# CA XOsoft™ High Availability for Windows

## Oracle Server Operation Guide

r12.5

# 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 9

About This Guide . . . . . 9
Related Documentation . . . . . 9
Server Requirements . . . . . 10
    Base Configuration . . . . . 10
    Log On Account Conditions . . . . . 11
    Servers Operating in a Workgroup . . . . . 11

## Chapter 2: Redirection Methods 13

How Redirection Works . . . . . 13
DNS Redirection . . . . . 13
Move IP Redirection . . . . . 14
    Add IP on the Master Server . . . . . 14
Scripts Redirection . . . . . 17

## Chapter 3: Creating and Using Scenarios 19

Create Oracle HA Scenario . . . . . 20
Scenario Properties . . . . . 26
Run the Scenario from Outside the Wizard . . . . . 28
Stop a Scenario . . . . . 30
View a Report . . . . . 30

## Chapter 4: Switching Over and Switching Back 33

How Switchover and Switchback Work . . . . . 34
Initiate Switchover . . . . . 35
Initiate Switchback . . . . . 37
Switchover Considerations . . . . . 40

## Chapter 5: Recovering Data 41

The Data Recovery Process . . . . . 41
How to Recover Oracle Data . . . . . 41
Setting Bookmarks . . . . . 42
Data Rewind . . . . . 43
Recover Active Server . . . . . 50

**Appendix A: Additional Information and Tips** 53

Spool Directory Settings . . . . . 53
Manually Recover a Failed Server - Move IP Address . . . . . 54

**Index** 55

# Documentation Changes

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

- Oracle Base Configuration (see page 10) -- This new topic has been added to augment the list of conditions provided in Base Configurations.
- DNS Redirection (see page 13) -- This new topic has been added to provide Redirection Method information specific to Oracle HA scenarios.
- Create Oracle HA Scenario (see page 20) -- This topic has been updated to include information specific to creating Oracle scenarios.
- How Switchover and Switchback Work (see page 34) -- This new topic has been added to provide information on how Oracle servers can be switched over and switched back.
- How to Recover Oracle Data (see page 41) -- This topic has been revised with additional information specific to Oracle scenarios and includes best practices.

# Chapter 1: Introduction

CA XOsoft High Availability (HA) is a high availability solution based on asynchronous real-time replication and automated application switchover and switchback to provide cost-effective business continuity for file servers and other application servers on both 32- and 64-bit Windows servers.

CA XOsoft HA offers push-button or fully automatic switchover of mission-critical servers over a LAN or WAN, server status monitoring, and integrated continuous data protection as a guard against data corruption, all in a system that sets the standard for ease of configuration and management.

When disaster strikes, whether in the form of a hurricane, a blackout, or far more likely, a virus attack or software or user error, your ability to respond well to the crisis can make the difference between a bright future and the end of your business. At the very least, you are likely to incur significant costs in lost business and, perhaps more importantly, lost confidence by your customers, investors, and other stakeholders. Disaster recovery planning (DRP) is not just about insurance. It is about maintaining your competitive edge. CA XOsoft HA is designed to give you that edge.

This section contains the following topics:

About This Guide (see page 9)
Related Documentation (see page 9)
Server Requirements (see page 10)

## About This Guide

This document describes how to implement a CA XOsoft HA solution for Oracle. Please review each procedure before you begin. It is essential that you have the appropriate resources and permissions to carry out each task.

## Related Documentation

Use this Guide along with the *CA XOsoft Installation Guide* and the *CA XOsoft User Guide*.

# Server Requirements

To implement CA XOsoft or CA XOsoft HA, refer to the appropriate list of requirements, depending on the server type you selected. These components are licensed separately. If you do not have the license required to access support for a specific server type, please contact Tech Support.

## Base Configuration

**Base Configuration**

- Two servers running Windows Server 2000, 2003, or 2008 with the same level of service packs and hot fixes installed.
- All IP addresses are statically assigned (DHCP-assigned IP addresses on the Master or Replica server are not supported)
- The protected server is not a domain controller or DNS server
- Both servers should reside in the same Active Directory forest and also be members of the same domain or trusted domain.

Both servers should reside in the same Active Directory forest and be members of the same domain or trusted domain.

**Notes:**

- CA XOsoft HA for Oracle does not support Microsoft cluster servers
- CA XOsoft HA for Oracle does not support Oracle RAC

To use 32-bit Oracle versions on 64-bit operating systems, you must first install the Oracle Client 11.x or later on the Oracle machine. This is required to successfully run Oracle HA scenarios.

Download the Oracle Client from the Oracle website:

http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/winx64soft.html
http://www.oracle.com/technology/software/tech/oci/instantclient/htdocs/winx64soft.html

Install the Instant Client Package - Basic in the current installation directory for the CA XOsoft Engine, or in one of the OS default PATH directories.

If you decide to enable Oracle support, you must provide the Oracle home path. If your Oracle server is installed without the 32-bit Oracle client library, then you must also provide the Oracle Instant Client Path.

## Log On Account Conditions

The CA XOsoft Replication and CA XOsoft HA Engine service must satisfy certain account conditions for successful communication with other components. If these requirements are not met, scenarios may not run. If you lack the permissions required, contact your local IS team.

- It is a member of the Domain Admins group. If the Domain Admins group is not a member of the built-in domain local group Administrators you must use an account that is.
- It is a member of the local machine Administrators Group. If the Domain Admins group is not a member, add the account manually.
- For servers in a workgroup, use the Local System account.
- Master and Replica servers must reside in the same Active Directory forest.

## Servers Operating in a Workgroup

For servers in a workgroup, set the CA XOsoft Engine service account to Local System (unless you have added any custom scripts that require higher level permissions). Oracle servers in a workgroup can use Redirect DNS only with DNS servers that allow non-secure updates. You can use Move IP, and custom redirection scripts normally, but the Switch Computer Name method is not supported.

# Chapter 2: Redirection Methods

This section contains the following topics:

How Redirection Works (see page 13)
DNS Redirection (see page 13)
Move IP Redirection (see page 14)
Scripts Redirection (see page 17)

## How Redirection Works

Each of the server types supported by CA XOsoft can be configured to use one or more redirection methods. You should enable redirection methods based on your environment and business needs. The supported redirection methods for Oracle follow:

## DNS Redirection

**Important!** This is the recommended redirection method for Oracle HA scenarios. During scenario creation, this setting is On by default.

DNS Redirection changes the Master server DNS to the Replica server IP address. Upon failure of the Master, the Replica server modifies the appropriate DNS record so that references to the Master server resolve to the Replica's IP address rather than the Master's IP address. This redirection method requires no network reconfiguration and works in LAN and WAN network configurations.

DNS redirection works only with A (host) type records and cannot update CNAME (Alias) records directly. However, if the CNAME record points to the modified A record, it is indirectly redirected.

Using the record that has the Master server's name is the default, however you can configure CA XOsoft HA to redirect any DNS A (host) record via the *Master's name in DNS* setting in the switchover properties tab.

# Move IP Redirection

Move IP redirection involves moving the Master server IP address to the Replica server.

This redirection method is preferred for Virtual Machine scenarios and is usable only in a LAN configuration in which the Master and Replica servers reside in the same network segment. In this configuration, switchover of the Master server causes the Replica to take over one or more of the IP addresses assigned to the Master server.

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

When using Move IP as the redirection method, you must first add IP addresses to the Master host. For more information, refer to the topic, Add IP on the Master Server.

## Add IP on the Master Server

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

**Important!** Perform the following only if you are using the Move IP redirection method.

**Add IP Address to Master Server**

1. Open the Control Panel and choose Network Connections.
2. Right-click Local Area Network and choose Properties.
3. Click Internet Protocol (TCP/IP) and then click the Properties button.
4. Click Advanced.

5. Click Add and enter an additional IP address (XO-IP).

   In the following screenshot, the XO-IP IP address is 192.168.220.23 and the current production server IP address is 192.168.220.111.

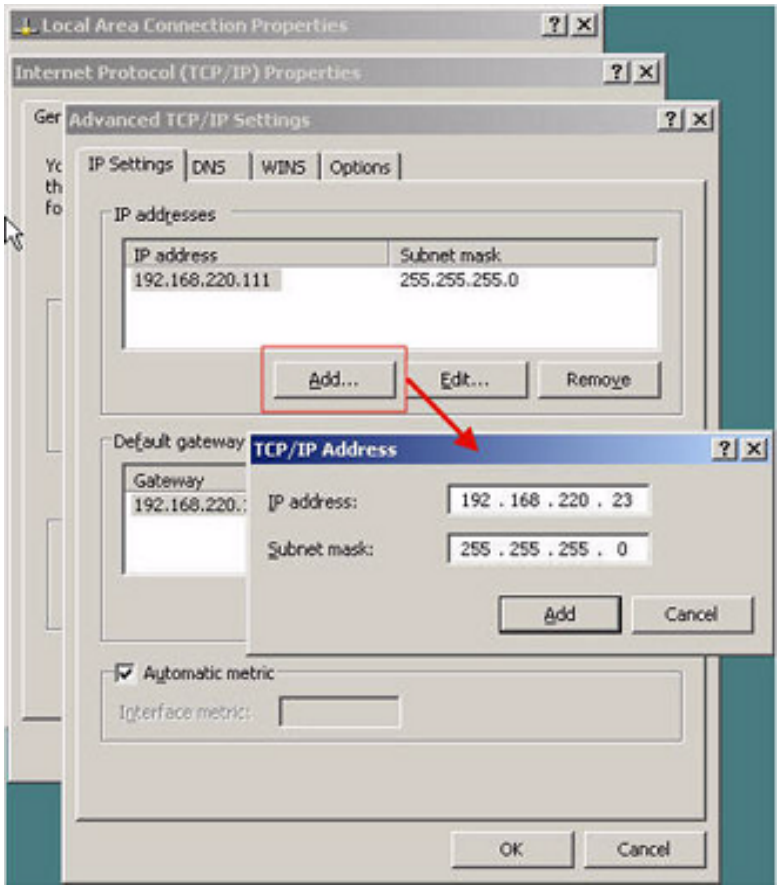


6. Click Add.
7. Click OK.
8. Click OK to exit the LAN settings.

After you add the IP to the Master, 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, from directly in the Wizard
- For existing scenarios, by modifying the master host name

The procedures for both ways follow.

## Add XO-IP to Existing Scenarios

Perform this procedure only if you are using the Move IP redirection method.

**To add the XO-IP to existing scenarios:**

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

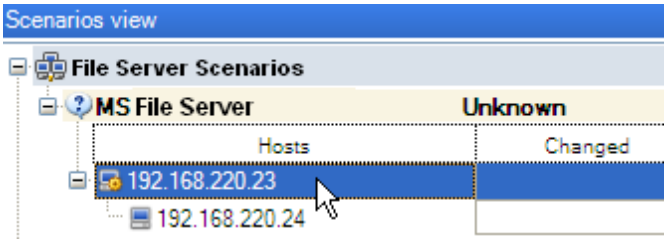


2. Right-click the Master and select **Rename** from the pop-up menu. Then, enter the XO-IP address.
3. On the Framework pane, select the **Switchover** tab and then select the Replica server as the switchover host.
4. Set the **Move IP** option to On. Ensure that the IP address under **Move IP**, **IP/Mask** 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.**

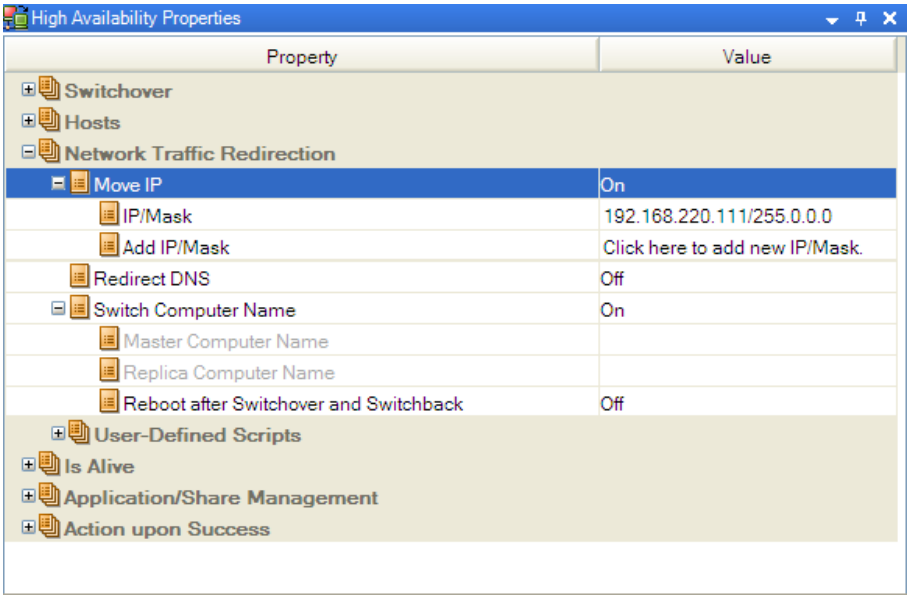

### Add XO-IP to New Scenarios

Perform this procedure only if you are using the Move IP redirection method.

During the initial run of the Scenario Creation Wizard, enter the XO-IP and Replica IP addresses instead of the server names.

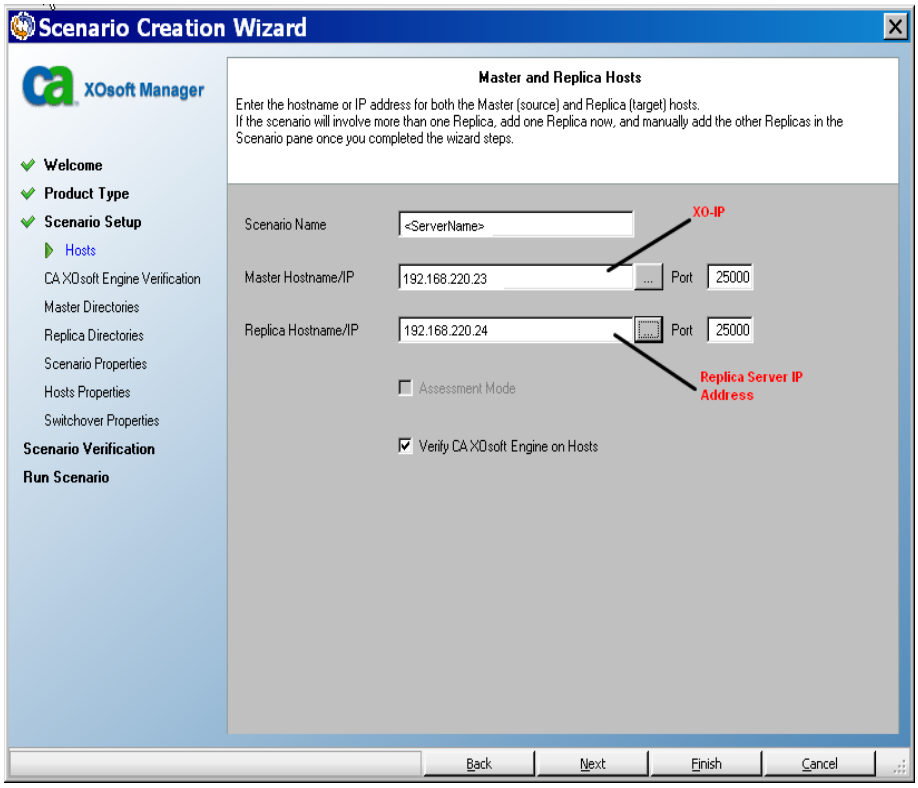


# Scripts Redirection

Custom Scripts: CA XOsoft HA can trigger custom scripts or batch files to perform the user redirection or any additional steps not covered by built-in methods. If the above methods are not appropriate or do not fully meet all requirements, please see the *CA XOsoft User Guide* for details on scripted redirection methods.

# Chapter 3: Creating and Using Scenarios

The following topics describe how to create and configure a replication scenario.

This section contains the following topics:

Create Oracle HA Scenario (see page 20)
Scenario Properties (see page 26)
Run the Scenario from Outside the Wizard (see page 28)
Stop a Scenario (see page 30)
View a Report (see page 30)

# Create Oracle HA Scenario

**To create an Oracle HA scenario**

1. Start CA XOsoft Manager. From the Scenario menu, choose New or click the New Scenario button.
2. From the Welcome dialog, choose Create a New Scenario, choose a scenario group and click Next.
3. From the Select Scenario and Product Type dialog, select Oracle, High Availability Scenario (HA), the desired Tasks on Replica, if any, and then click Next.

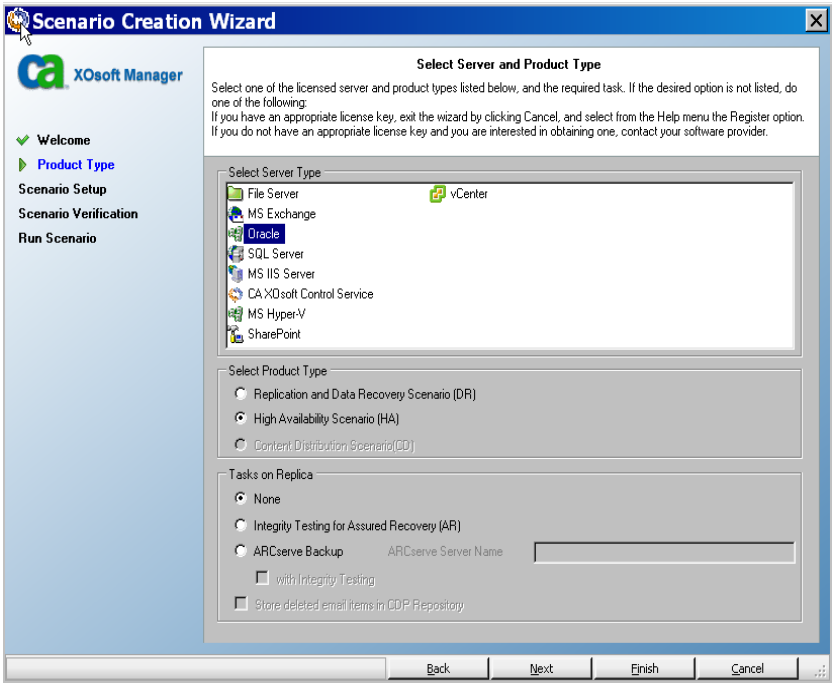

4. From the Master and Replica Hosts dialog, type a Scenario Name, enter the Hostname or IP Address and Port number for both Master and Replica servers, and then click Next.

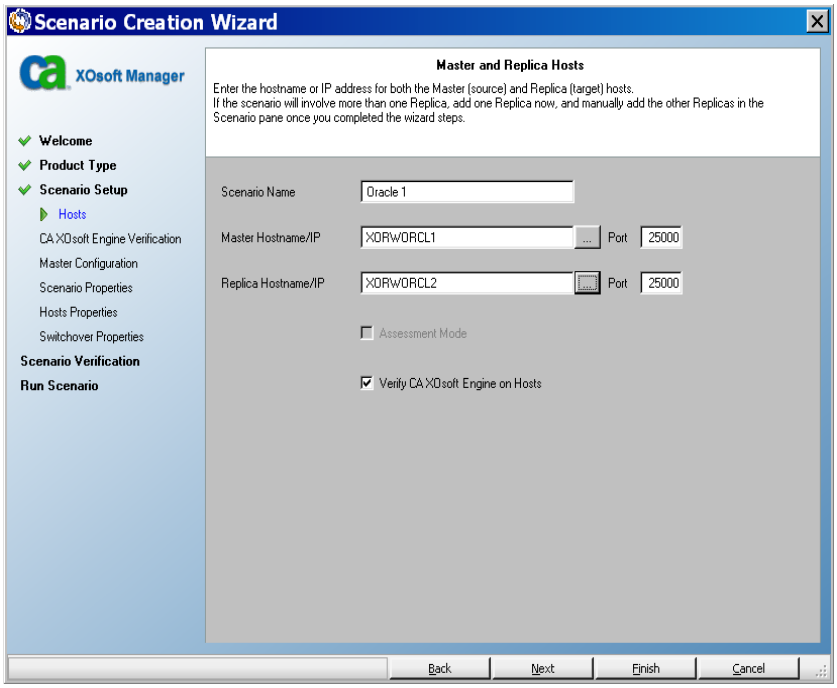


5. Wait for Engine Verification to complete. If prompted for User Credentials, enter the appropriate information. If the Engine service is not installed or not the correct version, click Install to upgrade it on one or both servers, then Verify Again. When complete, click Next.

   The Database Connection dialog opens.

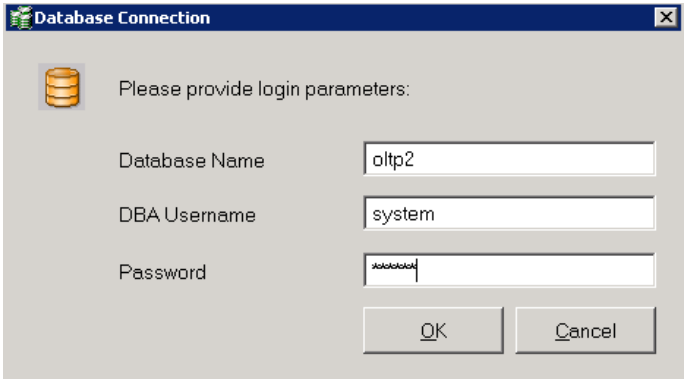

6. Enter the database name, DBA username and password and click OK. Enter credentials for an account with administrative privileges in the database so that CA XOsoft can query the Oracle configuration on the Master server.

   For the database name, use the database "connection string" configured in the Oracle client configuration file tnsnames.ora on the Master server.

   **Note:** If you are using a 32-bit Oracle version on a 64-bit operating system, you must enter the full database path in the Database Name field, in the following format:

   IP/Database_Name

   Example:

   172.16.80.10/DB1

   Enter the DBA Username and Password as usual and click OK.

7. When the Tablespaces for replication dialog appears, review the list of results auto-discovered on the Master Server. Clear and select databases for replication, as desired, enable the "Replicate new user created files in listed root directories" option and click Next.

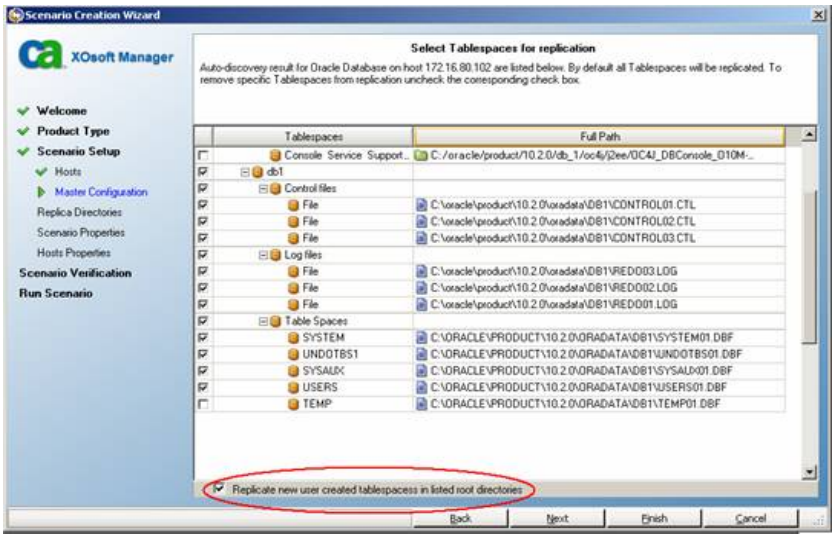


   Selecting this option automatically replicates tablespaces to Oracle databases as you add them. If this option is not selected, CA XOsoft HA replicates only the databases selected at configuration time.

8. From the Root Directories dialog, review the list of folders to be replicated and click Next.
9. From the Scenario Properties dialog, configure the settings that control the entire scenario. For more information, see the CA XOsoft User Guide. For Oracle scenarios, we recommend enabling the Replicate NTFS ACL option, under Replication, Optional Settings. Click Next.

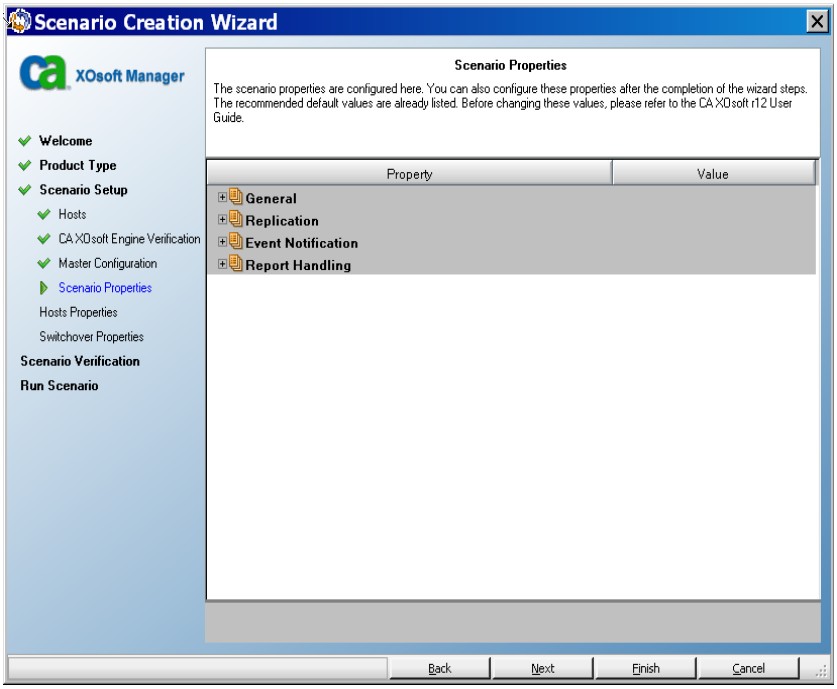

10. From the Master and Replica Properties dialog, configure the desired settings and then click Next.

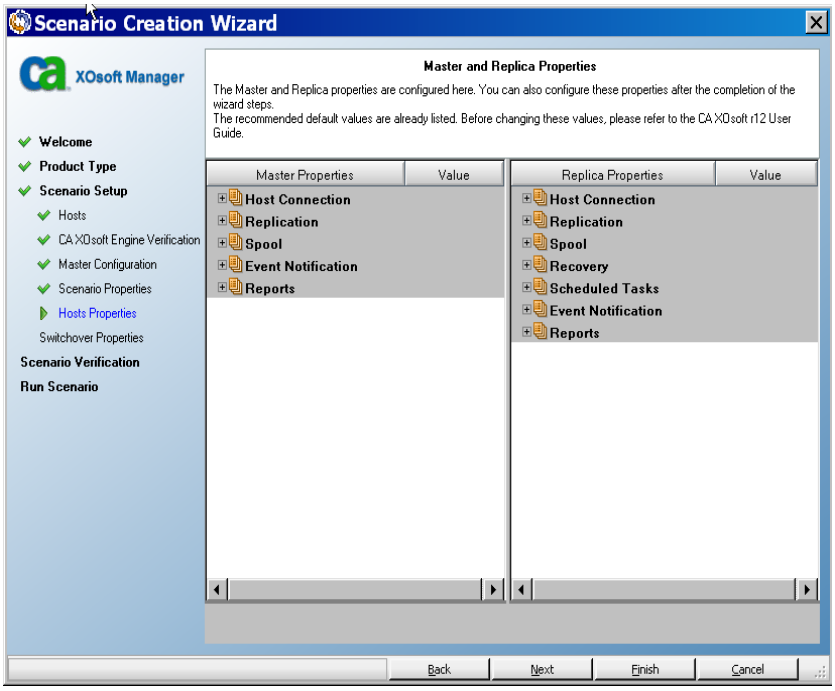

11. Wait for the Switchover Properties dialog to appear. Configure the redirection settings as described in IP Redirection, and then click Next.
12. From the Switchover Initiation dialog, configure the desired Initiation options and click Next. We recommend setting only one of these to Automatic.

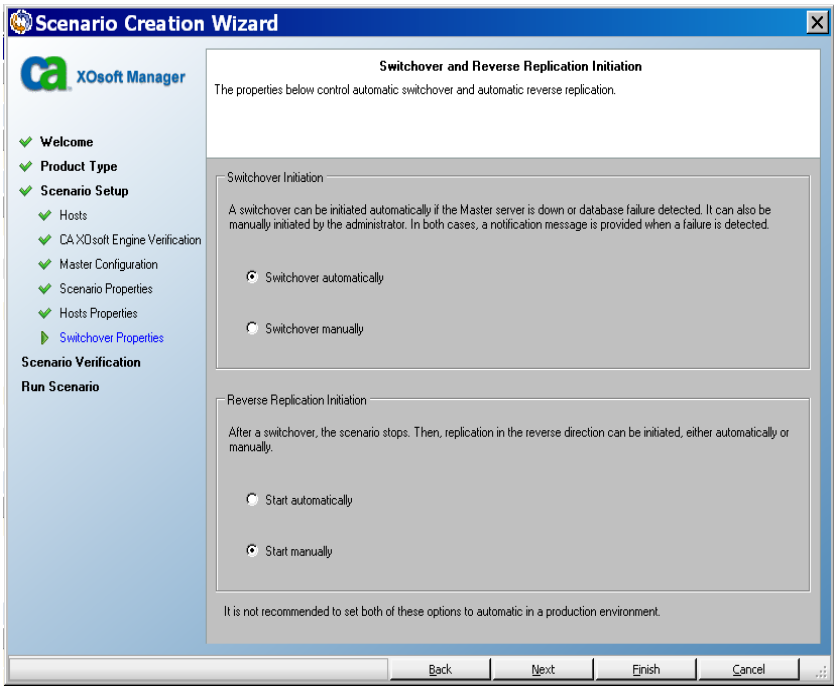


**Note:** If you selected Assured Recovery in the Tasks on Replica options, that dialog appears now. For more information, see the CA XOsoft User Guide.

13. From the Scenario Verification dialog, ensure no errors or warnings appear. While warnings do not prevent a scenario from running as errors do, both should be corrected before you proceed. Once successfully verified, click Next.
14. Click Run Now to start data synchronization or click Finish to run the scenario later.

# Scenario Properties

If you wish to change a scenario configured through the Wizard or configure additional settings, you can use the Properties pane to modify the scenario.

The Properties pane and its tabs are context-sensitive and change whenever you select a different node from a scenario folder. You must stop a scenario before configuring its properties. Certain values cannot be modified once set; they are noted.

For full details on configuring scenario properties and their descriptions, see the CA XOsoft User Guide.

Properties are organized into tabs on the CA XOsoft Manager Framework pane. The tabs displayed are based upon server type, CA XOsoft solution, and scenario status.

**Settings on the Root Directories tab**

Select a Master Server from the Scenario Pane. Double-click its Directories folder to add or remove Master Root Directories. Select or clear checkboxes next to folders, as desired, to include or exclude them. You may also edit directory names.

Select a Replica Server from the Scenario Pane. For each Master Root directory, you must specify a Replica Root directory. Double-click the Directories folder for the Replica server. Select or clear checkboxes next to folders, as desired, to hold the corresponding Master directory.

Directories with Oracle database files are automatically included by Oracle auto-discovery in the scenario at the scenario configuration step (see the topic, Create Oracle HA Scenarios). There is no need to add these directories manually. If you want to include additional directories in the scenario, you can add them using the Root Directories tab. We strongly recommend not including directories with Oracle binaries, DLLs and other Oracle instance directories that are not part of the database.

### Settings on the Properties Tab

#### Scenario Properties

These settings establish default behavior for the entire scenario.

- General properties -- cannot be changed once created
- Replication properties -- choose the replication mode (Online or Scheduled), synchronization values (File or Block, Ignore Files of Same Size/Type) and optional settings (Replicate NTFS Compress Attribute, Replicate NTFS ACL, Synchronize Windows Shares, Prevent Automatic Re-sync upon Error)
- Event notification properties -- specify a script to run or choose email notification
- Report Handling -- specify report settings, email distribution or script execution

#### Master and Replica Properties

These settings establish server properties on both Master and Replica. Some settings vary by server type.

- Host connection properties -- Enter the IP address, Port number and Fully Qualified Name of the Master and Replica
- Replication properties -- These properties differ for Master and Replica. See the User Guide for more information.
- Spool properties -- Set the size, minimum disk free size and directory path. See Spool Directory Settings (see page 53) for more information.
- Event notification properties -- specify a script to run or choose email notification
- Report properties -- choose synchronization or replication reports, specify distribution or script execution
- (Replica) Recovery properties -- set delay or data rewind properties

**Settings on the HA Properties Tab**

These settings control how switchover and switchback are performed

- Switchover properties -- choose automatic or manual switchover, and reverse replication settings
- Hosts properties -- specify the Master and Replica Fully Qualified Name
- Network Traffic Redirection properties -- choose Move IP, Redirect DNS, or User-defined scripts
- Is Alive properties -- set the heartbeat frequency and check method
- DB Management properties -- instructs CA XOsoft to manage shares or services on a database server
- Action upon Success properties -- defines custom scripts and arguments for use

If you added a new data or redo log file to the Oracle database, you need to update the scenario. Stop the scenario and rerun Oracle instance auto-discovery by double-clicking the instance name from the Root Directories tab.

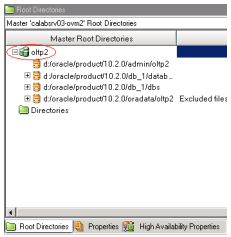


# Run the Scenario from Outside the Wizard

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, online replication starts automatically, continuously updating the Replica with all of the changes that occur on the Master.

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

**To run the scenario from outside the wizard**

1. From the Scenario pane, select the scenario you want to run.
2. Click **Run** on the Standard toolbar.

   Before initiating synchronization and replication, CA XOsoft verifies your scenario configuration. When verification completes successfully, CA XOsoft Manager displays the message: *Are you sure you want to run scenario "scenario_name?"* If problems are discovered, the top pane displays any warning and error messages resulting from verification.

   **Note:** Scenario Verification checks many different parameters between the Master and Replica servers to ensure a successful switchover. If any errors or warnings are reported you should not continue until they are resolved.
3. Correct errors before you continue. Errors are reported on the Event pane.

   **Note:** Replication of mount points succeeds 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.

   When no error is reported, the **Run** dialog appears and contains synchronization options.

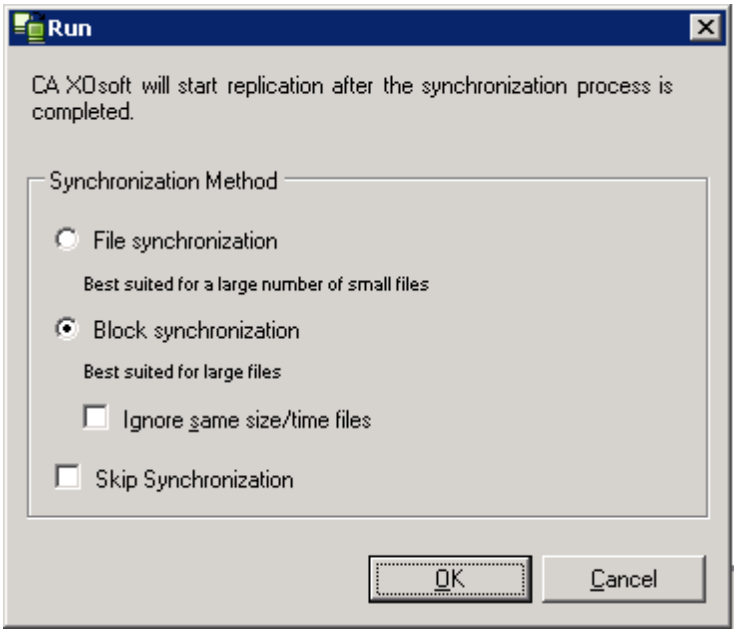


   **Note:** Do not use Skip Synchronization for any scenarios replicating a database.

4. If you have a large number of small files, select File Synchronization. If you have large files, select Block Synchronization. Select the Ignore same size/time files to skip the comparison of files with the same path, name, size and modification time, which are generally identical, to reduce synchronization time. You should enable the Skip Synchronization option only when you are certain the files on both Master and Replica are identical. The default selections are File Synchronization and Ignore files of same size/time option enabled.
5. Click the **OK** button. 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 window when the synchronization is complete: *All modifications during synchronization are replicated.*

At this point, the scenario is operational and active. By default, a Synchronization Report is generated when synchronization finishes. To view the report, refer to the topic, View a Report. You can also generate regular Replication Reports to monitor the replication process on each participating server. For more information, see the *CA XOsoft User Guide*.

# Stop a Scenario

**To stop a scenario**

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.

# View a Report

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.

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

### To view a report

**Note:** Though an Exchange report is shown for illustrative purposes, the steps and screens are similar regardless of scenario type.

1. To view a report, first you need to open the Report Center. There are two ways to open it:

   - On the Overview Page, click the **Report Center** link on the **Quick Start** pane on the left:

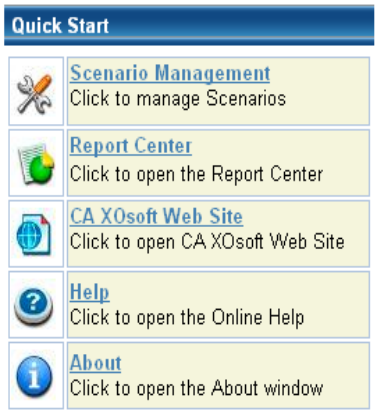


2. From the **Tools** menu, select the **Reports** option and then **Show Scenario Reports**.

   The Report Center opens in a new window:

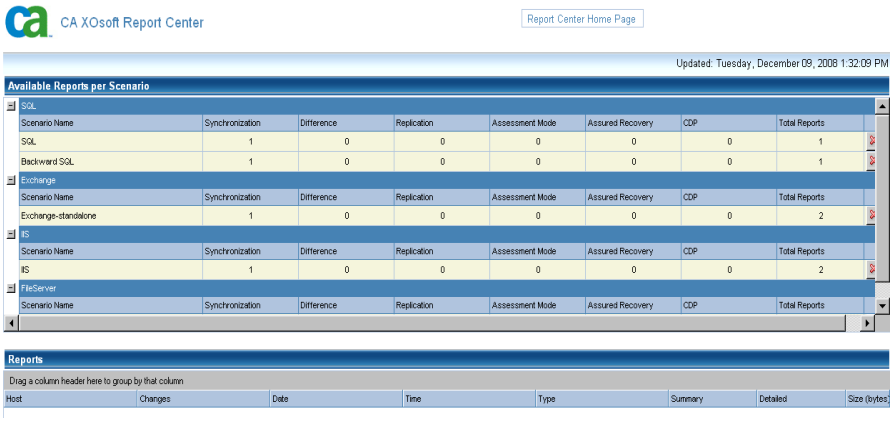


   The Report Center consists of two tables:

   - 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.

3. 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:

Drag a column header here to group by that column

| Host | Changes | Date | Time | Type | Summary | Detailed | Size (bytes) |
|---|---|---|---|---|---|---|---|
| XORWEXCH2K7-1 | UnKnown | Today | 03:29:37 | Assured Recovery | | | 811 |
| XORWEXCH2K7-1 | Changes found | 12/07/08 | 22:29:48 | Synchronization | | | 28415 |

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

The report you selected appears in a new window:

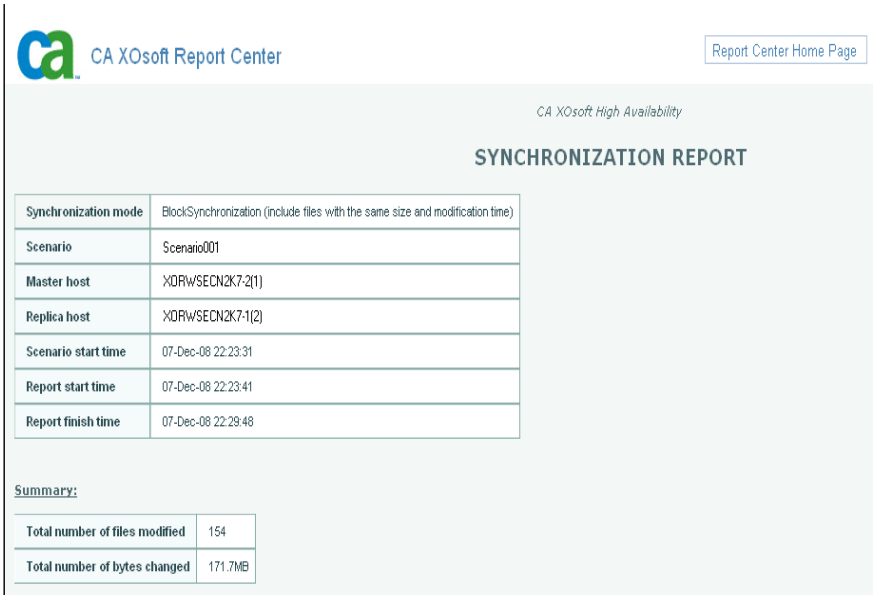

CA XOsoft Report Center

Report Center Home Page

*CA XOsoft High Availability*

SYNCHRONIZATION REPORT

| | |
|---|---|
| Synchronization mode | BlockSynchronization (include files with the same size and modification time) |
| Scenario | Scenario001 |
| Master host | XORWSECN2K7-2(1) |
| Replica host | XORWSECN2K7-1(2) |
| Scenario start time | 07-Dec-08 22:23:31 |
| Report start time | 07-Dec-08 22:23:41 |
| Report finish time | 07-Dec-08 22:29:48 |

Summary:

| | |
|---|---|
| Total number of files modified | 154 |
| Total number of bytes changed | 171.7MB |

# Chapter 4: Switching Over and Switching Back

*Switchover* and *Switchback* is the process in which active and passive roles are exchanged between the Master and Replica servers, so that if the Master is currently active, it changes to passive after Switchover passes the active role to the Replica. If the Replica is active, it changes to passive after Switchover passes the active role to the Master. Switchover can be triggered at the push of a button or automatically by CA XOsoft HA when it detects that the Master is unavailable, if you enabled the Perform Switchover Automatically option from the Switchover and Reverse Initiation dialog. When this option is Off, the system notifies you that the Master server is down so you can manually initiate switchover from the CA XOsoft Manager.

This section contains the following topics:

How Switchover and Switchback Work (see page 34)
Initiate Switchover (see page 35)
Initiate Switchback (see page 37)
Switchover Considerations (see page 40)

# How Switchover and Switchback Work

After the HA scenario starts running and the synchronization process is completed, the Replica checks the Master on a regular basis, by default every 30 seconds, to see if it is alive. There are three types of monitoring checks:

- **Ping** -- a request sent to the Master to verify that the Master is up and responding
- **Database check** -- a request that verifies the Oracle Server is running
- **User-defined check** -- a custom request you can tailor 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. Then, depending on the HA scenario configuration, CA XOsoft HA sends you an alert or automatically initiates a switchover.

When you created an HA scenario, you defined how you want the switchover to be initiated.

- If you selected the Initiate Switchover manually option from the Switchover and Reverse Replication Initiation page, perform a manual switchover. For more information, refer to the topic, Initiate Switchover (see page 35).
- 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 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.

Triggered (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 manually initiating the switchover by clicking the Perform Switchover button. Server ping response, application service status, and database connectivity are monitored. The timeout parameters are configurable and are more extensively covered in the *CA XOsoft HA User Guide*

When you created an HA scenario, you defined how you want the reverse scenario to be initiated.

- If you selected the Initiate Reverse Replication automatically option from the Switchover and Reverse Replication Initiation page, replication in the reverse direction (from Replica to Master) automatically begins after a switchover, if the original Master server is online.

- If you selected the Initiate Reverse Replication manually option, you need to perform switchback manually. If you performed switchover with the Master server offline, or if you stop a backward scenario and restart it again, in both cases, data from the Replica to the Master must be resynchronized before you can switch back to the Master.

When the Reverse Replication feature is off, to start reverse replication after a switchover has occurred, click the Run button. The benefit to this feature is, if both the master and replica servers were online and connected during switchover, resynchronization in the reverse direction is not required. Resynchronization involves comparing the data on the master and replica servers to determine which changes to transfer before real-time replication starts; this can take some time. If automatic reverse replication is turned on, and both servers were online during switchover, replication is reversed without the need for resynchronization. This is the one situation in which resynchronization is not required.

# Initiate Switchover

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

**Note:** Though the following steps show Exchange scenario screens as examples, the procedure is similar for all server types.

**To initiate manual switchover**

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

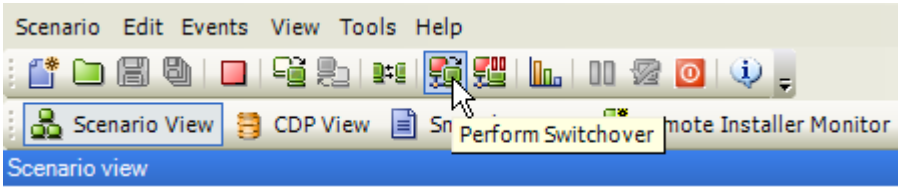


A confirmation message appears:

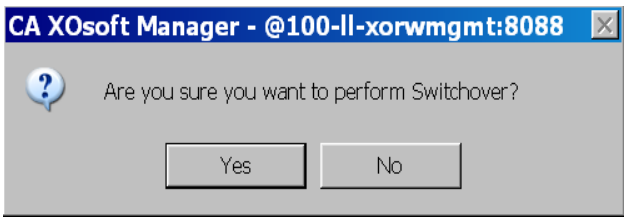

3. Click **OK** on the **Perform Switchover** confirmation message. This procedure initiates a switchover from the Master server to the Replica server:

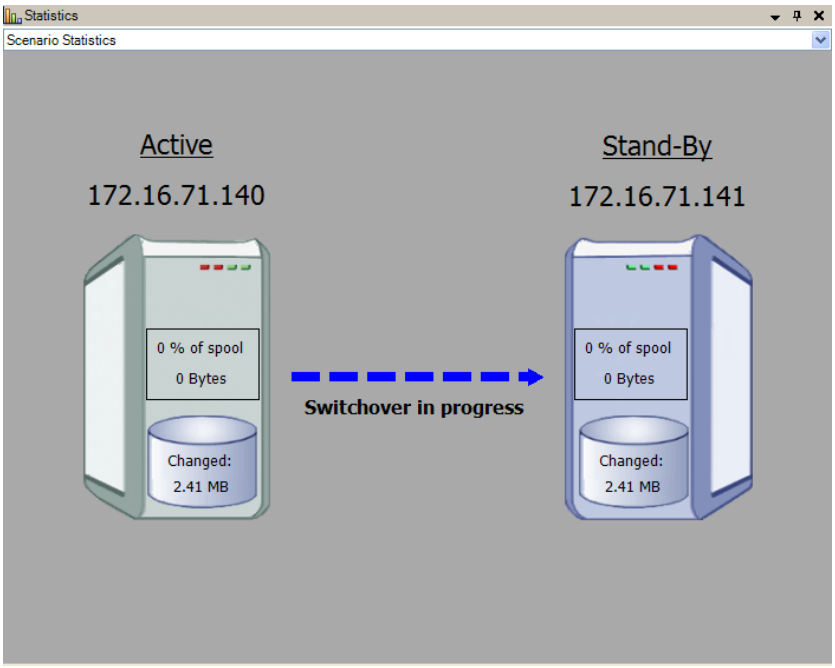


Detailed information about the switchover processes is located in the Events pane during switchover.

4. After the switchover is completed the scenario stops:

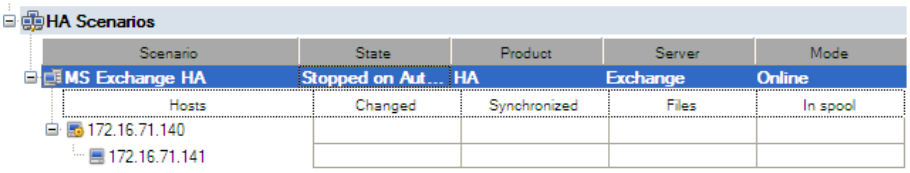


**Note:** The only case in which the scenario may continue to run after switchover is when **automatic reverse replication** is defined as **Start automatically**.

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

Now, the Master becomes the Stand-by server and the Replica becomes active server.

# Initiate Switchback

After a switchover is initiated, whether manually or automatically, at some point, you will want to reverse the server roles and make the original Master the active server again the Replica the standby server. Before you switch back the roles between servers, decide if you want to the data on the original Replica server to overwrite the data on the original Master. If yes, you must first perform a reverse scenario, called a backward scenario

**To initiate switchback**

1. Ensure that both Master and Replica servers are available on the network and that the CA XOsoft Engine is running.
2. Open the Manager and select the desired scenario from the Scenario pane.
3. Perform one of the following:
   - If the scenario is already running, skip directly to Step 4
   - If the scenario is not running, perform these steps and then go to Step 4:
     a. Click Run on the toolbar to start the scenario.

     CA XOsoft HA detects that a switchover has occurred and verifies its state and configuration. After verification completes, the Verification Results dialog appears, listing existing errors and warnings if detected, and prompting you to approve the running of the backward scenario. If desired, click the Advanced button to open an additional pane with detailed information about the hosts that participate in the scenario.

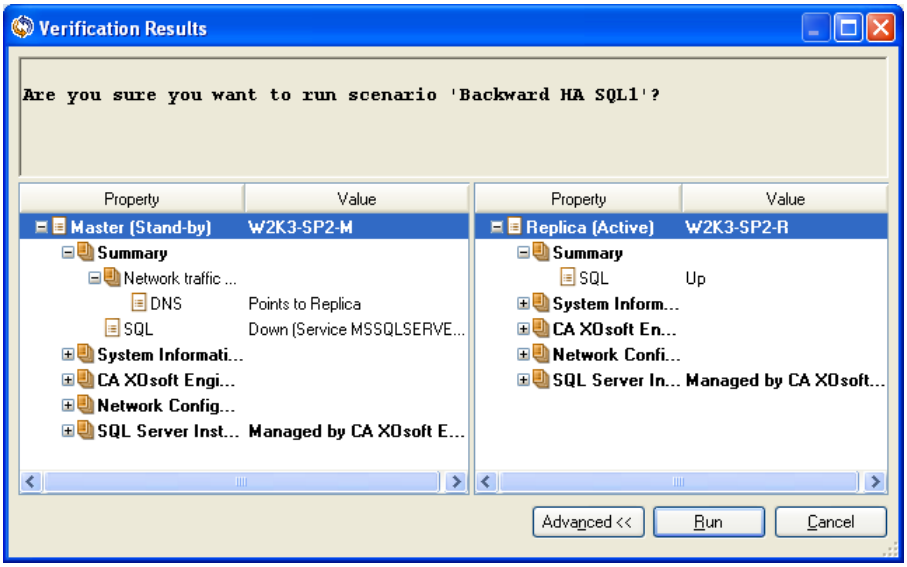

b. Select a synchronization method from the Run dialog and click OK to start resynchronization.

**Note:** For Oracle servers, choose Block Synchronization.

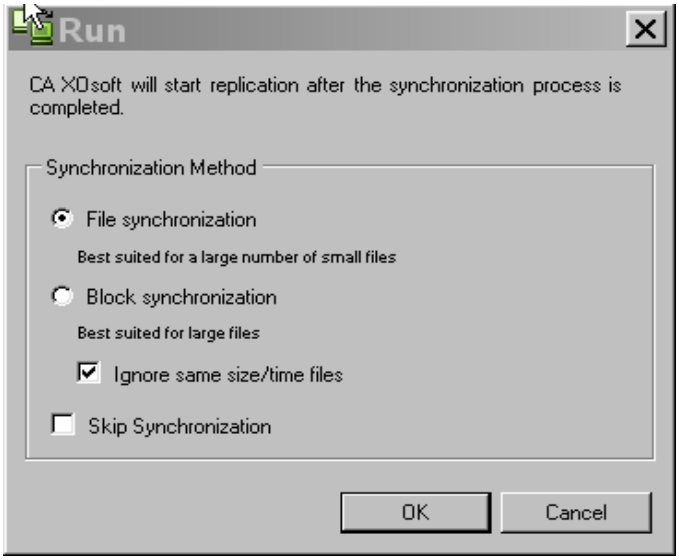


After resynchronization completes, you receive a message in the Event pane: All modifications during synchronization period are replicated. Now, replication from the active server to the standby server begins:

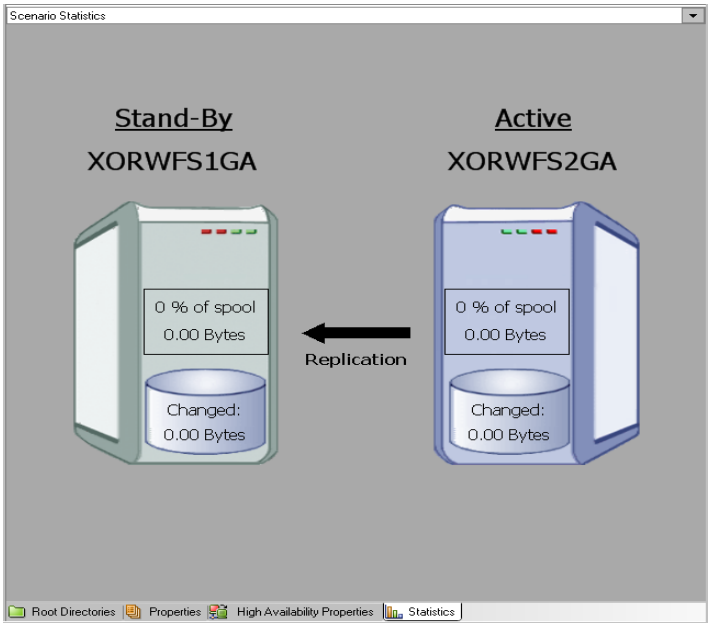

**Note:** You are now ready to reverse the roles between the Master and Replica servers.

4. Click Perform Switchover on the toolbar while the scenario is running to reverse the server roles. A confirmation message appears.
5. Click Yes to clear the message and start the switchback process.

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

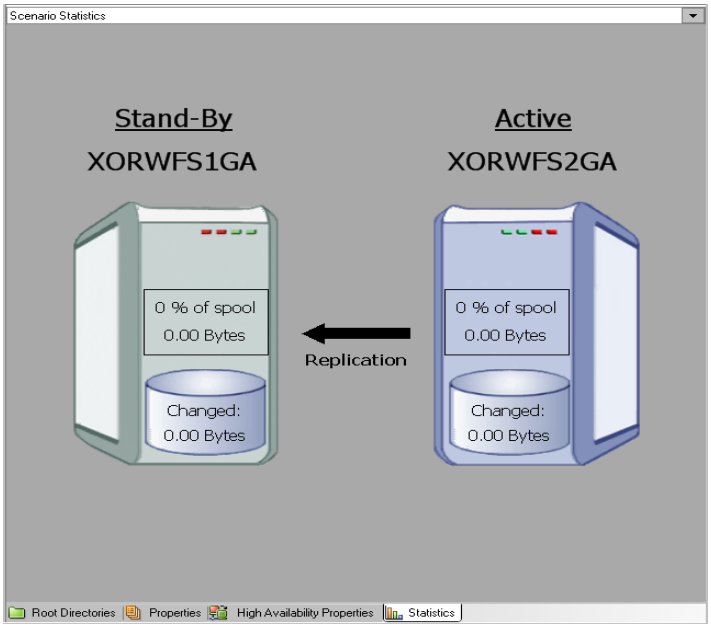


**Note:** The scenario will continue to run after the switchback when the Reverse Replication Initiation option is defined as Start Automatically.

You may now run the scenario again in its original (forward) state.

# Switchover Considerations

It is not recommended to set both the Switchover and Reverse Replication Initiation options to automatic in a production environment. While these options are individually beneficial it is best practice to set only one or the other to automatic. The reason for this recommendation is that automatic switchover allows CA XOsoft HA to trigger a switchover, after a failure is detected, without administrative involvement. If automatic reverse replication is also on, CA XOsoft HA may start to overwrite data on the failed production server before an administrator is able to assess the failure situation. Overwriting data on a failed server before an administrator can assess the situation may have undesirable consequences. Due to this possibility, setting both options to automatic is not recommended. Please choose either one or the other depending on what best fits your requirements. When only one option is used at a time these settings can be very beneficial and may be used safely.

# Chapter 5: Recovering Data

This section contains the following topics:

The Data Recovery Process (see page 41)
How to Recover Oracle Data (see page 41)
Setting Bookmarks (see page 42)
Data Rewind (see page 43)
Recover Active Server (see page 50)

## 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.

## How to Recover Oracle Data

Data recovery is covered in full detail in the CA XOsoft User Guide.

Because the **Recover Lost Data from Replica** procedure requires that you stop Oracle instances on both Master and Replica servers, it is not recommended and therefore, not included in this Operation Guide.

Instead, we recommend that you switch to the Replica, and run a backward scenario to synchronize data from the Replica back to the Master. In this manner, stopping the Oracle instance on the active Replica server is not required.

If you find that data on the Master is corrupted, you can use the Data Rewind process to restore it, if desired. That procedure is provided in this Guide.

# 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.

   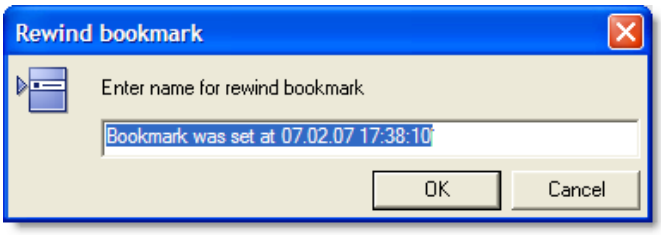
   

   The text that appears in the **Rewind bookmark** dialog will appear in the **Rewind Points Selection** dialog as the bookmark's name. The default name includes date and time.
3. Accept the default name, or enter a new name for the bookmark. It is recommended to give a meaningful name that will later 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.**

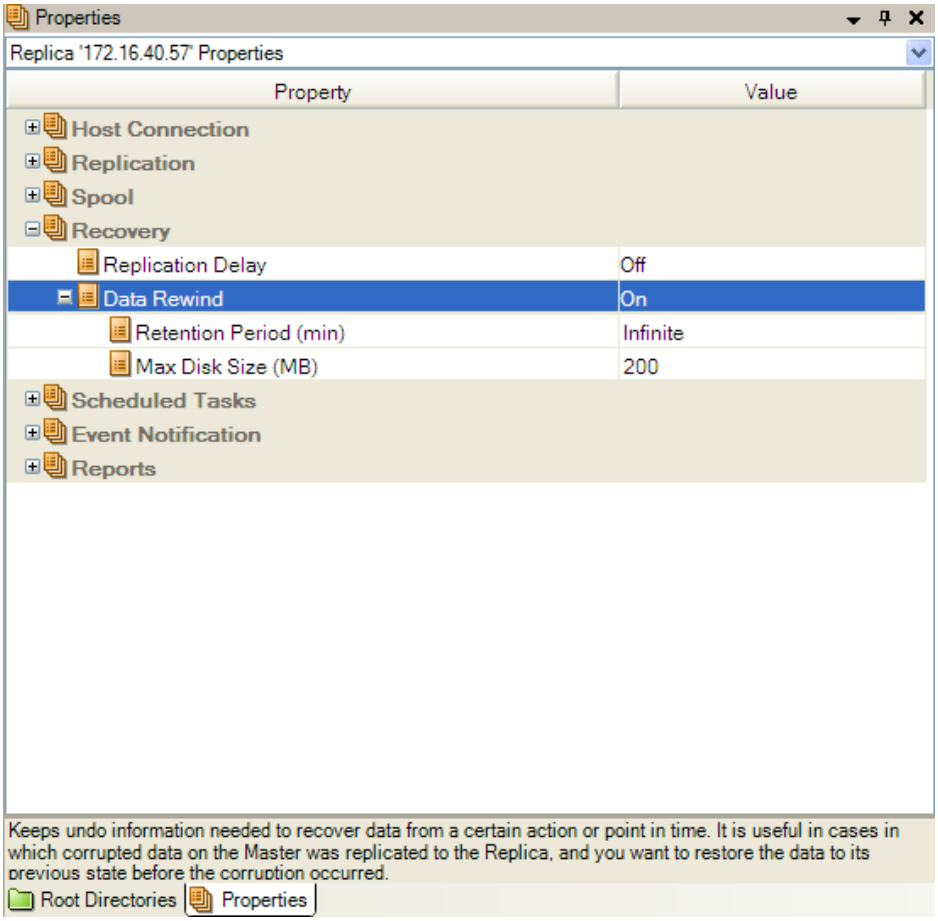

If this option is set to Off, the system will not register data rewind points. For more information about Data Rewind parameters (Retention Period, Max Disk Size), see the *CA XOsoft User Guide*.

**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.

**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.

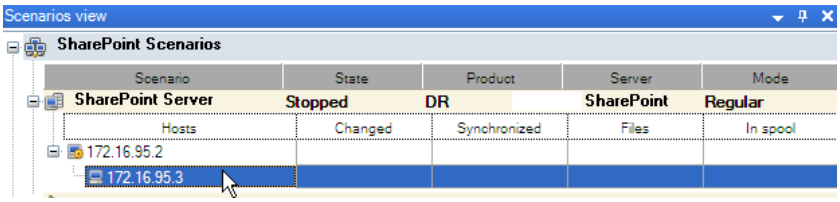

4. From the **Tools** menu, select **Restore Data**, or click the **Restore Data** button. If you are prompted for user credentials, enter the appropriate information and click OK.

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

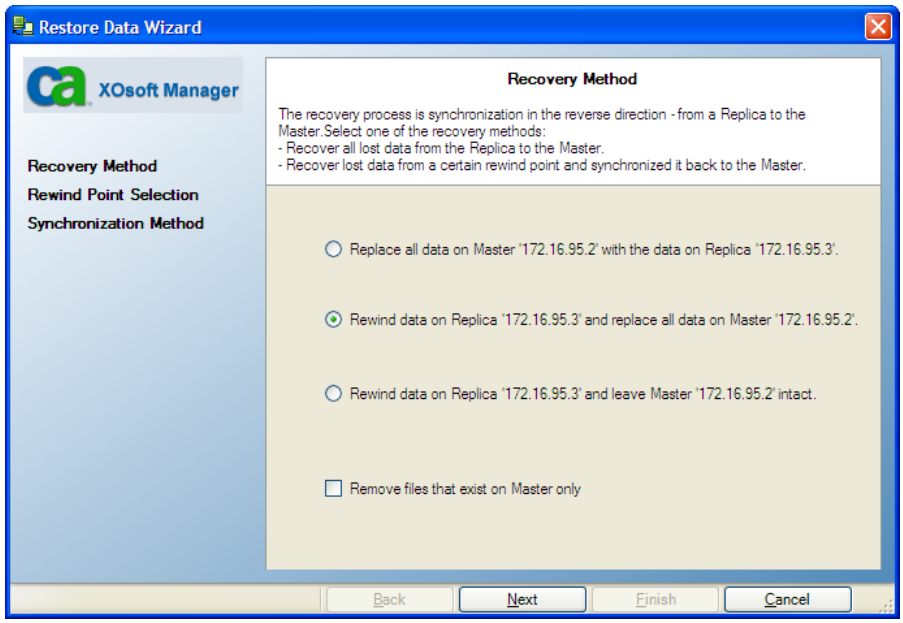

5. Select one of the Rewind data options, depending on whether you want the rewind data synchronized back to the Master (option 2) or left on the Replica only (option 3).

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

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

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

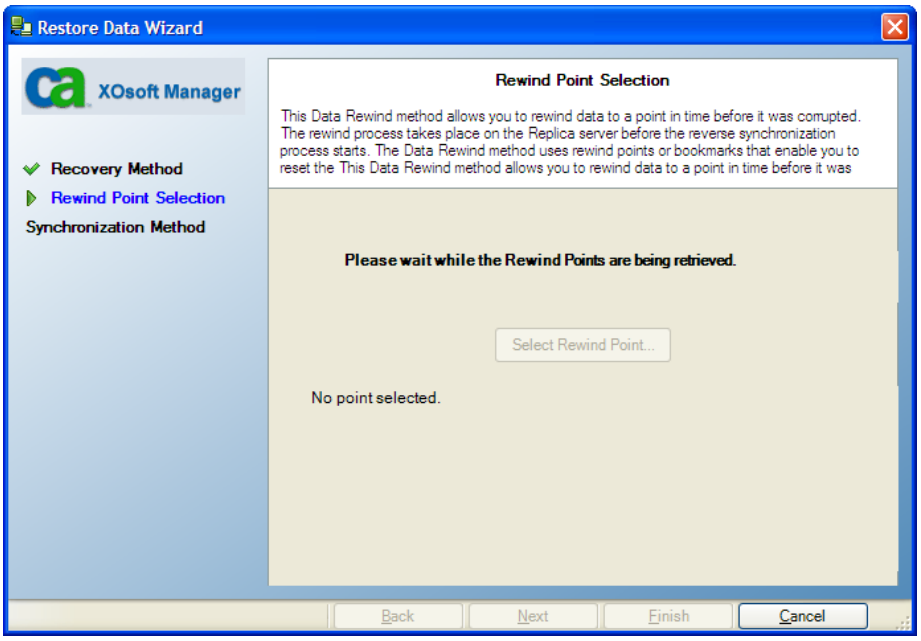


7. Wait until the **Select Rewind Point** button is enabled, and click it to view the existing rewind points.

The **Select Rewind Point** dialog opens.

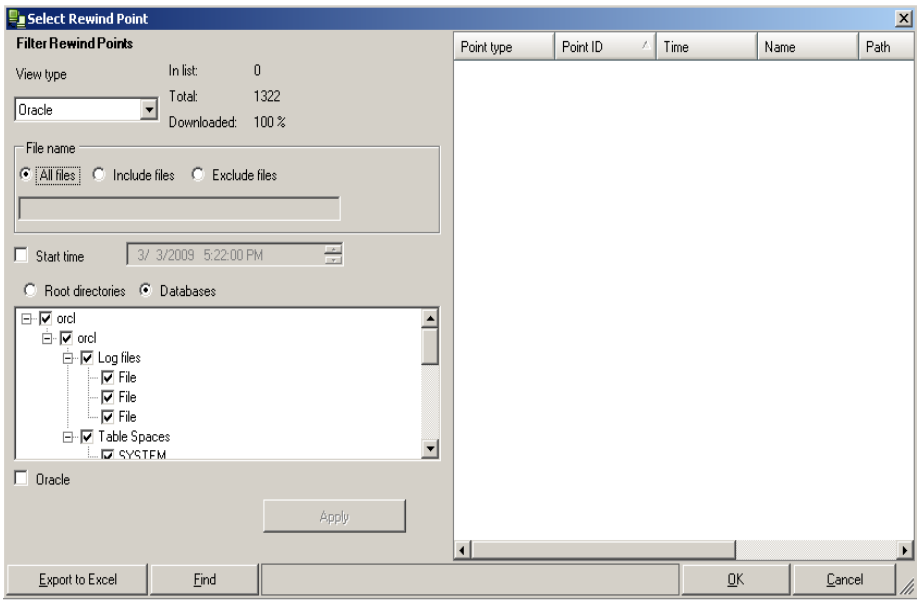


The **Select Rewind Point** dialog displays a list of all rewind points. These include modifications of folders and files that were automatically registered by the system and user-defined bookmarks.

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

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.

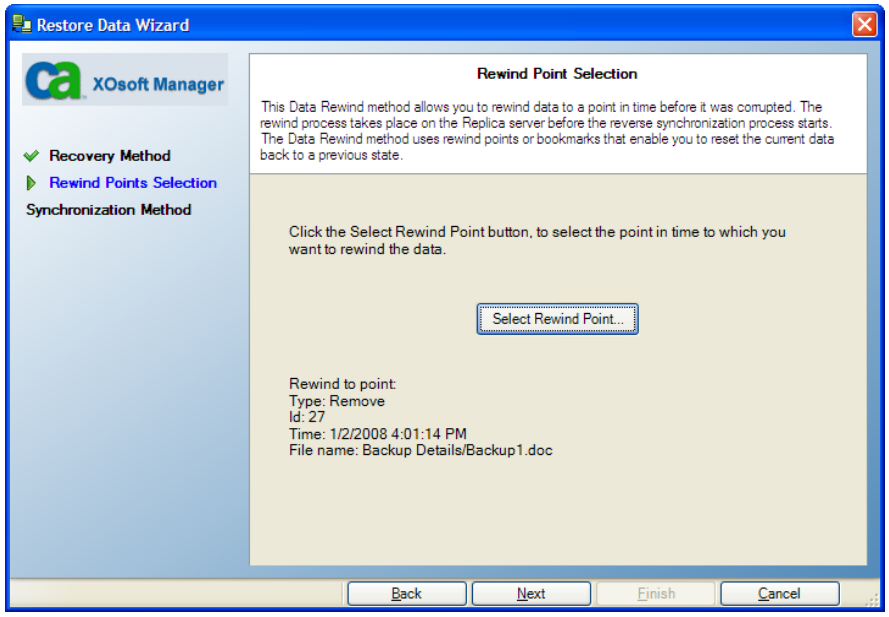


9. Click **Next**. The **Synchronization Method** page is displayed.

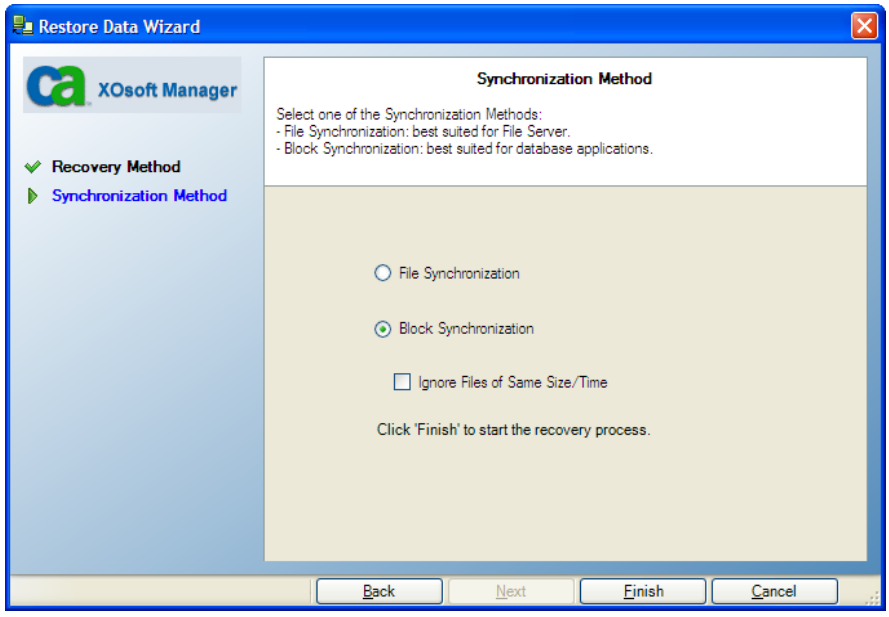

10. Select the **Block Synchronization** method and click **Finish**.

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

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

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

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

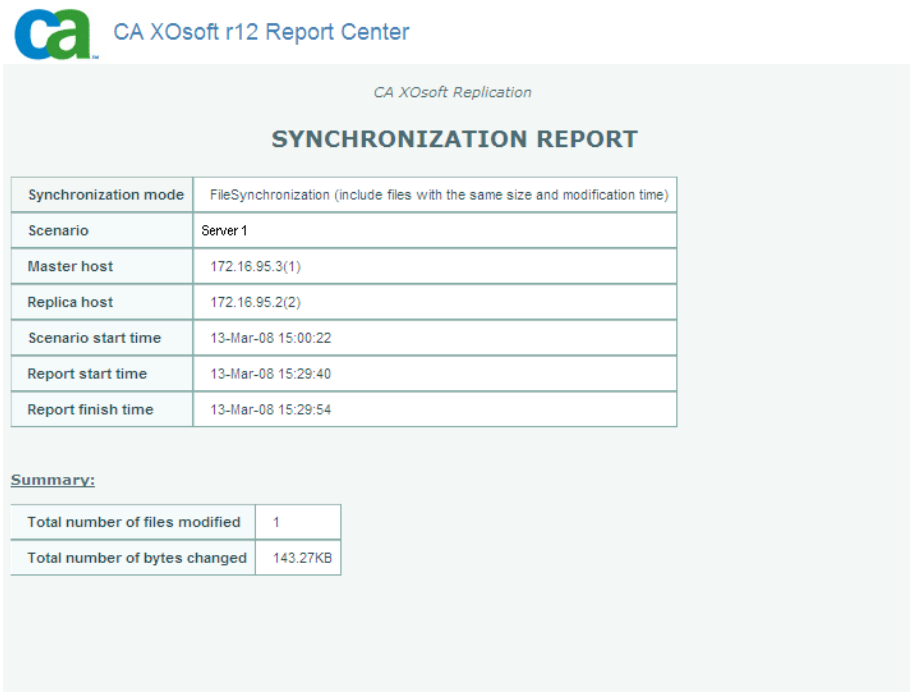

CA XOsoft r12 Report Center

*CA XOsoft Replication*

**SYNCHRONIZATION REPORT**

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

**Summary:**

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

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

# Recover Active Server

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

**Important!** While this option is the right choice in many situations, use it with caution. If used improperly data loss can occur. Normally, CA XOsoft HA 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 that then overwrites what may be a more current data set. When using Recover Active Server, CA XOsoft HA is forcing 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.

If Recover Active Server does not solve the problem, you can manually recover a server. For more information, refer to the topic, Manually Recover a Failed Server. (see page 54)

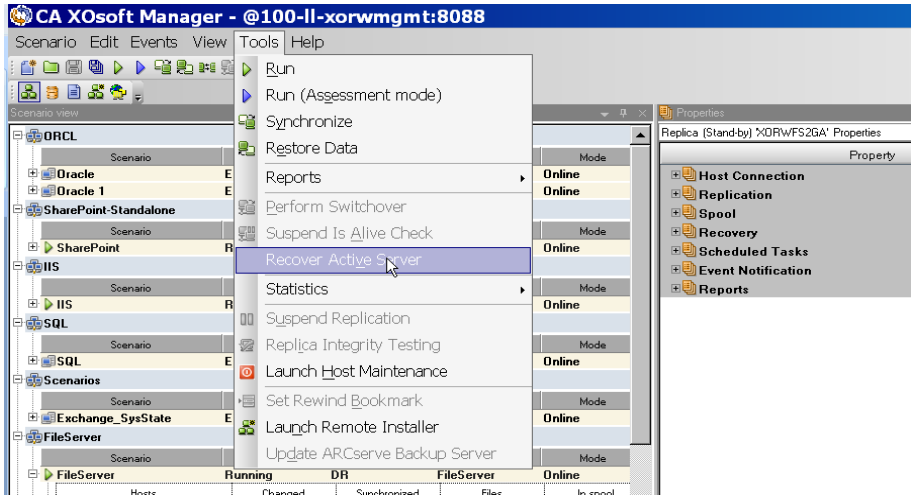

Select either *Make Master Active* or *Make Replica Active* depending onto which server you want to force the active role.

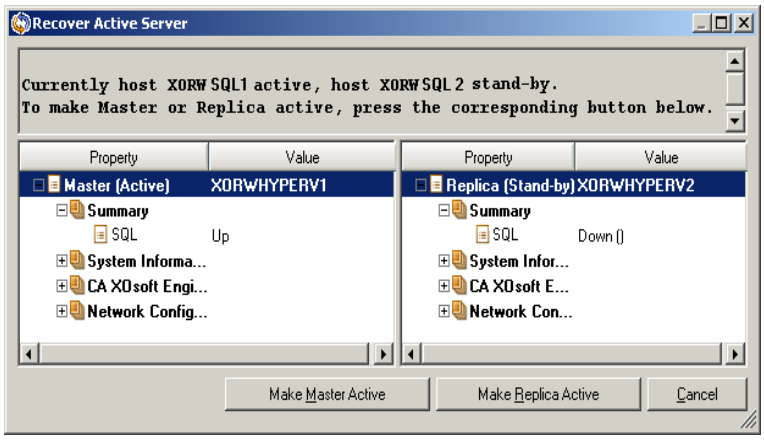


**Important!** If a legitimate switchover in a disaster situation occurs and users are redirected to the Replica server for any period of time, it is important to replicate all changes on the Replica back to the Master before making the Master server active. Using *Recover Active Server* in such a situation results in loss of data.

# Appendix A: Additional Information and Tips

This section contains the following topics:

Spool Directory Settings (see page 53)
Manually Recover a Failed Server - Move IP Address (see page 54)

## Spool Directory Settings

The CA XOsoft spool is a folder on disk where data to be replicated is backed up (spooled) if bandwidth is not sufficient to transfer the amount of changes in real-time. Data can spool due to temporary network disconnections, network congestion, or simply because the network bandwidth is not sufficient to transfer the amount of data changing over on the server. In addition to storing changes waiting on available bandwidth, spool space is also used as part of the normal synchronization process. Thus, some spool build up during synchronization is normal.

Place the CA XOsoft spool folder on a drive with relatively low use such as a dedicated volume or boot/system volume. Do not place the spool folder on a volume containing frequently accessed system (OS), user, or application data. Examples include volumes containing databases, shared files, or the system pagefile. By default, the spool folder is located in the tmp folder under the CA XOsoft installation directory. The spool parameters, located in the properties tab (on both master and replica) or set with the New Scenario Wizard, determines 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. For example, if you are replicating 50 GB of data on a server you should ensure that at least 5 GB of space is available for spool. Please note that this space is not pre-allocated.

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

**Note:** The CA XOsoft Spool Directory is not a pre-allocated space folder and will be used only if needed.

# 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 conficts.
2. From the TCP/IP properties dialog, remove the additional IP address.
3. Reboot the server and reconnect to the network.
4. If it is not already running, start the scenario from the CA XOsoft Manager. If automatic reverse replication 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. Perform a manual switchover to return the active role to the Master server. It is recommended that you do so outside of normal business hours.

# Index


## A

About This Guide • 9
Add IP on the Master Server • 14
Add XO-IP to Existing Scenarios • 16
Add XO-IP to New Scenarios • 17
Additional Information and Tips • 53

## B

Base Configuration • 10

## C

Contact CA • iii
contacting technical support • iii
Create Oracle HA Scenario • 7, 20
Creating and Using Scenarios • 19
customer support, contacting • iii

## D

Data Rewind • 43
DNS Redirection • 7, 13
Documentation Changes • 7

## H

How Redirection Works • 13
How Switchover and Switchback Work • 7, 34
How to Recover Oracle Data • 7, 41

## I

Initiate Switchback • 37
Initiate Switchover • 34, 35
Introduction • 9

## L

Log On Account Conditions • 11

## M

Manually Recover a Failed Server - Move IP Address • 50, 54
Move IP Redirection • 14

## R

Recover Active Server • 50
Recovering Data • 41
Redirection Methods • 13
Related Documentation • 9
Run the Scenario from Outside the Wizard • 28

## S

Scenario Properties • 26
Scripts Redirection • 17
Server Requirements • 10
Servers Operating in a Workgroup • 11
Setting Bookmarks • 42
Spool Directory Settings • 27, 53
Stop a Scenario • 30
support, contacting • iii
Switching Over and Switching Back • 33
Switchover Considerations • 40

## T

technical support, contacting • iii
The Data Recovery Process • 41

## V

View a Report • 30