

RDS

Directory Synchronization

Rollback Guide

Software Version 3.1.1

For Windows, Linux and UNIX operating systems

August 4, 2009

Copyright © 2009 RepliWeb® Inc., All Rights Reserved

The information in this manual has been compiled with care, but RepliWeb, Inc. makes no warranties as to its accuracy or completeness. The software described herein may be changed or enhanced from time to time. This information does not constitute a commitment or representation by RepliWeb and is subject to change without notice. The software described in this document is furnished under license and may be used and/or copied only in accordance with the terms of this license and the End User License Agreement.

No part of this manual may be reproduced or transmitted, in any form, by any means (electronic, photocopying, recording or otherwise) without the express written consent of RepliWeb, Inc.

Windows, and Windows XP are trademarks of Microsoft Corporation in the US and/or other countries. UNIX is a registered trademark of Bell Laboratories licensed to X/OPEN.

Any other product or company names referred to in this document may be the trademarks of their respective owners.

Please direct correspondence or inquiries to:

RepliWeb, Inc.
6441 Lyons Road
Coconut Creek, Florida 33073
USA

Telephone: (954) 946-2274
Fax: (954) 337-6424

Sales & General Information: info@repliweb.com
Documentation: docs@repliweb.com
Technical Support: <http://support.repliweb.com>
Website: <http://www.repliweb.com>

Table of Contents

1. INTRODUCTION	1
RDS ROLLBACK.....	1
2. RECORDING CHANGES	3
3. ROLLING BACK	5
4. USING CLI	6
ROLLBACK DEFINITION QUALIFIERS	6
ROLLBACK COMMAND.....	8

1. Introduction

The purpose of this document is to outline how RepliWeb RDS can be configured to effectively provide comprehensive deployment level recovery by rollback of content. Rollback can be an integral part of content deployment, allowing a system to return to a previous state if the need arises.

RDS Rollback

Rollback functionality provides a means of effectively undoing the effects of multiple replications on a target system to a user-specified 'point-in-time'. The benefit of such a system over a simple 'archiving' of specific files is its ability to specify a time relative to the targets (i.e. 8 hours ago) and systematically undo the effects of all scheduled or event-driven changes that were replicated.

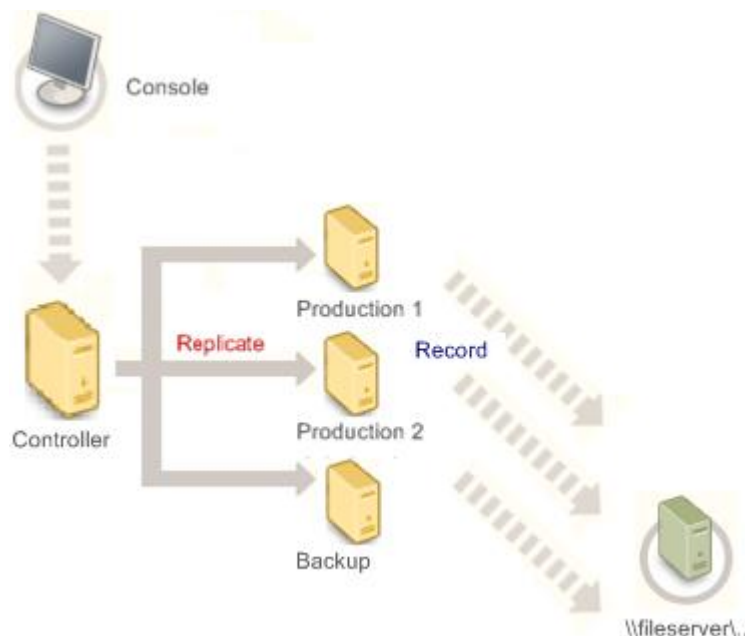


Figure 1: Sample Topology

Rollback functionality is employed through the 'recording' of all replication related changes on a target system, followed by potentially 'rolling back' the content at a later time. The recording phase occurs prior to any propagation of changes to a target system and includes the recording of creation, deletion and modification to files and directories (including ownership and permission changes). Recording can be defined to take place locally on each target system, or on a network-wide file server defined through a UNC share-point (i.e. \\nasstorage\rds_recording) – effectively providing unattended and efficient tape-less backup and restoration. System administrators also custom define the

recording history to a time relative to the present, for example, specifying that a replication should be able to rollback to a state 24 hours in the past. RDS will then only keep the changes that occurred in the last 24 hours, and automatically purge recorded changes that have expired.

2. Recording Changes

To record your changes, set the replication job to record all changes. The recording phase occurs during the replication, prior to any propagation of changes to a target system and includes the recording of creation, deletion and modification to files and directories (including ownership and permission changes).

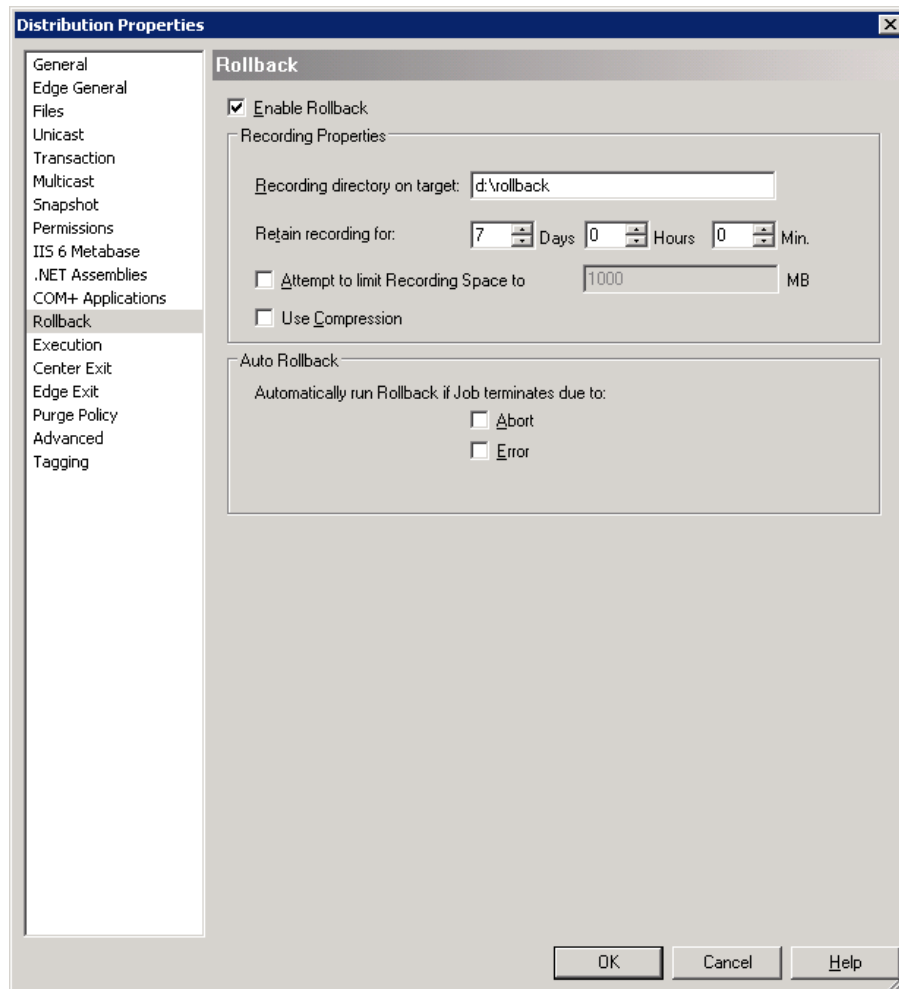


Figure 2: Rollback Definition

Set the job to record enable rollback recording on the target.

Recorded information includes:

- Files and directories created on the target
- Files and directories removed from the target
- Files and directories modified on the target (included modified permission, ownership and Windows ACLs)

NOTE: Recording only applies to data impacted by a replication as defined above – it does not record a complete copy of the target directory.

You will need to specify where the recorded data will be stored. The recording directory may be specified as a path on the target machine (`D:\rds_recordings`, `/var/rollback/`) or a network share (`\\nasstorage\rds_recording`).

Since the recording operation takes place on the target (i.e. a Satellite during an upload job, or the Controller during a download job), the recording directory must be accessible for read and write with the credentials used to access the target file system. This refers to the Satellite user defined in an upload job, or the Controller user (i.e. the username used to connect to the Console, or used in the command line) in a download job.

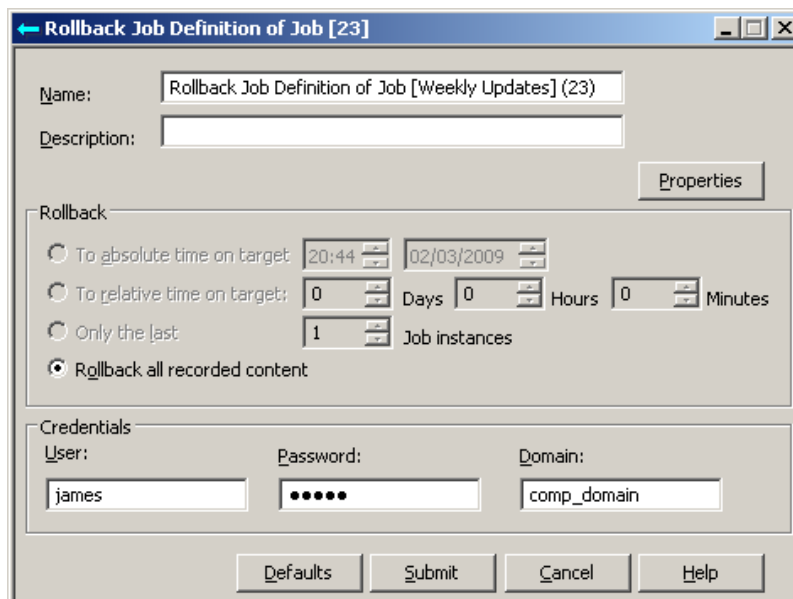
Next, define how long recorded information should be kept for. This directly relates to how far back a job can be rolled back. For example, a time of 18 hours will allow the job to undo all replication related changes that took place on the target system for the last 18 hours.

Recorded data can be kept in a compressed form. While preserving disk-space, this feature will increase recording and rollback times.

3. Rolling Back

A **Rollback** job may be submitted for jobs that had the **Enable Rollback** option checked in the **Rollback** tab. When the job is completed (successfully or not) using the right mouse menu or the main menu you can select to recover previously recorded data. The Rollback Job allows point-in-time recovery capabilities.

A rollback is similar to an undo function. It enables you to return content to its previously replicated state. Rollbacks are performed on target machine. You cannot roll back content on the source.



The screenshot shows a dialog box titled "Rollback Job Definition of Job [23]". It contains the following fields and options:

- Name:** Rollback Job Definition of Job [Weekly Updates] (23)
- Description:** (empty text box)
- Properties:** (button)
- Rollback:**
 - To absolute time on target: 20:44, 02/03/2009
 - To relative time on target: 0 Days, 0 Hours, 0 Minutes
 - Only the last 1 Job instances
 - Rollback all recorded content
- Credentials:**
 - User:** james
 - Password:** (masked with dots)
 - Domain:** comp_domain
- Buttons:** Defaults, Submit, Cancel, Help

Figure 3: Recovery

For **Scheduled** jobs, you can rollback to instances of the job.

Rolling back, (recovering data) is performed by selecting a point-in-time to return to, in addition to various job properties that can be set to define the rollback process. This includes standard job properties like pre/post commands, recovery options, etc.

NOTE: Jobs that use the Rollback functionality can be rolled-back only after the job completes (successfully or not). In order to rollback a job, the job cannot be running. To rollback an instance of a **Scheduled** job, the scheduled job itself needs to be aborted; it is not sufficient that the instance is completed.

The Rollback job can be set to perform Pre/Post commands both on the target and the source, as well as standard RDS Exit procedures like sending Email and reports and the completion of the Rollback job.

4. Using CLI

RDS Job rollback can also be performed using Command-Line.

Rollback Definition Qualifiers

Rollback functionality is employed through the 'recording' of all replication related changes on a target system, followed by a potential 'roll back' of the content at a later time. The recording phase occurs prior to any propagation of changes to a target system and includes the recording of creation, deletion and modification to files and directories (including ownership and permission changes).

The following qualifiers are used to define the properties of the recording phase of a job. To actually activate a rollback job, see the **Rollback Job Qualifiers** description below.

-rb_enable|norb_enable

Use this qualifier to enable rollback recording on the target.

Recorded information includes:

- Files and directories created on the target
- Files and directories removed from the target
- Files and directories modified on the target (included modified permission, ownership and Windows ACLs)

NOTE: Recording only applies to data impacted by a replication job as defined above – it does not record a complete copy of the target directory.

-rb_location=recording_folder

Specify where the recorded data will be stored. The recording location may be specified as a path on the target machine (D:\rds_recordings) or a network share (\\nasstorage\rds_recording)

Since the recording operation takes place on the target (i.e. a Satellite during an upload job or the Controller during a download job), the recording directory must be accessible for read and write with the credentials used to access the target file system. This refers to the Satellite user defined in an upload job, or the Controller user (i.e. the username used to connect to the Console, or used in the command line) in a download job.

This qualifier must be specified when using the **-rb_enable** qualifier.

-rb_retain=n in minutes

Specify in minutes how long recorded information should be kept for and directly relates to how far back a job can be rolled back. For example, a time of 18 hours will allow the job to undo all replication related changes that took place on the target system for the last 18 hours.

By not utilizing the `-rb_retain=n` qualifier, the default setting for retained recordings which is seven days will be used.

`-rb_compress|-norb_compress`

Using `-rb_compress` will store all recorded data in a compressed form. While preserving disk-space, this feature will increase recording and rollback times. The default is `-norb_compress`.

`-rb_enable_disk_space_limit|-norb_enable_disk_space_limit`

Use the `-rb_enable_disk_space_limit` to enforce space limits used for the job's recording during the Rollback process.

After the recording phase, the size of the recording repository is checked; if it is larger than allowed by the job configuration, R-1 will start deleting the oldest recordings until it reaches the size limit. The last recording is never deleted, even if it is larger than allowed on its own.

`-rb_max_rec_dir_size=dir_size` in MB

Set the rollback directory's maximum size. If the rollback directory passes the limit, R-1 will delete old recordings. (In the worst case all recordings but the last one will be deleted).

`-auto_rollback_on_abort|-noauto_rollback_on_abort`

Specify whether to perform automatic rollback in case the job aborts.

`-auto_rollback_on_error|-noauto_rollback_on_error`

Specify whether to perform automatic rollback in case the job encounters an error.

Example

The following command is used to submit a Replication job that records Rollback information on a network share drive.



```
>rds submit -controller=34.0.25.71 -controller_user=rwroot
-controller_pass=***
-replication_name=\"Web Site Update\" -satellite=152.163.12.2
-user=user_name -password=password
-source_dir=/usr/backup -target_dir=d:\data\content
-rb_enable
-rb_location=\\nasstorage\RDS_recording
RDSAPI-S-SUS, job <17103> successfully submitted
```

Rollback Command

This command is used to rollback a Replication Job. A **Rollback** job may be submitted for jobs that had the **Use Rollback** option checked in the `submit` command. When the job is completed (successfully or not) you can select to recover previously recorded data. The Rollback Job allows point-in-time recovery capabilities.

A rollback is similar to an undo function. It enables you to return content to its previously replicated state. Rollbacks are performed on target machine. You cannot roll back content on the source.

Rolling back, (recovering data) is performed by selecting a point-in-time to return to, in addition to various job properties that can be set to define the rollback process. This includes standard job properties like pre/post commands, recovery options, etc.

NOTE: Jobs that use the Rollback functionality can be rolled-back only after the job completes (successfully or not). In order to rollback a job, the job cannot be running. To rollback an instance of a **Scheduled**, the scheduled job itself needs to be aborted; it is not sufficient that the instance is completed.

Absolute Time Example:

The following command is used to Rollback a scheduled upload job to an absolute time. Data will be rolled back to 10:30 PM according to the Target machine (in this case the Satellite):



```
>rds rollback -job=3212 -replication_name="Rolling back"  
-description="Rolling back due to..." -controller=controller_name  
-controller_user=user_name -controller_password=password  
-user=user_name -password=password  
-rollback_to_time=2230  
DSAPI-S-SUS, job <4121> successfully submitted
```

NOTE: If you are using different credentials than those used in the original job, make sure there are sufficient access rights to the recording location and the target directory.

Relative Time Example:

The following command is used to Rollback a scheduled Upload job to a relative time. Data will be rolled back to its state 4 hours ago:



```
>rds rollback -job=212 -replication_name="Rolling back"
-description="Rolling back due to..." -controller=controller_name
-controller_user=user_name -controller_password=password
-user=user_name -password=password
-rollback_to_time="-4 hours"
RDSAPI-S-SUS, job <3721> successfully submitted
```

-job=job_id

Use this qualifier to specify the job that needs to be rolled back. Note that you can only roll back completed jobs.

-replication_name=rollback_job_name

The `-replication_name` qualifier defines the name of the replication rollback job.

-description=rollback_job_description

The `-description` qualifier provides a description of the replication rollback job.

-rollback_to_time=time_expression

Specify the date and time to recover from. This can be an absolute or a relative time.

NOTE: The time specified is calculated at the Target machine (Satellite in case of an Upload job, Controller if it's a Download job). It is never the Console time.

- **Absolute time on target** – Specify the exact time to roll back to. Format supports most standard time formats like: "ddmmyyyhhmm". For a complete list of supported time formats please refer to the RDS User Guide.

In case a rollback instance is not found at the exact time specify, RDS will roll back to the nearest recorded data, without passing the time specified.

- **Relative time on target** – Specify how long to roll back to, in days / hours / minutes. Here, too, the calculated time to roll back from is calculated on the target machine. For relative time format, please refer to the RDS User Guide.
- **Rollback all recorded content** – Using this option, the rollback process will bring the target directory to its state before the replication took place. To use this option, omit the `-rollback_to_time` qualifier from the command.

NOTE: For non-scheduled Replication jobs, this is the only option available.

For additional information, contact us at support.repliweb.com