

# Arcserve® Replication and High Availability

## Operation Guide for UNIX and Linux r16.5

arcserve®

Pre-release Document, only for reference

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by Arcserve 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 Arcserve. This Documentation is confidential and proprietary information of Arcserve and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and Arcserve governing your use of the Arcserve software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and Arcserve.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all Arcserve copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to Arcserve that all copies and partial copies of the Documentation have been returned to Arcserve or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, ARCSERVE 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 ARCSERVE BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF ARCSERVE IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is Arcserve.

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.

© 2015 Arcserve, including its affiliates and subsidiaries. All rights reserved. Any third party trademarks or copyrights are the property of their respective owners.

## Arcserve Product References

This document references the following Arcserve products:

- Arcserve® Replication
- Arcserve® High Availability (HA)

## Contact Arcserve

The Arcserve Support team offers a rich set of resources for resolving your technical issues and provides easy access to important product information.

<https://www.arcserve.com/support>

With Arcserve Support:

- You can get in direct touch with the same library of information that is shared internally by our Arcserve Support experts. This site provides you with access to our knowledge-base (KB) documents. From here you easily search for and find the product-related KB articles which contain field-tested solutions for many top issues and common problems.
- You can use our Live Chat link to instantly launch a real-time conversation between you and the Arcserve Support team. With Live Chat, you can get immediate answers to your concerns and questions, while still maintaining access to the product.
- You can participate in the Arcserve Global User Community to ask and answer questions, share tips and tricks, discuss best practices and participate in conversations with your peers.
- You can open a support ticket. By opening a support ticket online, you can expect a callback from one of our experts in the product area you are inquiring about.

You can access other helpful resources appropriate for your Arcserve product.



# Contents

---

<b>Chapter 1: Introduction</b>	<b>7</b>
Related Documentation .....	7
Supported Server Types .....	7
Server Requirements.....	8
<b>Chapter 2: Installing and Uninstalling Arcserve RHA</b>	<b>9</b>
Prerequisites .....	9
Install the Engine.....	10
Install the Engine into Non-Global Solaris Zones .....	10
Managing the Engine.....	11
Installing the Manager .....	11
Prepare Environment for HA with the GSSnupdate Script .....	13
Upgrading Arcserve RHA .....	15
Uninstalling Arcserve RHA.....	16
Uninstall Arcserve RHA .....	16
<b>Chapter 3: Redirection Methods</b>	<b>17</b>
<b>Chapter 4: Managing Scenarios</b>	<b>19</b>
Create a Replication Scenario .....	19
Create a High Availability Scenario.....	22
Start a Scenario .....	24
Stop a Scenario.....	26
UNIX/Linux Scenario Considerations.....	27
<b>Appendix A: Installed Files</b>	<b>29</b>
Files Installed on Red Hat and Novell SUSE Linux Enterprise .....	29
Files Installed on IBM AIX.....	30
Files Installed on Solaris .....	30
<b>Appendix B: Troubleshooting</b>	<b>33</b>
Unload xofs Drivers .....	33
The Moved IP is not Activated after Switchover .....	35
Oracle Listener Cannot Accept Client Connection After Failover.....	35



# Chapter 1: Introduction

---

This document is intended for experienced UNIX and Linux system administrators interested in implementing and deploying the Arcserve RHA solution in their environment. The document provides all the details necessary to install and uninstall the product, create Replication (DR) and High Availability (HA) scenarios, manage scenarios, and recover lost data.

This section contains the following topics:

[Related Documentation](#) (see page 7)

[Supported Server Types](#) (see page 7)

[Server Requirements](#) (see page 8)

## Related Documentation

Use this Guide along with the *Arcserve RHA Installation Guide*, the *Arcserve RHA Administration Guide*, and the *Arcserve RHA PowerShell Guide*.

## Supported Server Types

Refer to the Release Notes for Supported Server Types.

## Server Requirements

To implement Arcserve RHA, 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, contact Technical Support.

### For HA Scenario:

- Master and Replica must be running the same operating system with same level of service packs and hot fixes. For example, you cannot create an HA scenario from AIX to Solaris.

### For Oracle HA scenario:

- Two servers running supported UNIX/Linux Server with the same level of service packs and hot fixes installed.

**Note:** For a complete list of supported operating systems and applications, see the Arcserve RHA Release Notes.

- Install the same Oracle version on both servers, including service packs and hot fixes.
- The Oracle SID on the Master must match the Replica.
- The ID of the Oracle user (for example, "oracle" user) and Oracle groups (for example, oinstall or dba) must have identical IDs on the master and replica servers. For example, if the ID of "oracle" on the master is 300, then the ID on the replica must be 300.
- Ensure the path to the ORACLE\_HOME directory and the path to the database files are identical on both Master and Replica.

To minimize network traffic, Oracle temporary tablespace names and path are excluded from replication. Make sure that the Oracle database on the Replica server is configured with the same temporary tablespace names and path used on the master server.

**Note:** Arcserve HA for Oracle does not support Oracle RAC.



# Chapter 2: Installing and Uninstalling Arcserve RHA

---

This chapter explains how to install and uninstall the Arcserve RHA Engine and Arcserve RHA Manager.

This section contains the following topics:

[Prerequisites](#) (see page 9)

[Install the Engine](#) (see page 10)

[Managing the Engine](#) (see page 11)

[Installing the Manager](#) (see page 11)

[Prepare Environment for HA with the GSSnsupdate Script](#) (see page 13)

[Upgrading Arcserve RHA](#) (see page 15)

[Uninstalling Arcserve RHA](#) (see page 16)

## Prerequisites

Before the installation, verify the following prerequisite:

On Linux platforms, some packages required by the RHA engine may not be installed by default. You must install these packages before you install the RHA engine. The following command lists the packages required by the RHA engine.

```
rpm -qpR <rha-rpm-file>
```

For example, `rpm -qpR arcserve_rhel5_x86_64.rpm`

## Install the Engine

The Arcserve RHA installation bundle consists of packages for all supported platforms and a common installation script.

### To install the Engine

1. Become "superuser" (root).
2. Change the working directory to the directory that contains installation files.
3. Run the *install.sh* script.
4. Follow the instructions until the end of the installation.

You are prompted to accept the license agreement and asked if you want to create Arcserve RHA group and enable Oracle support.

### Notes:

- If you decide to enable Oracle support, you must provide the Oracle Owner, Oracle Home path, and Oracle Base path. The Oracle Owner user is required first and then the product retrieves the Home path and Base path through the Oracle Owner user environment. If the Home path and Base path cannot be found, then you must manually add them. If your Oracle server is installed without the 32-bit Oracle client library, then you must also provide the Oracle Instant Client path.
- If the Arcserve RHA package has been installed, you are prompted to reinstall it.
- To allow non-root users to manage scenarios, you must create the "caarha" group on your machine and ensure the group works with the supplementary group.

## Install the Engine into Non-Global Solaris Zones

For non-global zones on Solaris, you need to run the *configure.sh* script in each non-global zone.

**Note:** For non-global zones on Solaris 11, make sure you install the RHA engine on global zones before installing on non-global zones.

### To install the Engine into non-global Solaris zones

1. Run the *install.sh* script in global zones as the upper process.
2. Run the *configure.sh* script in each non-global zone after installation.

## Managing the Engine

After the installation, the Arcserve RHA Engine is automatically managed by the operating system: it is started during the operating system boot sequence, and stopped during the operating system shutdown process. These procedures are done automatically through `.rc` scripts.

However, if you need to manually stop the engine and then start it again, do the following:

### Linux

#### To start an Engine

```
/etc/init.d/ARCserveRHA start
```

#### To stop an Engine

```
/etc/init.d/ARCserveRHA stop
```

### Solaris

#### To start an Engine

```
/etc/init.d/ARCserveRHA start
```

#### To stop an Engine

```
/etc/init.d/ARCserveRHA stop
```

### AIX

#### To start an Engine

```
/opt/CA/ARCserveRHA/bin/ARCserveRHA.rc start
```

#### To stop an Engine

```
/opt/CA/ARCserveRHA/bin/ARCserveRHA.rc stop
```

## Installing the Manager

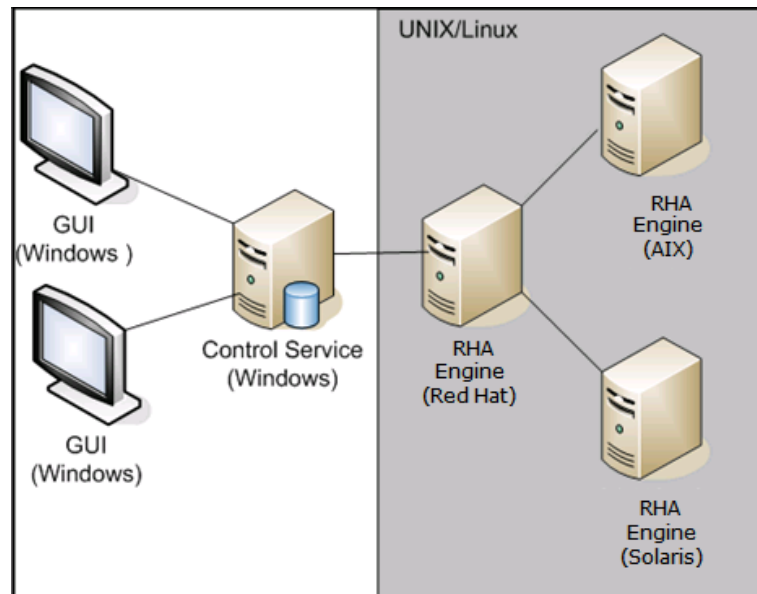
The Arcserve RHA Manager is installed only on Windows platforms. Installing Arcserve RHA components on a Windows platform is very straightforward. The installation package contains a file called *Setup.exe* that runs a standard MSI wizard.

- This (soft) installation does not require reboot or application shutdown.
- The required level of *instmsi.exe* is 2.0.2600.2 or higher. If you do not meet this minimum requirement, the Arcserve RHA installation automatically upgrades the Windows installer for you. However, upgrading the installer requires a reboot.

Standard prompts facilitate the installation. Your only major decision is on which servers to install the applications (Arcserve RHA Manager and Arcserve RHA PowerShell are installed together):

- Install Arcserve RHA Manager and the Control Service on any Windows Server computers that have network access to the machines that you intend to manage, as the following diagram shows.
- The default installation directory (INSTALLDIR) is: \Program Files\CA\ARCserve RHA. All executables, DLLs, and configuration files are located in INSTALLDIR.
- A Windows user running the Arcserve RHA Manager requires Read-Write permission to the installation directory.

Install the Arcserve RHA Engine on all servers participating in UNIX/Linux scenarios. For more information about installing the Arcserve RHA Manager, see the *Arcserve RHA Installation Guide*.



## Prepare Environment for HA with the GSSnsupdate Script

To run high availability scenarios in secure DNS Zone domains, you need to configure your UNIX or Linux server so that it is able to authenticate and change the DNS HOST A records for the Master and Replica servers in scenarios.

You may configure your UNIX/Linux environment for Arcserve RHA using the included script, written by PERL. Contact Support to obtain other versions of this script.

The following steps should be performed on all UNIX/Linux hosts in your environment.

### To prepare UNIX and Linux hosts for HA in secure DNS zone domains

1. Configure the Linux host as a Kerberos client. To do this, edit the file, `/etc/krb5.conf` and make the following changes:

```
[libdefaults]
default_realm = <DOMAIN name i.e. XOSOFT.ORG>

[realms]
<DOMAIN name> = {
  kdc = <DC Server fqdn>:88
  admin_server = <DNS Server fqdn>
  default_domain = <DOMAIN name i.e. XOSOFT.ORG>
}

[domain_realm]
.<domain name >= <DOMAIN name>          i.e. .xosoft.org =XOSOFT.ORG

[logging]
default = FILE:/var/log/krb5libs.log
kdc = FILE:/var/log/krb5kdc.log
admin_server = FILE:/var/log/kadmind.log

[libdefaults]
default_realm = XOLAB.COM

[realms]
XOLAB.COM = {
  kdc = server01 .xolab.com:88
  admin_server = server01 .xolab.com:749
  default_domain = XOLAB.COM
}

[domain_realm]
.xolab.com = XOLAB.COM

[kdc]
profile = /var/kerberos/krb5kdc/kdc.conf
```

2. Create a keytab file with `ktpass` on the domain controller you specified in "kdc" in the previous step. Log in under a domain admin account on the KDC.

**Note:** The `ktpass` utility may not be loaded on Windows by default. You may obtain it from the Windows Support Tools package.

```
ktpass -princ host/<a name you pick@DOMAIN> -mapuser <domain admin@DOMAIN> -pass <password>
-out c:\ <filename>.keytab -ptype KRB5_NT_PRINCIPAL
```

```
C:\>ktpass -princ host/xodemo@XOLAB.COM -mapuser testuser@XOLAB.COM -pass qazwsx
-out C:\xodemo.keytab -ptype KRB5_NT_PRINCIPAL
Targeting domain controller: calabdc01.xolab.com
Successfully mapped host/xodemo to testuser.
Key created.
Output keytab to C:\xodemo.keytab:
Keytab version: 0x502
keysize 48 host/xodemo@XOLAB.COM ptype 1 (KRB5_NT_PRINCIPAL) vno 4 etype 0x3 (DES
S-CBC-MD5) keylength 8 (0x8545b3195d835497)
Account testuser has been set for DES-only encryption.
C:\>
```

3. Transfer the keytab file <filename>.keytab to the Linux host securely.
4. On the Linux host, combine the keytab file into a single file called, /etc/krb5.keytab, using the ktutil program. You may need to install the Ktutil package first.

```
[root@server01-1x1 bin]# ktutil
ktutil: rkt /root/ <filename>.keytab
ktutil: wkt /etc/krb5.keytab
ktutil: list
ktutil: q
```

```
[root@server01 -1x1 ~]# /usr/kerberos/sbin/ktutil
ktutil: rkt ./xodemo.keytab
ktutil: wkt /etc/krb5.keytab
ktutil: list
slot KVNO Principal
-----
1 5 host/test@XOLAB.COM
ktutil: q
[root@calaberv07-1x1 ~]#
```

5. Ensure the Linux host time is synchronized to NTP server or to Microsoft DNS server.
6. Copy nsupdate-gss.pl PERL script and install the set of PERL libraries required by the script. These PERL sources can be found at <http://search.cpan.org> or other RPM.

```
perl-Digest-HMAC
perl-Net-IP
perl-Net-DNS
perl-GSSAPI
perl-Krb5...
```

7. Change Kerberos keytab principal name in the nsupdate-gss.pl script located in /opt/CA/ARCServe RHA/ to match the name you chose above, for example, host/xodemo.

```
#####
my $sys_retcode = system("kinit -k host/xodemo");
if ($sys_retcode != 0) {
    print "Failt to create and cache Kerberos ticket through 'kinit'\n";
    exit 1;
}
```

```
#####
```

8. Run the following test to ensure the environment is ready and the script can securely update DNS A records.

```
./nsupdate-gss.pl 192.168.88.1 shagu01-RHEL5u3-x64 xosoft.org --add_ips=192.168.88.21 --ttl=60
--del_ips=192.168.88.31
```

9. Define the nsupdate-gss.pl script in the ws\_rep.cfg file located in the install directory /opt/CA/ARCserveRHA/bin to perform a secure update against Microsoft DNS. Remove the # in front of the "GSSNsupdateScript = " line, as shown below.

```
#####
```

```
# Script for secure update dynamically against MS DNS server on UNIX/Linux hosts
```

```
GSSNsupdateScript = "[INSTALLDIR]/scripts/nsupdate-gss.pl"
```

```
#
```

```
# User can visit arcserve support, or
```

```
# define the script by shell or PERL with the arguments matching
```

```
# GSSNsupdateScript NAME_SERVER HOST DOMAIN [options]
```

```
# Options:
```

```
# --add_ips=IPS target IPs for A records to be added
```

```
# --ttl=TTL TTL for the added A records
```

```
# --del_ips=IPS target IPs for A records to be removed
```

10. Stop and restart the Engine:

```
/etc/init.d/ARCserveRHA stop
```

```
/etc/init.d/ARCserveRHA start
```

11. Repeat this procedure for the other host.

## Upgrading Arcserve RHA

Consider the following before upgrading to this release:

- Stop any running scenarios.
- You do not need to manually uninstall the previous version. The Installation process removes the prior version automatically.
- Copy old scenarios to the machine running this version of the Arcserve RHA Management GUI. Once copied, you can import them into this version of Arcserve RHA using the Scenario, Import menu selection. Scenarios may be located at:
  - UNIX: /opt/CA/ARCserveRHA/bin/ws\_scenarios
  - Windows: Program Files (x86)/CA/ARCserve RHA/Manager/ws\_scenarios

## Uninstalling Arcserve RHA

Before uninstalling Arcserve RHA you should stop all running scenarios and verify that there are no directories mounted by xofs.

To verify there are no directories mounted by xofs, make sure that the `/etc/xofs_mnttab` file is empty. If it is not empty, see *Unload xofs Drivers* (see page 33).

**Note:** You can skip the verification process by rebooting your computer at the end of uninstall procedure.

## Uninstall Arcserve RHA

The following procedure can be used to uninstall Arcserve RHA on all supported UNIX and Linux platforms.

### To uninstall Arcserve RHA

1. Become "superuser."
2. Make sure that all the replication scenarios have been stopped.
3. Run the following uninstall script:

```
/opt/CA/ArcserveRHA/bin/uninstall.sh
```

You are prompted to confirm the uninstall procedure. For example: *Uninstall Arcserve RHA (y/n)*

4. Type `y` and press Enter.
5. [Optional] Reboot.



# Chapter 3: Redirection Methods

---

Arcserve High Availability (Arcserve 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.

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)
- Switch the server hostname

**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.

Redirection method selection is based on the requirements of the application being protected; certain methods may not apply to a particular scenario. For more information, see the application-specific Operation Guide.

If you use Switch Hostname as the redirection method in UNIX/Linux scenarios, you should also set DNS redirection.

## **Move IP redirection Method for Red Hat Enterprise Linux 6.x and CentOS 6.x**

To use the Move IP Redirection method for Red Hat Enterprise Linux 6.x and CentOS 6.x, perform the following steps:

1. Turn off the NetworkManager service and start the native network service on both the master and replica server. For more information on how to turn off the NetworkManager server, see the respective operating system user guide.
2. On the master, manually configure alias IP addresses by adding `ifcfg-ethModule:<n>` configuration files.



# Chapter 4: Managing Scenarios

---

This section describes how to create, start, and stop scenarios using the Manager. You can also use PowerShell to manage your scenarios. For more information about using PowerShell to create a scenario, see the *Arcserve RHA PowerShell Guide*.

**Important!** Arcserve RHA is best suited for replicating data files. It is not recommended to replicate executable files and libraries since doing so can impact application performance. During scenario creation, you may exclude such files from replication.

Effective with this release, you can use High Availability scenarios to protect your UNIX environment.

This section contains the following topics:

[Create a Replication Scenario](#) (see page 19)

[Create a High Availability Scenario](#) (see page 22)

[Start a Scenario](#) (see page 24)

[Stop a Scenario](#) (see page 26)

[UNIX/Linux Scenario Considerations](#) (see page 27)

## Create a Replication Scenario

Protecting UNIX and Linux machines using Arcserve RHA requires scenarios that identify the Master and Replica servers, as well as properties that control scenario behavior. You can create File Server or Oracle scenarios, as needed. Auto-discovery is used to protect only one Oracle instance in one scenario. If you want to protect more than one Oracle instance in one scenario, use the File Server scenario type. Scenario creation is a wizard-based process and the steps for creating File Server and Oracle scenarios are quite similar. Where differences exist, they are clearly noted. For more detailed instructions, see the appropriate Operation Guide.

**Note:** Root directories cannot be intersected in different scenarios; that is, you cannot create a second Replication scenario using the same master directory and root directory used for your first Replication scenario.

**Note:** RHA cannot protect the block device mounted to the RHA root directory after you run the scenario.

Use your web browser to connect to Arcserve RHA using the Control Service, log onto the portal page and then run the Management GUI. If this is the first time you have run the GUI, it is installed automatically using ClickOnce™ technology. The following is the URL of the portal page:

`http://<ControlServiceHost>:8088/entry_point.aspx`

1. Select Scenario, New from the Manager, or click the New Scenario button on the toolbar.

The New Scenario Wizard opens.

2. Select Create New Scenario, and click Next.

The Select Server and Product Type dialog opens.

3. Select the required scenario options, as follows:

- a. From the Select Server Type list, select the type of scenario you want to create. For UNIX/Linux replication, only File Server and Oracle scenarios are supported.
- b. From the Select Product Type options, select Replication and Data Recovery Scenario (DR).

4. Click Next.

The Master and Replica Hosts dialog opens.

5. Enter the names or IPs of the Master and Replica servers you want to protect, and click Next.

The Engine Verification dialog opens. Wait while the Engine version status is verified. If the current version is outdated, Arcserve RHA asks if you want to update it now. You can install the latest version or continue. When verification is complete, click Next.

**Note:** If the scenario type is Oracle, you are prompted for user credentials. You should enter credentials for an account with administrative privileges in the database so that Arcserve RHA can query the Oracle configuration on the Master server.

6. Enter the Oracle instance name, Oracle DBA and user name/password and click OK.

- In **File Server** scenarios, the Master Root Directories dialog opens. Select the files, directories, or both to replicate from the Master and click Next to continue to the Replica Root Directories screen.
- In **Oracle** scenarios, the Tablespaces for Replication dialog opens in place of the Master Root Directories dialog. Arcserve RHA provides a list of results auto-discovered on the Master server. Select or clear databases for replication as needed and enable the option Replicate new user-created files in listed root directories. Click Next to continue to the Replica Root Directories screen.

The Replica Root Directories dialog opens. Depending on the type of scenario you are creating, the Replica Root Directories dialog shows a list of directories for the Oracle database or File Server.

7. Select the files, directories, or both on the Replica in which the data that you want to replicate will be stored, and click Next.

**Note:** If the same directory structure exists on the Replica server, the wizard automatically selects it as the directory to which to replicate.

The Scenario Properties dialog opens.

8. Configure the Master and Replica properties. Typically, the default values are sufficient. You can modify all the settings in this pane after the scenario is created.

**Note:** The spool is a folder on the disk where data to be replicated is temporarily stored (that is, spooled). The spool parameters, located in the Properties tab on both the Master and Replica, or set with the New Scenario Wizard, 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. See the *Arcserve RHA Administration Guide* for more information.

The Master and Replica Properties dialog opens.

9. Set the Data Rewind option to *On* on the Replica Properties pane if you want to enable recovery of lost data from the Replica using rewind points. For more information about using Data Rewind, see the topic Data Rewind.
10. Click Next when you are satisfied with the Master and Replica properties.

The Scenario Verification dialog opens and the wizard verifies the scenario configuration.

11. Click Next if the scenario is verified successfully, otherwise if errors appear click the Back button and correct the scenario settings.

After the scenario is verified, the Scenario Run page opens where you are prompted to run the scenario. Running the scenario starts the data synchronization process.

12. Click Run Now to run the scenario, or click Finish to finish scenario creation and run the scenario later.

The synchronization process starts.

#### For non-global zones on Solaris

Create multiple scenarios in the global zone and then manage all scenarios from there. Most steps are same except the following steps:

1. Enter the global zone IP address as the master IP address.
2. Use the following commands to get the zone and zone root directory information:

```
/usr/sbin/zoneadm list -vi  
zonecfg -z <zonename> info
```

3. While setting the root directory, manually add the non-global zone root directory to the path that you want to protect. In this case, you cannot use auto discovery for the applications such as Oracle.

## Create a High Availability Scenario

Create UNIX HA scenarios as you would any other HA scenario, using the Scenario Creation Wizard.

For information about creating an HA scenario for an Oracle database, see the *Arcserve RHA Oracle Server Operation Guide*.

### To create a new UNIX HA Scenario

1. Start Arcserve RHA Manager. Select File, Create, New Scenario or click the New Scenario button.

The Welcome dialog opens.

2. Click Create a New Scenario and type a Scenario Group Name, or select one from the list and then click Next.

The Select Server and Product Type dialog opens.

3. Select File Server and ensure the High Availability Scenario (HA) product type is selected. For Tasks on Replica, select None and then click Next.

The Master and Replica Hosts dialog opens.

4. Type a Scenario Name, enter the Hostname or IP Address and Port number for both the Master and Replica servers, enable the Verify Arcserve RHA Engine on Hosts option, and then click Next. You may be prompted for user credentials. If so, enter the appropriate credentials and click OK.

The Engine Verification dialog opens if you enabled engine verification.

5. Wait for verification to complete and click Next to continue when it is complete.

The Master Root Directories dialog opens, showing the list of directories identified on the Master.

6. Click Next.

The Replica Root Directories screen opens.

7. Set the Replica's directory and click Next.

The Scenario Properties screen opens.

8. Configure the properties as needed and click Next. For more information about configuring scenario properties, see the *Arcserve RHA Administration Guide*.

The Master and Replica Properties screen opens.

9. Configure the properties as needed and click Next. For more information about configuring Master and Replica properties, see the *Arcserve RHA Administration Guide*.

The Switchover Properties screen opens.

10. Configure the switchover properties as needed and click Next.

**Note:** Recommended default values are already defined. For more information about configuring switchover properties, see the *Arcserve RHA Administration Guide*.

1. The Switchover and Reverse Replication Initiation screen opens.

11. Make your selections for the following options to control automatic switchover and automatic reverse replication and click Next.

#### **Switchover Initiation**

Specifies whether to initiate a switchover automatically or manually. Automatic switchover is initiated when the Master server is down or a database failure is detected. Manual switchover is initiated by the Administrator. In both cases a notification message is provided when a failure is detected.

#### **Reverse Replication Initiation**

Specifies whether to initiate reverse replication automatically or manually. After a switchover, the scenario stops and reverse replication is initiated.

**Note:** It is recommended that you do **not** set both of these options to automatic in a production environment.

The Scenario Verification screen opens.

12. Wait while the Scenario Verification process completes.

If Scenario Verification lists any errors, you must resolve them to continue. If any warnings are listed, you should also resolve them to successfully continue. After making changes, click Retry to repeat verification.

13. Click Next.

The Scenario Run dialog opens.

14. Click Run Now to start synchronization and activate the scenario. Click Finish to run the scenario later.

### For non-global zones on Solaris

Create multiple scenarios in the global zone and then manage all scenarios from there. Most steps are same except the following steps:

1. Enter the global zone IP address as the master IP address.
2. Use the following commands to get the zone and zone root directory information:


```
/usr/sbin/zonectl list -vi  
  
zonectl -z <zonename> info
```

3. While setting the root directory, manually add the non-global zone root directory to the path that you want to protect. In this case, you cannot use auto discovery for the applications such as Oracle.

## Start a Scenario

You can start a scenario using the Manager.

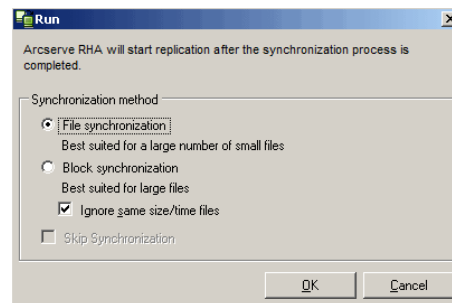
### To start a scenario

1. Select the scenario that you want to run from the Scenario pane.
2. Click Run  on the toolbar.

A Verification Results dialog opens prompting you to approve running the scenario.

3. Click Run. Use the Advanced button to display scenario details of the Master and Replica.

The Run dialog opens.



**Note:** When you start UNIX-based scenarios, you cannot skip file/block synchronization.



4. Choose File synchronization and click OK.

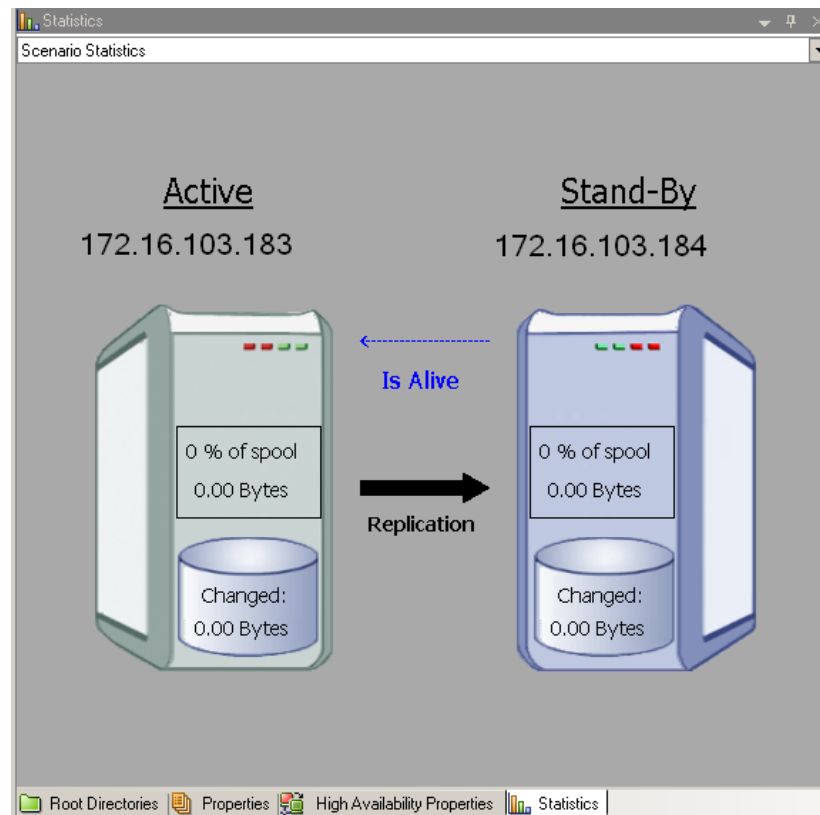
**Note:** If you are running a scenario for an Oracle server, remove the check from the Ignore same size/time files check box and select Block Synchronization.

The Manager now indicates that the scenario is running by the green play symbol to the left of the scenario:

The screenshot shows a window titled "Scenario view" with a tree view on the left and a table on the right. The tree view shows "Scenarios" expanded to "FileServer 1", which has a green play button icon. The table below shows the details of the running scenario.

Scenario	State	Product	Server	Mode
FileServer 1	Running	DR	FileServer	Online
Hosts				
	Changed	Synchronized	Files	In spool
172.16.0.10	0.00 Bytes	22.48 MB	162	0.00 Bytes
172.16.0.11	0.00 Bytes	20.33 MB	150	0.00 Bytes

Once a scenario is running, a Statistics tab appears (on the bottom of the far right pane):




In addition, a Report is generated by default once synchronization occurs.

## Stop a Scenario

You can stop a scenario using the Manager.

### To stop a scenario

1. Click the Stop button  on the toolbar.  
A confirmation message appears prompting you to approve stopping the scenario.
2. Click Yes.

The scenario stops.

**Note:** After stopping the scenario, the Manager no longer shows the green play symbol to the left of the scenario and the Statistics tab is no longer available.

## UNIX/Linux Scenario Considerations

Consider the following when creating UNIX/Linux scenarios:

- A directory can be present in only one scenario.
- For Network File Sharing (NFS), V4 is not supported. You should install the Engine on the NFS server and make sure the exported directory resides in the root directory.
- NFS support is not provided with SUSE 11.
- For Windows-to-UNIX replication, Windows ACLs are lost.
- For UNIX-to-Windows replication, UID, GUID, symbolic links, and hard links are lost.
- File names on Windows are not case-sensitive, so conflicts are possible.
- Root directories cannot be intersected in different scenarios; that is, you cannot create a second Replication scenario using the same master directory and root directory used for your first Replication scenario.
- HA parameters on a UNIX/Linux host:
  - System Information: The AD Domain Controller and MS Cluster properties are not supported so the values are always set to *No*. For the DNS Server property, only the BIND server process *named* is checked; if *named* is running, then the DNS Server property is set to *Yes*, otherwise it is set to *No*.
  - Network Configuration: NetBIOS name is not supported.
- Hard links are now supported. Hard links cannot be created between different file systems or between root directories on the same file system.
- Running scenarios concurrently from different non-global zones on Solaris is not supported. The work-around is to create multiple scenarios from the Solaris global zone and managing all scenario from there.

**Note:** Make sure that all directories in local zones are accessible from the Solaris global zone.



# Appendix A: Installed Files

---

During the Arcserve RHA installation, different files are installed for different operating systems, as described in the following sections.

This section contains the following topics:

[Files Installed on Red Hat and Novell SUSE Linux Enterprise](#) (see page 29)

[Files Installed on IBM AIX](#) (see page 30)

[Files Installed on Solaris](#) (see page 30)

## Files Installed on Red Hat and Novell SUSE Linux Enterprise

On Linux platforms, the following files are installed:

Installed File	Description
/opt/CA/ArcserveRHA/bin/ws_rep	The Arcserve RHA Engine operates in user mode as a daemon. Its primary responsibility is to collaborate with the xofs kernel module (filtering file system), in order to track file system changes and to propagate these changes according to the replication scenario.
/opt/CA/ArcserveRHA/kernel/fs/xofs.*	A proprietary file system - <i>xofs</i> . Implemented in the form of a loadable kernel module. The primary purpose of <i>xofs</i> is to track file system changes and to notify the Engine of these changes. It is loaded during system start up (through /etc/init.d/Arcserve RHA). <b>Note:</b> <i>.up</i> is for uniprocessor while <i>smp</i> is for symmetric multiprocessor.
/etc/init.d/ArcserveRHA	Startup script used to start up and shut down the Engine.
/etc/pam.d/ws_rep	Needed by Arcserve RHA to authenticate management connections from the GUI.
/opt/CA/ArcserveRHA/bin/README	Arcserve RHA README File
/opt/CA/ArcserveRHA/bin/ws_rep.cfg	Arcserve RHA configuration file
/opt/CA/ArcserveRHA/bin/uninstall.sh	Uninstalls the software.

## Files Installed on IBM AIX

On AIX platforms, the following files are installed:

Installed File	Description
/opt/CA/ARCserveRHA/bin/ws_rep	The Arcserve RHA Engine operates in user mode as a daemon. Its primary responsibility is to collaborate with the xofs kernel module (filtering file system) in order to track file system changes and to propagate these changes according to the replication scenario.
/opt/CA/ARCserveRHA/bin/xofs.ext	A proprietary file system - <i>xofs</i> . Implemented in the form of a loadable kernel extension. The primary purpose of xofs is to track file system changes and to notify the Engine of these changes. It is loaded during system start up (through /opt/CA/ARCserveRHA/bin/ARCserveRHA.rc).
/opt/CA/ARCserveRHA/bin/xoctl	Auxiliary utility (file system helper). Preloads xofs during system start-up.
/opt/CA/ARCserveRHA/bin/xoumount	Auxiliary utility ( <i>umounts xofs</i> ). Analog of standard <i>umount</i> command line utility. <b>Note:</b> The standard <i>umounts</i> command will not work for xofs in the current release of Arcserve RHA.
/opt/CA/ARCserveRHA/bin/ARCserveRHA.rc	Startup script used to start up and shut down the Engine.
/opt/CA/ARCserveRHA/bin/uninstall.sh	Uninstalls the software.

## Files Installed on Solaris

On Solaris platforms, the following files are installed:

Installed File	Description
/opt/CA/ARCserveRHA/bin/ws_rep	The Arcserve RHA Engine operates in user mode as a daemon. Its primary responsibility is to collaborate with the xofs kernel module (filtering file system) in order to track file system changes and to propagate these changes according to the replication scenario.
/usr/kernel/fs/xofs and /usr/kernel/fs/sparcv9/xofs	A proprietary file system - <i>xofs</i> . Implemented in the form of a loadable kernel module. The primary purpose of xofs is to track file system changes and to notify the Engine of these changes. Loaded on demand during replication scenario start up.
/opt/CA/ARCserveRHA/bin/uninstall.sh	Uninstalls the software.

/opt/CA/ARCserveRHA/bin/configure.sh      Configures ws\_rep in non-global zones.  
(zones)





# Appendix B: Troubleshooting

---

The following information is provided to help you troubleshoot Arcserve RHA scenarios for UNIX/Linux:

- This version of Arcserve RHA provides an "uninject" feature that allows xofs to unload automatically, even when there are some open files during a scenario stop operation.

This section contains the following topics:

[Unload xofs Drivers](#) (see page 33)

[The Moved IP is not Activated after Switchover](#) (see page 35)

[Oracle Listener Cannot Accept Client Connection After Failover](#) (see page 35)

## Unload xofs Drivers

If some of the directories were inadvertently left under xofs control during the uninstall procedure (for example, the scenario was running and the directory was mounted by xofs), the uninstall procedure cannot unload the xofs driver from your system.

In this case, you can either reboot the computer or unload the xofs driver manually.

### To unload xofs drivers

1. Check if there are any xofs mount points using the following command:

```
#cat /etc/xofs_mnttab
```

2. Stop all the processes that hold your directory. Use the *fuser* command appropriate for your platform to discover processes that have open files in your directory:

#### AIX and Solaris

```
#fuser -c <dir_from_xofs_mnttab>
```

#### Linux

```
#fuser -u <dir_from_xofs_mnttab>
```

3. Use the following `umount` command for the directory discovered in step 1.

```
#umount <dir_from_xofs_mnttab>
```

4. Check that no process is using the xofs driver and manually unload it. Follow the appropriate instructions for your platform:

### **AIX**

Run as Root the following command to check whether xofs is loaded into the kernel:

```
echo lke | kdb| grep xofs
```

Run as Root the following command to unload the xofs driver:

```
/opt/CA/ARCserveRHA/bin/xoctl u /opt/CA/ARCserveRHA/bin/xofs.ext
```

### **Solaris**

Run the following command to check whether the driver is loaded into memory:

```
modinfo|grep xofs
```

Run the following command to manually unload the xofs driver:

```
modunload -i <xofs ID>
```

### **Linux**

Run the following command to verify that the reference counter of the xofs driver is 0:

```
/sbin/lsmmod|grep xofs
```

Run the following command to manually unload the xofs driver:

```
/sbin/rmmod xofs
```

## The Moved IP is not Activated after Switchover

### Valid on Linux

#### Symptom:

After a successful switchover, when I use Move IP, the moved IP (on the new active server) does not seem to be activated. Although, the operating system shows the IPs are activated and a local ping works as well, I cannot access the IP address from outside the server.

#### Solution:

The solution is to manually run the `ifup` command for the moved IP on the new active server after switchover. For example, `ifup eth0:1`.

Optionally, you can automate this by creating a script and running this script using the RHA user interface after switchover.

To run the script, navigate to HA Properties, Action upon Success, User-Defined Script and then provide the script name with the full path.

The following is an example of a script.

```
#!/bin/bash
ifup eth0:1
```

## Oracle Listener Cannot Accept Client Connection After Failover

### Valid on Linux

#### Symptom:

When I use Move IP only, then by default, the Oracle listener on Replica does not accept the client connection after failover.

#### Solution:

This is because the Oracle listener explicitly listens to the specified IP and port. When the moving IP is moved to Replica, the connection is set up for the original production IP and is not accepted by the Oracle listener. The solution is to configure the listener to listen to the IP that you want to move.