - Term • 14
- Root directories
 - auto-discovery • 123
 - edit • 108
 - filter • 109
 - remove • 109
 - Replica • 125
 - selecting • 105
- Run
 - after Reboot, automatically • 139
 - dialog • 67
 - mode • 70
- Run Mode • 70
- Run Reverse Replication Scenario • 195
- Run Reverse Replication Scenario after Switchover • 224
- Run Script after Synchronization
 - on Replica • 170
- Run Script before Synchronization
 - on Replica • 170
- Run Script upon Trigger File Creation
 - on Replica • 170
- Running
 - DR scenario from wizard • 46
 - HA scenario, from Wizard • 195
 - replication • 67
 - scenario • 67
 - synchronization • 73
- Running the Replication Process • 67

S

- Save Scenarios • 40, 129
- Saving scenarios • 129

- Scenario
 - Assured Recovery, creating • 267
 - backward • 195
 - creating using wizard • 46
 - defining Master and Replica, manually • 104
 - deleting • 129
 - Disaster Recovery, defining • 46
 - exporting • 130
 - for CDP Repository • 305
 - graphical view • 67
 - group • 45
 - High Availability defining • 195
 - ID • 139
 - importing • 131
 - operations • 129
 - pane ,customizing • 37, 82
 - properties, setting • 137
 - removing • 129
 - replication mode • 139
 - Report settings • 144
 - running from wizard • 46
 - saving • 129
 - stopping • 71
 - term • 14
 - type, defining • 46
 - verification • 46
- Scenario Operations • 129
- Scenario Pane • 83
- Scenario properties
 - configuring • 137
 - event notification • 143
 - General • 139
 - optional settings • 139
 - Replication • 139
 - report handling • 144
 - Schedule synchronization • 145
 - setting • 137
 - understanding • 46, 138, 195, 267, 305
- Schedule Replication Suspension • 77, 174
- Schedule Synchronization • 73, 79, 140, 145, 277
- Schedule the Bandwidth Limit • 11, 171, 176
- Scheduled Assured Recovery mode • 276
- Scheduled replication mode • 22, 139
- Scheduled Tasks • 174
- Scheduler • 286
- Scheduling
 - advanced • 145, 148
 - Assured Recovery • 276
 - replication • 139

-
- suspension • 77, 174
 - synchronization • 145
 - Search for Deleted Items and Retrieve Them • 329
 - Select Master Directories and Their Contents for Replication • 104, 105, 110
 - Select Registry Keys for Synchronization • 117, 118, 119
 - Select Replica Root Directories • 54, 104, 125, 340
 - Select the Name and Storage Location of the Synchronized Registry Keys • 117, 121
 - Selecting
 - Master directories and their contents • 105
 - Replica root directories • 125
 - scheduling bandwidth limit • 170, 176
 - Set a Schedule for Automatic Synchronization • 146
 - Set Advanced Schedules • 148
 - Set the System State Protection Schedule • 154, 155
 - Set Up a User Group • 352
 - Set User Rights • 354
 - Setting
 - Assured Recovery properties • 284
 - bookmarks • 186
 - High Availability properties • 221
 - Master properties • 164, 165
 - Replica properties • 164, 170
 - Scenario, properties • 137
 - Setting Assured Recovery Properties • 174, 280, 283, 284
 - Setting Bookmarks • 186
 - Setting CDP Scenario Properties • 11, 313
 - Setting High Availability Properties • 221
 - Setting Master and Replica Properties • 57, 163, 201, 252, 273, 309, 344
 - Setting Scenario Properties • 137
 - Setting the Super User Group • 355
 - Setting Up Snapshot Creation • 290
 - Shadow CopiesSee VSS Snapshots • 289
 - Simultaneous synchronization and replication • 22
 - Simultaneous Synchronization and Replication • 22
 - Skipping Synchronization • 67
 - Snapshot Management window • 294
 - SnapshotsSee VSS Snapshots • 289
 - Spool • 57, 167, 173, 201, 252, 273, 309, 344
 - directory for CDP Repository • 305
 - directory in Master • 167
 - directory in Replica • 173
 - max size in Master • 167
 - max size in Replica • 173
 - min disk free size in Master • 167
 - min disk free size in Replica • 173
 - of Master, settings • 167
 - of Replica, settings • 173
 - SSL, opening the Overview page using • 31
 - Stacking panes • 38
 - Standard Toolbar • 40
 - Standby to Active Redirection Script • 225
 - Start DB • 286
 - Start DB, in AR test • 286
 - Starting
 - replication • 67
 - scenario • 67
 - State information • 83
 - lost connection • 83
 - State Information • 83
 - Statistics • 83
 - pane • 84
 - refresh • 89
 - refresh rate • 89
 - Statistics Tab • 84
 - Stop DB • 288
 - Stop Replication • 71
 - Stop upon Error • 139
 - Stopping
 - database on Run • 170
 - replication • 71
 - scenario • 71
 - Storage groups and mailbox stores
 - Creating • 195
 - removing • 195
 - updating • 195
 - Store System State Protection Properties • 155, 156, 158
 - Supported application and database servers • 19
 - Supported Application and Database Servers • 19
 - Suspend Replication • 41, 75
 - Suspend Replication Manually • 76
 - Suspension
 - activating • 75
 - how it works • 24
 - manual • 76
 - resume replication after • 77
 - schedule • 77, 174
-

- term • 14
- Switch Computer Name • 225
- Switchback • 211
 - initiating • 211
 - understanding • 193
- Switching Back the Control Service Roles • 245, 258, 262, 263
- Switching Over and Switching Back • 193
- Switching the Roles of the Active and Standby Control Services • 258
- Switchover • 41, 208
 - active and standby host • 193, 223
 - Active to Standby Redirection Script • 225
 - add IP/Mask • 225
 - automated • 224
 - automatic initiation, defining • 195
 - automatic or manual, defining • 195
 - host name, defining • 224
 - hosts, setting • 224
 - Identify Network Traffic Direction Script • 225
 - initiating • 208
 - Is Alive check • 229
 - Move IP • 225
 - NetBIOS name for • 225
 - perform automatically, define • 224
 - properties • 221, 224
 - Properties page • 195
 - Reboot after, setting • 225
 - Redirect DNS • 225
 - redirection methods, define • 225
 - reverse replication scenario initiation, defining • 195
 - Run Reverse Replication Scenario after, define • 224
 - Standby to Active Redirection Script • 225
 - Switch Computer Name • 225
- Switchover Properties • 224
- Synchronization • 40, 73
 - advanced scheduling • 148
 - automatic • 22
 - automatic, defining • 139
 - automatic, preventing upon error • 139
 - automatic, scheduling • 145
 - Block • 20, 139
 - by user request, defining • 139
 - dialog • 73
 - excluding dates from schedule • 147
 - File • 139
 - filter • 21

- hours • 146
- how it works • 19
- keep deleted files on Replica during • 170
- manual • 73
- method, selecting for run • 67
- methods • 20
- reportsee Synchronization report • 168
- run script after on Master • 165
- run script after on Replica • 170
- run script before on Master • 165
- run script before on Replica • 170
- scheduling • 145
- skipping • 67
- solution • 19
- starting • 73
- suspend replication during • 75
- term • 14
- types • 139
- Synchronization Filter • 21
- Synchronization Methods • 20
- Synchronization report
 - detailed • 98
 - generating, define for Master • 168
 - summarized • 97
 - viewing • 97
- Synchronization Report • 97
- Synchronize Registry Keys • 11, 52, 105, 117, 141, 183, 189, 340
- Synchronize Windows Shares • 139
- Synchronizing hours • 146

T

- Templates
 - creating • 60
 - creating scenario using • 63
 - using • 60
- Terms and abbreviations • 14
- Terms and Abbreviations • 14
- Testing Assured Recovery and Managing VSS Snapshots • 265
- Testing Assured RecoverySee Assured Recovery • 275
- The CA XOssoft Manager • 82
- The CA XOssoft Overview Page • 81
- The Data Recovery Process • 181
- The Switchover and Backward Scenario Processes • 245, 258, 259, 260
- Toolbar buttons • 40
- Toolbars • 40

Trigger File, run script upon creation
 on Master • 165
 on Replica • 170
TTL, DNS • 225

U

Understanding Assured Recovery Properties • 273, 286
Understanding CDP Scenario Properties • 310, 315, 319
Understanding High Availability Properties • 202, 223, 253
Understanding Master Properties • 93, 165
Understanding Replica Properties • 93, 170
Understanding Scenario Properties • 56, 93, 138, 200, 251, 272, 309, 342
Understanding the CDP Repository • 297
Understanding the CDP Repository Statistics and Reports • 319
Understanding the Content Distribution Solution • 333
Understanding the Control Service Scenario • 243
Understanding the High Availability System and the Switchover and Switchback Procedures • 193
Understanding VSS Snapshot Properties • 291, 293
Universal Shadow Storage Volume VSS Snapshots
 Universal Shadow Storage Volume • 293
Unmounting VSS Snapshots • 295
User-defined scripts
 for Switchover redirection • 225
Using Templates • 60
Using the CDP Repository • 28, 48, 297
Using the Content Distribution Solution • 11, 333
Using the Manager • 239
Using the Master Cluster • 239
Using VSS Snapshots • 289

V

Verification
 Engine on Hosts • 46
 HA scenario results dialog • 195
 scenario • 46
View Events in a Separate Window • 90
View Incoming Events • 91

Viewing
 events • 90
 events in separate window • 90
 incoming events using pop-up window • 91
 Manager screen options • 37
 reports • 94
 VSS Snapshots • 294
Viewing a Report • 59, 69, 94, 281, 313, 323, 346
Viewing and Arranging the Manager Screen • 37
Viewing and Managing Snapshots • 294
Viewing Options • 37
Viewing Snapshots • 42, 294
Viewing Toolbar • 42
VSS Snapshots
 configuring creation • 289
 creation, setting up • 290
 Management window • 294
 Managing • 295
 Max Storage Size per Volume • 293
 Preferred Number of Snapshots to Keep • 293
 properties • 293
 using • 289
 viewing • 294
 window • 294

W

Warnings, before running • 195
Weekly Activity, scheduling synchronization • 139
Wizard
 Scenario creation • 46