

Replay 4 Replication

Technical Overview

This document contains a technical overview of the Replay Replication facility new to Replay 4. Please review this document before implementing replication.



Copyright 2009 by AppAssure, Inc.

ALL RIGHTS RESERVED.

No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means — graphic, electronic, or mechanical, including photocopying, recording, taping, Web distribution or information storage and retrieval systems — without the written permission of the publisher.

For permission to use material from this publication, contact

AppAssure Software, Inc.
1925 Isaac Newton Square, Suite 440
Reston, VA 20910
+1 703-547-8686
info@appassure.com

CONTENTS

Introduction.....	4
How it works.....	4
Terms	4
Snapshots.....	5
Replication	5
Replication Implementation Details	6
Replication Topologies	6
1 : 1 Configuration	6
Many : 1 Configuration	7
Many: Many Configuration	8
Replication Schedules	8
Recovery Point Copy and Consume	9
Authentication between Replay Cores	9

TECHNICAL OVERVIEW

Introduction

This document describes the Replication feature of Replay 4. Replication is designed to replicate recovery points between Replay Cores in an efficient and safe manner to enable off-site backup and off-site disaster recovery.

Replay Replication option can be enabled or disabled on a per protected server basis. This option is available in single server and multi-server implementations. The cool thing about Replay Replication is that it replicates the compressed and deduplicated recovery points over the WAN so its bandwidth efficient saves storage on the LAN and at the DR location.




Figure 1 - Off-Site Disaster Recovery

How it works

This section describes how Replay Replication works by first defining our terminology followed by a review of our approach.

Terms

- **Protected Server** – A Windows workload, physical or virtual, that is protected by a Replay Agent. Supported environments include Windows® Server 2008 R2, Windows® 2008 Server, Windows® 2003 Server, Windows 7, XP, Vista.
- **Replay Core**– processes and stores compressed and deduplicated recovery points as incremental forever images of the protected servers. Cores can be installed on the protected

server or on a dedicated server or worker virtual machine. This role performs many functions including creating virtual standby environments, application validation, restores and replication.

- **Recovery Points** – Represents a point-in-time image of the protected server. They are stored as compressed and deduplicated files on storage available to the Replay Core, DAS, NAS, SAN.
- **Replication Source** – The server from where the replication of the recovery points is initiated. Replication is always performed between two Replay Cores.
- **Replication Target** – The server that is the destination of the replicated recovery points.

Snapshots

Replay 4 creates point-in-time snapshots of the protected servers and they are stored as compressed and deduplicated recovery points in a directory accessible from the Replay Core. The recovery points are maintained on a per protected server basis in a directory structure as follows:

```
Drive:\TevRepository\ProtectedServerName1
  recoverypoint4Incremental
  recoverypoint3Incremental
  recoverypoint2Incremental
  recoverypoint1Base
```

```
\\ShareName\TevRepository\ProtectedServerName2
  recoverypoint4Incremental
  recoverypoint3Incremental
  recoverypoint2Incremental
  recoverypoint1Base
```

Replication

Replication copies the recovery points from a replication source to replication target on a per protected server basis. As recovery points are replicated, all recovery point files and registry settings are replicated to the secondary Replay Core, such that the loss of the replication source won't mean the loss of all recovery points and protection settings. In the case of a Replay Core outage, the Replay Core can be rebuilt, and resynchronized with a mirror, though until resync is complete the primary Replay Core won't be able to accept new snapshots.

Here are some facts about Replay Replication:

- As new snapshots are taken on the replication source, they will be replicated to replication target.
- Replication is performed on a per protected server basis.
- Recovery points are replicated oldest-first, to ensure the replication target always has a valid and mountable epoch chain.
- Recovery point changes due to rollup will not be replicated reducing the amount of data that needs to be transferred or the WAN. The replication target will perform its own rollup on the replicated recovery points.

- Rollup and replication will not conflict with one another.
- It is assumed that the replication target is on a different subnet and behind a different firewall than replication source, thus the replication target will only communicate with the agents protected by replication source for rollback purposes..
- The replication target is able to perform all normal Replay restore functions such as BMR, instant rollback, VM export, even if the replication source is down.
- If Exchange is being protected, Exchange and system DLLs from the replication source will be propagated to the replication target so features such as force mountability check, MR, and P2V are available on the replication target.
- Mismatched versions between replication source and replication target are detected.
- The replication source and replication target can be any combination of supported operating systems.
- In the event of a link failure or transfer error, when the link comes back up, the transfer is resumed from where it was interrupted.
- All recovery points that are replicated are verified for integrity.

Replication Implementation Details

- Internally, the Replay Cores communicate over HTTP. The target listens on port 8080 for HTTP connections from the source. All communications are initiated by the source to the target.
- Replay uses the Adler-32 checksum algorithm to verify the integrity of transferred files, and detect when a partially-transferred file can be resumed without data loss.
- To determine what files to replicate, Replay chooses all source files not yet on the target, and those files on both which have not yet been rolled up but have different metadata between the two sides.
- Each file is transferred in a separate HTTP PUT operation. Once an entire epoch is transferred, it is loaded into Replay and available for use. Thus, some replicated recovery points are available for use even before all recovery points have replicated.

Replication Topologies

Replay support various replication topologies to meet your business needs.

1 : 1 Configuration

The 1 : 1 configuration is that standard configuration and is useful for protection of a single server or group of servers from 1 site to another.

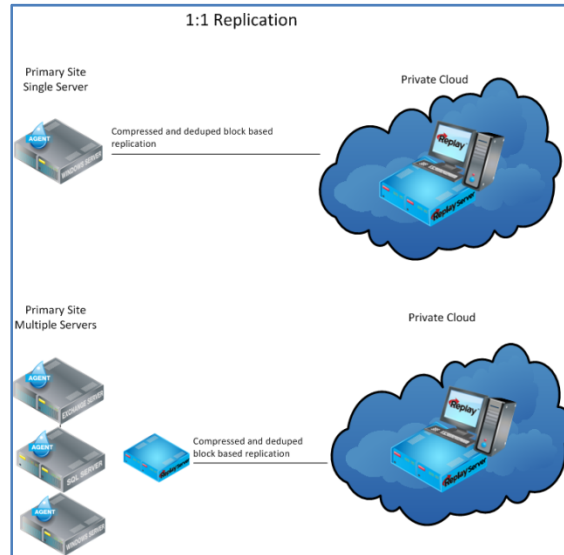


Figure 2

Many : 1 Configuration

The Many : 1 configuration is useful for protection of remote offices from a centralized location; multiple replication sources can replicate to one replication target as shown in the figure below.

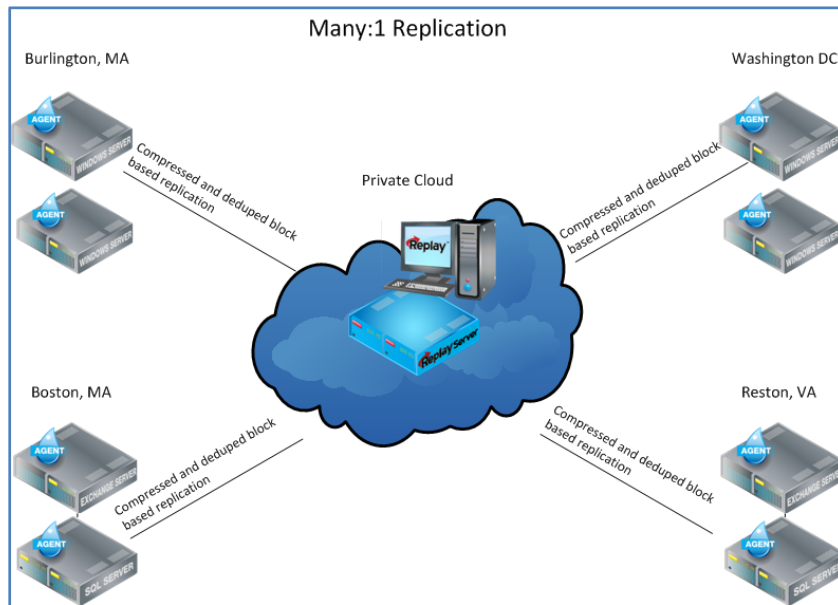


Figure 3

Many: Many Configuration

Replay Replication supports replication on a per protected server basis. This means that you can replicate different protected server's recovery points to different replication targets in different locations as shown in the figure below.


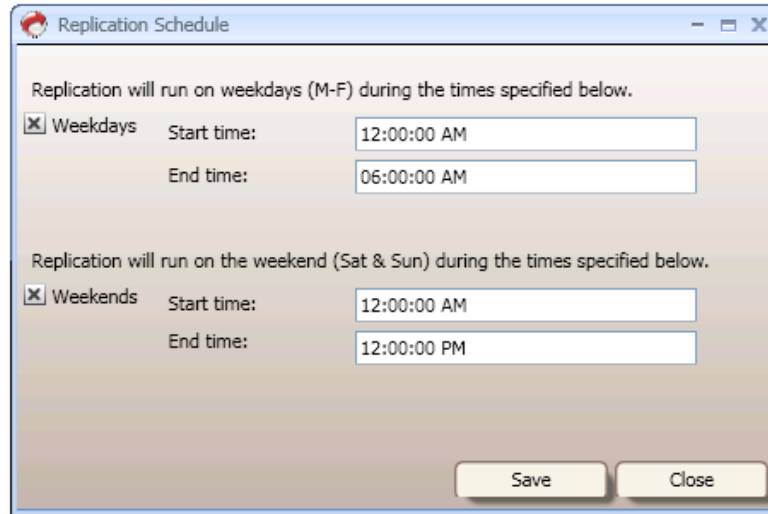


Figure 4

Replication Schedules

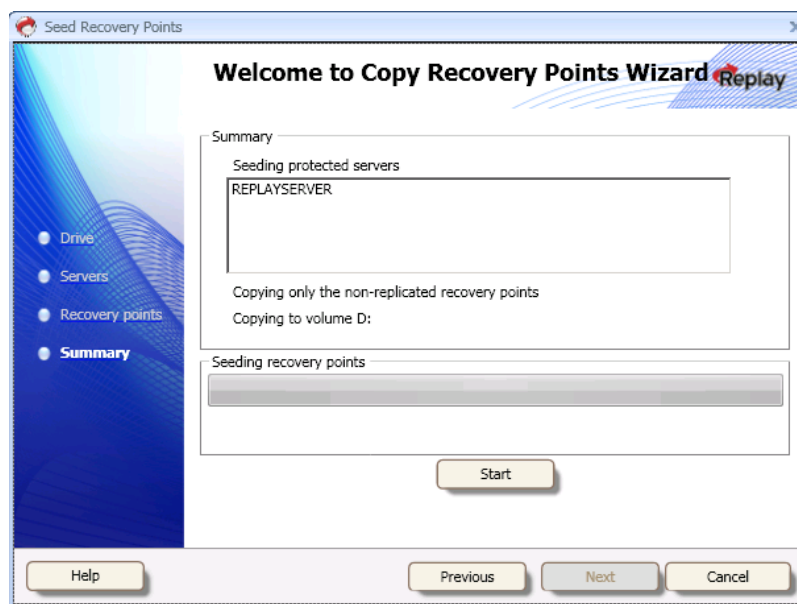
Replication can be scheduled to run during certain times during the day. Any time outside of this window, replication will be suspended. The schedule is specified as a start and stop time for weekdays, and a start and stop time for weekends.

- Users can, on an ad-hoc basis, pause a currently running replication for an hour, a day, or until resumed
- Users can configure VM export, recovery point export, or continuous rollback on the replication target for recovery points replicated from the replication source.



Recovery Point Copy and Consume

The copying feature is useful when the WAN link is insufficient to move large amounts of data. It copies the existing recovery points including the base image from the replication source to a local removable device. Once the copy is complete, the removable drive can be inserted in the replication target and consumed. The replication target is seeded with the recovery points from the replication source including the base image. The replication will resume by replicating the new recovery points only.



Authentication between Replay Cores

Replay uses NTLM authentication over HTTP. Once authenticated between the replication source and the replication target, Replay generates a one-time-use key, which it exchanges via RSA over HTTP. Each request sent to the replication target over HTTP is authenticated with a separate one-time key. Thus

authentication is hardened against Replay attacks, but potentially vulnerable to a man-in-the-middle attack since we have no way of authenticating the identity of the remote server.