



# **Hitachi Dynamic Replicator - Scout Host Administration Guide**

Copyright © 2009 Hitachi, Ltd., Hitachi Data Systems Corporation, ALL RIGHTS RESERVED

Notice: No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd. and Hitachi Data Systems Corporation (hereinafter referred to as "Hitachi Data Systems").

Hitachi, Ltd. and Hitachi Data Systems reserve the right to make changes to this document at any time without notice and assume no responsibility for its use. Hitachi, Ltd. and Hitachi Data Systems products and services can only be ordered under the terms and conditions of Hitachi Data Systems' applicable agreements. All of the features described in this document may not be currently available. Refer to the most recent product announcement or contact your local Hitachi Data Systems sales office for information on feature and product availability.

This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.

Hitachi is a registered trademark of Hitachi, Ltd. in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd. in the United States and other countries.

All other trademarks, service marks, and company names in this document are properties of their respective owners.

## Contents


<b>Part 1: Introduction to Scout .....</b>	<b>9</b>
1 Introduction to Scout.....	10
1.1 Overview.....	10
1.2 Scout Components.....	12
1.2.1 CX Server (Configuration Server).....	12
1.2.2 Process Server (PS).....	13
1.2.3 Volume Replication –VX Agent .....	14
1.2.4 File Replication – FX Agent .....	18
1.2.5 CDP Retention Option.....	20
<b>Part 2: Protecting data using Scout .....</b>	<b>21</b>
2 Using the Management Console.....	22
2.1 Introduction to the Management Console.....	22
2.2 Logging into the CX-CS User Interface.....	23
2.2.1 CX Dashboard.....	24
2.2.2 CX-CS User Interface Components.....	28
2.2.3 License Management .....	29
3 The VX Agent.....	38
3.1 Volume Protection .....	39
3.1.1 When to choose VX Replication .....	39
3.1.2 Profiling.....	41
3.1.3 Differences between Profiler and Plain VX Replication .....	44
3.1.4 Preparing Target Volume for Replication.....	44
3.1.5 Before setting up Replication Pairs.....	48
3.1.6 Setting Volume Replication .....	49
3.1.7 Microsoft Clustered Volume Replication.....	59
3.1.8 Editing a Replication Pair .....	63
3.1.9 Stop or Break a Replication.....	73
3.1.10 Pause Replication .....	74
3.1.11 Force Delete.....	75
3.2 Additional notes for Solaris.....	76
3.2.1 Protecting zpool on Solaris .....	76
3.2.2 Protecting Solaris 10 with zones.....	78
4 The FX agent.....	79
4.1 File Protection.....	80
4.1.1 When to use FX Replication.....	81
4.1.2 FX Execution .....	86
4.1.3 Creating User-defined FX Templates .....	87
4.1.4 Things to Remember while Setting FX Job Options .....	98
4.1.5 Replicating Files through FX .....	100
4.1.6 Group Scheduling Modes .....	105
4.1.7 Editing Existing Job Groups and Jobs .....	108
4.1.8 FX error codes, causes, and solutions.....	110
<b>Part 3: Recovery .....</b>	<b>113</b>
5 Recovery Operations.....	114
5.1 Data Consistency .....	114

5.1.1	Generating consistency markers/ Tags.....	114
5.1.2	Automating Tag Generation through FX.....	114
5.2	Recovery.....	115
5.2.1	When to Use Recovery .....	115
5.2.2	Determine the Time to Recover.....	117
5.2.3	Recovery Snapshots .....	119
5.2.4	Scheduled Snapshots .....	131
5.2.5	Performing Target Volume Rollback.....	139
5.2.6	Recovery Pair Status .....	142
5.2.7	Target Drive Rollback Status .....	143
<b>Part 4: Monitoring</b>	<b>.....</b>	<b>145</b>
6	Checking Status.....	146
6.1	Protection Status .....	147
6.1.1	Protection Status -Volumes.....	147
6.1.2	Protection Status - Files .....	150
6.1.3	Protection Status - Snapshots .....	153
6.2	Possible situations while replication.....	154
6.2.1	Resync required set to Yes .....	154
6.2.2	Resizing source volume.....	155
6.2.3	Throttling.....	159
6.2.4	Disk Space Warning Threshold.....	160
6.3	Agent Heartbeat.....	161
7	Configuring Alerts.....	162
7.1	Trap Listeners.....	162
7.2	Adding Users and Email Alerts.....	165
7.2.1	Types of Users .....	165
7.2.2	To add an Administrator or User.....	165
7.2.3	Mail server settings.....	173
7.3	Agent Settings .....	173
7.3.1	Agent settings.....	173
7.3.2	Process Server .....	174
8	Log Management .....	175
8.1	Host – Log details .....	175
8.2	CX - Log Details .....	176
8.3	Download Logs.....	177
8.4	Audit log: .....	177
8.5	CX Log Rotation.....	179
8.6	Clear FR Logs by Date.....	179
8.7	Retention Reserve Space Settings .....	180
8.8	Installers .....	180
9	Analyzer and Trending.....	181
9.1	Trending.....	181
9.1.1	When is Trending Used.....	181
9.1.2	Network Trending .....	182
9.1.3	Bandwidth Report.....	183
9.1.4	Health Report.....	185
9.1.5	Detailed Trending .....	194
9.1.6	Analyzer .....	194

<b>Part 5: Other Administration Tasks.....</b>	<b>197</b>
10 Configuration settings .....	198
10.1 CX Backup/Restore .....	198
10.2 License Management.....	199
10.2.1 License Validity and Expiry.....	199
10.2.2 Capacity Calculator.....	200
10.2.3 Upgrading a License .....	201
10.2.4 CX Cluster .....	202
10.2.5 Configure Remote CX.....	206
10.2.6 Differences between CX HA and Remote CX features .....	211
10.2.7 Process Server Failover.....	213
10.2.8 Versions & Patches.....	214
10.3 Process Server Traffic Load Balancing.....	215
10.4 User Documents.....	216
10.5 RX Settings.....	216
<b>Part 6: Protecting Applications.....</b>	<b>217</b>
11 Application Support .....	218
11.1 Application support .....	218
11.2 Exchange Server Support.....	220
11.2.1 Protection .....	220
11.2.2 Recovery Types .....	221
11.3 SQL Server .....	222
11.3.1 Protection .....	222
11.3.2 Recovery types.....	222
<b>Part 7: Command Line Tools.....</b>	<b>223</b>
12 Console Tools.....	224
12.1 Issuing Consistency Tags.....	225
12.1.1 VACP Command Line Options (Windows).....	225
12.1.2 VACP Command Line Options (Linux/ Solaris).....	235
12.2 CX UI v/s cdpcli .....	238
12.3 CDPCLI Interface.....	239
12.3.1 When is cdpcli used .....	239
12.3.2 Cdpcli on windows.....	239
12.3.3 Cdpcli on Linux/ Solaris.....	261
12.4 Generating reports through bwreport.pl.....	281
12.4.1 Daily report.....	281
12.4.2 Weekly report .....	282
12.4.3 Monthly and yearly report.....	282
12.4.4 Custom report.....	282
12.5 ICAT .....	283
12.5.1 Introduction to ICAT .....	283
12.5.2 How ICAT works .....	284
12.5.3 Before you begin.....	284
12.5.4 Configure.....	285
12.5.5 Resume mode .....	290
12.6 Application Support.....	291
12.6.1 Application.exe.....	291
12.6.2 DNS.exe .....	293

12.6.3	Winop.exe.....	294
12.6.4	Tweaking FailoverServices.conf.....	300
<b>Part 8: Performance Tuning .....</b>		<b>301</b>
13	Tuning FX.....	302
13.1	Tuning the FX for Faster Execution.....	302
13.1.1	Options for Faster Execution .....	302
13.2	Auto Delete FX Logs on Source / target.....	304
14	Tuning VX.....	305
14.1	Source side Driver Tweaks .....	305
14.1.1	Working Process.....	305
14.1.2	Registry Values.....	306
14.2	For production environments with aggressive data writes .....	307
14.3	Tunable for Solaris.....	308
14.3.1	Common Tunable.....	308
14.3.2	Volume Specific Tunable.....	309
14.4	Tuning target side agent .....	310
14.5	Tuning TCP Window Sizes.....	316
15	Tuning CX.....	317
15.1	Using a Lower Encryption.....	317
15.2	TCP Window tuning for Linux based CX .....	317
16	Bandwidth Shaping .....	318
16.1	Preparing for bandwidth shaping .....	318
16.2	Creating bandwidth shaping pairs.....	319
16.2.1	Policy Conflicts.....	322
16.3	How Policies are enforced .....	323
16.4	Effect of Stopping VX Replication Pairs on the Bandwidth Policies .....	323
16.5	Advanced Configuration .....	323

## Conventions

Keywords, command buttons and other such fields are enclosed in “ ” while being bold (for example, to denote  “**Next**” is used).

Inputs for commands and Variables are shown in *Italics*

File names and paths are shown in **bold**

Commands are shown in **Courier new font**

Mandatory keywords, arguments and inputs are enclosed within < >.

Optional arguments are enclosed in [ ]



### Notes:

Contain suggestions or tips.



### Caution:

Contains critical information

## References

Although this document been designed sequentially, you may choose to skip to sections that are of interest. Other documents which you may want to refer are “**Installation Guide**” for installing and configuring Hitachi Dynamic Replicator - Scout. If your requirement is specific to applications like Exchange, SQL server, Oracle etc then refer the solution documents. Also refer to the “**Troubleshooting Guide**” for possible issues and their workarounds.

## Scope of document

This document primarily deals with the CX user interface and functionalities. The “[recovery](#)” section does elaborate on recovery operations through the CX user interface. Console based recovery is also explained in this document.

## Target audience

This document is intended for Hitachi Dynamic Replicator - Scout administrators, Hitachi Dynamic Replicator partners, sales, and engineering teams.

## How this document is organized

This document is divided into eight parts.

### **Part 1: Introduction to Scout**

Contains introduction to the Scout, components of Scout, CDP retention option technology etc.

### **Part 2: Protecting data using Scout**

Chapters 2, 3 and 4 are included in this section. Topics covered are:

- Introduction to the CX User interface
- Volume level replication and
- File level replication

### **Part 3: Recovery**

All recovery operations that can be performed through the CX user interface are covered in this section.

### **Part 4: Monitoring**

Chapters under this section deal with monitoring backup process, health of the application etc.

### **Part 5: Other administration tasks**

Administrative tasks such as adding or deleting agents to the CX server, backing up the CX server configuration, monitoring any applied patches etc are covered in this section.

### **Part 6: Application support**

An introduction to the application support is given in this section. Refer to solution documents for each application for detailed steps in protecting and recovering applications.

### **Part 7: Command line tools**

This section covers console tools, their syntaxes, and examples. Both windows and Linux example commands are shown in black and white backgrounds respectively. Issuing consistency markers on production server and performing recovery operations through the DR server's console are covered in this chapter.

### **Part 8: Performance Tuning**

Scout comes with certain default settings ideal for production environment. However, these values can be tweaked for better performance considering available resources.



#### **Notes:**

Refer to the installation guide for platform specific installation with examples



# Part 1: Introduction to Scout

This chapter introduces the Scout and its components. A high level overview of Scout is discussed followed by Scout components and their architecture. At the end of this chapter you would

- Understand Scout components and architecture
- Understand CDP retention option technology

# 1 Introduction to Scout

## 1.1 Overview

Hitachi Dynamic Replicator's Scout is an application-aware business continuity solution that combines enterprise-class disaster recovery and advanced continuous data protection (CDP) in a single product. Scout is ideally suited for implementing consolidated real time backups and remote replication based on CDP technology (Continuous Data Protection).

CDP technology provides the capability to rollback/rewind the system to a point back in time without any data copies or restores. Data changes are captured at block level on production volumes thus effectively reducing the RPO.

Components of Scout include host agents (VX or FX) on the production and backup servers along with a dedicated server (CX).

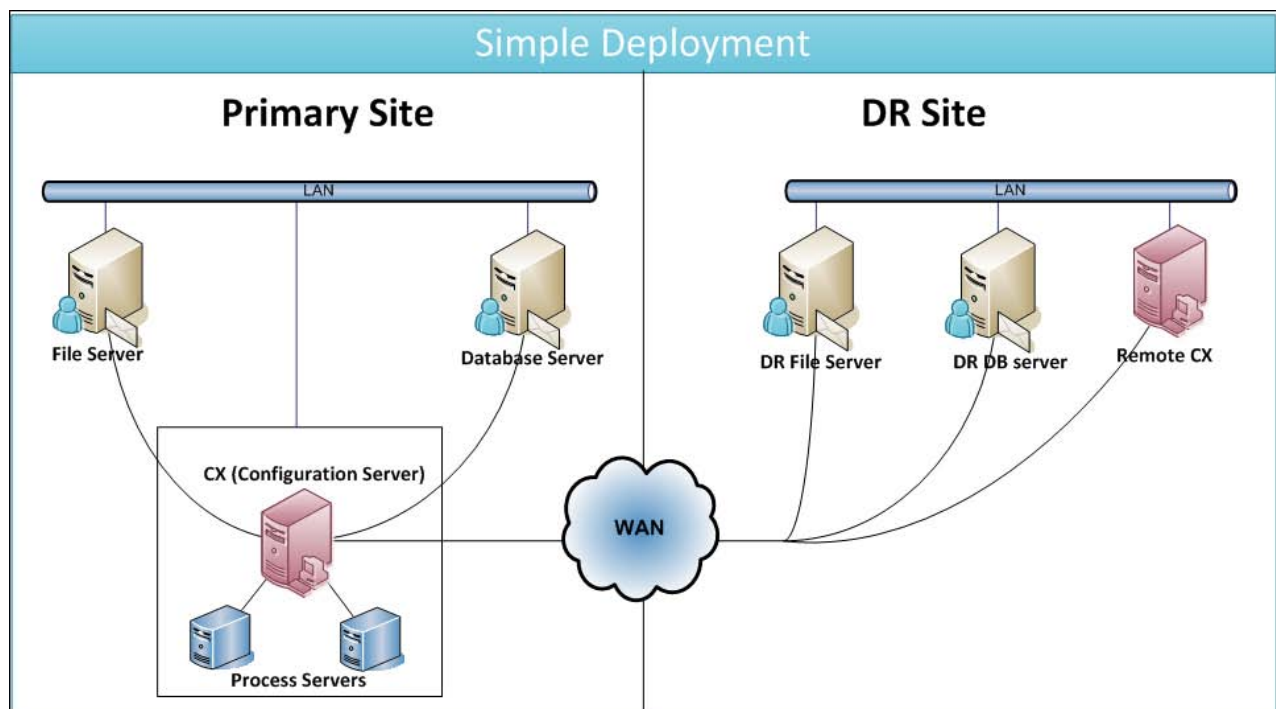


Figure 1: A typical Scout deployment

### The CX server (configuration Server) (CX-CS)

Operations such as setting replication pairs, performing recovery operations, License Management, generating statistical reports, log management and status reporting are performed through the CX server's user interface (CX UI).

CX server resides within the same LAN as that of the source host and provides a central web based interface through which the CX server is administered, while target(s) are usually placed over a WAN or a LAN depending on the type of recovery policy in place.

**Process Server (CX-PS)**

The process server is deployed within the same LAN as of the CX server. Each replication pair is assigned a process server which takes care of all the offload activities.

**Host Agents:**

A host agent can either be a VX agent (for volume level protection) or a FX agent (for file and directory protection). Hosts are offloaded by intelligent agents (FX or VX), which do not compete for resources so that business applications are given priority.

**VX Agent:**

VX agents are volume replication agents and perform block level replication. VX agent has a small footprint on the host and moves only byte level differences based on data changes. Features like snapshots (virtual and physical) can be taken on the target host without disrupting the replication.

**FX agent:**

FX agents are file replication agents and perform file and folder level replication. Each of these replications may be scheduled to run at a particular time. FX jobs are used in combination with the VX agent as part of the application solution.

**RX Server:**

The objective of the RX server is to consolidate backup administration from one single interface rather than depend on individual CX UI to monitor several CX servers. You will be able to monitor a group of CX servers on one interface thus simplifying monitoring and saving time. Some of the advantages of using the RX are:

- Centralized monitoring ability for all CX servers.
- Dashboard with centralized and consolidated reports from all CX servers including CX health, bandwidth usage, alerts and license statistics.
- Multi-tenant architecture

## 1.2 Scout Components

### 1.2.1 CX Server (Configuration Server)

The CX server lies in the center of the Production, DR and Process servers and forms the nucleus of the Scout solution. All the administrative tasks such as backup, monitor, restore etc are performed through the CX server's web interface. Additionally the CX server also generates reports, graphs, email and trap alerts.

CX Server may be installed on a wide range of platforms, please refer to the installation guide for platform support and detailed installation steps.

As the picture shows below, all the process servers, VX and FX hosts communicate with the CX server on a constant basis. The process server (PS) is used to offload the CX server for VX replication pairs. You may choose to install the PS and CX within the same machine; however for better scalability a separate PS is preferred.




Figure 2

A list of services that run on the CX server as given in the table below

Table 1

Service	Purpose
Timeshot manager monitor services	Required to monitor health of replication activity, dispatches mails, generates graphs from rrd/log files sent by PS, monitor agent licenses etc.
Apache	Required for agents to post information about replication activity. Serves as a control path for a replication activity.
Mysql, Scheduler	MYSQL DB stores all the metadata information. Scheduler is used for setting up FX jobs on configuration server.

## 1.2.2 Process Server (PS)

Process servers are either Windows or Linux machines deployed along with the CX servers to handle all the offload activity such as FTP for data transfer; generate log files for graphs, compress data on its way to the DR servers, etc. The PS is assigned to a replication pair while configuring a volume replication. Each PS may handle many VX replication pairs depending on the available resources. The PS is in constant communication with the CX server, production and DR servers.

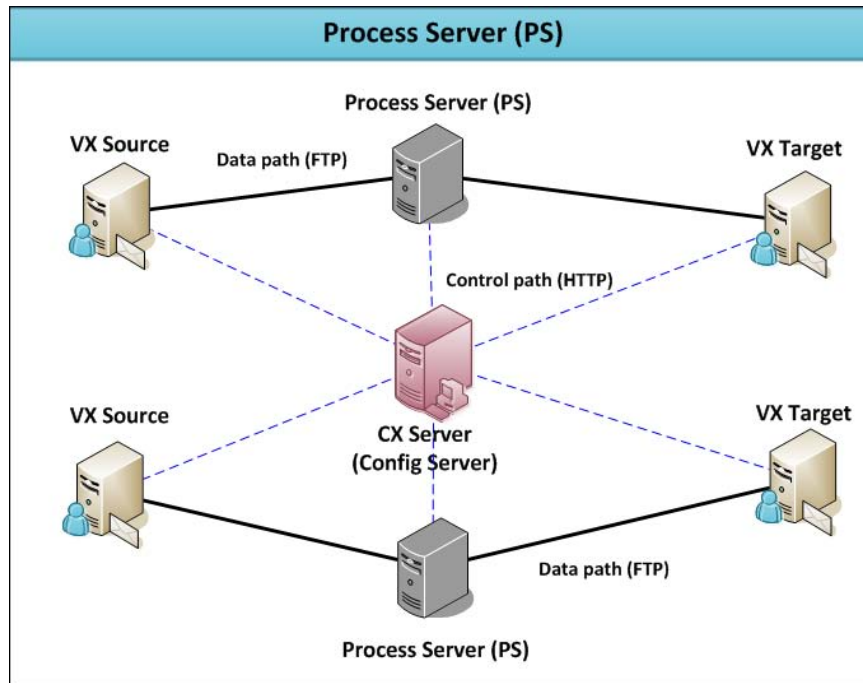


Figure 3

A list of services that run on the PS as given in the table below

Table 2

Service	Purpose
Timeshot manager monitor service	Required to register the process server with Configuration server, monitor health of replication activity from source/target to process server.  Generate rrd/log files for bandwidth trending/data transfer activity.
Volume synchronization processes	Required to process differentials/resync files sent by agent
Proftpd Service	Required to facilitate FTP transfer from source / Target to PS

### 1.2.3 Volume Replication –VX Agent

The VX agent is a software component installed on the production machine and its corresponding DR machine to enable volume level replication between them. The VX agent constantly communicates with the CX server and the respective process server to perform continuous backup. The VX agent protects the source host at a volume level by continuously replicating the protected source volume to the target.

Once you set a replication pair the replication pair goes through stages. The first stage is called “**Initial Sync Step 1**” where all the data on the disk is replicated to the corresponding target. The second step is “**Initial Sync Step 2**” where all the data changes occurred during step 1 are replicated to the target. The third stage is “**Differential sync**” where all the writes are intercepted by the VX agent and replicated to the target. Differential Sync is a continuous process. Since volume replication is done at a block level the file system of the target changes to that of the source.

#### VX components

Once the VX agent is installed a service named “**svagents**” is created. This service runs two threads “**dataprotection.exe**” and “**s2.exe**”.




Figure 4: VX agent composition while acting as source.

The components required by the VX agent differ depending on the role played by the VX agent, i.e. source VX or target VX. Vacp consistency tags, s2.exe and the driver will not be used on the target VX agent. The target VX agent uses cache manager and dataprotection.exe to update the data changes from the source




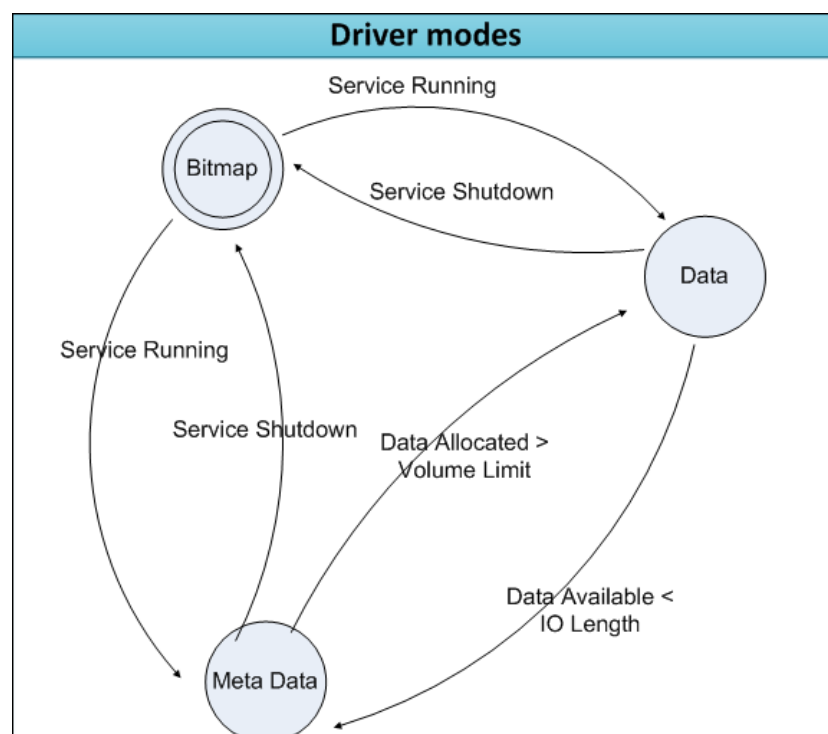
Figure 5

**Driver:** The driver is a low level component that intercepts disk writes and passes them over to S2.exe. The driver operates in three different modes depending on the data change rates and available resources

**Table 3**

	Data Mode	Meta Data mode	Bitmap mode
Data capture	The driver captures the offset, length and the data itself.	The driver captures the offset and length of the changed data.	The driver captures data changes and stores them within a sparse file. Each bit in this sparse file represents the chunk of data on the volume.

The following picture shows how the driver shifts modes to adapt to large data change rates or to handle resources.



**Figure 6**

**Dataprotection.exe:** On the production server dataprotection.exe is responsible for replicating all data on disk to the DR server. It is used while the replication pair is in **"Initial Sync Step 1"**. On the target server dataprotection.exe is responsible for both replication and recovery.

**S2.exe:** This process runs only on the production server and starts along with dataprotection.exe. S2.exe works in sync with the driver to replicate real time writes happening to the production volume.

**VACP Consistency Tags:** Also called as bookmarks, consistency events or consistency tags. These consistency tags are generated either manually or through scheduled jobs on the production machine

through the “**vacp.exe**” command line utility found under the VX agent installation path. These tags act as markers on the disk to which a recovery operation may be performed.

Cache manager (target VX): The cache manager is used by the target VX agent to move data from the CX-process server, while the dataprotection.exe applies these data changes to the target volume and then deletes the files from the cache store. The retention logs on the target host are also updated by dataprotection.exe in the process. The picture below shows the differentials on CX-process server being moved to the cache store on the target host and then applied to the target volume.

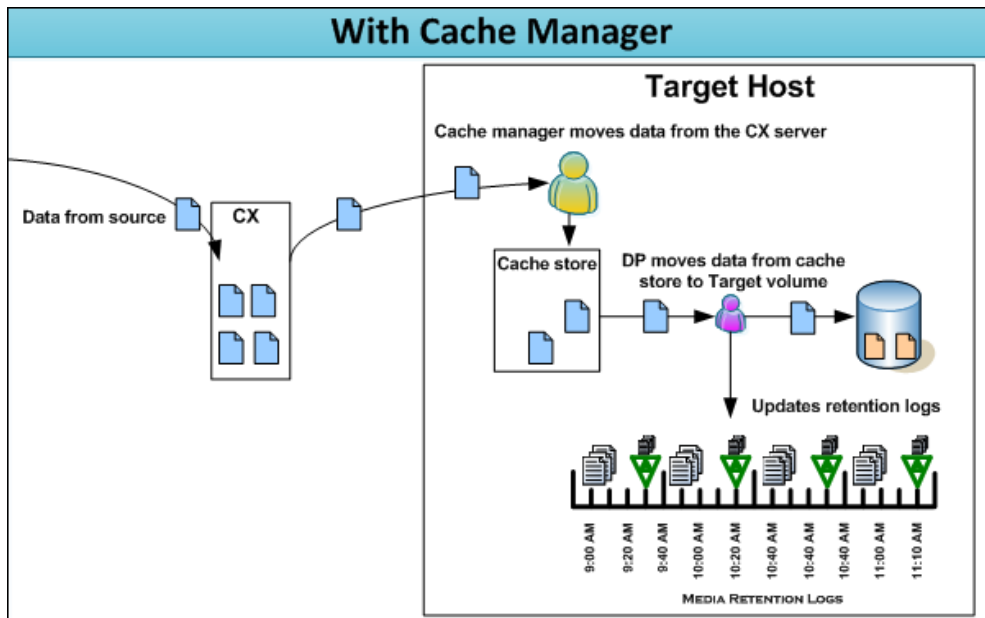


Figure 7

The amount of differential data on the CX-process server and on the target host may be seen through the CX-configuration server's UI on the "[Protection Status](#)" screen on page 147.



### VX Agent Communication

The source and target VX agents never communicate with each other directly. All communication between the source and the target VX agents happen through the CX-process server. This means that the source and target VX agents are unaware of each other. However when a “**Direct Copy**” is performed, the VX agent performs replication within the same host

The VX agents communicate configuration and status information to the CX-configuration server over the HTTP protocol. The default out of the box configuration uses HTTP over standard port 80. The source and target VX agents use FTP/FTPS as the data transfer protocol to send and receive data from the assigned CX-process server. The target VX agents additionally can open connections to port 873 of the CX-process server when using offload resync.




Figure 8:

HTTP and FTP/FTPS protocols were explicitly chosen as the communication protocols due to their firewall friendly nature. Almost all stateful firewalls come with rules predefined for both HTTP and FTP/FTPS, i.e. there is no need to open a range of ports specifically.

## 1.2.4 File Replication – FX Agent

Scout FX File Replication allows:

- Asynchronously replicating complete file systems from Production server (source machine) to DR Server (target machine).
- Scheduling jobs through the CX-configuration server's interface involving the source and or the target host.

The FX agent is used in combination with the VX agent for protecting applications like SQL server, Exchange server etc.

To install FX agent log in as an administrator or any account with admin privileges. Refer to the installation guide for detailed installation steps.

### FX Components

Once the FX agent is installed a service named “**frsvc**” is created on windows platform and on other platforms the service name will be “**svagents**”.




Figure 9:

### Scheduler

The FX scheduler is a CX-configuration server's user interface component through which you may schedule a particular FX job to execute on a repetitive basis. To learn more about scheduler refer to the section [Group Scheduling modes](#) on page 105.

### Pre Script

Prescript is an option within the FX execution. By using the prescript option a script file (batch file or a shell script) can be passed on to the FX job that is executed before the actual FX replication begins.

### Post Script

Post script is again optional and can accept a script file which executes after the FX replication is complete. To learn more about FX execution sequence refer to the section [FX execution](#) on page 86.

### File replication

This component performs file or directory replication. You will be presented with many options on the CX-CS UI to tweak the replication behavior.

### FX Agent Communication

The FX agents communicate configuration and status information to the CX-CS server over HTTP protocol. The source and target FX (unlike the VX) agents communicate directly to each other. The data transfer occurs directly between the source and target FX agents without involving the CX-PS. The data transfer protocol by default, is a single socket connection to port 874 of the source or target. The port number can be customized under the FX “Job Options”.




Figure 10:

## 1.2.5 CDP Retention Option

### What is CDP Retention Option?

CDP retention Option technology enables you to roll a volume back in time or take a snapshot of a protected volume back in time. CDP retention option is specific for volume replication pairs. This option can either be set while setting up the replication pair or can be enabled or disabled at a later point.



Figure 11: CDP retention option on the CX-CS UI

### What happens when CDP retention option is enabled?

If this option is enabled, a screen is displayed where you are prompted to enter location of retention logs, and the type of policy to manage the logs. The VX agent on the target host intercepts each write received from the CX-process server then updates the target volume(s) accordingly and generates corresponding retention logs.

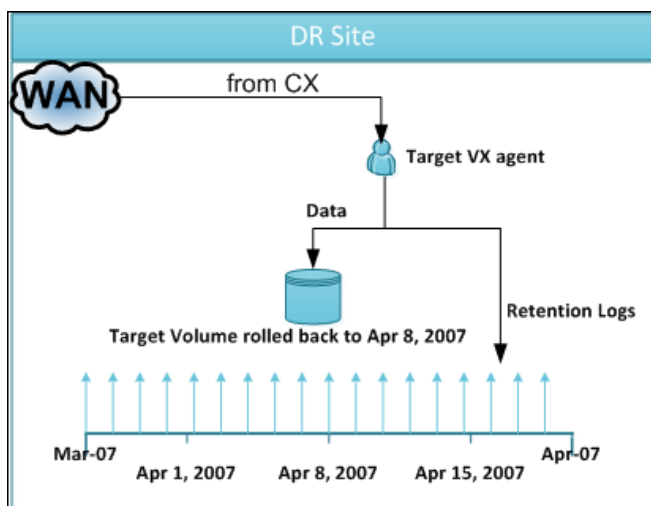


Figure 12: CDP retention option logs being generated

These retention logs store information about data changes within a time period. This time frame is referred to as the retention window. Any rollback or back in time snapshot can be made only within this retention window.

Retention log files grow in size to accommodate data changes happening on the production server. There are three kinds of retention policies that control retention log size time-based, space-based and a combination of both. Depending on the type of policy enforced the retention window is maintained by discarding older data changes within the retention log files to make room for new data changes. To learn more refer to the section [Types of CDP retention option policies](#) on page 57

# Part 2: Protecting data using Scout

File Protection and Volume Protection are explained in this section.

## **Chapter 2**

In this chapter, we look into the CX user interface, and explore basic operations. At the end of this chapter you would be able to

- Log into the CX interface
- Identify CX UI components with their respective functionalities

## **Chapter 3**

Chapter 3 deals with setting up volume replication pairs, CDP retention option, and types of retention policies, and editing a volume existing replication pair.

## **Chapter 4**

The FX agent is introduced in this chapter. FX agent is used to replicate files or folders and also used for data consistency jobs.

## 2 Using the Management Console

### 2.1 Introduction to the Management Console

The CX user interface (management console) is web based and therefore has to be accessed through a web-browser. Scout User Interface is best viewed with a resolution of 1024 X 768 and supports Microsoft Internet Explorer 6+ (recommended), and Firefox 1.0+

Through the CX-CS UI you may perform a variety of continuous backup operations at both Volume and file levels. Given below is a high level overview of the CX-CS UI capabilities.

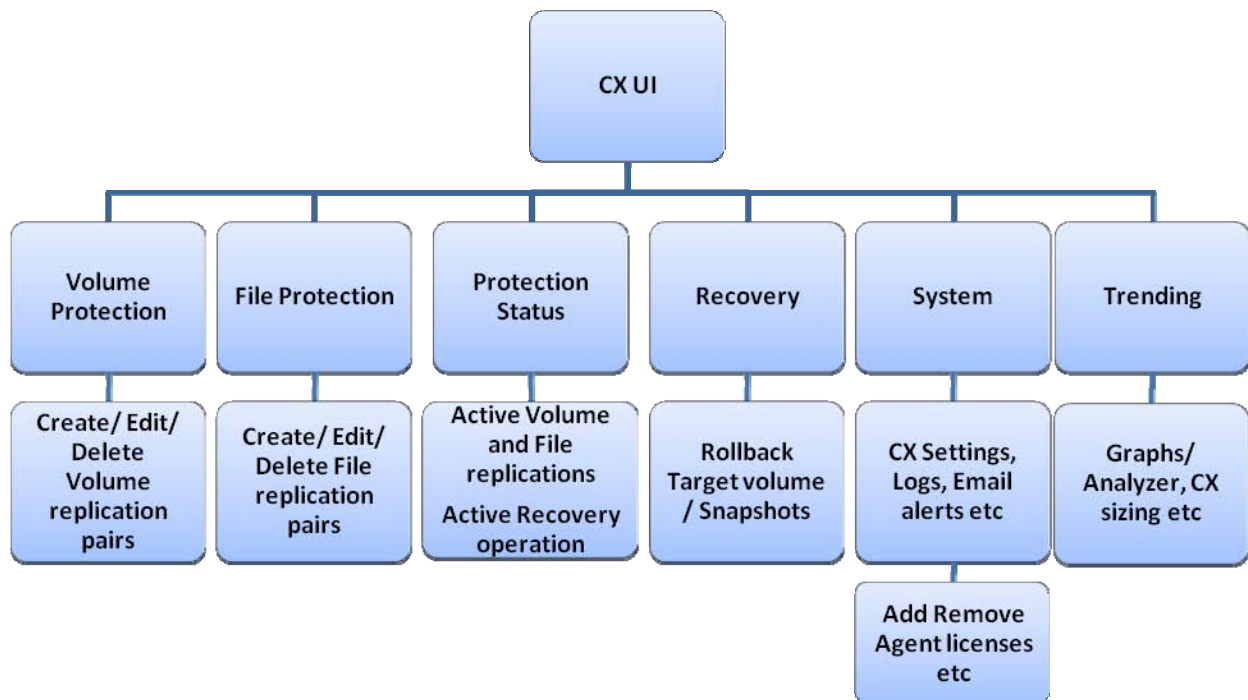


Figure 13:

Through the CX-CS UI you would be able to

- Set File and Volume replication pairs
- Monitor status of replication pairs and recovery operations
- Define bandwidth policies specific to Scout
- Download Logs
- Set email alerts
- Obtain graphs, view health reports.
- Perform recovery operations (target volume rollback, snapshots etc)



#### Caution:

Linux-based CX server relays a loop back address if a fully qualified domain name is not used, so while working with the CX UI always use a fully qualified domain name.

## 2.2 Logging into the CX-CS User Interface

**Step 1.** To access the CX-CS UI, open the browser, and type the CX-configuration server's IP address. For evaluation, enter the user id and password as given below.

User id: "admin"

Password: "password"

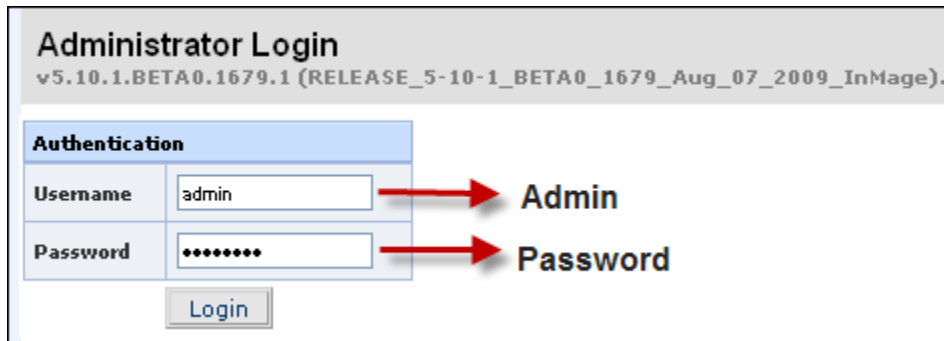
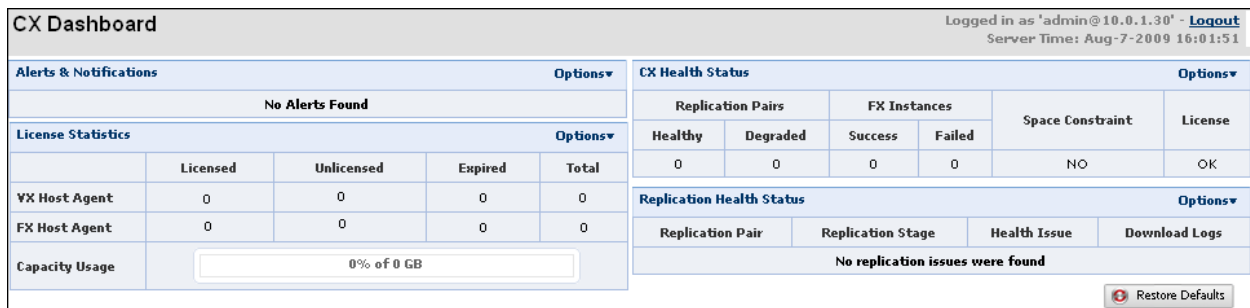


Figure 14: Logging in

**Step 2.** The "CX Dashboard" screen opens up on successful authentication. The CX Dashboard shows four default panels. When you login for the first time, the dashboard screen is blank as shown below.



Alerts & Notifications					Options▼
No Alerts Found					

License Statistics					Options▼
	Licensed	Unlicensed	Expired	Total	
VX Host Agent	0	0	0	0	
FX Host Agent	0	0	0	0	
Capacity Usage	0% of 0 GB				

CX Health Status						Options▼
Replication Pairs		FX Instances		Space Constraint	License	
Healthy	Degraded	Success	Failed			
0	0	0	0	NO	OK	

Replication Health Status				Options▼
Replication Pair	Replication Stage	Health Issue	Download Logs	
No replication issues were found				

Restore Defaults

Figure 15: After logging in

**Step 3.** After logging in for the very first time, assign license to VX and FX agents as per requirement. To assign license refer to Section [2.2.3 License Management](#) on page [27](#).

## 2.2.1 CX Dashboard

The CX dashboard appears as soon as you login to the user interface. You may also click on the home icon on the top right hand side of the CX UI. This screen presents a comprehensive view of the CX server status such as number of replication pairs, license status, degraded VX or FX replications, alerts etc.

There are four docks in this screen. You may rearrange these docks as desired. Each of the dock contains “Options” specific to it. Through the options menu you can delete the dock, refresh it , minimize or maximize it and even enable auto refresh.

The options menu lists four options as shown in the picture.



Figure 16: Options Tab

- **Refresh:** is used to refresh the dock, each time a dock is refreshed it picks up new data.
- **Settings:** To set auto refresh at regular intervals. To set auto refresh options for any of the dock click “Settings”, this will open the Auto Refresh UI as shown below. Check the “**Auto Refresh**” box, and choose refresh interval from “**Refresh Interval**” drop down menu. It will refresh the required window after specified interval.

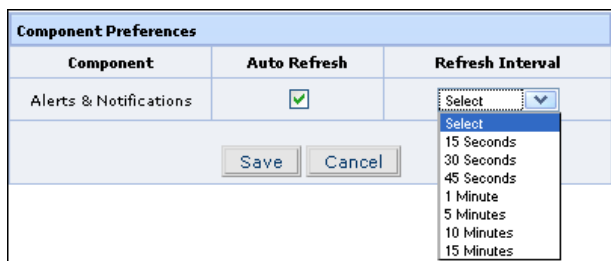


Figure 17

- **Min/Max:** is used to minimize or maximize the dock.
- **Remove:** is used to remove the dock from the screen.

To load defaults for all docks click the “**Restore Defaults**” on dashboard on the bottom of the screen.



### Notes:

The rearrangement of “Dashboard Docks” or configuration of “Option” are system specific. The corresponding information is stored as browser cookies.

“Restore Default” tab helps to resume the default rearrangement of dashboard docks. By default the “Auto Refresh” option not checked in.



## Alerts & Notifications

All alerts and notifications are displayed in this dock. For the list of alerts and notifications refer to the [To add an Administrator or User](#) section on page 165.

Alerts & Notifications		Options ▾
2009-08-19 13:15:25	<b>Agent Down</b> imits111.dev-domain.net : FX Agent has been down for more than 900 seconds. Jobs Affected : [GroupId]:1 [JobId]:1 [Source:Path]=>[Target:Path]; imits111.dev-domain.net-/tmp/source=>[imits111.dev-domain.net-/tmp/target] [JobId]:2 [Source:Path]=>[Target:Path]; imits111.dev-domain.net-/tmp /source=>[imits111.dev-domain.net-/tmp/target] [GroupId]:2 [JobId]:3 [Source:Path]=>[Target:Path]; imits111.dev-domain.net-/tmp/source=> [imits111.dev-domain.net-/tmp/target] (Number of occurrences: 31040 times in last 24 hours)	
2009-08-19 13:15:25	<b>Agent Down</b> LOCALHOST.LOCALDOMAIN : VX Agent has been down for more than 900 seconds. No replication pairs set for the host. (Number of occurrences: 15143 times in last 24 hours)	
2009-08-19 13:15:15	<b>Capacity Threshold Exceeded</b> For license 30chandu_per-cx-1 Utilize capacity is greater than threshold capacity (Number of occurrences: 10700 times in last 24 hours)	

Figure 18

## License Statistics

This dock shows the number of licensed unlicensed and expired license for both VX and FX agents. In case of capacity based license also it shows percentage of capacity usage.

License Statistics				Options ▾
	Licensed	Unlicensed	Expired	Total
VX Host Agent	5	0	0	5
FX Host Agent	1	0	0	1
Capacity Usage	<div><div>84.77% of 2 GB</div></div>			

Figure 19

## CX Health Status

This dock shows the Replication Pairs, FX jobs, RPO Statistics, Space Constraint on the CX server and the state of the CX License.

CX Health Status						Options ▾	
Replication Pairs		FX Instances		RPO Statistics		Space Constraint	License
<div><div></div>4</div>	<div><div></div>7</div>	<div><div></div>0</div>	<div><div></div>0</div>	<div><div></div>4</div>	<div><div></div>7</div>	NO	OK

Figure 20

- All VX replication pairs are shown under “**Replication Pairs**” column, healthy pairs are shown as green and degraded pairs are shown as orange.
- FX Replication pairs are shown under “**FX Instances**” column healthy pairs as green and degraded pairs as orange.
- RPO Statistics shows number of healthy and degraded pairs.
- Space Constraint shows, if any source constraint exist for CX sever.
- License column displays information regarding license of VX and FX.

## Replication Health Status

This dock lists all the VX replication pairs which are degraded on this CX server. You may click on any of the replication pair to view further details. The “**Replication Stage**” shows the status of the replication pair while the “**Health Issue**” shows the reason.

Replication Health Status			Options ▾	
Replication Pair	Replication Stage	Health Issue	Download Logs	
imits111.dev-domain.net : /dev/mapper/ /varun-v1 -> imits111.dev-domain.net : /dev/mapper/varun-v2	Resyncing (Step I)	Not Progressing	Source	Target
imits111.dev-domain.net : /dev/mapper/ /varun-v3 -> imits111.dev-domain.net : /root/volume5	Resyncing (Step I)	Not Progressing	Source	Target
imits111.dev-domain.net : /root/volume6 -> imits111.dev-domain.net : /root/volume4	Resyncing (Step I)	Not Progressing	Source	Target

Figure 21

Health Issues can degrade a replication pair. Given below is the list of health issue. For possible workarounds refer to Table 17 on page 169.

- “**RPO Violation**”: Occurs when the RPO has exceeded the threshold
- “**Resync Required set to Yes**”: When VX replication pairs are in “Differential Sync” with the “Resync Required” field set to Yes
- “**Throttled**”: When the VX replication pair is throttled
- “**Disconnected from Source**”: When source volume is disconnected from the CX-CS server
- “**Not Progressing**”: This occurs when there is no differential data from the source host for the last half hour.
- “**Pause Pending**”: A VX replication is being paused
- “**Paused**”: A VX replication has paused
- “**Source License Expired**”: production server’s license has expired
- “**Target License Expired**”: DR server’s license has expired
- “**Delete Pending**”: When a VX replication is being stopped
- “**Deletion failed**”: When a deleting a VX replication failed

## 2.2.2 CX-CS User Interface Components

The CX UI is divided into two major areas; the left hand side is the command pane containing six icons, each with a specific functionality

1. Volume Protection
2. File Protection
3. Protection Status
4. Recovery
5. System and
6. Trending

Clicking on any of them changes the area to the right also known as the view pane.

**Volume Protection:** Used to set volume replication pairs, edit existing volume replication pairs, or even remove a volume replication.

**File Protection:** Used to set file replication pairs, edit, or delete existing file replication pairs

**Protection Status:** Used to monitor activities such as file replication, volume replication and recovery operations

**Recovery:** Used to perform recovery operations such as snapshots and target volume rollback.

**System:** Used to set or remove licenses to agents, download log files, set bandwidth shaping etc.

**Trending:** Used to view a wide variety of graphs such as compressed data, network trending, etc.

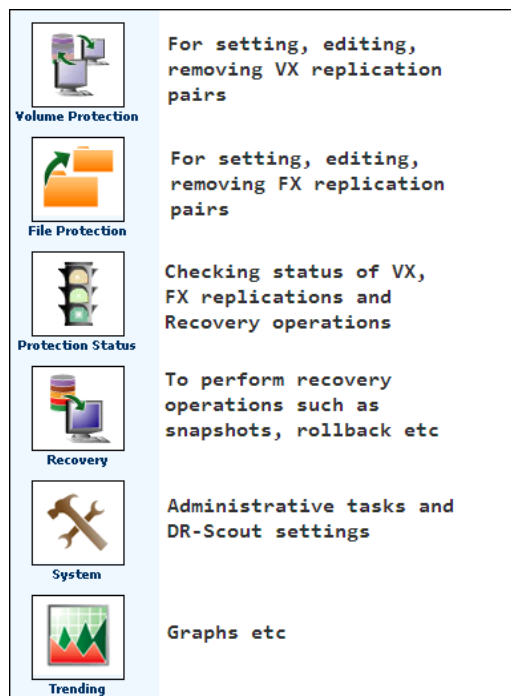


Figure 22:

### 2.2.3 License Management

After logging in to the CX UI for the first time the CX UI is blank. To perform any backup or recovery operation, assign appropriate license to the agent, only then the agent can be used. There are two types of licenses,

- Capacity based license and
- Host based license

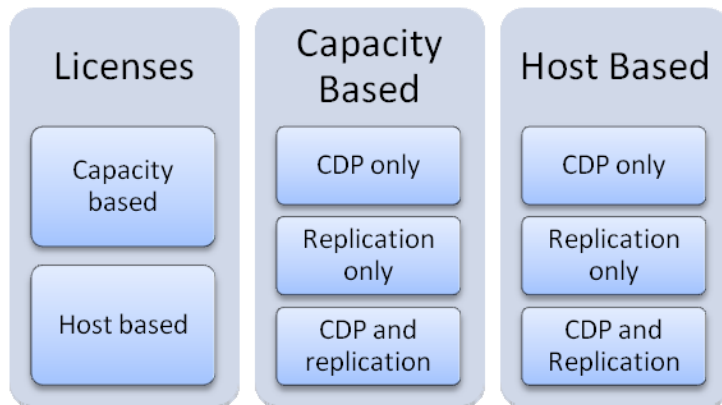


Figure 23:

#### Capacity-based License

A group of hosts are assigned the same license and the combined amount of data protected among these hosts is restricted by the license. For instance, if the license has a capacity of 1 GB it can be assigned to “n” number of hosts; however the combined source volume size of all these hosts cannot exceed the 1 GB limit. Capacity based license is specific to VX agent only.

#### Host-based License

Each host is assigned with a unique license and does not have any limit on the amount of data replicated. FX agent can be assigned a host based license only.

Table 4: Differences between host based and capacity based license

Feature	Host based	Capacity based license
Hosts	One license per host	“N” number of hosts
Amount of data protection	No restriction	Restricted by the license
Expiry	Independent for each host	When CX license expires

## Further Classification

Both host based and capacity based license are further classified into five types

- CDP only
- Replication only
- CDP and replication both
- 1-N replication (dependent) and
- CX HA (dependent)

When using a “**CDP only**” license, features like secure transport and compression would be disabled and while using a “**Replication only**” license, CDP retention option feature is disabled. However while using a CDP and Replication license, secure transport, compression and CDP retention option are enabled.

**Table 5: Features unique for CDP only and Replication only licenses**

Feature	CDP Only license	Replication Only license
Secure Transport	Disabled	Enabled
Compression	Disabled	Enabled
CDP retention option (Includes Recovery Snapshots Features – Rollback, Event and Consistency based Snapshots)	Enabled	Disabled

1-N replication license works in combination with CDP only, Replication only or both CDP and replication license. Without which a 1-N replication cannot be set.

CX HA license works in combination with CDP only, Replication only or CDP and replication license. To learn more about CX HA, refer to the section [CX Cluster](#) on page 202.

## Assigning Licenses

You can assign and release licenses for VX and FX agents through the CX UI. Every agent should be assigned a valid license before it can take part in any replication job. The Scout licensing page can be accessed by navigating to the “**License Management**” tab under the “**System**” menu.

To upgrade to a newer license, upload the license file through the licensing page and the upgrade is automatically done by the CX server.

When you install the CX server, the UI is blank. Navigate to the license management page (accessible through “**System**” -> “**License Management**”). Here you will see all the VX and FX agents pointed to this CX-CS server.

For an agent to take part in a replication you will first need to assign a valid license to it. To assign a license to the agent you will need to upload a license file to the CX-CS server.

System: License Management  
v5.10.1.GA.1701.1 (RELEASE\_5-10-1\_GA\_1701\_Aug\_29\_2009\_InMage).

Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Sep-1-2009 10:50:35

[Monitoring](#) | [Bandwidth Shaping](#) | [CX Settings](#) | [Agent Settings](#) | [License Management](#) | [Agent Heartbeat](#) | [Process Server Traffic Load Balancing](#) | [Process Server Failover](#) | [Remote](#)  
[CX](#) | [Versions and Patches](#) | [User Documents](#) | [Logs](#) | [Install](#) | [RX Settings](#)

License Details

[License file not uploaded]

Host Mac Address

00:12:3f:65:d3:7e

[Capacity Calculator](#)

License Upload

Filename	Action
<input type="text"/>	<div><div>Browse...</div><div>Upload</div></div>

Unlicensed Hosts

Server	License
<input type="radio"/> IMITS088 [vx,fx]	<input type="radio"/> Capacity License
<input type="radio"/> W2K3E321145PAS1 [vx,fx]	<input checked="" type="radio"/> Host License
<input type="radio"/> W2K3E321145PAS2 [vx,fx]	<div><input type="text"/></div>

Set License

Reset

Licensed Hosts

Server	Agent Type	License	Protected Capacity (GB)	Migrate License
<input type="radio"/> InMageProfiler	vx	6335794F-9FFB-42c5-9DF4-85A8E59D0B24	-	
Total:			0	

Release License

Reset

\* Cluster Node

Figure 24: License Management



### Notes:

Upload a new license file when the older license expires.

## License Upload

**Step 1.** Click “Browse”. The Choose File window appears.

**Step 2.** Select the valid license file and click “Upload”. You will be prompted for confirmation, click “ok”

License Details	
<b>[License file not uploaded]</b>	
Host Mac Address	00:12:3f:65:03:7e
<a href="#">Capacity Calculator</a>	
License Upload	
Filename	Action
C:\Documents and Settings\... Browse...	Upload

Figure 25

Once the license is uploaded, the “License Details” shows all the information about the license. Now that the CX server has a license, proceed to assign licenses to the agents (VX and FX)

License Details	
License Upload Date	Sep 01, 2009
License Version	1.2
Customer Name	InMage_Customer
Partner Name	InMage
Host Mac Address	00:12:3f:65:03:7e
License Mac Address	svsHillview
Expiration Date	Permanent
License Count	19 (10 vx, 9 fx)
License Capacity (GB)	1000
<a href="#">Capacity Calculator</a>	
License file uploaded successfully.	
License Upload	
Filename	Action
Browse...	Upload

Figure 26



### Notes:

The old license file is backed up under the directory “/home/svsystems/etc” on CX in case you accidentally uploaded an invalid license file. You can restore it by gaining shell access to the CX-CS.



## Unlicensed Hosts



Unlicensed Hosts	
Server	License
 W2K3E321145PAS1 [vx,fx]	<input checked="" type="radio"/> Capacity License <input type="radio"/> Host License
 W2K3E321145PAS2 [vx,fx]	<div>License: vx, InMage Sy-stem-sinc-cx01-licensetemp1</div>
<div>Set License Reset</div>	

Figure 27:

**Step 3.** Select the type of license, either capacity or host based license. Then select the appropriate VX or FX agent to which you want to assign the license.

**Step 4.** Select an available license from the License drop-down box.

**Step 5.** Click on “Set License”. The page will refresh and add the new VX or FX agent to the list of licensed hosts.

**Step 6.** While assigning the capacity license observe that the license still appears in the license column unlike the host license.



### Notes:

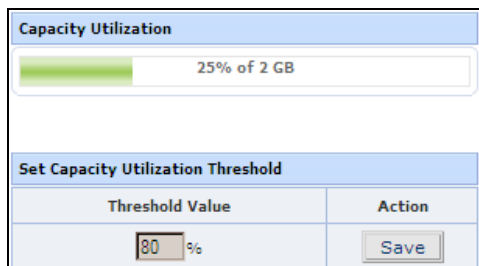
All VX licenses begin with VX as suffix as part of the license string.

An unlicensed host also sends log messages to the CX server like “host down” which may be ignored.

While replication is in progress refrain from using disk management tools on target volume

## Capacity utilization

This shows the amount of capacity consumed by replication pairs. Once the “**Set Capacity Utilization Threshold**” limit is reached, email alerts are sent out (if email alerts are configured). To learn more about email alerts refer to the section [Adding users and email alerts](#) on page 165.



Capacity Utilization	
25% of 2 GB	
Set Capacity Utilization Threshold	
Threshold Value	Action
80 %	<button>Save</button>

Figure 28:

## Licensed Hosts

If a host is no longer required it can be removed from the UI by releasing its license. It will appear under unlicensed hosts (to remove from unlicensed hosts as well, access the agent configuration window and change the IP address of the CX, or uninstall the VX agent)

The Scout licensing page can be accessed, by navigating to the “**License Management**” tab under the “**System**” menu

Licensed Hosts					
	Server	Agent Type	License	Protected Capacity (GB)	Migrate License
	InMageProfiler	vx	6335794F-9FFB-42c5-9DF4-85A8E59D0B24	-	
	W2K3E321145PAS2	vx	InMageSy-stem-sinc-vxSE-licensetemp1 [replication & cdp]	0	<a href="#">Assign Capacity License</a>
	W2K3E321145PAS2	fx	InMageSy-stem-sinc-fxFR-licensetemp6	-	
	W2K3E321145PAS1	fx	InMageSy-stem-sinc-fxFR-licensetemp7	-	
	W2K3E321145PAS1	vx	InMageSy-stem-sinc-cx01-licensetemp1 [replication & cdp]	0	<a href="#">Assign Host License</a>
Total:				0	
<div><button>Release License</button><button>Reset</button></div>					

Figure 29 Releasing Assigning Licenses

**Step 7.** Select the host and agent whose license you want to release.

**Step 8.** Click on “**Release License**”. This will free the license from its agent and add it to the list of free licenses.

**Step 9.** To move an agent from host based to capacity based license, click on the link “**Assign Capacity License**”, similarly to move from a capacity based license to a host based license, select the desired host and click on “**Assign Host license**”

## License Validity and Expiry

A license file is bound to a single CX box. The licensing module validates this by checking the MAC address of one of the interfaces on the host as specified by the customer, with the information in the license file. If there is a mismatch, none of the VX or FX jobs will work. VX agents at source will be throttled, FX agents will not receive their job configuration information and unlicensed hosts will not be allowed to receive a license. A MAC mismatch is alerted to the user under license details and so is a license expiry.

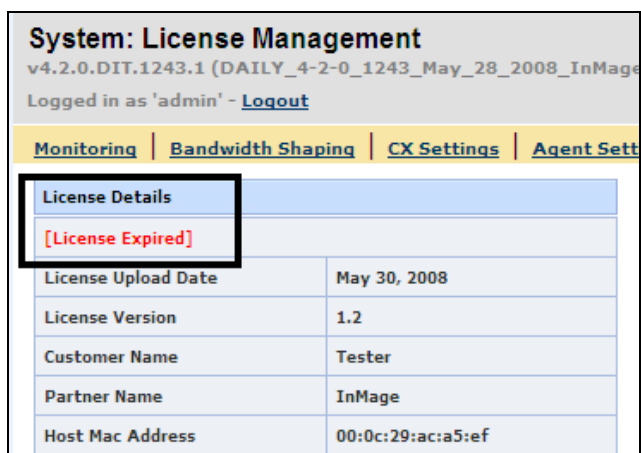


Figure 30: License expired for Capacity-based License

The licensing scheme also supports time limited evaluation licenses. Each license has its own expiration date, after which, the license is unusable. If the license of an agent participating in a replication expires, the replication will come to a halt. License expiration for licensed agents is alerted to the user through red highlighted errors in the license management UI as shown.

You will be alerted about the expiry of the license for a particular agent by a red highlighted error in the Licensed Hosts window as shown below.

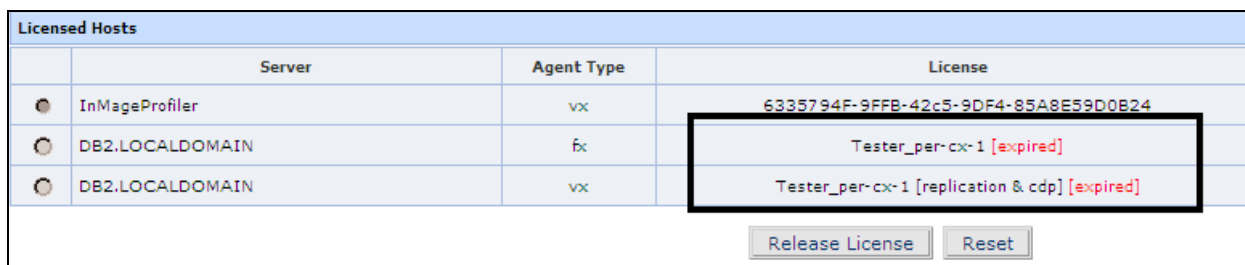


Figure 31 License Expiration Alerts

## Capacity Calculator

A capacity calculator is shown under “**License Details**”. It is used to add up the desired volume capacity. To move to a higher capacity license, identify the list of volumes that are to be protected through the capacity calculator

License Details	
License Upload Date	Sep 01, 2009
License Version	1.2
Customer Name	InMage_Customer
Partner Name	InMage
Host Mac Address	00:12:3f:65:03:7e
License Mac Address	svsHillview
Expiration Date	Permanent
License Count	19 (10 vx, 9 fx)
License Capacity (GB)	1000
<a href="#">Capacity Calculator</a>	

Figure 32

Select the list of volumes and the “**Expected capacity growth**” value as a percentage of the total capacity of selected volumes. The result will be displayed on the screen. You may then make a request for a higher capacity license

Capacity Calculator	
Select the devices you want to protect	<a href="#">Check All</a>   <a href="#">Uncheck All</a>
<input type="checkbox"/> SQL2K8SRC-64	
<input type="checkbox"/> W2K3-TGT-64BIT	
<input type="checkbox"/> PROD-SERV	
<input type="checkbox"/> E:\Mount_prod	
<input checked="" type="checkbox"/> E	
<input checked="" type="checkbox"/> F	
<input checked="" type="checkbox"/> G	
Expected capacity growth	<input type="text" value="5"/> %
<input type="button" value="Calculate"/>	

Figure 33

## Upgrading a License

If a capacity based license is uploaded over an existing host based license, you can upgrade to the new license without disrupting any backup or recovery operations.

On the CX CS UI, click on “**System -> License Management**” then scroll down to “**Licensed Hosts**”. All VX agents with host based licenses will show an “**Assign Capacity License**” link under the Upgrade License column. To upgrade the VX license to capacity based license, just click on the link and the new license is assigned to the agent without disrupting any operation.

Licensed Hosts					
	Server	Agent Type	License	Protected Capacity (GB)	Migrate License
	InMageProfiler	vx	6335794F-9FFB-42c5-9DF4-85A8E59D0B24	-	
	W2K3E321145PAS2	fx	InMageSy-stem-sinc-fxFR-licensetemp6	-	
	W2K3E321145PAS2	vx	InMageSy-stem-sinc-vxSE-licensetemp1 [replication & cdp]	0	<a href="#">Assign Capacity License</a>
	W2K3E321145PAS1	fx	InMageSy-stem-sinc-fxFR-licensetemp7	-	
	W2K3E321145PAS1	vx	InMageSy-stem-sinc-cx01-licensetemp1 [replication & cdp]	0	<a href="#">Assign Host License</a>
Total:				0	
<div>Release LicenseReset</div>					

Figure 34

### 3 The VX Agent

In this chapter you will learn to use the Volume replication agent through the CX-configuration server's user interface. Basic examples in this chapter explain in detail to setup replication pairs, editing settings of an existing replication pair, and even stopping a replication pair. At the end of this chapter you would be able to:-

1. Set Replication Pairs
  - From windows production server to Windows DR server
  - Backup a Microsoft cluster environment through VX agent
2. CDP retention option Technology
  - Type of retention policies
  - Policies
  - To enable or disable CDP retention option etc
3. Edit an Existing Replication Pair
  - Enable or disable CDP retention option
  - Expose a target volume to the DR Server without stopping a replication pair
  - Other settings like forcing a Resync, enabling or disabling encryption.

## 3.1 Volume Protection

The VX agent service is set to run every time the host starts. You can check the status of the agent service through “**services.msc**” on windows host. For Linux hosts, access the console and navigate to the “**bin**” folder under VX agent installation path. To check the status of the VX agent service use the command “**./status**” to stop the service use the command “**./stop**” and to start use “**./start**”. The agent service is very critical and needs to run at all times.

### 3.1.1 When to choose VX Replication

Volume protection is used to replicate data at a volume level, since Scout VX agent performs block level replication it is ideal for protecting enterprise applications or volumes with heavy data change rates.

To setup VX replication pairs

4. Install VX agent on the production and DR Servers
5. Point both VX agents to the same CX-configuration server
6. Assign appropriate licenses to both the agents
7. Ensure that the VX agents are up and running on both the production and DR servers
8. Ensure the VX driver is not bypassed

Through “**Volume Protection**” you can

9. Set Volume replication pairs
10. Make changes to a replication pair (enable CDP retention option or disable it, expose target volume etc)
11. Break or pause a replication pair

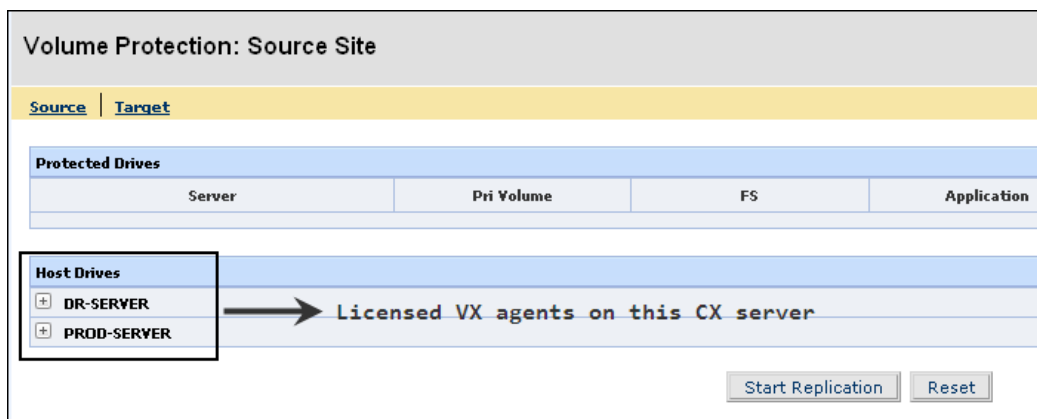


Figure 35: Volume Protection

Click on the “**volume protection**” the “**Source Site**” is displayed. The source site contains the list of VX hosts pointed to this CX-CS server. Expand any of the host to display a list of volumes under it. On a cluster environment “**Cluster Group(s) Volumes**” shows below the hosts, expanding those displays the cluster volumes. Refer the section “[Microsoft clustered volume replication](#)” on page 59 to set replication on a cluster environment.

Cluster Group(s) Volumes							
	Server	Pri Volume	FS	Application	Capacity (Bytes)	Free Space (Bytes)	Replication Status
	Cluster:EXCLUSTER, Group:EXVIRTSERV Servers:EXCLUSTER1	K,L		Unknown	0	0	Inactive
	Cluster:EXCLUSTER, Group:Group 3 Servers:EXCLUSTER1	M		Unknown	0	0	Inactive
	Cluster:EXCLUSTER, Group:Cluster Group Servers:EXCLUSTER1	N		Unknown	0	0	Inactive
	Cluster:EXCLUSTER, Group:Group 0 Servers:EXCLUSTER1	Q		Unknown	0	0	Inactive
<div> Start Replication Reset </div>							

Figure 36:



#### Notes:

The target volume should be either equal or larger than its corresponding source volume

All the data on the target volume will be lost once the replication pair is set

While the replication is in progress the target volume will always be locked unless you manually unhide the target volume

The file system on the target volume changes to that of the source volume

Replication goes through three stages:-

Resync Step 1: All the data on the source volume is replicated to the target volume in this stage

Resync Step 2: All the data changes accumulated during resync step 1 will be replicated to the target volume

Differential Sync: A continuous process where all data changes are intercepted by the VX agent and then transferred to the target host.



### 3.1.2 Profiling

#### Why is InMageProfiler Used?

A production server is “profiled” to gain valuable insights to information such as data change rates, data compressibility, required bandwidth storage requirements etc.

#### How does profiler work?

After installing a VX agent on the production server (or a host that is to be profiled), a CX-CS and PS servers are introduced within the same LAN, and a replication pair is set from the CX-CS UI with the production volume as the source and the target being “**InMageProfiler**”. This enables the CX-CS to track all data changes on the production volume(s) on a continuous basis.

To attain higher levels of accuracy, profiling should span for at least two weeks, ideally spanning the monthly boundary. This allows for gathering a statistically significant amount of data

After Profiling DR planners can answer questions like:-

- What is the total storage capacity required for Disaster Recovery of select applications?
- What is the bandwidth required for a near zero RPO? What is the bandwidth reduction due to compression?
- Does the currently provisioned bandwidth suffice for a continuous DR implementation?
- What is the storage required for desired retention window?
- How are the data changes distributed throughout the day, week or month?
- What is the bandwidth requirement for a certain RPO?

## How to profile a volume?

**Step 10.** Once appropriate license is assigned to the production server agent, the agent appears on the other CX screens enabling you to perform backup operations. Click on **“Volume Protection”**, select the volume that is to be profiled, and click on **“Start Replication”**.

Volume Protection: Source Site

Source

Target

Protected Drives

Server	Pri Volume	FS	Application	Capacity (Bytes)	Frees Space (Bytes)

Host Drives

DR-SERVER

PROD-SERVER

Server	Pri Volume	FS	Application	Capacity (Bytes)	Free Space (Bytes)	Replication Status
PROD-SERVER	E ( Logistics )	NTFS	Unknown	536870400	531504640	Inactive
PROD-SERVER	F ( Human Resources )	NTFS	Unknown	536870400	531504640	Inactive
PROD-SERVER	G ( Information Technology )	NTFS	Unknown	536870400	531504640	Inactive
PROD-SERVER	H ( Engineering )	NTFS	Unknown	1073740800	1065686016	Inactive
PROD-SERVER	I ( New Volume )	NTFS	Unknown	1073740800	1065686016	Inactive

Start Replication

Reset

Logged in as 'admin@10.0.1.30' - Logout

Server Time: Aug-11-2009 19:18:55

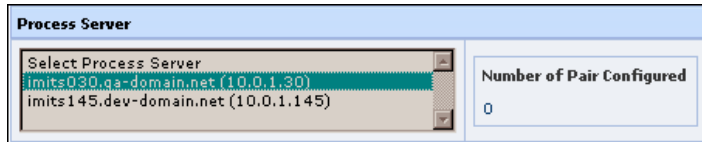
Figure 37: Selecting Volume for Profiling

**Step 11.** The **“target site”** page opens up, select **“InMageProfiler”**, and scroll down to select the **“Process Server”**.

Host: PROD-SERVER Drive: E  Capacity: 536870400				
Select a target volume				
	Server	Volume	Capacity (Bytes)	Free Space (Bytes)
+ DR-SERVER				
+ PROD-SERVER				
	InMageProfiler	P	81920000	81920000

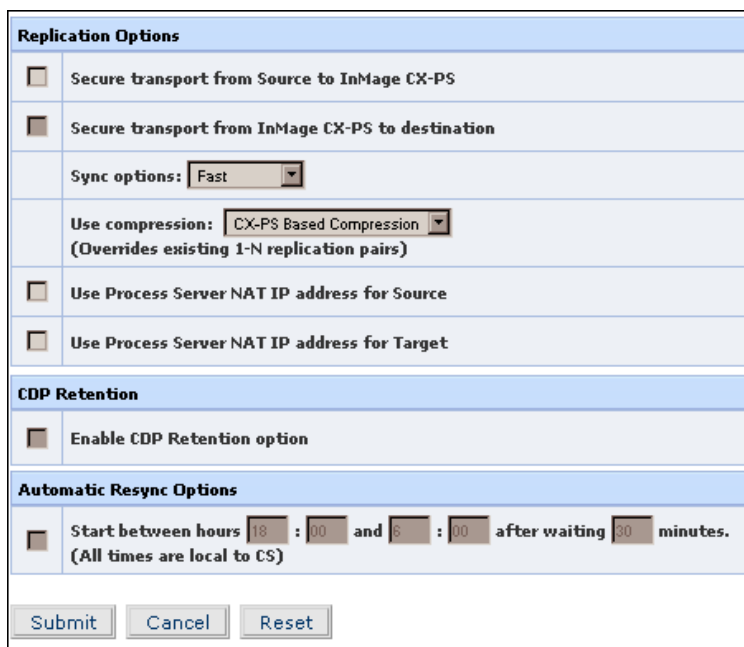
Figure 38: Target as InMage Profiler

**Step 12.** Select the desired “**Process Server**” that will handle this replication pair.



**Figure 39**

**Step 13.** Most of the “**Replication options**” are disabled because the volume is being profiled, scroll down, and click on “**Submit**”. Unlike a normal replication which starts from “**Resync Step 1**”, the profiling pair starts from “**Differential Sync**” mode.



**Figure 40:**

**Step 14.** While a volume is being profiled it cannot act as a source volume for any other target. You may monitor the status through the “**Protection Status**” screen. Observe that the replication pair starts directly from the “**Differential Sync**”

Protection Status									
Logged in as 'admin@10.0.1.30' - <a href="#">Logout</a> Server Time: Aug-11-2009 19:21:52									
25 Pairs <span>▼</span> <span style="float: right;">Page 1 of 1</span>									
Volume Protection Status									
Server	Volume	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)		Resync progress	RPD	Status	Resync Required
				On CX-PS	On Target				
PROD-SERVER->InMageProfiler	E (Logistics) -> p	N/A	N/A	0	N/A	N/A	0	Differential Sync	NO <span>⊕</span>

**Figure 41:**

### 3.1.3 Differences between Profiler and Plain VX Replication

Table 6:

	InMageProfiler	Plain VX replication
Target	CX-CS server	Another host with VX agent
Replication Status	Starts and remains in Differential Sync	Resync step 1-> Resync step 2-> Differential Sync
Data	Only differential data is replicated	All data including data on disk is replicated
Objective	To size the production environment	To protect data.

### 3.1.4 Preparing Target Volume for Replication

#### Preparing Windows volumes for replication

Ensure that you disable “write caching on the disk” on all windows DR servers (target host)

After installing agents on the source and target (on windows platform), you should next optimize the source volumes. This step is optional and may be omitted, but by doing this, the bootstrap process and subsequent replication will achieve better compression and thus complete sooner. You only need to optimize the source volumes that you intend to replicate. Other volumes should not be optimized.

To optimize the source volume, you should invoke the executable file “**volopt.exe**” (found in the agent installation folder), the volume optimization tool. This is recommended when the source volume has considerable amount of free space. This reduces the duration of the initial sync and result in higher compression levels. However, if the source volume is low on free space, this step can be ignored.

This tool takes a single drive letter as an argument. The time it takes to complete is proportional to the amount of unused space on the volume. You may estimate the time remaining by looking at the progress indicator.

This tool creates a file on the specified volume; this file takes up all the free space on the volume leaving only 1 GB (by default) free space and then deletes the file automatically. This helps the VX agent to perform a relatively faster initial sync. Alternatively, you can also specify the name of the file, and amount of free space to be left unused. Multiple volopt processes may be run in parallel. Contention for the disk controller may reduce the expected savings in time.

The target volume on the remote host does not require optimization. It suffices to just perform a full format (as opposed to a quick format) of the target volume before starting replication. Ensure that you do not intend to use the target volume for any purpose other than replication, since all data will be lost during the format.

## Preparing Linux volumes for replication

Linux mount accepts a label to indicate a block special device containing a file system. If you opt to use these labels to identify the replicated devices on the source system, they may result in a conflict of names on the target. An example is when both the source and the target systems address the root volume by label and then you choose to replicate the root volume on to another device on target.

Since Label is part of file system, it too gets replicated and now the target has two devices with the same label. In this case, the boot process gets confused about which one to use. This can be the case for any volume with a Label. But `"/root"` and `"/boot"` label conflict could result in a system not coming up on reboot.

Here is an example of labels being used with entries in `"/etc/fstab"`.

<code>LABEL=/boot</code>	<code>/boot</code>	<code>ext3</code>	<code>defaults</code>	<code>1</code>	<code>2</code>
<code>LABEL=/root</code>	<code>/</code>	<code>ext3</code>	<code>defaults</code>	<code>1</code>	<code>1</code>



### Notes:

This situation does not arise when you are not using labels for any of the replicated devices. If otherwise you can resolve this in two ways.

### Solution 1

Use unique labels on the target system. This approach ensures that source and target have no common labels. Since each file system has its own command to set and get labels, Refer to their documentation.

**Example:** `e2label` for `ext3` and `resiserfstune` for `reiserfs`.

### Solution 2

Stop using labels and use explicit device names. This approach solves the problem by identifying the block devices with complete path names, e.g. `"/dev/sda1"`, `"/dev/sda2"` etc.

Here are a few files that use labels at boot time. So ensure that these are updated with either the unique labels or device names depending on the approach taken.

Example of files using Labels:

`"/boot/grub/grub.conf"` (RedHat)

`"/boot/grub/menu.lst"` (SUSE)

`"/etc/fstab"`

Examples:

**Solution 1:** The procedure below explains how to set unique labels.

Sample of `“/etc/fstab”` on target before the change

```
# This file is edited by fstab-sync - see 'man fstab-sync' for details
LABEL=/          /                ext3      defaults    1 1
LABEL=/boot      /boot          ext3      defaults    1 2
```

Steps involved in preparing the target volume (UNIX/ Linux specific)

**Step 15.** Read current labels based on the device names

```
[root@rut-flt-tgt ~]# e2label /dev/sda1
/boot
[root@rut-flt-tgt ~]# e2label /dev/sda2
/
```

**Step 16.** Change the labels

```
[root@rut-flt-tgt ~]# e2label /dev/sda1 /target_boot
[root@rut-flt-tgt ~]# e2label /dev/sda2 /target_root
```

**Step 17.** Modify the `“/etc/fstab”` file to reflect the new labels just set.

```
# This file is edited by fstab-sync - see 'man fstab-sync' for details
LABEL=/target_root /                ext3      defaults    1 1
LABEL=/target_boot /boot          ext3      defaults    1 2
```

**Step 18.** Modify the file `“/etc/grub.conf”`. Before changing the label

```
title Red Hat Enterprise Linux AS (2.6.9-42-test.ELsmp) root (hd2,0)
kernel /vmlinuz-2.6.9-42.ELsmp ro root=LABEL=/ initrd /initrd-2.6.9-
42.ELsmp.img
```

**Step 19.** Replace the old labels with ones just created.

```
title Red Hat Enterprise Linux AS (2.6.9-42.ELsmp) root (hd0,0)
kernel /vmlinuz-2.6.9-42.ELsmp ro root=LABEL=/target_root initrd
/initrd-2.6.9-42.ELsmp.img
```

**Solution 2:** The procedure below explains how to use device path names. Modified “/etc/fstab”.

```
. # This file is edited by fstab-sync - see 'man fstab-sync' for
details
/dev/sda2      /                ext3    defaults    1 1
/dev/sda1      /boot           ext3    defaults    1 2
```

Modified “grub.conf”

```
title Red Hat Enterprise Linux AS (2.6.9-42.ELsmp) root (hd0,0)
kernel /vmlinuz-2.6.9-42.ELsmp ro root=/dev/sda2 initrd /initrd-
2.6.9-42.ELsmp.img
```

## Preparing Solaris volumes for replication

Before installing the unified agent on Solaris machine, ensure that you create the partition from the cylinder 1 rather than the cylinder 0 on both the production and DR servers. The replication will not complete when partition is made from cylinder 0.

### 3.1.5 Before setting up Replication Pairs

VX supports volume level replication for Windows and Linux partitions. Due to this cross platform support you may replicate from and to any of them. Heterogeneous combinations like Windows source to Linux target and Linux source to windows target are quite possible. Setting a replication pair is similar across all of them. Additionally, you may replicate from and to mount points as well. Recovery operations can be performed from the CX-CS UI and also through the command line interface. More of command line interface can be found under “[cdplici](#)” section on page 239.

Replications can be set across platforms, however recovery operations differ. In this section, you will see how to setup replication pairs.




Figure 42:

#### Windows VX Agent

Windows VX agent gets installed under “C:\program files\InMage systems” by default. You may change this while installing the VX agent. Once installed, a reboot is required for the VX agent to fully function. Then switch to the CX-CS UI where the VX agent is pointed to and assign a license.



#### Caution:

On windows machines, ensure that you disable “**write caching on the disk**” on the DR server (target host)

On Solaris platforms, always use the “**shutdown**” command instead of “**reboot**” as the later is not graceful and would also mark all the Solaris replication pairs for a resync.

Example of commands not to use on Solaris machines:

- **Reboot - shutdown -y -g0 -i6**
- **Shutdown - shutdown -y -g0 -i0**
- **Power off - shutdown -y -g0 -i5**

To avoid resync across shutdown/bootup for vxfs file-system, download the “**vx\_freeze\_thaw**” binary from the Hitachi Data Systems Web Portal and place the binary under the appropriate directory depending upon the operating system version.

- Solaris 9, place the binary under “**/usr/bin**” directory.
- Solaris 10: place the binary under “**/lib/svc/method**” directory.



### 3.1.6 Setting Volume Replication

Once volume optimization is complete on the source volumes, format the corresponding target volume (equal or larger than the source volume in size), and then you may start replicating the source volume to the target volume. Scout supports cross platform volume replication, so you can set volume replication pair from Window source to Window target, Window source to Linux target, Linux source to Window target, and Linux source to Linux target. The process for volume replication across platform is same, so the volume replication steps explained below is with the help of window source to window target environment, which can be used for cross platform replication.

#### Step 20. Choose Source Volume

To begin replication, you will need to map the source volume onto the target volume from the CX-CS UI. Navigate to “**Volume Protection->Source Site**”. Expand the source host to select a volume that is to be replicated and click on “**Start Replication**” at the bottom of the screen.

Volume Protection: Source Site

Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Aug-11-2009 19:24:10

Source | Target

Protected Drives

Server	Pri Volume	FS	Application	Capacity (Bytes)	Frees Space (Bytes)
--------	------------	----	-------------	------------------	---------------------

Host Drives

☐ DR-SERVER  
☒ PROD-SERVER

	Server	Pri Volume	FS	Application	Capacity (Bytes)	Free Space (Bytes)	Replication Status
	PROD-SERVER	E ( Logistics )	NTFS	Unknown	536870400	531504640	Inactive
	PROD-SERVER	F ( Human Resources )	NTFS	Unknown	536870400	531504640	Inactive
	PROD-SERVER	G ( Engineering )	NTFS	Unknown	536870400	531504640	Inactive
	PROD-SERVER	H	NTFS	Unknown	1073740800	1065686016	Inactive
	PROD-SERVER	I ( New Volume )	NTFS	Unknown	1073740800	1065686016	Inactive

Start Replication

Reset

Figure 43: Selecting the source

Ensure that the selected volume is not offline (See in the “Replication Status” column.). The “**application**” field in source site and target site shows the applications installed on those volumes, for now Scout supports MS Exchange 2003 and MS SQL 2000 (Scout is aware of these applications and provide greater support for these applications such as application consistency, failover, fallback etc).

## Step 21. Choose Target Volume

After clicking “Start Replication” button in the above screen you will be prompted to choose target volume. Choose the required volume from the listed volumes.

Host: PROD-SERVER Drive: E Capacity: 536870400					
Select a target volume					
	Server	Volume	Capacity (Bytes)	Free Space (Bytes)	In Use?
<input type="checkbox"/> DR-SERVER					
<input checked="" type="radio"/>	DR-SERVER	E ( New Volume )	536870400	531483648	NO
<input type="radio"/>	DR-SERVER	F ( New Volume )	536870400	531483648	NO
<input type="radio"/>	DR-SERVER	G ( New Volume )	1073740800	1065665536	NO
<input type="radio"/>	DR-SERVER	H ( New Volume )	1073740800	1065665536	NO
<input type="checkbox"/> PROD-SERVER					
<input type="radio"/>	InMageProfiler	P	81920000	81920000	NO

Figure 44: Selecting the target

Target volumes that are smaller than the source volume size will be grayed out. The target volume should be larger or the exact same size as of the corresponding source volume.

Select the WAN volume to which the replicated data will be written. In this screen, you complete the source – target mapping.



### Caution:

While replication is in progress, the target volume will be locked and cannot be accessed by any other program. Windows 2008 prompts for a format of the target volume which should be ignored.

All the data in the target volume will be overwritten, with no recourse for recovery

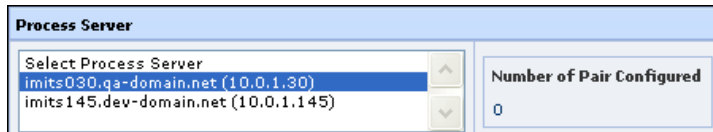
The target volume should never be the boot volume of the remote host, or the volume where the system's paging file is located. Overwriting these volumes leads to an unbootable remote system

## Step 22. Set Replication Options

### Selecting Process Server

All the process servers pointed to the CX-CS server are listed here; you may choose a desired process server which will handle the entire offload activities specific to this replication pairs. For better scalability you may point more process servers to the CX-CS server.

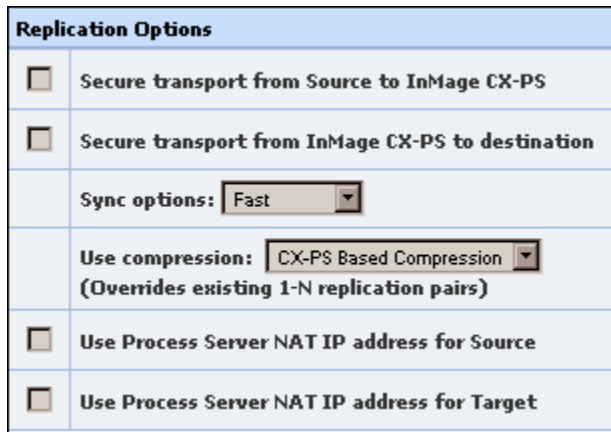
Select the process server and scroll down to set the “**Replication Options**”



Process Server	
Select Process Server	
imits030.ga-domain.net (10.0.1.30)	▲
imits145.dev-domain.net (10.0.1.145)	▼
Number of Pair Configured	
0	

Figure 45: Process Server

### Setting Replication Options



Replication Options	
<input type="checkbox"/>	Secure transport from Source to InMage CX-PS
<input type="checkbox"/>	Secure transport from InMage CX-PS to destination
Sync options:	Fast
Use compression:	CX-PS Based Compression (Overrides existing 1-N replication pairs)
<input type="checkbox"/>	Use Process Server NAT IP address for Source
<input type="checkbox"/>	Use Process Server NAT IP address for Target

Figure 46

- “**Secure Transport from Source to InMage CX**”: Check this option to encrypt data transmission from source host to CX-PS.
- “**Secure transport from InMage CX to destination**”: Check this option to encrypt data transmission from CX-PS to remote target.

Refer to the section [Using a lower encryption](#) on page 317 to switch to a lower encryption.



#### Notes:

However, it is important to remember, the encryption can have a certain degree of performance penalties compared to unencrypted transmissions. The performance impact shall be limited to CX-PS and Target.

If CX-PS and the target are on the same system then do not enable compression for better performance

## Sync options

There are three sync options Fast, Offload and Direct copy

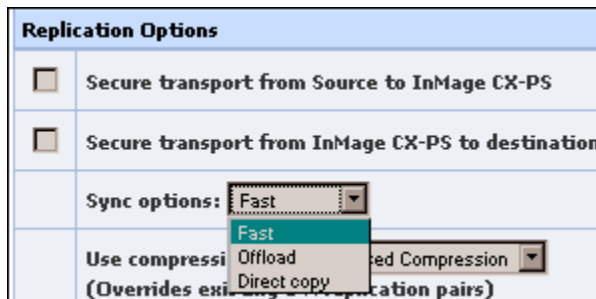


Figure 47

- **“Fast”**: This option performs a faster Resync at the cost of source host’s CPU resources.
- **“Offload”**: On the other hand performs a slower Resync with relatively less source host CPU cycles.
- **“Direct Copy**: Direct Sync is used when both source and target volumes reside on the same machine. This is mainly used for migrating data between similar or heterogeneous storage hardware, such as data migration between DAS to SAN etc. Data can be quickly migrated without any impact on the network, since the data gets replicated directly from the source disk to target disk. Only the differentials are sent to CX-PS and then back to target disk. A checksum of the data is not computed during initial Resync to minimize the CPU impact on the host. Also, number of threads used for direct sync is configurable, giving better control on the speed of the initial resync. This allows you to throttle down or up your I/O usage. Direct copy is used to reduce the RPO value by physically transporting the target volume to a remote host and then reusing the same volume for its corresponding source. Any disk controlled by production server including a different storage subsystem may be used as a target volume here, thus suitable for data migration between arrays, or even DAS to SAN etc. CPU usage for **“Direct Copy”** is expected to be cumulatively lower than a normal Resync in most cases since a checksum is not performed. However, the I/O is performed much faster, and could as well compete with the production I/O.

## Use Compression

The drop down menu for Use Compression provides three options such as **“No Compression”**, **“CX-PS Based Compression”**, **“Source Based Compression”**. Choose first one to replicate data without compression. Choose second one for data compression by CX and choose third one for source based data compression.

Replication Options	
<input type="checkbox"/>	Secure transport from Source to InMage CX-PS
<input type="checkbox"/>	Secure transport from InMage CX-PS to destination
Sync options: <span>Fast</span>	
<input type="checkbox"/>	Use compression: <span>CX-PS Based Compression</span> (Overrides existing)
<input type="checkbox"/>	Use Process Server

Figure 48



### Caution:

Do not attempt [cdpcli command](#) based recovery operations while the replication pair is in **“Reync”** mode. Recovery should only be performed on replication pairs in **“Differential Sync”**

## Process Server NAT IP address

At times the process servers are placed over different networks from the production or DR servers as shown in the picture below. The agent cannot communicate with the process servers in this case. This issue is addressed by assigning a NAT IP address for the process server through the CX-CS UI.

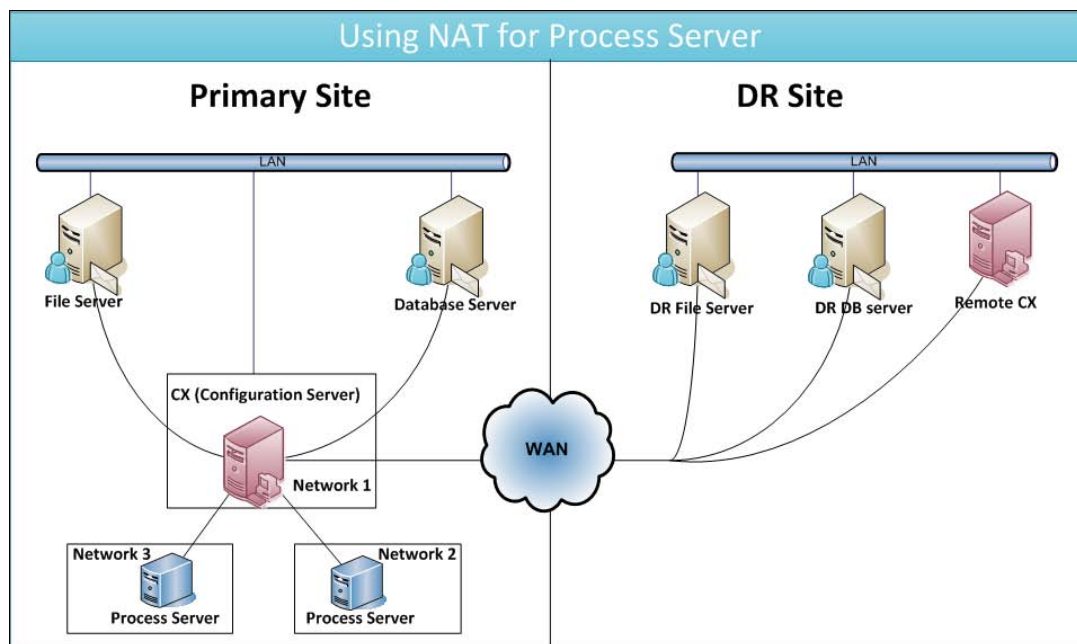


Figure 49

To enable the source agent to communicate with the process server on a different network enable the option **“Use Process Server NAT IP address for source”**. Similarly to enable target hosts to communicate with process servers on different networks, enable the option **“Use Process Server NAT IP address for Target”**.

<input type="checkbox"/>	Use Process Server NAT IP address for Source
<input type="checkbox"/>	Use Process Server NAT IP address for Target

Figure 50

You may assign a NAT IP to a process server under the **“System -> Agent settings”** screen. Select the desired process server then enter a NAT IP and click on **“Change Settings”**

Process Server		
	IP Address	NAT IP Address
<input checked="" type="radio"/>	imits030.qa-domain.net (10.0.1.30)	10.1.145.38
<input type="radio"/>	imits145.dev-domain.net (10.0.1.145)	

[Change Settings](#)

Figure 51

### CDP Retention Option

By default “**Enable CDP retention option**” is enabled. This gives roll back capabilities to the replication pair. You may choose to disable to perform a plain volume replication. Refer to the [CDP retention option](#) section on page 20 to learn about CDP retention option functionality.

CDP Retention	
<input checked="" type="checkbox"/>	Enable CDP Retention option

Figure 52: CDP retention option

### Automatic Resync Options

Automatic resync is used when a replication pair is required to address data inconsistencies automatically. During replication if there is any inconsistency from either of the agents (source or target) a “**resync required**” (can be seen after a replication is set under protection status) field is set to “**Yes**” indicating a force resync to be done at a later time (requires manual intervention).

Automatic Resync Options	
<input checked="" type="checkbox"/>	Start between hours <input type="text" value="18"/> : <input type="text" value="00"/> and <input type="text" value="6"/> : <input type="text" value="00"/> after waiting <input type="text" value="30"/> minutes. (All times are local to CS)
<input type="button" value="Submit"/> <input type="button" value="Cancel"/> <input type="button" value="Reset"/>	

Figure 53:

When the “**Resync Required**” field is set to “**Yes**” for a replication pair, automatic resync is performed when you configure “**Automatic Resync Options**” without any manual intervention. The replication pair waits for a certain period of time (by default it is 30 minutes as shown above) and then it performs a forced resync within the time frame specified, this ensures data consistency and minimizes manual intervention.

Click “**Submit**” to continue.

**Step 23.** The “CDP Retention option” generates retention log files on the target host. This screen requires you to define a policy how these logs should be handled.

**Volume Protection: Retention Options**  
 Logged in as 'admin' - [Logout](#)

Pair Details			
Server	Pri Volume	Remote Server	Volume
PROD-SERVER	E	DR-SERVER	E

Retention Logging Policy			
Retention Policy	Roll-backward		
Retention Log Size	0.00 (MB)	Current Retention Log Size	0.00 (MB)
Unused Space	256.00 (MB)		
Retain changes upto	256 MB (Cannot be less than 256 MB)		
Retain changes upto the (time)	(Days) (hrs.)		
On insufficient disk space	<input checked="" type="radio"/> Purge older retention logs <input type="radio"/> Pause differentials		
Log data directory	G:\retention_logs (Eg:- K:\log_data) G,H are drives suggested for storing rollback log files.		

Configure Threshold for Alerts	
Alert when disk space utilization reaches	80 %

Suggested Volumes For Retention Logs	
Volume	Available Space (MB)
G ( New Volume )	760
H ( New Volume )	760

**Figure 54:**



**Notes:**

A Resync required is set to “Yes” when the CX server detects any inconsistency on the target volume.

CDP retention option logs are created freshly after each Resync

Email alerts will be sent out for the above two scenarios if configured accordingly.

Ensure that retention logs are stored on a NTFS partition (if target is a windows system) or a native Linux file system if the target is a Linux system.

Ensure enough free space for retention logs



## Types of CDP Retention Option Policies

You may handle these retention logs through three different modes:

- Space-based policy
- Time-based policy
- Composite policy (blend of both)

**Space-Based policy:** Space-based retention policy ensures that all latest changes amounting to at least 256 MB (or the amount of memory you choose to enter) are stored into retention logs, thus enabling a rollback to any point in time within the memory limit. Sometimes this can span over months (less writes to disk), or lasts for few minutes (aggressive writes to source volume). Enter the space (minimum of 256 MB) and then enter the location of the retention logs (recommended drives will be shown on the UI). Ensure that there is

- Enough free space on the retention log volume and
- It is an NTFS volume (if on windows)
- It is a Linux native file system (if on Linux)

**Time-Based Policy:** To opt for this policy, do not enter the memory limit; just enter the time in days, and hours for which you wish to retain changes. All the changes that fall in the specified time gets stored into the retention logs and the space is adjusted automatically. The space occupied for retention logs under the time-based policy is shown as 0 on the CX UI since there is no space restriction here. You may observe this under the “**volume protection -> Target -> View Details**” screen

**Composite Policy:** Using time and space together is called composite policy. The matrix is shown below

**Table 7: CDP Retention Option Policy**

	Memory (for e.g. 256 MB)	Time (for e.g. 1 day)	Result
Condition	Memory threshold hit	Time threshold not hit	256 MB worth logs stored (will not overshoot allocated space)
Condition	Memory threshold not hit	Time threshold hit	Logs for specified time (1 day) are available despite having free space (Will not overshoot time limit)

If CDP retention option is not enabled the replication starts immediately (once you click on “**Submit**” under “**replication options**”). Retention log folders need to be excluded from virus scanners to avoid rise in RPO time.



### Notes:

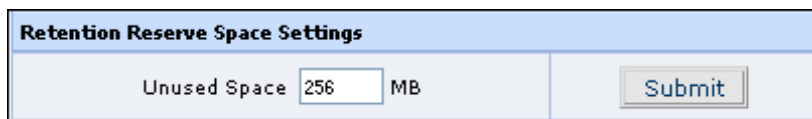
Retention logs are pruned periodically to make space for newer data changes in favor of older data changes. When logs about older changes are pruned, virtual snapshots which fall into the same time range are deleted.

The Fields in CDP retention option are:

**Retention Policy:** Displays the type of policy in place

**Retention Log Size:** A minimum of 256 MB is required

**Unused space:** Amount of free space to be left out while allocating retention logs, this can be configured through “System -> Agent Settings -> Retention Reserve space settings”, Enter the required amount of space to be left unused on the retention volume and then click on “Commit”



Retention Reserve Space Settings	
Unused Space	<input type="text" value="256"/> MB
<input type="button" value="Submit"/>	

Figure 55:

**Retain changes up to:** A maximum limit of retention logs for this replication pair, as and when there are more changes to be accommodated, older changes are removed to accommodate new changes. Filling in days and hours rather than size will make it a time based retention policy.

**On Insufficient disk space:** On low disk space on the retention log volume, the retention engine deletes older logs to make space for new changes, this is the default behavior. Instead you may choose to stop applying differentials to the target until free space is available.

**Log Data Directory:** This is where the retention logs are stored as files; always use a NTFS volume to store retention logs.



#### Notes:

The volume containing CDP retention option logs will not be displayed on the CX-CS UI. This is to safe guard the retention logs from accidental deletion.

CDP retention option logs will always be stored on the Target host.

Retention logs cannot be stored on a FAT volume so be sure to give a non FAT volume for retention logs while on windows platform

While storing retention logs on a Linux target host, FAT, NTFS and FUSE file systems should not be used, always use a native file system to store retention logs

You should avoid creating retention log on network drives or other drives that may become unavailable. Select a volume where you can guarantee the disk space selected will always be available.

On Linux target hosts, ensure that the retention log volume has its entry in the “/etc/fstab” so that it is mounted every time when the target is rebooted. When this entry is not there in the “/etc/fstab” then the retention log volume will not be mounted after a reboot and new retention logs will be created using the “home” space. To avoid loss of retention logs the entry must also be made in “/etc/fstab”.

**Step 24.** Click on “**Submit**” to start the replication pair. On successful volume replication pair setting, the pair will be displayed in “**Protection Status**” as shown below. The status of the replication start from “**Resyncing (Step I)**” then moves to “**Resyncing (Step II)**” before reaching “**Differential Sync**”

Protection Status									
<div> <div>25 Pairs</div> <div> <div>Page 1 of 1</div> </div> </div>									
Volume Protection Status									
Server	Volume	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)		Resync progress	RPO	Status	Resync Required
				On CX-PS	On Target				
PROD-SERVER->DR-SERVER	E ( Logistics ) -> E	0	0	0	0	N/A	3.73 minutes	Differential Sync	NO

**Figure 56: Protected Drives**

### 3.1.7 Microsoft Clustered Volume Replication

Clustering is about high availability and Scout has the ability to work-in a Microsoft Clustering Environment

#### Microsoft Clustering

A Server cluster is a collection of independent servers that together provide a single and highly available platform for hosting applications

#### Server Cluster nodes

Server Cluster nodes can be either active or passive.

- **Active Node.** When a node is active it is actively handling requests
- **Passive Node.** When a node is passive, it is idle, on standby waiting for another node to fail/ Multi-node clusters can be configured using different combinations of active and passive nodes.
- **If an active node fails and there is a passive node available,** application and services running on the failed node can be transferred to the passive node. Since the passive node has no current workload, the server should be able to assume the workload of another server without any problems (providing all servers have the same hardware configuration).
- **If all servers in a cluster are active and a node fails,** the applications and services running on the failed node can be transferred to another active node. Since the server is already active, the server will have to handle the processing load of both systems. The server must be sized to handle multiple workloads or it may fail as well.

The following are the sequence of steps to use Scout in a Microsoft clustering environment

12. Install and configure Microsoft clustering (Ensure proper cluster setup)
13. Failover all active instances of applications to one single node (Active)
14. Install VX agent on the passive node (without any active instances of applications)
15. Reboot each node after installation
16. Failback to the node with VX agent and repeat the process on other nodes
17. Setup replication pairs as required.

Scout VX agent detects volumes that are part of a Microsoft Cluster. You will notice the same on

## “Volume Protection Page: Source Site”

**Step 25.** Click on “Volume Protection”, the “Source site” shows “Cluster Group(s) Volumes”, expand the list, select the appropriate cluster, and then click on “Start Replication”.

Cluster Group(s) Volumes							
	Server	Pri Volume	FS	Application	Capacity (Bytes)	Free Space (Bytes)	Replication Status
<input checked="" type="radio"/>	Cluster:EXCLUSTER, Group:EXVIRTSEV Servers:EXCLUSTER1	K,L		Unknown	0	0	Inactive
<input type="radio"/>	Cluster:EXCLUSTER, Group:Group 3 Servers:EXCLUSTER1	M		Unknown	0	0	Inactive
<input type="radio"/>	Cluster:EXCLUSTER, Group:Cluster Group Servers:EXCLUSTER1	N		Unknown	0	0	Inactive
<input type="radio"/>	Cluster:EXCLUSTER, Group:Group 0 Servers:EXCLUSTER1	Q		Unknown	0	0	Inactive

Figure 57

**Step 26.** This opens up the next screen, select the required volume, and then click on “Next”.

**Volume Replication: Cluster Setup**

Cluster: EXCLUSTER  
Cluster Group: EXVIRTSEV

Configured Cluster Drives			
Primary Server	Primary Volume	Remote Server	Remote Volume
<input type="button" value="Finish"/> <input type="button" value="Remove"/> <input type="button" value="Reset"/> <input type="button" value="Cancel"/>			

Unconfigured Cluster Drives						
	Server	Pri Volume	FS	Capacity (Bytes)	Last YX Sentinel Heartbeat Time	Replication Status
<input checked="" type="radio"/>	EXCLUSTER1,EXCLUSTER2	K	NTFS	1071627264	0000-00-00 00:00:00	Inactive
<input type="radio"/>	EXCLUSTER1,EXCLUSTER2	L	NTFS	1071627264	0000-00-00 00:00:00	Inactive

Figure 58

**Step 27.** This opens up the “Target Site” screen, select the target volume, and then scroll down to set the “Replication Options”.

Cluster: EXCLUSTER  
Cluster Group: EXVIRTSEV  
  
Drive: K  
Capacity: 1071627264

Select a target WAN volume					
	WAN Server	Volume	Capacity (Bytes)	Free Space (Bytes)	In Use?
<input 5"="" data-kind="parent" type="button" value="+&lt;/input&gt;&lt;/td&gt;&lt;td data-cs="/> EXCLUSTER1					
<input 5"="" data-kind="parent" type="button" value="+&lt;/input&gt;&lt;/td&gt;&lt;td data-cs="/> EXCLUSTER2					
<input type="button" value="−"/>	BAKP-SERV				
<input checked="" type="radio"/>	BAKP-SERV	G (New Volume)	1073740800	1065644032	NO

Figure 59

**Step 28.** Select the PS and check the “Enable CDP retention option” and then click on “Submit”.

Process Server	
Select Process Server	Number of Pair Configured
ITGTR5U2-64.inmage.in(10.0.164.73)	2
w2k3-32PPS(10.0.227.227)	
SR5U2-64.inmage.in(10.0.248.73)	

**Figure 60**

Replication Options	
<input type="checkbox"/>	Secure transport from Source to InMage CX-PS
<input type="checkbox"/>	Secure transport from InMage CX-PS to destination
Sync options:	Fast
Use compression:	CX-PS Based Compression
(Overrides existing 1-N replication pairs)	
<input type="checkbox"/>	Use Process Server NAT IP address for Source
<input type="checkbox"/>	Use Process Server NAT IP address for Target

CDP Retention	
<input checked="" type="checkbox"/>	Enable CDP Retention option

Automatic Resync Options	
<input type="checkbox"/>	Start between hours 18 : 00 and 6 : 00 after waiting 30 minutes. (All times are local to CS)

Submit	Cancel	Reset
--------	--------	-------

**Figure 61:**

**Step 29.** In this screen, you can define the type of retention policy for this particular replication pair.

Enter the retention values, for spaced-based, enter the amount of space that can be used for retention logs, (minimum space is 256 MB and it should not be on a FAT file system). The appropriate drives will be suggested on the UI. Once done setting up the retention policy, click on “**Submit**” to start the replication pair.

**Volume Protection: Retention Options**  
 Logged in as 'admin' - [Logout](#)

Pair Details			
Server	Pri Volume	Remote Server	Volume
EXCLUSTER1,EXCLUSTER2	K	BAKP-SERV	G

Retention Logging Policy			
Retention Policy	Roll-backward		
Retention Log Size	0.00 (MB)	Current Retention Log Size	0.00 (MB)
Unused Space	1.00 (MB)		
Retain changes upto	<input type="text"/> MB (Cannot be less than 256 MB)		
Retain changes upto the (time)	<input type="text"/> (Days)	<input type="text"/> (hrs.)	
On insufficient disk space	<input checked="" type="radio"/> Purge older retention logs <input type="radio"/> Pause differentials		
Log data directory	<input type="text"/> h:\logs (Eg:- K:\log_data) H drive is suggested for storing rollback log files.		

Configure Threshold for Alerts	
Alert when disk space utilization reaches	<input type="text"/> 80 %

**Figure 62**

**Step 30.** The last screen appears, click on “**Finish**” under “**Configured Cluster Drives**” to start the replication pair.

**Volume Replication: Cluster Setup**  
 Cluster: EXCLUSTER  
 Cluster Group: EXVIRTSERV

Configured Cluster Drives				
	Primary Server	Primary Volume	Remote Server	Remote Volume
	EXCLUSTER1,EXCLUSTER2	K	BAKP-SERV	G

**Figure 63**

**Step 31.** You can check the status of the replication pair under “**Protection Status**”. The replication pair starts from “**Resync Sync step 1**” then proceeds to “**Resync Step 2**” and finally reaches “**Differential Sync**”.



**Notes:**

For replicating a System volume (boot volume replication) refer the “boot volume replicating” solution document

### 3.1.8 Editing a Replication Pair

At times you may need to edit a replication pair's settings such as

- Enable or disable CDP retention option for an existing replication pair
- Break or stop a replication pair
- Pause a replication pair
- Make a target volume visible
- Resync an existing replication pair
- Increase threshold for Resync and Differential Sync.
- Enable or disable compression for an existing replication pair
- Enable or disable encryption for an existing replication pair etc

To edit replication pair settings navigate to **"Volume Protection → Target"** then select the replication pair that you want to edit then click on **"View Details"**.



Volume Protection: Target Site							
<div>Source   Target</div>							
Configured replication pairs							
	Server	Pri Volume	Remote Server	Volume	Last VX Agent Heartbeat Time	Profiling Mode	RPO Threshold
	PROD-SERVER	E (Logistics)	DR-SERVER	E	2009-08-11 19:42:11	NO	30
	PROD-SERVER	F (Human Resources)	DR-SERVER	F	2009-08-11 19:42:11	NO	30
<div> <a href="#">View Details</a> <a href="#">Apply Changes</a> <a href="#">Stop Replication</a> <a href="#">Pause Replication</a> <a href="#">Force Delete</a> <a href="#">Reset</a> </div>							

Figure 64:

This opens **"Replication Statistics"** screens. The replication statistics screen is divided into three tabs.

#### "Statistics" tab:

The statistics screen is read only screen where you can see **"Pair Settings"**, **"pair details"**, and **"Retention Settings"**.

Volume Protection: Replication Statistics

Source | Target

Statistics | Reports | Settings

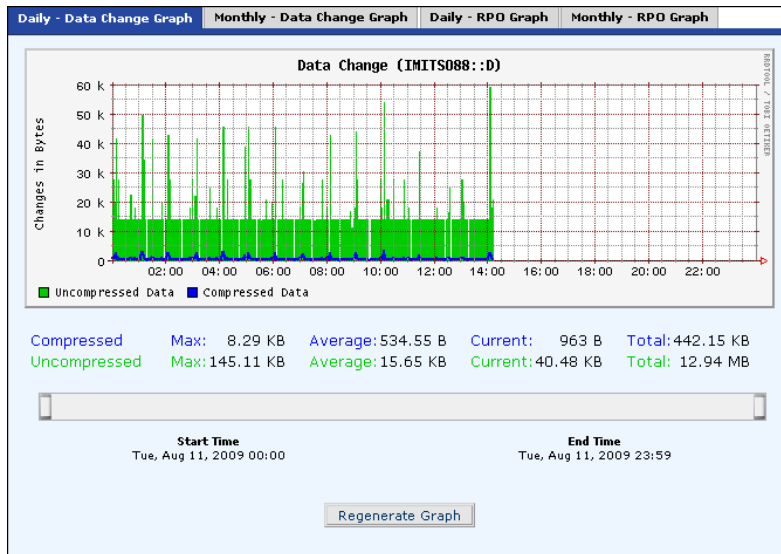
Pair Details										
Server		Primary Volume	Remote Server		Target Volume	Process Server	Replication Pool	Fast Resync Unmatched %	Agent Log	
PROD-SERVER		E ( Logistics )	DR-SERVER		E	imits030.qa-domain.net [ 10.0.1.30 ]	8	4.88	N/A	
Pair Settings										
Visible	Visible Drive Mode	Profiling Mode	Secure CX-PS to Destination	Secure Source to CX-PS	Resync Mode	RPO Threshold	Replication Pool (1-24)	Resync Files Threshold (MB)	Differential Files Threshold (MB)	Compression Enable
No	N/A	No	No	No	Fast	30	8	2048	8192	CX-PS Based
Retention Settings										
Retention	Retention Log size limit (in MB)		Retention Time limit		Log data directory		Disk Space Threshold (%)	Unused Space (in MB)		On insufficient disk space
Enabled	256.00		Not Configured		G:\retention_logs\9616a68f8f		80	256.00		Purge older logs

**Figure 65:**

Additionally you may also see the data change graphs and RPO graphs on daily and monthly basis.

### Daily – Data change Graph

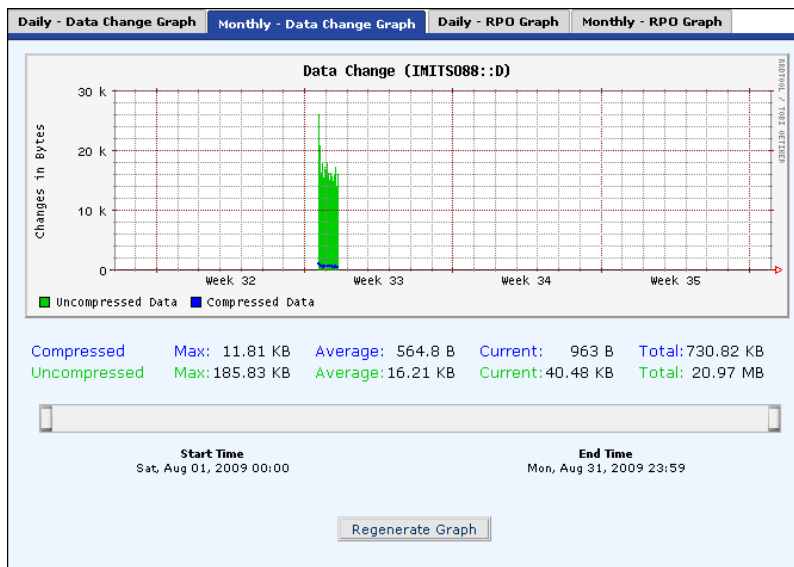
This graph shows data changes on the Y axis and the corresponding time on the X axis. You may use the slider bar to close in on a desired time frame and then click on **“Regenerate Graph”** for the new graph.



**Figure 66**

### Monthly – Data change Graph

Click on the **“Monthly- Data Change Graph”** to get a monthly view. The Y axis shows the amount of data and the X axis shows marks time in weeks.



**Figure 67**





### Daily – RPO Graph

Click on the “**Daily – Data Change Graph**” to see the RPO graph. The Y axis shows the RPO in minutes and the X axis shows the corresponding time. Additionally you may use the sliders to close in on a desired period and then click on “**Regenerate Graph**” to get the new graph.




Figure 68

### Monthly- RPO Graph

Click on the “**Monthly–RPO Graph**” for the monthly RPO graph. Here the Y axis shows the RPO in minutes and the X axis marks time in weeks. You may again use sliders to close in on a desired time and “**Regenerate Graph**”




Figure 69



## “Settings” Tab

Pair details is read only and no values can be modified through this

Volume Protection: Replication Settings

Source

Target

Statistics

Reports

Settings

Pair Details							
Server	Primary Volume	Remote Server	Target Volume	Process Server	Replication Pool	Fast Resync Unmatched %	Agent Log
PROD-SERVER	E ( Logistics )	DR-SERVER	E	imits030.qa-domain.net [ 10.0.1.30 ]	8	4.88	N/A

Figure 71:

Pair settings can modify the following settings

Pair Settings											
Visible	Visible Drive Mode	Resync	Profiling Mode	Secure CX-PS to Destination	Secure Source to CX-PS	Resync Mode	RPO Threshold	Replication Pool (1-24)	Resync Files Threshold (MB)	Differential Files Threshold (MB)	Compression Enable
<input type="checkbox"/>	<input checked="" type="radio"/> Read Only <input type="radio"/> Read-Write	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Fast	30	5	2048	8192	CX-PS Based

Figure 72

18. **“Visible”**, **“Visible Drive Mode”**: Target drive visible (read only or read write mode).
19. **“Resync”**: Force resync. This field is used to resync a replication pair manually.
20. **“Profiling mode”**: Profile the source volume to predict resource requirements.
21. **“Secure CX-PS to destination”**: Enable or disable encryption between CX server and target.
22. **“Secure source to CX-PS”**: Enable or disable encryption between source and CX server.
23. **“Resync Mode”**: Toggle between fast and offload resync.
24. **“RPO Threshold”**: If RPO increases beyond this limit, email alerts are sent to the configured email id.
25. **“Replication Pool”**: The tmanagerd thread that handles this replication
26. **“Resync files Threshold”**: Cache folder size on CX server while replication is in resync. If this limit is exceeded then email alerts are sent to the configured id. If the value is set to zero then replication pauses and the replication will not be throttled while in resync mode.
27. **“Differential Files threshold”**: Cache folder size on CX server while replication is in differential sync. If this limit is exceeded then email alerts are sent to the configured id. By default, the value is 8192 MB (8 GB). This can be customized to suit your environment. Once this threshold is reached, the replication pair is throttled. During throttling, the source stops its transfers to the CX server until all the accumulated differentials drain from the CX server to the target host. If the value is set to zero then replication pauses and the replication will not be throttled while in differential sync.
28. **“Compression Enable”**: Compression mode (either CX-PS based or host based).

Pair Settings											
Visible	Visible Drive Mode	Resync	Profiling Mode	Secure CX-PS to Destination	Secure Source to CX-PS	Resync Mode	RPO Threshold	Replication Pool (1-24)	Resync Files Threshold (MB)	Differential Files Threshold (MB)	Compression Enable
<input checked="" type="checkbox"/>	<input checked="" type="radio"/> Read Only <input type="radio"/> Read-Write	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Fast	30	3	2048	8192	CX-PS Based
<div>Restart Resync</div> <div>Accept Changes</div> <div>Reset</div>											

Figure 73: Observe that the Windows Replication Pair does not show “mount point” column

The target volume can be made visible under read only or read write mode by enabling the “Visible” option. This in turn enables the “Visible drive Mode”, select the Read only or Read-Write as desired and click on “Accept Changes” to confirm changes or “Reset” to cancel changes made. For a Linux target, you will need to enter a mount point (persists even after a reboot)

Pair Settings												
Visible	Visible Drive Mode	Mount Point	Resync	Profiling Mode	Secure CX to Destination	Secure Source to CX	Resync Mode	RPO Threshold	Replication Pool (1-24)	Resync Files Threshold (MB)	Differential Files Threshold (MB)	Compression Enable
<input type="checkbox"/>	<input type="radio"/> Read Only <input type="radio"/> Read-Write		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Fast	0	8	2048	8192	CX Based
<div>Restart Resync</div> <div>Accept Changes</div> <div>Reset</div>												

Figure 74: Mount Point will be enabled once you check the “Visible” check box

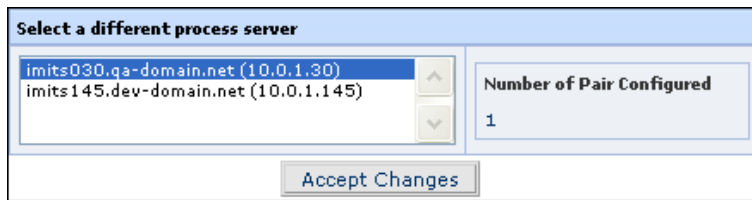
If the target is visible then a red exclamation mark (!) is seen beside the “Server” under “Volume Protection -> Target side”

<div>Statistics</div> <div>Reports</div> <div>Settings</div>							
Pair Details							
Server	Primary Volume	Remote Server	Target Volume	Process Server	Replication Pool	Fast Resync Unmatched %	Agent Log
PROD-SERVER (!)	E (Logistics)	DR-SERVER	E	imits030.qa-domain.net [ 10.0.1.30 ]	8	4.88	N/A
A target volume for this drive is visible. Please lock it before adding a new pair.							

Figure 75:

### Select a different process server

You may change the process server of a VX replication pair here. Select the desired process server and click on “Accept Changes”. This will move the replication pair to the new process server.

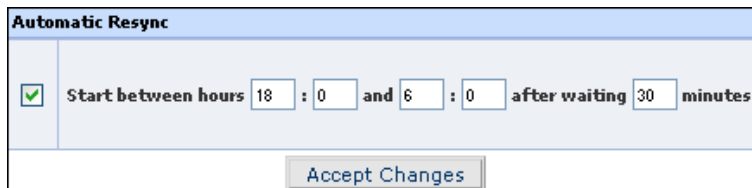


The dialog box has a title bar "Select a different process server". Inside, there is a list box containing two entries: "imits030.ga-domain.net (10.0.1.30)" and "imits145.dev-domain.net (10.0.1.145)". The first entry is selected. To the right of the list box is a label "Number of Pair Configured" with the value "1". At the bottom center is a button labeled "Accept Changes".

Figure 76

### Automatic Resync Options

Automatic resync options can be edited, enabled or disabled for an existing replication pair.



The dialog box has a title bar "Automatic Resync". Inside, there is a checkbox with a green checkmark. To its right is the text "Start between hours 18 : 0 and 6 : 0 after waiting 30 minutes". At the bottom center is a button labeled "Accept Changes".

Figure 77: Edit Automatic Resync

## “Retention Settings” (Retention policy)

Through retention settings screen you can change the following settings

29. Enable or disable retention settings
30. Increasing the size of retention logs
31. Alert when disk space utilization reaches some %
32. Move retention logs from one location to another location on the target server (DR server)

Retention settings have two command buttons “Edit” and “Disable Retention”. Clicking on “Disable Retention” disables CDP retention option for the VX replication pair.

Click on the same button to enable CDP retention option. (If retention is disabled, the caption changes to “Enable CDP retention option”.) Clicking on “Edit” will open up “Retention logging policy”.

Retention Settings						
Retention	Retention Log size limit (in MB)	Retention Time limit	Log data directory	Disk Space Threshold (%)	Unused Space (in MB)	On insufficient disk space
Enabled	256.00	Not Configured	H:\log\ea92838c75	80	256.00	Purge older logs
<a href="#">Edit</a>				<a href="#">Disable Retention</a>		

Figure 78:

## Retention Logging Policy

Retention Logging Policy			
Retention Policy	Roll-backward		
Retention Log Size	256.00 (MB)	Current Retention Log Size	1.91 (MB)
Unused Space	256.00 (MB)		
Retain changes upto	256 MB  (Cannot be less than 256 MB)		
Retain changes upto the (time)	0 (Days)	0 (hrs.)	
On insufficient disk space	<input checked="" type="radio"/> Purge older retention logs <input type="radio"/> Pause differentials		
Log data directory	H:\log\ea92838c75	<a href="#">Edit</a> J,K are drives suggested for moving retention log files.	
Configure Threshold for Alerts			
Alert when disk space utilization reaches			80 %
<a href="#">Save</a> <a href="#">Cancel</a>			

Figure 79:

You may change the alert settings when the retention disk reaches a threshold through the “Alert when disk space utilization reaches”. This is set to 80% by default indicating that if the disk is 80% full then an email alert is sent (this again requires email alerts to be enabled)

There are two ways to define the retention logs

33. Specify a certain size (not less than 256MB)
34. Time-based (not less than 1 hour)
35. Composite policy ( a combination of time based and space based policy)

Time-based policy retains only the latest changes as mentioned in the policy  
 Space-based policy retains changes as long as logs do not exceed the specified size.  
 A composite policy however uses both approaches where you can define both size and time.

**“On insufficient disk space”** on the retention volume: Either older retention logs are deleted to make space for new changes (default behavior) or differentials are paused (based on the option selected).

**“Log Data Directory”**: You may change the Retention log size, Retention time limit, Alert settings and even move retention logs through this screen. Click on the **“Edit”** link to change the location of the retention logs on the target host. You will be asked for a confirmation before moving the retention logs. The corresponding replication pair will be paused while moving the retention logs and will resume once the retention logs are moved.

## Retention Settings

Retention settings can either be edited (increase or decrease size) using the **“Edit”** button or retention can be disabled by clicking on **“Disable Retention”**. By doing this older retention logs cannot be used for recovery operations and may be deleted manually.

Retention Settings						
Retention	Retention Log size limit (in MB)	Retention Time limit	Log data directory	Disk Space Threshold (%)	Unused Space (in MB)	On insufficient disk space
Enabled	256.00	1 day	H:\log\ea92838c75	80	256.00	Purge older logs
			<a href="#">Edit</a>	<a href="#">Disable Retention</a>		

Figure 80: Pair with CDP retention option Enabled

Retention Settings						
Retention	Retention Log size limit (in MB)	Retention Time limit	Log data directory	Disk Space Threshold (%)	Unused Space (in MB)	On insufficient disk space
Enabled	0.00	1 day	H:\log\ea92838c75	80	256.00	Purge older logs
			<a href="#">Edit</a>	<a href="#">Disable Retention</a>		

Figure 81: Pair with CDP retention option Disabled

For time based retention, the **“Retention Log Size limit”** is shown as 0 indicating that there is no space restriction for retention logs here.



### Notes:

Before Enabling / Disabling CDP retention option for a replication pair, stopping a replication pair modifying retention log folder and performing a target volume rollback always unmount all virtual snapshots corresponding to that replication pair



### 3.1.9 Stop or Break a Replication

Select the desired replication pair and click on **“Stop Replication”** to stop the replication. You will be prompted for confirmation, click on **“OK”** to continue to break the replication pair. On clicking **“OK”** You should see a **“Stop Replication Options”** appear on the same screen.

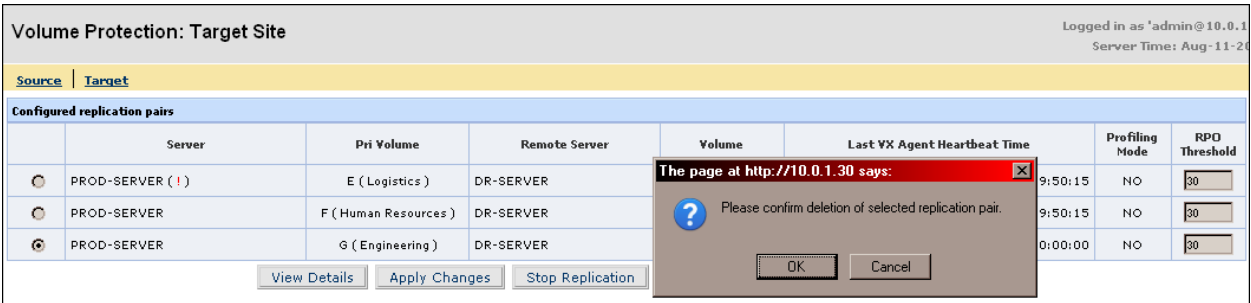


Figure 82: Stop replication, Apply changes, Reset

**“Stop Replication Options”** contains **“Clean CDP retention logs”** and **“Unlock target drive”**. You may select both of them to delete retention logs corresponding to this replication pair and unlock the target volume. However it is not mandatory to select either of the options if you intend perform both these tasks manually at a later time. Click on **“Submit”** to break the replication pair.

The target volume may be unlocked at a later time through the cdpcli.exe utility on the target host. Refer to the section [CDPCLI interface](#) on page 239. Finally, the retention logs may be removed manually to free up disk space. Once the replication pair is stopped, it will no longer appear on the CX-CS UI.




Figure 83:



#### Notes:

While breaking a replication with a Linux target host, you will also be prompted to mount the target volume.

When the pair is in Resync, the “unlock target drive” option will be disabled

When a Solaris source volumes is unmounted while replication, you will need to manually delete the replication pair through the CX UI.

The replication pair may take a few moments before it is deleted from the CX-CS UI, you may notice this from the “**Protection Status**” screen as shown below

Volume Protection Status										
Server	Volume	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)		Resync progress	RPO	Status	Resync Required	View Details
				On CX-PS	On Target					
PROD-SERVER->DR-SERVER	G ( Engineering ) -> I	0	0	0	0	0 %	0.55 minutes	Resyncing (Step 1) [Deletion pending]	YES	
PROD-SERVER->DR-SERVER	F ( Human Resources ) -> F	0	0	0	0	N/A	0.57 minutes	Differential Sync	NO	
PROD-SERVER->DR-SERVER	E ( Logistics ) -> E ( Logistics )	0	0	0	0	N/A	1.78 minutes	Differential Sync (target visible)	NO	

Figure 84

“**Apply changes**” button is used to change the RPO Threshold.

“**Reset**” button resets all values back. (Default value of RPO threshold is 0).



#### Notes:

If the Source volume has been used in other replication pair then, there shall be no new entry listed, however, the entry can be found under the “Volume Protection ->Target Site” or under “Protection Status”.

### 3.1.10Pause Replication

Rather than stopping the replication pair, you may choose to pause it at times. Click on “**Volume Protection – Target Site**” then select the replication pair you desire to pause and click on “**Pause Replication**”, you will be prompted for a confirmation, click on “**OK**” and the replication pair is paused.

Volume Protection: Target Site							
<div> <div>Source</div> <div>Target</div> </div>							
Configured replication pairs							
	Server	Pri Volume	Remote Server	Volume	Last Vx Agent Heartbeat Time	Profiling Mode	RPO Threshold
	PROD-SERVER (1)	E ( Logistics )	DR-SERVER	E ( Logistics )	2009-08-11 19:53:16	NO	<input type="text" value="30"/>
	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	2009-08-11 19:53:16	NO	<input type="text" value="30"/>
	PROD-SERVER	G ( Engineering )	DR-SERVER	I	2009-08-11 19:53:16	NO	<input type="text" value="30"/>
<div> <div>View Details</div> <div>Apply Changes</div> <div>Stop Replication</div> <div>Pause Replication</div> <div>Force Delete</div> <div>Reset</div> </div>							

Figure 85

While the replication pair is paused, others will continue to backup your data. Pause replication is applicable only for the selected pair. While you are replicating one source volume to multiple target volumes, you will be prompted either to pause the replication for the selected target host or for all target hosts.

You may resume the replication pair by selecting the replication pair and clicking on “**Resume Replication**”. Observe that the “**Pause replication**” button will change to “**Resume Replication**”.

### 3.1.11 Force Delete

“Force Delete” is similar to that of “Stop Replication” the following are the differences between them

Table 8

Feature	Stop Replication	Force Delete
Deleting CDP retention logs	You will be presented with an option on the CX UI to delete or retain CDP retention logs	CDP retention logs will not be deleted
Unlocking the target volume	You will be presented with an option to unlock the target volume when the pair is not in resync	Target will remain locked, you may unlock this through the <a href="#">cdpcli command line utility</a> on the target host
Cache directory cleanup	Background process	Does not clean cache directory
Deleting virtual snapshots	Background process	Does not unmount virtual snapshots

When a stop replication fails, or when you intended to forcefully delete a replication pair, navigate through the “Volume Protection” -> “Target” then select the desired replication pair and click on “Force Delete”. You will be asked for a confirmation if the “Stop Replication” fails before, click on “Ok” to continue with force delete

Volume Protection: Target Site

Logged in as 'admin@10.0.1'

Server Time: Aug-11-20

Source | Target

Configured replication pairs

	Server	Pri Volume	Remote Server	Volume	Last VX Agent Heartbeat Time	Profiling Mode	RPD Threshold
	PROD-SERVER ( 1 )	E ( Logistics )	DR-SERVER	E ( Logistics )	2009-08-11 19:53:16	NO	<input type="text" value="30"/>
	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	2009-08-11 19:53:16	NO	<input type="text" value="30"/>
	PROD-SERVER	G ( Engineering )	DR-SERVER	I	2009-08-11 19:53:16	NO	<input type="text" value="30"/>

View Details

Apply Changes

Stop Replication

Pause Replication

Force Delete

Reset

Figure 86

## 3.2 Additional notes for Solaris

### 3.2.1 Protecting zpool on Solaris

- To protect a zpool, you will need to replicate all the underlying disks/partitions/logical volumes of the zpool to the DR server.
- Solaris allows creating partitions starting from block 0. This is where the disk's metadata information is stored. If this partition is used as target of a replication or destination of a physical snapshot, this block 0 gets overwritten thus making the disk unusable. It is suggested to avoid partitions that start at block 0 being used in above mentioned cases
- To recover a zpool, you can either use the physical snapshot or a virtual snapshot approach. In both approaches all the replication pairs associated with protecting zpool needs to be recovered to a common time.
- Since only one zpool can be active in a directory, ensure that there is no existing zpool of same name/guid on the DR server and delete any old virtual snapshots.
- **Virtual snapshot approach:** Recover all replication pairs associated with the zpool to a common time and then issue the "**zpool import -d /dev/vs/dsk <zpool name>**" command to recover the zpool on the DR server.
- We recommend that you use the "**inmzpool**" command line utility found under the VX agent installation directory instead of the zpool command. "**inmzpool**" is a command line utility that should be used to perform all kinds of operations when virtual snapshots are used for zpools. You can pass the same arguments as that of "**zpool**" as "**inmzpool**" is a wrapper around the "**zpool**" command



#### Caution:

When using the virtual snapshot approach, always use the read-write snapshot because zpool does not support RO devices and import may fail randomly. However while importing you can use the "read-only" mount option to protect against any writes or use the "read-write" mount option to allow writes (By default Zpool import mounts in read-write mode)

On Solaris platforms, always use the "**shutdown**" command instead of "**reboot**" as the later is not graceful and would also mark all the Solaris replication pairs for a resync

Only zpools made out of full disk, partitions and logical volumes can be supported. ZFS can support even a regular file based vdevs which are not supported for protection. Only zpools based on volumes are supported

- **Physical snapshot approach:** The process slightly differs when you opt for a physical snapshot approach. When the target and snapshot devices are under the same directory then the import will panic and may result in a system crash. To avoid this physical snapshots are taken on another directory. The table below shows the list of allowed volumes compatible with the target volume type.

**Table 9**

Source Volume	Target Volume	Physical snapshots
Sun volume (/dev/md/dsk) Or Veritas volume (/dev/vx/dsk) Or Solaris partition (/dev/dsk)	Solaris partition (/dev/dsk)	Sun volume (/dev/md/dsk) Or Veritas volume (/dev/vx/dsk)
		Veritas volume (/dev/vx/dsk) Or Solaris partition (/dev/dsk)
	Veritas volume (/dev/vx/dsk)	Sun volume (/dev/md/dsk) Or Solaris partition (/dev/dsk)

- Create physical snapshots for all replication pairs associated with the protecting zpool to a common time.
- Use the below command to bring the zpool online on the DR server.  
`"zpool import -d <physical snapshot device path or virtual snapshot device path> <name of the zpool>"`



**Notes:**

When new components are added to the zpool, they have to be replicated as well.

Vacp is not supported for recovering a zpool by extension no event based recovery is supported for zpool devices

To replicate a complete disk beneath the zpool you will need to replicate the s0 partition of the corresponding disk as zpool uses the EFI labeling.

As a best practice it is advisable to start using the partition from cylinder 50



**Caution:**

Zpool with full disk members with VTOC labels are not supported.

Zpools imported using virtual snapshot devices will be destroyed on

- VX pair delete:
- Virtual snapshot delete
- Retention pruning

### 3.2.2 Protecting Solaris 10 with zones

On Solaris 10 machines the agent software is installed on the global zone. Check the below list to before setting VX replication pairs

- Select the target machine which has similar or better hardware resources ((CPU, Memory, Disk Space etc) than its corresponding source Solaris machine.
- You will need to share the VX install location to all the zones on the source Solaris machine.
- Acquire the list of all zones on the source Solaris machine that are to be protected. Export each zone to an individual zone configuration file. The zone configuration file contains details about the zone.
- You will then need to read the zone configuration file to get the location of the zone data and directories exported to the zone.
- Identify the partitions, mount points from the zone configuration and create corresponding partitions and mount points on the target server. Replicate the volumes identified in each of the zone configuration file.
- You will need to place the zone configuration files on one of the source volumes so that a copy of it is maintained on the target Solaris machine. The zone configuration files should be updated whenever there are changes to any zone resources. So that the zone configuration is always updated on the target system.
- You will need to configure consistency jobs on each of the replication to enable event based recovery operations.

## 4 The FX agent

In the previous chapter we have seen the Volume replicating agent, in this chapter we will focus on the File replication agent (FX agent).

After this chapter you will be able to

- Set FX replication
- Understand FX job options and settings
- Understand FX execution sequence for both pull and push mode
- Default FX templates that are shipped and a brief description
- Create user defined templates
- Understand various scheduling modes
- Edit an existing FX job

## 4.1 File Protection

On Windows servers, the FX agent is installed as a software service that is configured to start automatically as part of the system startup. The agent service shows up as “**InMage Scout FX Agent**” in windows Services GUI. The command “**services.msc**” can be used to display the Services GUI.

On Unix/Linux hosts, the FX agent is installed as a daemon process. As part of the agent installation, you can choose for or against an automatic start of FX agent daemon during system startups. The **start**, **stop** and **status** scripts in the installation directory can be used to start, stop and display the status of FX agents respectively.

After this, all the action is by the FX agent software. The agent installation folder is configurable during install and defaults to “**/usr/local/InMage/FX**” and the binaries take about 6MB of disk storage.



### Notes:

FX agent service needs to be running with domain user privileges on windows.

Ensure that the “**config.ini**” file under the FX installation path contains valid domain and user names. In most of the cases the “**root**” user belongs to the “**root**” group.

When the root user does not belong to the root group edit the “**config.ini**” file to change the line **inmsyncGID = <name of the group to which the root user belongs>**. This ensure that the FX agent has valid user privileges

The FX agent locks the destination directory while replication so we recommend that you do not use the system drive (c:\) as the target destination



4.1.1 When to use FX Replication

FX replication is used to replicate a group of files or folders on a scheduled basis while keeping the properties intact. FX replication can be done across platforms. FX agent is used in conjunction with the VX agent to provide application support.

If VX agent is to “Volume Replication” then FX agent is to “File Protection”. A blank “File Protection” screen is shown below.




Figure 87:

Fields found on File Protection main screen are

**Filters:** Show the application name (given during configuring the FX replication pair), this is specially used when working with many FX jobs.




Figure 88:

“Application Per Page” and “Jobs Per Page”: Enter the required number of jobs to be displayed per page and click on “Ok”




Figure 89:

Clicking on the Arrows allows you to sort the jobs more efficiently.



Figure 90:

Click the “**New Job Group Wizard**” to open “File Protection” wizard. Click the “Add Job” to start a file replication pair.

Clicking on the “**Manage Templates**” will display the list of available FX templates; you may also choose to create user defined templates here

You may choose to delete default templates (not recommended). FX templates are very helpful while performing failover and failback. They are used in conjunction with the VX agent. Scout suit actively supports enterprise applications such as Exchange, SQL Server etc. Each of the template calls a script which you may want to alter based on your environment.

File Replication Templates	
	Template ID
<input type="radio"/>	BES Failover
<input type="radio"/>	CX Backup For Linux
<input type="radio"/>	CX Backup For Windows
<input type="radio"/>	CX DB Sync for Linux
<input type="radio"/>	DNS Failback
<input type="radio"/>	DNS Failover
<input type="radio"/>	ESX Linux Guests Discovery
<input type="radio"/>	ESX Windows Guests Discovery
<input type="radio"/>	Exchange 2007 Consistency
<input type="radio"/>	Exchange 2007 Consistency Validation

**Figure 91**

The following are the list of FX templates

**Exchange Discovery:** Discovers exchange installation, its configuration and volumes it spans

**Exchange failover without retention:** Exchange 2003 failover based on latest consistent point

**SQL Failover without retention:** Performs SQL 2000 for replication pairs configured without CDP retention option

**DNS failover:** Performs a DNS failover to the target host

**DNS failback:** Performs a failback to the source host

**SQL Planned Failover:** Performs a SQL 2000 planned failover

**SQL Unplanned failover:** Performs an unplanned failover of SQL 2000 database

**Exchange Consistency Validation:** Validates Exchange logs and database on the source and outputs the result to FX log

**Exchange log rotation:** Rotates logs on exchange (source side)

**SQL Discovery:** Discovers SQL 2000 database configurations and volumes involved

**SQL Consistency:** Issues consistency tags for SQL 2000

**Exchange consistency:** Issues consistency tags for Exchange 2003

**Exchange planned failover:** Performs failover of exchange 2003 to a latest consistent point

**Exchange unplanned failover:** Performs an unplanned failover of exchange

**SQL 2005 Discovery:** Discovers SQL 2005 installation, configuration and volumes it spans

**SQL 2005 consistency:** Issues consistency tags on the source SQL 2005 server

**SQL 2005 planned failover:** Performs a failover of SQL 2005 to the latest tag (which the script issues on the source and failover is made to this tag)

**SQL 2005 unplanned failover:** Performs SQL 2005 failover to a latest consistency tag.

**SQL 2005 failover without retention:** Performs a SQL 2005 failover for a replication pair configured without retention.

**BES Failover:** Performs Blackberry Exchange Server failover

**Exchange 2007 Consistency:** Issues exchange 2007 consistency tags on the source

**Exchange 2007 planned failover:** Performs exchange 2007 failover to a consistent point

**Exchange 2007 unplanned failover:** Performs exchange 2007 failover to a latest consistent point

**Exchange 2007 Consistency validation:** Validates exchange logs and database on the source, the result of this validation can be seen under the FX log

**Exchange 2007 failover without retention:** Performs exchange 2007 failover for replication pairs set without CDP retention option enabled.

**Windows File share information:** Replicates a windows network share to a target where it remains as a network share

**Linux File share information:** Replicates a Linux network share to a target where it remains as a network share

**Oracle FX Source -> CX:** Backing up oracle step 1

**Oracle FX CX -> Target:** Backing up Oracle step 2

**Oracle FX Target-> CX:** Backing up Oracle step 3 (Refer to the solution document for detailed steps involved in backing up oracle)

**CX Backup for Linux:** Takes backup of entire CX

**CX Backup for Windows:** Takes backup of entire windows

**CX DB Sync for Linux:** This FX template is internally used by CX HA and remote CX features. This performs CX database sync between two CX servers.

**ESX Linux Guests Discovery:** This template is used to discover Linux guests on a given ESX server.

**ESX Windows Guests Discovery:** This template is used to discover windows guests on a given ESX server.

**FileServer Consistency:** This template is used to issue consistency tags for file servers

**FileServer Discovery:** This FX template is used to perform

**FileServer Failover without Retention:** Template used to perform a file server failover when replication pairs are set without CDP retention.

**FileServer Planned Failback:** FX template is used to perform a failback of file server.

**FileServer Planned Failover:** This is used to perform a planned file server failover

**FileServer Unplanned Failover:** This is used to perform an unplanned file server failover.

**Linux file share information:** Used to replicate file share information on a Linux server

**MySQL Consistency:** Used to issue consistency tags on MySQL volumes.

**MySQL Planned Failover:** performs planned MySQL failover

**MySQL Unplanned Failover:** Performs unplanned MySQL failover

**Oracle(Unix/Linux) Consistency:** This job is used to issue consistency tags on Linux based oracle volumes

**Oracle(Unix/Linux) Planned Failover:** Used to perform Linux based Oracle planned failover

**Oracle(Unix/Linux) Unplanned Failover :** Used to perform Linux based Oracle unplanned failover

**Oracle\_Consistency\_Windows:** Used to issue consistency tags on oracle volumes (windows based)

**Sharepoint Consistency:** Used to issue consistency tags for SharePoint volumes

**Sharepoint Planned Failover:** Used to perform SharePoint planned failover

**Sharepoint Unplanned Failover:** Used to perform SharePoint unplanned failover

**Windows file share information:** This FX template is used to replicate windows share information

### 4.1.2 FX Execution

FX job executes in a pull mode by default. This can be changed through the FX job options, under miscellaneous options. There is a difference between both modes.

#### Pull mode




Figure 92:

#### Push mode

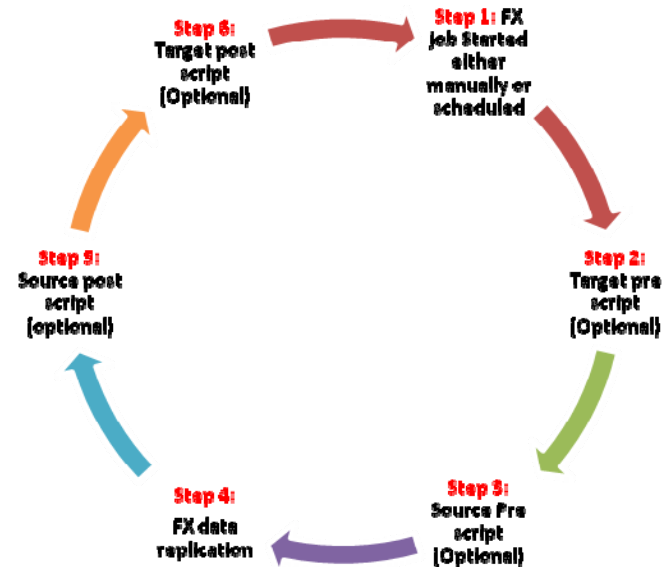


Figure 93:

### 4.1.3 Creating User-defined FX Templates

An FX template stores job options, scheduled time and applies them to a job when you choose the template at the time creating a replication pair.

To create a new FX template,

**Step 32.** Navigate to “**File Protection → Manage Templates**”. Enter the name for the new template and then click on “**Create**”

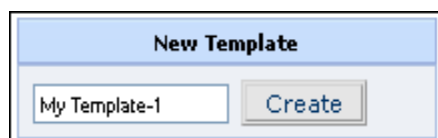


Figure 94:

**Step 33.** Schedule the template when it has to execute. This can be scheduled in different ways

- “**Run Now**”: Executes the job (to which this template has been applied) to execute once its configured
- “**Run At**”: specific time (runs only once)
- “**Run On Demand**”: Needs to be executed manually
- “**Run Every**”: Loops every given amount of time
- “**Daily At**”: Runs at a given time everyday
- “**Weekly On**”: Runs weekly at a given time. To learn more refer to the section [Group scheduling modes](#) on page 105.

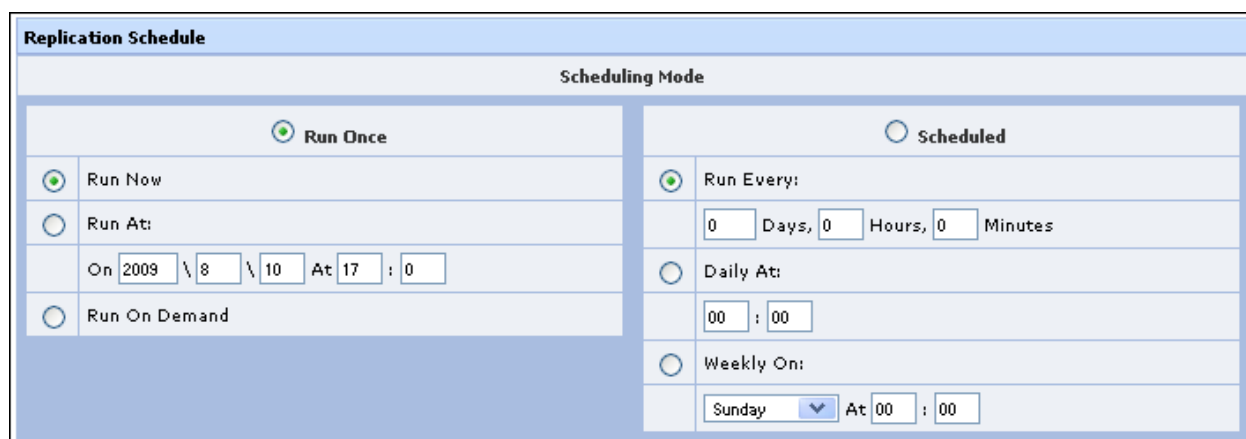


Figure 95:

**Step 34.** Then enter the “**Source**” and “**Target**” directory (this is where the templates are placed) then scroll down and set the “**Job Options**”.

Directory	
Source: <input type="text" value="D:\test"/>	Target: <input type="text" value="H:\test copy"/>

**Figure 96:**

**Step 35.** Choose the required “**File/Directory Options**”

Options	
File/Directory Options	
<input checked="" type="radio"/>	Copy the source directory to a subdirectory of the target directory
<input type="radio"/>	Copy the contents of the source directory directly into the target directory
<input checked="" type="checkbox"/>	Always perform checksum
	Checksum block size: <input type="text" value="8192"/>
<input type="checkbox"/>	Whole files (no incremental checks)
<input type="checkbox"/>	Create backup files
	Backup Directory: <input type="text"/>
	Backup Suffix: <input type="text"/>
<input checked="" type="checkbox"/>	Compress files

**Figure 97:**

- “**Copy the source directory to a subdirectory of the target directory**”: This replicates all the source files and folders into a sub folder on the target. (Default).
- “**Copy the contents of the source directory directly into the target directory**”: This copies all the files and folders directly into the target host
- “**Always perform checksum**”: Performs 128 bit md4 checksum to ensure data consistency. (Default selected with block size 8192).
- “**Whole Files (no incremental checks)**”: Performs a full backup, overwrites older files on the target.
- “**Create backup files**”: This would enable the options “**Backup Directory**” and “**Backup Suffix**”. Enter the complete directory path and the suffix.
- “**Backup Directory**”: This will backup existing files on the target to the specified folder with specified “**backup suffix**” and then the files will be replaced on the target (keeping the backup intact). However when the “**Push mode**” is selected , the backup directory is created under the target folder
- “**Compress files**”: This option compresses the files during transfer to save bandwidth (Default)



### Step 36. “Inclusion/Exclusion Options”

Inclusion/Exclusion Options	
<input checked="" type="checkbox"/>	Update only (Do not overwrite newer files)
<input type="checkbox"/>	Only update files that already exist at the destination
<input type="checkbox"/>	Ignore files that already exist at the destination
<input type="checkbox"/>	Ignore files with same size and timestamp at destination
<input type="checkbox"/>	Ignore files with same size
<input type="checkbox"/>	Exclude files matching pattern: <input type="text"/> Separated by ;
<input type="checkbox"/>	Include subset of exclude list matching pattern: <input type="text"/> Separated by ;

Figure 98:

- **“Update only (Do not overwrite newer files)”**: Updates only older files and does not overwrite newer files (Default). If a same (old) file exists on the target, it will be updated as required rather than overwriting.
- **“Only update files that already exist at the destination”**: If a file exists on the source and the target then by enabling this option the FX agent updates the file but will not create new files on the target.
- **“Ignore files with same size and timestamp at destination”**: Files with same size and timestamp will not be updated on the target.
- **“Ignore files with same size”**: Ignores the files on the target with the same size irrespective of the timestamp
- **“Exclude files matching pattern”**: Ignores files with specified extension. Does not copy files with specified pattern from source to target. If an exception has to be made then check the option below
- **“Include subset of exclude list matching pattern”**: Any exception from the above option has to be entered here.

For example,

If all the files with extension .txt has to be excluded (not to be copied to the target) then check the **“Exclude files matching patterns”** and enter **“\*.txt”**. If an exception has to be made for **“something.txt”** then **“something.txt”** has to be placed under **“Include subset of exclude list matching pattern”**

## Include / Exclude options further explained

The FR exclude option allows you to selectively exclude certain files from the list of files to be transferred. Multiple exclude patterns can be specified by separating them with a semicolon. The exclude and include patterns specified in the job options allow for flexible selection of which files to transfer/skip. An ordered list of include/exclude options specified is maintained, and when a filename is encountered, a check against each pattern is made in turn. The first matching pattern is acted on. If it is an exclude pattern, then that file is skipped. If it is an include pattern then that filename is not skipped. If no matching include/exclude pattern is found then the filename is not skipped.

The patterns can take several forms. The rules are explained below:

- If the pattern starts with a / then it is matched against the start of the filename, otherwise it is matched against the end of the filename. Thus `"/myfile"` would match a file called `"myfile"` at the base of the tree. On the other hand, `"myfile"` would match any file called `"myfile"` anywhere in the tree.
- If the pattern ends with a / then it will only match a directory, not a file, link or device.
- If the pattern contains a wildcard character from the set `*?` [Then expression matching is applied using the UNIX shell file name matching rules. Otherwise a simple string match is used.
- If the pattern includes a double asterisk `"**"` then all wildcards in the pattern will match slashes, otherwise they will stop at slashes.
- If the pattern contains a / (not counting a trailing /) then it is matched against the full filename, including any leading directory. If the pattern does not contain a / then it is matched only against the final component of the filename.
- If the pattern starts with `"+"` (a plus followed by a space) then it is always considered an include pattern, even if specified as part of an exclude option. The `"+"` part is discarded before matching.
- If the pattern starts with `"-"` (a minus followed by a space) then it is always considered an exclude pattern, even if specified as part of an include option. The `"- "` part is discarded before matching. The `+/-` rules are most useful in exclude lists, allowing you to have a single exclude list that contains both include and exclude options.

Below are some exclude/include examples:

- Exclude `"*.o"` would exclude all filenames matching `*.o`
- Exclude `"/myfile"` would exclude a file in the base directory called `myfile`
- Exclude `"myfile/"` would exclude any directory called `myfile`
- Exclude `"/myfile/*/bar"` would exclude any file called `bar` two levels below a base directory called `"myfile"`
- Exclude `"/myfile/**/bar"` would exclude any file called `bar` two or more levels below a base directory called `"myfile"`
- Include `"*/"; *.c"`; exclude `"**"` would include all directories and C source files
- Include `"myfile/"; "myfile/bar.c"`; exclude `"**"` would include only `"myfile/bar.c"` (the `"myfile/"` directory must be explicitly included or it would be excluded by the `"**"`)



### Notes:

The double quotes around the exclude/include patterns in the examples are purely for explanatory purposes and are to be discarded while specifying the patterns in the job options page.

### Step 37. “File Deletion Options”

One of the things regarding file deletion options that is that they may be disabled through the use of a specific key “**EnableDeleteOptions**” in the FX agent configuration file for UNIX agents and registry for Windows agents. This key value pair is set to the default value of 3255 (decimal) enabling all the delete options. Given below is a picture shows that the “**EnableDeleteOptions**” is set to 0, there by disabling the delete options.

File Deletion Options	
<input type="checkbox"/>	Enable file deletion options
<input type="checkbox"/>	Delete files at destination that do not exist at source
<input type="checkbox"/>	Delete excluded files on the source
<input type="checkbox"/>	Delete files after transfer only
<input type="checkbox"/>	Keep partially transferred files

Figure 99:

- “**Enable file deletion options**”: Checking this option will enable the options below it.
- “**Delete files at destination that do not exist at source**”: This deletes any additional files on the target that no longer exist at the source.
- “**Delete excluded files on the source**” : Excluded files will be deleted on the source
- “**Delete files after transfer only**”: Copies all source files to the target and then deletes the target files that do not exist on the source
- “**Keep partially transferred files**”: Maintains partial files

### Step 38. “Link Options”

Link options deal specially with symbolic links. Each of the option has its own operation

Link Options	
<input checked="" type="radio"/>	Copy symbolic links as symbolic links
<input checked="" type="radio"/>	Copy contents of symbolic links
<input checked="" type="radio"/>	Copy links outside the source tree
<input checked="" type="radio"/>	Ignore links outside the source tree

Figure 100:

- **“Copy symbolic links as symbolic links”**: Symbolic links are copied directly (not the actual data but the data path is copied)
- **“Copy contents of Symbolic Links”**: The actual data is copied rather than the symbolic links themselves
- **“Copy links outside the source tree”**: If the symbolic link is within the source directory (the one being replicated) and the actual data (the one symbolic link points to) is outside the source directory that will be copied as well.
- **“Ignore Links outside the Source tree”**: Symbolic link is within the directory being replicated and the actual data is outside the directory, then selecting this option will ignore the actual data.

### Step 39. “File Detail Options”

File Detail Options	
<input checked="" type="checkbox"/>	Preserve permissions
<input type="checkbox"/>	Preserve owner (root only)
<input type="checkbox"/>	Preserve group
<input type="checkbox"/>	Preserve devices (root only)
<input checked="" type="checkbox"/>	Preserve times

Figure 101:

- **“Preserve permissions”**: Maintains file attributes and their respective permissions on the target
- **“Preserve owner (root only)”**: Preserves file ownership attributes; this option is available to the super-user only.
- **“Preserve group”**: Preserves group attributes
- **“Preserve devices”**: Preserves devices after copy to the target, this option is available to the super-user only.
- **“Preserve times”**: Maintains times (edited, created etc) of the source files on the target

Step 40. “Secure Shell Options”

Secure Shell Options	
<input type="checkbox"/>	Enable Encryption
Encryption Key Type	<span>RSA</span>
Encryption Cipher	<span>NONE</span>

Figure 102:

- **“Enable Encryption”**: If the sshd service is not running on both the source and target or if the replication is within the same host then this option will be disabled. Checking this box will enable the options under it.
- **“Encryption Key Type”**: there are 2 types of types here RSA or DSA. Although DSA takes time it gives the best security.
- **“Encryption Cipher”**: Select the appropriate cipher for encryption

You can configure each job to be encrypted with specific encryption key type and cipher type independently. Editing the job will create new encryption keys under the installation folder.

Supported versions of ssh are **“OpenSSH server versions 3.8”** (or above) and **“SecSH server versions 3.2”** (or above). Ensure that the ssh binaries are installed, else the Secure Shell Options will be disabled. Open ssh binaries can be obtained from [www.openssh.com](http://www.openssh.com). Although Open ssh is not shipped with FX agent, it can be obtained separately from Hitachi Data Systems.



**Caution:**

FX agent service should be running with domain user privileges for encryption to work. Encryption will not work if the user name (the one FX agent starts with) is present in multiple domains.

Step 41. "Feedback Options":

Feedback Options	
<input type="radio"/>	Very high verbosity
<input type="radio"/>	High verbosity
<input type="radio"/>	Medium verbosity
<input checked="" type="radio"/>	Low verbosity
<input type="radio"/>	Very low verbosity
<input type="radio"/>	Quiet
<input checked="" type="checkbox"/>	Show file transfer statistics
<input checked="" type="checkbox"/>	Show progress during transfer

Figure 103:

- "Very high verbosity": Logs everything
- "High verbosity":
- "Medium verbosity": gives debug information
- "Low verbosity":
- "Very low verbosity"
- "Quiet": No logging
- "Show file transfer statistics": Displays statistics in "Protection Status"
- "Show progress during transfer": Displays statistics even while performing a transfer in "Protection Status"

## Step 42. “Miscellaneous Options”

Miscellaneous Options	
<input type="checkbox"/>	Create temporary files in: <input type="text"/>
<input checked="" type="checkbox"/>	Do not cross filesystem boundaries
<input checked="" type="checkbox"/>	IO timeout in seconds: <input type="text" value="9600"/>
<input checked="" type="checkbox"/>	Specify port number to use: <input type="text" value="874"/>
<input type="checkbox"/>	Limit bandwidth to <input type="text" value="0"/> KB/s
<input type="radio"/>	Push data from source to target (implies lower CPU load on source)
<input checked="" type="radio"/>	Pull data from source to target
	CPU throttle <input type="text" value="0"/>
	Send RPO alert if <input type="text" value="0"/> minutes passed
	Send E-mail alert if <input type="text" value="5"/> minutes passed without job progress
	Pre execution script pathname <input type="text"/>
	Post execution script pathname <input type="text"/>
	Pre execution script pathname (destination) <input type="text"/>
	Post execution script pathname (destination) <input type="text"/>
	Catch All job modifier <input type="text"/> <i>for power users only</i>

Figure 104:

- **“Create temporary files in”**: Creates temporary files in the specified directory (target) before transferring. However then the **“push mode”** is selected these temporary files are created under a directory placed under the target directory.
- **“Do not cross file system boundaries”**: Backup restricted to file system limitations
- **“IO timeout in seconds”**: IO timeout
- **“Specify port number to use”**: port used by FX agent
- **“Limit bandwidth to”**: Limits bandwidth usage within the specified range.
- **“Push data from source to target”**: Source host pushes data to the target, faster process but requires more of CPU resources
- **“Pull data to target from source”**: slower process, the target host pulls data from the source host. To learn more about FX execution while in pull or push mode refer the section **“FX execution”** on page 86.
- **“CPU throttle (source)”**: Maximum amount of CPU resource to be used.
- **“Send RPO alert if”**: Notifies all CX administrators about the job crossing RPO bounds when the actual RPO exceeds the specified RPO threshold. Note that the administrator must have e-mail notifications enabled.
- **“Send E-mail alert if”**: Specify the time ( in minutes) to receive email alerts if the FX job did not make any progress.
- **“Pre execution script pathname”**: The source side script executes first before the replication
- **“Post execution script pathname”**: The source side script executes after the replication is complete
- **“Pre execution script pathname (destination)”**: Target side script executes first at the target then the replication takes place

- **“Post execution script pathname (destination)”**: Target side script executes after the replication is done. Pre and post scripts also support arguments.  
For windows the syntax is: \ "c:\post Test.bat" <arg without spaces> \ "<arg with spaces>" (with the quotes)
- **“Catch all Job modifier”**:

The keys used in this field are preceded by “- -”(double -). If an invalid argument is used then the job will not execute and return an exit code of 1.

**Table 10: Catch all options**

Key	Effect
-A	Replicates ACL’s wherever ACL’s are supported.(Windows without domain )
-p	Replicates permissions on Unix (and between heterogeneous platforms like windows to Linux etc.)
-S /- -sparse	Handle sparse efficiently.
- n	Dry run, shows what would have been transferred
- - super	Preserves file and folder permissions on windows
- - inplace	Used to overcome error 12 which occurs due to low temporary directory space



**Notes:**

When the temporary directory cannot accommodate a larger incoming file the FX job ends with an error code of 12. To resume the job, use the “--inplace in the “catch all” field.

While performing a drive to drive replication, ensure that the target drive is double the source size else the job will fail with error code 12. Again you may use the “--inplace” in the catch all field to overcome this

An exit code 23 along with data compression and sync compression having value as NA indicate an error. Data compression and sync compression have non zero values and an exit code of 23 is thrown because of user permissions on windows (Access Control List) and mounting volume without a flag on Linux.

File Protection Status							
Filter	Job Description	Application	Status	Source Host	Source Directory	Target Host	Target Directory
<input type="button" value="Set"/> <input type="button" value="Clear"/>	<input type="text"/>	Select ▼	Select ▼				
<input type="checkbox"/>	firstjob...	runevery	Completed	R5U232BITTGT.INIMAGE.IN	/home/21	R5U232BITSRC.INIMAGE.IN	/home
More Details	Start Time	End Time	Last Update Time	Data Compression			
Log Trending	2008-12-22 02:20:01	2008-12-22 02:21:23	2008-12-22 02:21:23	N/A			

**Figure 105:**






**Caution:**

FX agent service must be restarted for the Job to take effect of the changes in the port settings when editing existing job. If the Daemon is not restarted the Jobs will fail to run.

To preserve permissions to the target you can use the “**Catch All**” option (all platforms)



**Figure 106: Entering Catch all options**

The File systems where the ACL's are not supported will throw error 23 while replicating the typical ACL related errors will show up in the job logs like

```
set_acl:  
sys_acl_set_file(D000000640000015E/D0000006400000188/D000000640000018B/D000  
000640000018C, SMB_ACL_TYPE_DEFAULT): No error
```

Or

```
set_acl:  
sys_acl_set_file(D000000640000015E/D0000006400000188/D000000640000018B/D000  
000640000018C, SMB_ACL_TYPE_DEFAULT): Operation not supported
```

Because FX also replicates a file along with its ACL permissions it is imperative that

**On windows platforms:** The source files and folders should be owned by a domain user (and not a local user)

#### 4.1.4 Things to Remember while Setting FX Job Options

Certain File Replication option combinations are known to cause unexpected behaviors. The known options and actual behaviors are documented in this section.

- **“Include and exclude options”**: If a file replication job contains rules set to exclude a set of files and to include a subset of those excluded files, and then include option will be ignored. All files under the exclude option will be excluded.
- **“Include and exclude options with other rules”**: Under the following set of rules: Keep files at target host that do not exist at source, exclude replication of certain files at source, and delete files at target (if present) that were excluded from replication at source; The behavior is as desired except that, the files that do not exist at the source host are deleted at the target along with the files excluded from replication.
- **“Symbolic links outside source tree”**: When the options to copy symbolic links as symbolic links and ignore links outside the source tree are set, replication of the symbolic links pointing to files outside the source tree will be skipped.
- **“Illegal filenames across platforms”**: Certain special characters (or names) are allowed on certain platforms while they are illegal on others. For instance, while replicating files from a UNIX source to Windows target, files at the source with names that may be illegal on the Windows platform will not be replicated, the same applies to paths as well. While entering log paths or file names, ensure that illegal names are not used. This is applicable for both VX and FX agents
- Scout CX residing on Linux and Scout FX File Replication agent for Linux, both, may start a daemon process on port 873 by default. This port number is not configurable in Scout CX. If Scout FX agent and CX server co-exist on the same box, a file replication job will not initiate if port 873 is in use by CX for volume replication. This can be circumvented, by configuring the file replication job to use a specified port. This port can be specified in the file replication job options page.
- FX agent service should always be running with domain user privileges.


Miscellaneous Options	
<input type="checkbox"/>	Create temporary files in: <input type="text"/>
<input checked="" type="checkbox"/>	Do not cross filesystem boundaries
<input checked="" type="checkbox"/>	IO timeout in seconds: <input type="text" value="9600"/>
<input checked="" type="checkbox"/>	Specify port number to use: <input type="text" value="874"/>
<input type="checkbox"/>	Limit bandwidth to <input type="text" value="0"/> KB/s
<input type="radio"/>	Push data from source to target (implies lower CPU load on source)
<input checked="" type="radio"/>	Pull data from source to target

Figure 107: Configuring FX job port number

- The FR job options allow for pre and post execution scripts to be run. Each time a job starts, it attempts to run the pre execution script if any. Notice that if the script invokes a system reboot, this will reoccur in a loop as the script will be run each time by a new instance of the FR job at startup. If it

is required that the system reboots only once, the necessary logic would have to be placed in the script itself.

- The Scout File Replication agent allows for replication to or from hosts using NAT IP addresses. A typical scenario where this feature could be used is while trying to replicate files from a primary server to a remote target that accesses external networks through a firewall/router or vice versa. In this scenario, you will have to configure the agent to be aware of the NAT address of the host (i.e. address of the host as seen by external networks). The NAT IP address can be set by accessing the agent configuration utility, and checking “**Enable fixed NAT IP address**” and entering the NAT IP address in the field below. Note that a VPN connection to the primary site network would be required on the remote target so the agent can communicate with the Scout CX server. Further, router policies and/or Windows firewall policies may need to be configured to allow access to the specified FX agent daemon port (configurable in the CX UI job options page), especially while running a replication job in the “**push mode**”. The default port used by the FX agent daemon is 874.



**Figure 108: Configuring NAT address**

The FX agent locks the destination directory while replication so we recommend that you do not use the system drive (c:\) as the target destination

## 4.1.5 Replicating Files through FX

### 4.1.5.1 Replicating one or more files

Scout provides a way to replicate a single file/ file types or only selected files rather than replicating the whole folder. The following is an example to replicate one or more files (or file types)

**Step 43.** Open CX UI, click on “**File Protection**”, click on “**New job group wizard**”, and then click on “**Add job**”. Enter the folder under which the file is located at source, select the destination, and enter the same path then click on “**Next**”.

The screenshot shows the 'File Protection Wizard: Replication Pair' window. It has a title bar and a main area with several sections. The 'Replication Hosts' section is at the top, followed by 'Application Name' (set to 'Single file replication') and 'Job Description' (set to 'Registry'). Below these are two columns: 'Source' and 'Destination'. Each column has a 'Host' section with radio buttons for 'DRServer [Windows]' and 'PRODUCTION [Windows]'. The 'Directory' section for both columns is set to 'e:\none'. At the bottom, there is a 'Template not selected' dropdown menu and 'Next ->' and 'Cancel' buttons.

Figure 109:

**Step 44.** This opens up “**File Protection Wizard: Options**”, scroll down for “**Inclusion/Exclusion Options**” and check “**Exclude file matching pattern**”. Enter “**\*.\***” (to exclude all files and folders), check “**Include subset of exclude list...**”, and enter the name of the file. If more than one file has to be replicated enter the names of the files separated by a semi colon. The rest of the process is similar to setting up a normal FX job.

The screenshot shows the 'File Protection Wizard: Options' window, specifically the 'Inclusion/Exclusion Options' section. It contains several checkboxes and text fields. The 'Update only (Do not overwrite newer files)' checkbox is checked. Below it are three unchecked checkboxes: 'Only update files that already exist at the destination', 'Ignore files that already exist at the destination', and 'Ignore files with same size and timestamp at destination'. Further down, the 'Exclude files matching pattern:' checkbox is checked, with the text field containing '\*.\*'. The 'Include subset of exclude list matching pattern:' checkbox is also checked, with the text field containing 'protect.reg'. Both text fields are followed by the text 'Separated by ;'.

Figure 110:

#### 4.1.5.2 Creating Jobs and Job Groups

**Step 45.** Click on the “File Replication” menu. The Replication Job Groups window appears.

The screenshot shows the 'File Protection' interface. At the top, it says 'Logged in as 'admin' - Logout' and 'Server Time: Nov-5-2008 02:20:36'. Below this is a header for 'Applications' with a filter set to 'Application Name' and a 'Sort by' dropdown. There are input fields for 'Application Per Page' and 'Jobs Per Page', both set to 5, with a 'Go' button. The main content area states 'No Replication Applications To Display'. At the bottom, there are two buttons: 'New Job Group Wizard' and 'Manage Templates'.

**Figure 111 :** File replication new job group

**Step 46.** Click on “New Job Group Wizard” to create a new job group. Click on “Add Job” to proceed

The screenshot shows the 'Replication Jobs' section. It has a header with 'Replication Jobs' and a sub-header 'Logged in as 'admin' - Logout'. Below this is a table with the following columns: 'Application Name', 'Source Host', 'Source Directory', 'Target Host', and 'Target Directory'. The table is empty, and there is a message 'No jobs added yet'. At the bottom, there are two buttons: 'Cancel' and 'Add Job'.

**Figure 112:** Creating a new job group

**Step 47.** Enter the “Application name”, “Job Description”, then select the source host as “Source” and target host as “Destination”. Enter the source and target directory then click on “Next” to proceed.

The screenshot shows the 'File Protection Wizard: Replication Pair' dialog box. It has a header with 'File Protection Wizard: Replication Pair' and a sub-header 'Logged in as 'admin' - Logout'. Below this is a section for 'Replication Hosts'. It contains fields for 'Application Name' (Shark DB) and 'Job Description' (Shark DB replication). There is a table with two columns: 'Source' and 'Destination'. Each column has a 'Host' row and a 'Directory' row. The 'Source' column has 'PROD-SERV [Windows]' selected for the host and 'c:\Shark\_db' for the directory. The 'Destination' column has 'BAKP-SERV [Windows]' selected for the host and 'c:\Shark\_db' for the directory. At the bottom, there is a 'Template not selected' dropdown and two buttons: 'Next ->' and 'Cancel'.

**Figure 113:** Adding a file replication job

**Step 48.** This opens up the “FX Job Options” screen. Refer to the [Creating User-defined FX Templates](#) section on page 87 for detailed explanation of each field.

**Step 49.** Scroll down and click on “Finish” to continue.

<input type="checkbox"/>	Limit bandwidth to <input type="text" value="0"/> KB/s
<input type="radio"/>	Push data from source to target (implies lower CPU load on source)
<input checked="" type="radio"/>	Pull data from source to target
	CPU throttle (source) <input type="text" value="0"/>
	Send RPO alert if <input type="text" value="0"/> minutes passed
	Send E-mail alert if <input type="text" value="5"/> minutes passed without job progress
	Pre execution script pathname <input type="text"/>
	Post execution script pathname <input type="text"/>
	Pre execution script pathname (destination) <input type="text"/>
	Post execution script pathname (destination) <input type="text"/>
	Catch All job modifier <input type="text"/> for power users only
<input <="" <input="" td="" type="button" value=" Cancel "/>	

**Figure 114**

**Step 50.** Click on “Set Schedule” to determine when to run the job.

### File Protection

Logged in as 'admin' - [Logout](#)

**Group Schedule**

Schedule Type	Schedule Time
No schedule set	

**Replication Jobs**

	Application Name	Source Host	Source Directory	Target Host	Target Directory
Run order 1					
<input type="radio"/>	Shark DB	PROD-SERV	c:\Shark_db	BAKP-SERV	c:\Shark_db

**Figure 115: Scheduled Job Group**

102

**Step 51.** Specify desired option for the FX job to execute. Then click on “Set Schedule”.

**File Protection Wizard: Scheduling**  
Logged in as 'admin' - [Logout](#)

**Replication Schedule**

**Scheduling Mode**

☒ **Run Once**

☐ Run Now

☐ Run At:

On 2008 \ 11 \ 5 At 3 : 0

☒ Run On Demand

☐ **Scheduled**

☒ Run Every:

0 Days, 0 Hours, 0 Minutes

☐ Daily At:

00 : 00

☐ Weekly On:

Sunday At 00 : 00

**Set Schedule** **Cancel**

**Figure 116:**



**Notes:**

If no hosts are listed under “Replication Hosts”, you will need to assign an FX license to the hosts. License assignment to hosts can be performed on clicking License Management menu under the System tab.

If a job is scheduled to run at regular intervals, all further schedules are skipped until the executing FX job is completed.

The FX job should always be run with domain user privileges on windows else the agent window shows an error “frsvc could not open service manager”

**Step 52.** You are returned back to the previous screen. Here you may choose add more jobs to the same group.

The screenshot shows the 'File Protection' interface. At the top, it says 'Logged in as 'admin' - Logout'. Below this is a 'Group Schedule' section with a table for 'Schedule Type' and 'Schedule Time'. The 'Schedule Type' has options 'Once At' and 'On Demand'. The 'Schedule Time' has a 'Set Schedule' button. Below this is a 'Replication Jobs' section with a table. The table has columns: 'Application Name', 'Source Host', 'Source Directory', 'Target Host', and 'Target Directory'. There are two jobs listed: 'Shark DB' and 'second db'. Below the table are buttons for 'Details', 'Remove', 'Cancel', 'Add Job', and 'Finish'.

Group Schedule	
Schedule Type	Schedule Time
Once At	On Demand

Set Schedule

Replication Jobs					
	Application Name	Source Host	Source Directory	Target Host	Target Directory
Run order 1					
	Shark DB	PROD-SERV	c:\Shark_db	BAKP-SERV	c:\Shark_db
Run order 2					
	second db	PROD-SERV	c:\second_db	BAKP-SERV	c:\second_db

Details Remove Cancel

Add Job

Finish

**Figure 117: Job group with added job**

You can add more jobs similarly by clicking “**Add Job**”, or delete created jobs at this stage by selecting the job and clicking “**Remove**”. If you click on “**Cancel**” file replication main screen appears without adding the job group and jobs. If you click on “**Finish**” job group is added and jobs to the database and start them based on the scheduling mode.



## 4.1.6 Group Scheduling Modes

A job group can be scheduled to run once or run recursively. This section describes the six ways of scheduling a job group.

36. **“Run Once”, “Run Now”**: Runs job in the job group immediately after creation.

The image shows a 'Replication Schedule' dialog box. It has a title bar 'Replication Schedule' and a sub-header 'Scheduling Mode'. There are two main sections: 'Run Once' and 'Scheduled'. The 'Run Once' section is active, indicated by a selected radio button. It contains three options: 'Run Now' (selected with a radio button), 'Run At:' (with a date and time picker set to 'On 2008 \ 1 \ 11 At 8 : 0'), and 'Run On Demand'. The 'Scheduled' section is inactive, indicated by an unselected radio button. It contains four options: 'Run Every:' (with a date and time picker set to '0 Days, 0 Hours, 0 Minutes'), 'Daily At:' (with a time picker set to '00 : 00'), 'Weekly On:' (with a day selector set to 'Sunday' and a time picker set to 'At 00 : 00'), and 'Run On Demand'. At the bottom of the dialog are two buttons: 'Set Schedule' and 'Cancel'.

Figure 118: Run once, run now

37. **“Run Once”, “Run At”**: Schedules the job group to run once at a specified date and time, with a resolution up to 1 minute.

The image shows a 'Replication Schedule' dialog box. It has a title bar 'Replication Schedule' and a sub-header 'Scheduling Mode'. There are two main sections: 'Run Once' and 'Scheduled'. The 'Run Once' section is active, indicated by a selected radio button. It contains three options: 'Run Now' (unselected), 'Run At:' (selected with a radio button, with a date and time picker set to 'On 2008 \ 1 \ 11 At 8 : 0'), and 'Run On Demand'. The 'Scheduled' section is inactive, indicated by an unselected radio button. It contains four options: 'Run Every:' (with a date and time picker set to '0 Days, 0 Hours, 0 Minutes'), 'Daily At:' (with a time picker set to '00 : 00'), 'Weekly On:' (with a day selector set to 'Sunday' and a time picker set to 'At 00 : 00'), and 'Run On Demand'. At the bottom of the dialog are two buttons: 'Set Schedule' and 'Cancel'.

Figure 119: Run once, run at scheduling mode

Hours are specified using a 24 hour scale.

38. **“Run Once”, “Run on Demand”**: Jobs in the job group are run on demand by using the job group start/stop buttons on file replication main page.

The screenshot shows the 'Replication Schedule' dialog box. The 'Scheduling Mode' section has two tabs: 'Run Once' (selected) and 'Scheduled'. Under 'Run Once', there are three options: 'Run Now', 'Run At:', and 'Run On Demand'. The 'Run At:' option is selected, showing a date and time picker set to 'On 2008 \ 1 \ 11 At 8 : 0'. At the bottom, there are 'Set Schedule' and 'Cancel' buttons.

Figure 120: Run once, run on demand scheduling mode

39. **“Run Scheduled”, “Run Every”**: Schedules the job group to run as per the entered value. Job runs for the first time and job runs based on the entered value. For example, you enter two days, job runs immediately, and thereafter, it runs for every two days interval.

The screenshot shows the 'Replication Schedule' dialog box. The 'Scheduling Mode' section has two tabs: 'Run Once' and 'Scheduled' (selected). Under 'Scheduled', there are three options: 'Run Every:', 'Daily At:', and 'Weekly On:'. The 'Run Every:' option is selected, showing a date and time picker set to '0 Days, 20 Hours, 0 Minutes'. At the bottom, there are 'Set Schedule' and 'Cancel' buttons.

Figure 121: Run scheduled, Run Every Scheduling Mode

40. **“Run Scheduled”, “Daily At”**: Schedules the job group to run once a day at a specified time.

**File Protection Wizard: Scheduling**

**Replication Schedule**

**Scheduling Mode**

☐ Run Once

☒ Scheduled

☐ Run Now

☐ Run At:

On 2008 \ 1 \ 11 At 8 : 0

☐ Run On Demand

☐ Run Every:

0 Days, 20 Hours, 0 Minutes

☒ Daily At:

00 : 00

☐ Weekly On:

Sunday At 00 : 00

Set Schedule Cancel

Figure 122: Run scheduled, Daily At Scheduling Mode

41. **“Run Scheduled”, “Weekly On”**: Schedules the job group to run once a week at a specified day and time, with resolution up to 1 minute.

**Replication Schedule**

**Scheduling Mode**

☐ Run Once

☒ Scheduled

☐ Run Now

☐ Run At:

On 2008 \ 1 \ 11 At 8 : 0

☐ Run On Demand

☐ Run Every:

0 Days, 20 Hours, 0 Minutes

☐ Daily At:

00 : 00

☒ Weekly On:

Sunday At 1 : 00

Set Schedule Cancel

Figure 123: Run scheduled, Weekly On Scheduling Mode



#### Notes:

Note that with respect to recurring File Replication jobs, once an active instance completes its run, it is removed and a new instance is created for the next run. This instance will bear null start/end times and its status will report “Not Started” indicating that it is waiting to be scheduled.

## 4.1.7 Editing Existing Job Groups and Jobs

**Step 53.** In order to edit settings of an existing job, go to the File Replication UI main page, select the group you wish to edit, and click on “**Details**”.

The screenshot shows the 'File Protection' interface. At the top, it says 'Logged in as admin' and 'Logout'. The server time is 'Jan-11-2008 07:11:27'. Below this is a table of applications. The 'Others' group is selected, showing a table with columns: Job Description, Status, Source Host, Source Directory, Target Host, Target Directory, RPO, Scheduled Type, GID, JID, Job Order, Scheduled Start Time, and Trending. The table contains one row for 'others...' with status 'Not started...'. Below the table are buttons for 'Stop', 'Start', 'Details', and 'Remove'. At the bottom, there are buttons for 'New Job Group Wizard' and 'Manage Templates'.

Job Description	Status	Source Host	Source Directory	Target Host	Target Directory	RPO	Scheduled Type	GID	JID	Job Order	Scheduled Start Time	Trending
others...	Not started...	DRSERVER	e:\two	PRODUCTION	e:\two	N/A	on demand	3	5	0	0000-00-00 00:00:00	N/A

Figure 124: Editing an existing File Job group



### Notes:

If a job is edited in a job group then older job description is removed and the protection status page shows ongoing jobs including previously executed FX jobs

**Step 54.** The File Replication group page shows up. Next, click on “**Edit**” to change the group settings.

The screenshot shows the 'File Replication Group Page'. It has two main sections: 'Group Schedule' and 'Replication Jobs'. The 'Group Schedule' section has a table with columns: Schedule Type, Schedule Time, and Next Run Time. The 'Replication Jobs' section has a table with columns: Application Name, Job ID, Job Order, Status, Source Host, Source Directory, Target Host, Target Directory, Last Update Time, and % Complete. Below the tables are buttons for 'Show Job Options', 'Edit', and 'Back'.

Schedule Type	Schedule Time	Next Run Time
On Demand	N/A	0000-00-00 00:00:00




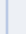
Application Name	Job ID	Job Order	Status	Source Host	Source Directory	Target Host	Target Directory	Last Update Time	% Complete
Single file replication	4	0	Not started	PRODUCTION	e:\one	DRSERVER	e:\two	2008-01-11 07:11:29	0
Others	5	0	Not started	DRSERVER	e:\Two	PRODUCTION	e:\two	2008-01-11 07:11:32	0

Figure 125: File Replication Group Page



**Step 55.** The File Replication groups edit page shows up. Click on “**Set Schedule**” to modify the existing group schedule. This brings up the scheduling page where the mode can be changed as described in the scheduling section.

Group Schedule						
Schedule Type		Schedule Time				
Once At		On Demand				
<input type="button" value="Set Schedule"/>						

Replication Jobs						
		Application Name	Source Host	Source Directory	Target Host	Target Directory
<i>Run order 1</i>						
		Single file replication	PRODUCTION	e:\one	DRSERVER	e:\two
<i>Run order 2</i>						
		Others	DRSERVER	e:\Two	PRODUCTION	e:\two
<input type="button" value="Details"/> <input type="button" value="Remove"/> <input type="button" value="Cancel"/>						
<input type="button" value="Add Job"/>						
<input type="button" value="Finish"/>						

**Figure 126: File Replication group edit page**

In addition to changing the group schedule, an existing job can be deleted by selecting the job and clicking on the “**Remove**” button. An existing job can also be modified by selecting the job and clicking on the “**Details**” button which will bring up the job wizard page where the source/target hosts, directories and job options can be changed as described in the earlier sections. Furthermore, new jobs can also be added to the existing group by clicking on “**Add Job**”, which brings up the job wizard. To change the job run order click on the  or . If jobs are in the same run order they will run in parallel. If two jobs are in the different run levels they will run sequentially. Run order 1 jobs will be executed first whereas jobs in run order 2 will run after the jobs in the run level 1 are completed.

You can delete an existing job by selecting the job and clicking on the “**Remove**” button. Job group definition can be completed by pressing “**Finish**” button.



**Notes:**

To “**Finish**” a job group, you must enter both job and its schedule. Otherwise the “**Finish**” option will not appear.

#### 4.1.8 FX error codes, causes, and solutions

Table 11

Error Codes	Description	Causes	Troubleshooting
1	Syntax or usage error	Invalid filter mentioned	Edit the job options and correct the filters under Include/ Exclude Options
		Invalid options in catch all job modifier	Edit the Job Options and correct the <b>“Catch all”</b> job modified under Miscellaneous options.
3	Errors selecting input/output files, folders	Destination is not a directory	Edit the FX job and set a valid target directory
		Unable to get details of a file or a directory	Verify the file system on which the files exist is available for use.
5	Error starting client-server protocol	A file daemon.conf may be corrupted. This can be found under the FX installation path (for non windows) And for windows daemon.conf can be found under FX installation folder\application data.	Restart the FX job.
10	Error in Socket IO	Source Directory does not exist	Edit the job settings and set a valid source directory
		Failed to connect to remote FX agent.	Check the FX agent service and start the service if not already started.  A firewall might be blocking the FX agent’s port (default 874). Unblock the FX ports from the firewall.  Port might already be in use by another application. Stop the FX agent service and verify if the port is in use.
11	Error in file I/O	Failed to open include/exclude file.	Include/Exclude option creates a file on the target (in case of pull mode). The error occurs if this file cannot be accessed.

			<p>Ensure enough free space and set read write permissions to the parent folder.</p> <p>For windows the parent folder is FX installation folder\application data</p> <p>And for non-windows the parent folder is FX installation folder.</p>
		Unable to create destination folder	Ensure FX agent service has permissions to create the directory in the target path.
		No free space available	FX installation folder may be running out of space. Move the log files to a different volume to make free space.
12	Error in rsync protocol data stream	Connection failure	<p>Ensure source and target FX services are up and running</p> <p>Ensure inmsync process is running on the source (in case of pull mode)</p> <p>Check the network connectivity.</p>
		Unable to write to the socket.	<p>Too many FX jobs are competing for the same resource such as sockets or disk space at the same time.</p> <p>Schedule the jobs to run such that they do not overlap with each other.</p>
14	Error in IPC code	Unable to create a pipe or create a child process	System policy restricting the FX agent to create child processes.
20	Received SIGUSR1 or SIGINT	FX agent service or its child process was stopped abruptly.	Identify if any user or application has accidentally stopped the process and restart the FX agent service.
23	Partial transfer due to error	Some of the files /file attributes are not transferred completely.	<p><b>Windows:</b></p> <p>The FX agent service should be started with domain user privileges</p> <p>Then FX job option “<b>Catch all job modifier</b>” should have the input value as --super</p> <p><b>In general</b></p> <p>Source FX agent should have permissions to read the files from the source and write them on the target.</p>
24	Partial transfer due to vanished source files/	This is a warning message indicating file deletion on the source	Run the job again to maintain data consistency.

		.folder during FX execution.	
30	Timeout in data send/receive	No activity from the FX job for 9600 secs will result in a timeout. This is used to detect network outage	While working with large number of files, there might be idle time during the FX replication. Either increase the idle time or Split the job to smaller jobs.
-255	Job Stopped by user	Job stopped through CX UI	Re run the job to replicate the data.
		FX service is down while job is in progress.	Start the FX service.
		Pre or Post script returned a non zero value.	Check the script for possible errors.
-127	<b>Inmlimit:</b> child terminated unexpectedly	Inmlimit is a child process for FX, which again spawns inmsync. This error is shown when inmsync exits unexpectedly	Check the job log for errors and restart the FX service.



# Part 3: Recovery

This chapter deals with Recovery through the CX user interface. In this chapter you will find examples for each of the recovery operation.

After this chapter, you would be able to

- Issue basic consistency marker
- Perform a target volume rollback
  - Based on time or an event
- Perform snapshot operation
  - Physical
  - Based on time or an event
  - Virtual read only, virtual read write
  - Based on time or an event
- Scheduled snapshots
  - Based on time and or an event

## 5 Recovery Operations

### 5.1 Data Consistency

To ensure data consistency, consistency markers (user defined events) are issued on the production volume(s). This is done through vacp.exe under the VX agent installation folder. To learn more about vacp.exe refer to the section [Issuing consistency tags](#) on page 225.

#### 5.1.1 Generating consistency markers/ Tags

Consistency markers are of three types: file system, user defined, and application.

File system tags flush all the volume data in the memory to the disk and then issue a marker which then gets replicated to the DR server's volume (Target volume).

**User Defined Tags:** When a user defined tag is issued a file system tag is generated first and then a user defined tag (with a custom name) is generated. These tags are used to identify important events such as marking month end jobs etc.

**Application Tags:** The application is paused then all the application data in the memory is flushed on to the disk and a marker is issued.

All consistency markers are stored within the CDP retention option log files.

#### 5.1.2 Automating Tag Generation through FX

By using the Pre or Post script in the FX job options, a script can be set to execute at regular intervals. Enclosing a set of commands within this script, you may automate issuing consistency tags at regular intervals.

## 5.2 Recovery

In this section you can find recovery operations such as snapshots and target volume rollback.

### 5.2.1 When to Use Recovery

Recovery operations are performed either to recover from an outage or to validate backups (on the target volume).

While a volume replication is active, the target volume is locked and cannot be accessed. To view the contents of the target volume, you may either choose to unhide the target volume which will pause the replication and increase the RTO or you may choose to take a snapshot of the target volume (which will not disturb the replication). A snapshot is a copy of the target volume.

Snapshots are broadly classified into two types

- Point in time snapshots
- Recovery snapshots

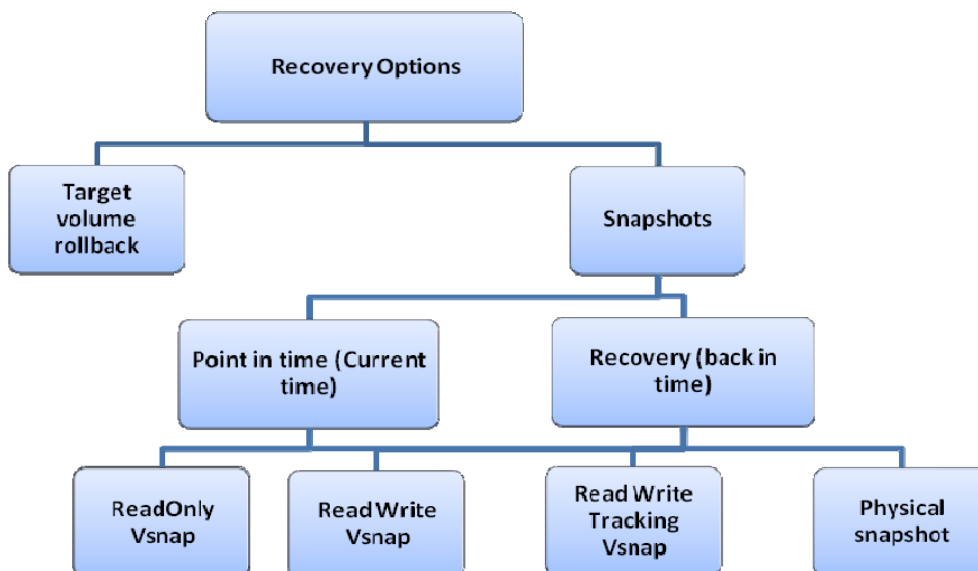


Figure 127: Types of recovery operations

**Point-in-time Snapshots:** Gives you the copy of the target volume as of that instance (the time when you are taking the snapshot). CDP retention option is optional since the snapshot generated will always be of the latest state of the target volume. This can be performed either through the command line interface or through the CX UI (through scheduled snapshots)

**Recovery Snapshots:** Can be taken based on time or an event at some point back in time and this requires CDP retention option to be enabled. Again, it can be performed from the command line interface or through the CX UI.

Recovery and point in time snapshots can both be taken on Physical or a virtual volume.

Taking a snapshot on a physical volume requires the intended snapshot volume it to be equal or larger than the target volume (in the replication pair).

**Virtual snapshots** are virtual in nature. They require minimal system resources and are faster in loading or unloading. Virtual snapshots are again of three types

- Read only (CX UI and CLI)
- Read Write (CX UI and CLI)
- Read Write tracking (CLI only)

As the name indicates read only snapshots are for informative purposes and are not capable of retaining writes on to them. Read write virtual snapshots on the other hand do retain writes on to them, this is done by maintaining an archive log on some part of the local disk as specified (either from UI or through cdpcli) . Read write tacking virtual snapshots goes a step forward; this is especially useful if a new virtual snapshot has to be updated with the writes of an unmounted virtual snapshot. All these snapshots can also be scheduled (scheduled snapshots)

The following table shows the list of recovery operations that you may perform through the CX UI. Observe that all the recovery operations may be performed based on time or event.

**Table 12**

Recovery Operations	Event based	Time based	Multiple Pairs
Virtual Snapshots	✓	✓	✓
Physical Snapshots	✓	✓	✓
Rollback	✓	✓	✓
View recovery range	✓	✓	✓
Scheduled snapshots	✓	✓	✓

## 5.2.2 Determine the Time to Recover

**Step 56.** To perform a snapshot or a rollback to some point back in time three steps are followed as shown below.

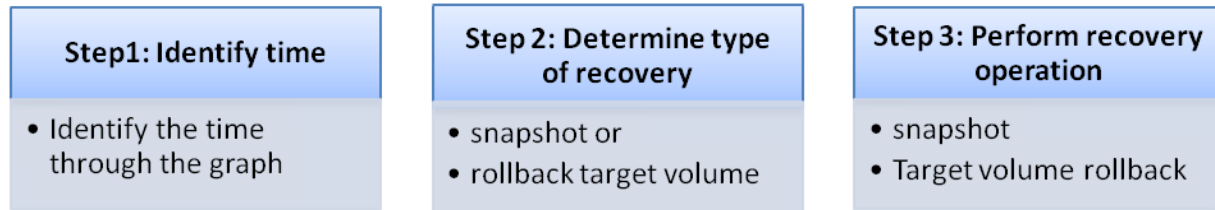


Figure 128

**Step 57.** To identify the time, Click on “**Recovery**” and select the desired replication pair and click on “**View Recovery Range**”.

Volume Recovery : Recovery Snapshots

Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Sep-2-2009 14:35:40

[Recovery Snapshots](#) | [Scheduled Snapshots](#)

[Recover](#) | [Monitor Recovery](#) | [Monitor Rollback](#)

Source Host: All | Target Host: All | Volume: | Search

1-3 of 3 Records | List 3 Records/Page | Page 1 of 1

	Server	Pri Volume	Remote Server	Volume	Replication Pool
<input checked="" type="checkbox"/>	PROD-SERVER	E ( Logistics )	DR-SERVER	E	2
<input type="checkbox"/>	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	3
<input type="checkbox"/>	PROD-SERVER	G ( Engineering )	DR-SERVER	G	4

[Recover](#) | [Rollback](#) | [View Recovery Range](#) | [Reset](#)

Figure 129



### Notes:

You may also choose multiple replication pairs to perform a snapshot, rollback or even view the recovery range.

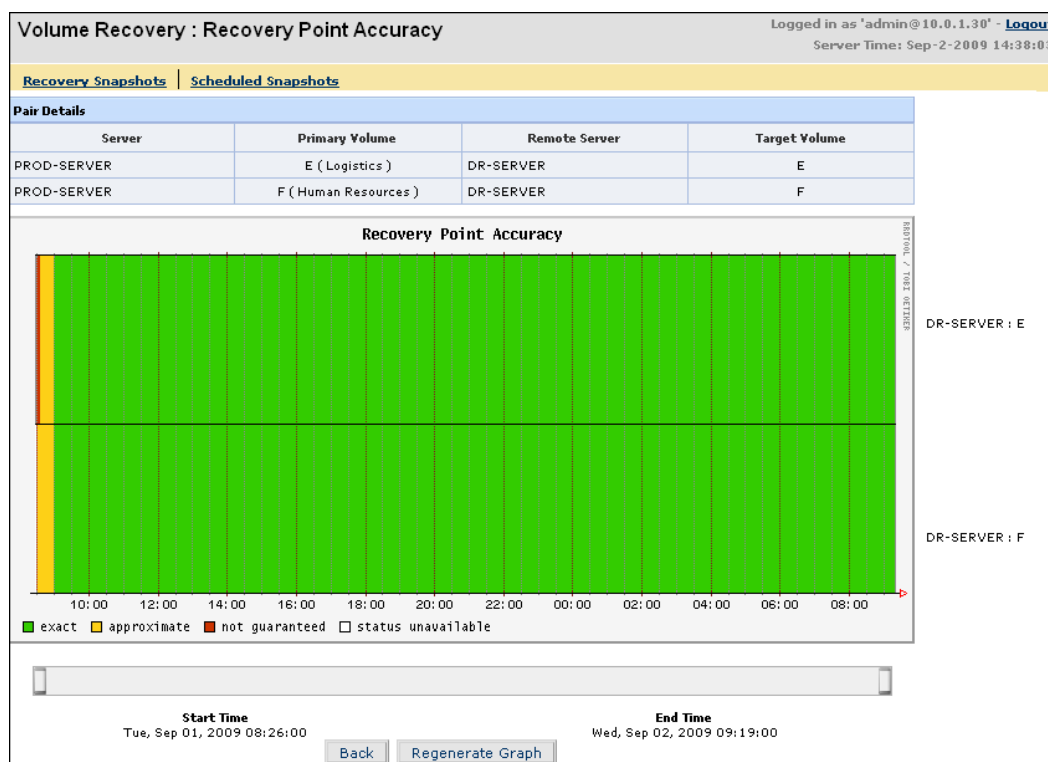
While performing an event based snapshot or rollback for more than one replication pair, you will be shown only the common tags (or events) or common time range between the replication pairs.

This opens the “**Recovery Point Accuracy**” screen. The graph below the “**Pair Details**” shows the health of retention logs. The green areas indicate data replication occurred in data mode and this is considered to be suitable for performing any recovery operations. You should be able to see a bar with two sliders, the left marker indicates the start time of the retention logs and the right marker indicates the end of the retention log time range.

Red and yellow areas indicate that the time frame may not be suitable for performing a recovery operation to it.

White areas indicate that the status is not available.

You may move the markers on the slider to close in on a time frame, then click on “**Regenerate Graph**” to further narrow down a consistent point.



**Figure 130:**

After closing in on a time, you may proceed to perform any of the recovery operation.



#### Notes:

You may also select multiple replication pairs to see their recovery point accuracy. Based on the replication pair you select you get that many sliders. Multiple Tabs appear when multiple replication pairs are selected below “Recovery Point”.

## 5.2.3 Recovery Snapshots

### 5.2.3.1 Virtual Snapshots

#### Overview

From the usage standpoint, virtual snapshots are similar to recovery based snapshots the only major difference being the target volumes are different.

From the functional standpoint

- Virtual Snapshots do not require any physical volumes. These volumes can be loaded and unloaded when required.
- Virtual snapshots consume relatively less time than that of a normal recovery snapshot. This is due to the fact that virtual snapshot process does not perform any data copying as opposed to the normal recovery snapshot which performs volume copy. Virtual snapshot uses retention logs/COW data to service read requests from the virtual snapshot volume. Write operations on the virtual snapshot volume shall be stored in a private file of the virtual snapshot volume

These unique characteristics of the virtual snapshots make it a very valuable tool for recovery process.

Some of the scenarios where these snapshots can be of great help are:-

- It is relatively easy to mount read-only application like reporting servers, standby servers, verification scripts etc. by using virtual snapshots volumes.
- When you need to do a roll back to a point –in-time, but not sure about the time to roll back.
- Virtual Snapshot can be used to browse the data across different points in time horizon without having to do a full recovery snapshot, thus helping to zero-in to the desired point in a quicker way.
- Once the recovery point has been determined, all it remains is either recover only the corrupted files/folders or resort to complete fully recovery snapshot for the determined recovery-point.
- If we would like to view the replicated target volume, one way is to make the target volume visible. But, this would stop applying the differentials to the target volume. Other option is to create point-in-time virtual snapshot of the target volume. In this way, differentials are still applied to the target volume, but the same time, we were able to view the target volume at that particular instance in time.

Other characteristics include

- Like Recovery Snapshots, CDP retention option need to be enabled for creating virtual snapshots
- There can be more than one virtual snapshot for replication pairs.
- Virtual Snapshots can be created as “**Read Only**” or “**Read Write**”. All the data written to virtual snapshot volumes is stored in separate log files specific to virtual snapshots.
- Can create virtual snapshots with read-write and tracking enabled. This would allow tracking of writes in the virtual volume. The tracked data of a virtual volume can be applied on to another virtual volume/ physical volume.
- Virtual Snapshots are available even after a reboot.
- Virtual Snapshots Volumes are automatically dismounted, if the recovery point used for the virtual snapshot does not comply with the specified retention policies. i.e. you can find the older virtual snapshot volumes removed as the retentions logs get pruned along with time.

### **Creating a Virtual snapshot**

To create a virtual snapshot you need a replication pair (with or without CDP retention option) and a CDP license. Virtual snapshot creates a virtual volume and mounts it. This can be accessed through the operating system interface (either windows or Linux).

Creating Virtual Snapshots is a two step process:-

- Map Generation Phase: Where CDP engine create a Meta data for the Virtual Snapshot based on recovery criteria. In case of point-in-time virtual snapshot, an empty map is created.
- Mounting Phase: The CDP engine mounts the Virtual Snapshot Engine to specified drive letter or mount point

Virtual snapshots should not be created on Removable media (floppy drives, CD drives, network drives etc). In case of a virtual volume the maximum space required for it is not more than the size of the target volume. Virtual snapshots are again of three types

42. Read only
43. Read Write
44. Read Write tracking



### 5.2.3.2 Creating a Virtual Snapshot

To create a virtual recovery snapshot, click on **“Recovery”** on the CX UI. Select the desired servers that correspond to the replication pair and then click on **“Recover”** button.

Volume Recovery : Recovery Snapshots

Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Sep-2-2009 14:35:40

[Recovery Snapshots](#) | [Scheduled Snapshots](#)

[Recover](#) | [Monitor Recovery](#) | [Monitor Rollback](#)

Source Host: All | Target Host: All | Volume: | Search

1-3 of 3 Records | List | 3 Records/Page | Page 1 of 1

	Server	Pri Volume	Remote Server	Volume	Replication Pool
<input checked="" type="checkbox"/>	PROD-SERVER	E ( Logistics )	DR-SERVER	E	2
<input type="checkbox"/>	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	3
<input type="checkbox"/>	PROD-SERVER	G ( Engineering )	DR-SERVER	G	4

[Recover](#) | [Rollback](#) | [View Recovery Range](#) | [Reset](#)

Figure 131:

A next screen **“Recovery Snapshots Options (Multiple Pairs)”** appears. Virtual snapshot may be created based on time or an event. Also, you can select multiple pairs and click **“Recover”** to create a snapshot (either physical or virtual) based on a common time or common consistency tag.

## Time-based virtual snapshot on windows

**Step 58.** To create time-based virtual snapshots select an option “**Using Time**” under “**recovery options**”. The “**Recovery Point**” section is changed to accommodate time.

Volume Recovery: Recovery Snapshots Options (Multiple Pairs) Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Sep-2-2009 14:41:38

**1** Snapshot Options **2** Review

Server	Primary Volume	Remote Server	Target Volume	Replication Pool
PROD-SERVER	E ( Logistics )	DR-SERVER	E	2
PROD-SERVER	F ( Human Resources )	DR-SERVER	F	3

**Recovery Options**

Logs Available From 2009/09/01 08:26:49:919 (GMT) To 2009/09/02 09:24:14:381 (GMT)

Recovery Based On

☒ Using Time ☐ Using Application consistency & Event based

Figure 132: Selecting “Using Time”



### Notes:

Event based snapshots have higher levels of data consistency than time based snapshots  
While performing a Physical snapshot, ensure that the physical disk is larger than the replicated volume.

**Step 59.** Enter the desired time between the retention range and select the “**Drive type**” as “**Virtual**”

**Recovery Point**

Logs Available From 2009/09/01 08:26:49:919 (GMT) To 2009/09/02 09:24:14:381 (GMT)

2009 \ 9 \ 2 At 09 : 24 : 14 : 381 [Recovery Point Accuracy>>](#)

**Drive Type**

☒ Virtual ☐ Physical

Figure 133

**Step 60.** While performing multiple read/write virtual snapshots, you will be shown “**Bulk Vsnap Options**”. Rather than specify write log path and mount points for each of the virtual snapshot, you may type in a common directory and common mount point. Each virtual snapshot will create its own directory under the specified “**Parent Directory for Vsnap Logs**”. Click on “**Next**” to continue.

Bulk Vsnap Options				
Read/Write	Parent Directory for Vsnap Logs		Parent Directory for Mountpoints	
<input checked="" type="checkbox"/>	H:\WRW			

Virtual Drive Options				
Target Volume	Read/Write	VSNAP Write Log Path	Mount Point	Virtual Device
DR-SERVER:E	<input checked="" type="checkbox"/>	H:\WRWlog1 J,I,H,D are drives suggested for storing Data log files.		B ▼
DR-SERVER:F	<input checked="" type="checkbox"/>	H:\WRWlog2 J,I,H,D are drives suggested for storing Data log files.		U ▼

**Figure 134**

**Step 61.** You should see the “**Review**” screen, click on “**Finish**” to create the snapshots.

**Volume Replication:Recovery Snapshot Options** Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Sep-2-2009 14:58:45

1  
Create/Modify  
Export Options

2  
**Review**

Server	Primary Volume	Remote Server	Snapshot Volume	Recovery Based On	Recovery Point	Drive Type
DR-SERVER	E	DR-SERVER	B	Time	2009/9/2 09:24:14:381	Virtual
DR-SERVER	F	DR-SERVER	U	Time	2009/9/2 09:24:14:381	Virtual

**Figure 135**

**Step 62.** Observe the status of the snapshot on the “**Recovery**” screen. You may also delete the snapshot from this screen.

**Volume Recovery : Recovery Snapshots** Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Sep-2-2009 15:09:08

Recovery Snapshots
Scheduled Snapshots

Recover
Monitor Recovery
Monitor Rollback

**Source Host**

**Target Host**

**Volume**

1-3 of 3 Records      List  Records/Page      Page  of 1

<input type="checkbox"/>	Server	Pri Volume	Remote Server	Volume	Replication Pool
<input type="checkbox"/>	PROD-SERVER	E ( Logistics )	DR-SERVER	E	2
<input type="checkbox"/>	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	3
<input type="checkbox"/>	PROD-SERVER	G ( Engineering )	DR-SERVER	G	4

**Figure 136**

For a windows target mount points are supported only on NTFS file system. The directory should be empty, if the directory does not exist then a fresh directory is created. The size of a read write virtual snapshot is equal to the source volume; exceeding this limit will result in data write failed on Virtual snapshot. If you are attempting to take a virtual snapshot on a Linux platform, you have to enter only **“Mount point”** since there will not be any drive letters.



**Notes:**

Event based recovery options will require consistency tags to be issued on the source volume. This can be done either through an FX job or through the vacp command line tool. To learn more about issuing consistency tags refer the section [“issuing consistency tags”](#) on page 225 in this document.

Creating a virtual snapshot will always create certain meta data files.

For Read only Virtual snapshots: If CDP retention option is enabled for the replication pair. These metadata files will be stored under the retention log path

If retention is disabled for the replication pair. A path (data log path) has to be explicitly mentioned to store these files. This is applicable for both CX UI and through the CLI interface.

For Read Write Virtual snapshots: Since a data log path has to be mentioned for storing data about writes, the meta data files will be stored under the same path.

Recovering multiple pairs

Time based :  
When you are recovering multiple pairs based on time, click on the “Recovery Point Accuracy” under the “Recovery Point”

Recovery Point			
Logs Available	From 2009/08/25 14:12:26:356 (GMT)	To 2009/08/26 06:24:54:264 (GMT)	
2009	\ 8	\ 26 At 06 : 24 : 54 : 264	<a href="#">Recovery Point Accuracy&gt;&gt;</a>

Figure 137: Recovery Point Accuracy for time based recovery

You should see a new screen listing the selected replication pairs, and a corresponding graph. You can use the slider to close in on the desired time and to regenerate the graph. Click “Submit” to return to the previous screen where the time is filled up.

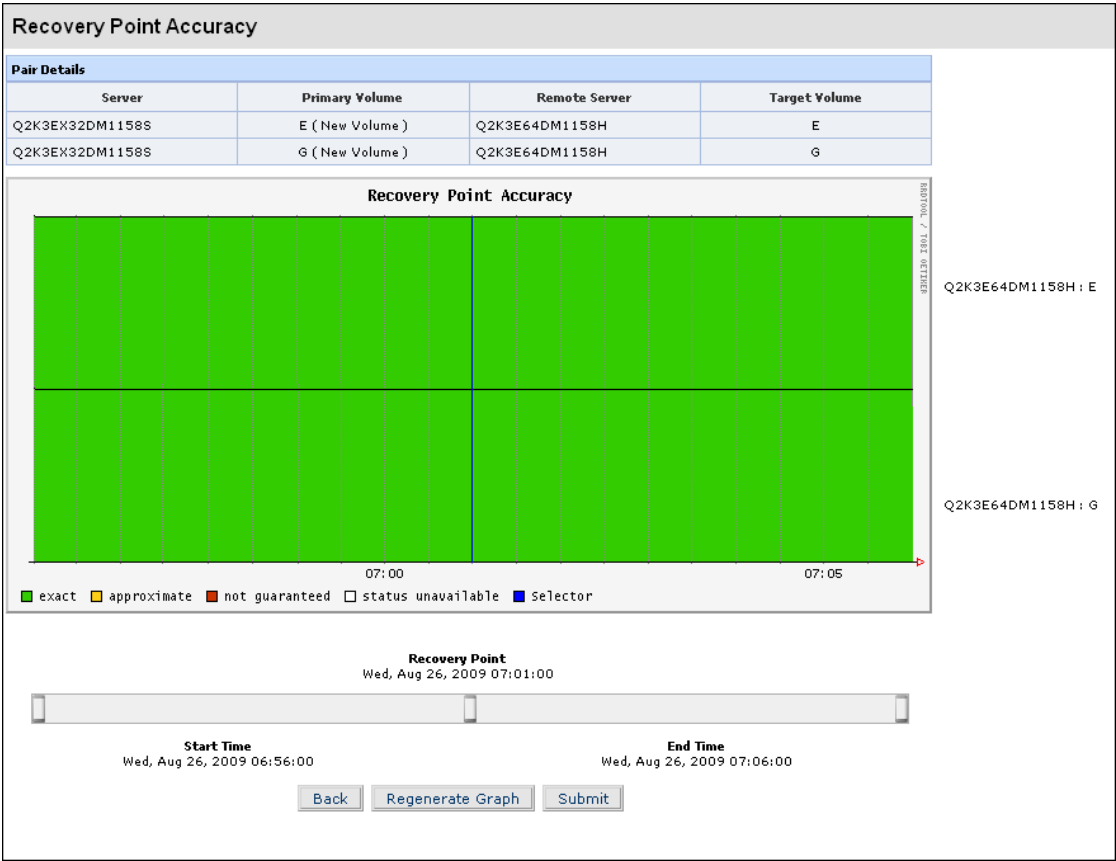


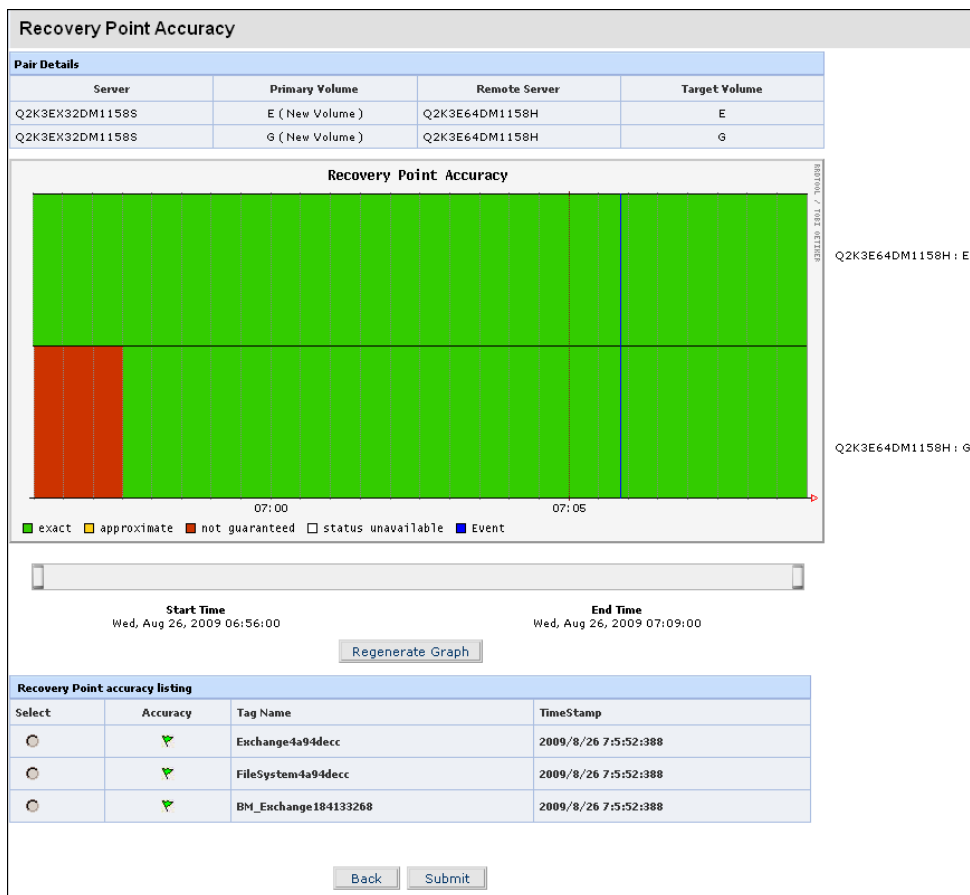
Figure 138

Event based:

When you are performing event based snapshot for a single or group of replication pairs, you can click **“Recovery Point Accuracy List”** under **“Recovery Tag”** to list the common time range, tag accuracy and corresponding graph.

**Figure 139: Recovery Point Accuracy for event based recovery**

You may use the slider to close in on the desired consistency tag. Select the desired tag and click **“Submit”** to return to the previous screen



**Figure 140**

We recommend that you always choose the “Exact” Recovery point indicated in green



**Notes:**

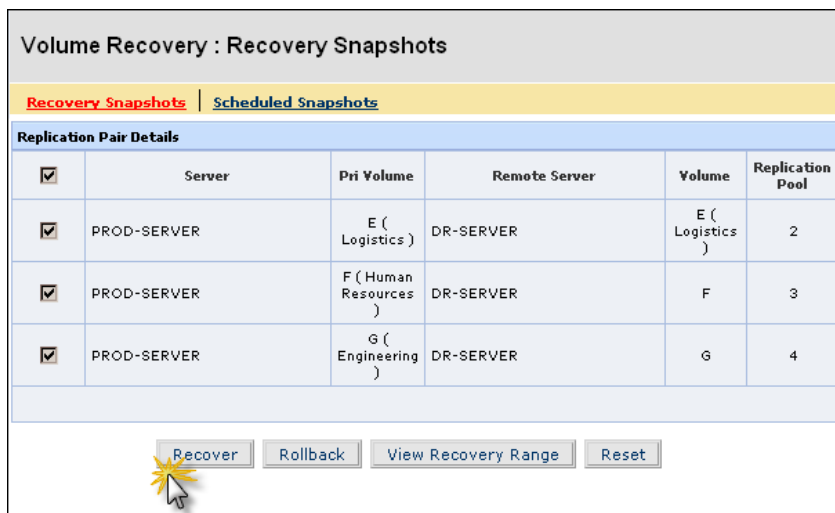
**Recovery Point Accuracy Range Listing:**

A tag will be "Approximate" when an I/O occurs at the time of a bitmap read operation by the driver on the source side. Refer to the section “[Tuning VX](#)” on page 305 for a detailed list of driver tweaks

**Event based snapshots have higher levels of data consistency than time based snapshots**

## Event-based Virtual Snapshot

**Step 63.** To create an event based virtual snapshot navigate to recovery in the UI, then select a server, and click on “**Recover**” button.

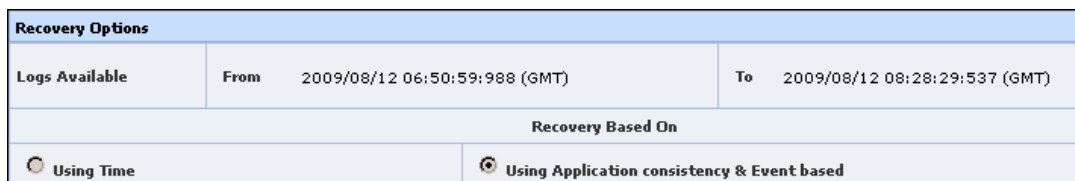


Volume Recovery : Recovery Snapshots					
<a href="#">Recovery Snapshots</a>   <a href="#">Scheduled Snapshots</a>					
Replication Pair Details					
<input checked="" type="checkbox"/>	Server	Pri Volume	Remote Server	Volume	Replication Pool
<input checked="" type="checkbox"/>	PROD-SERVER	E ( Logistics )	DR-SERVER	E ( Logistics )	2
<input checked="" type="checkbox"/>	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	3
<input checked="" type="checkbox"/>	PROD-SERVER	G ( Engineering )	DR-SERVER	G	4

[Recover](#) [Rollback](#) [View Recovery Range](#) [Reset](#)

Figure 141

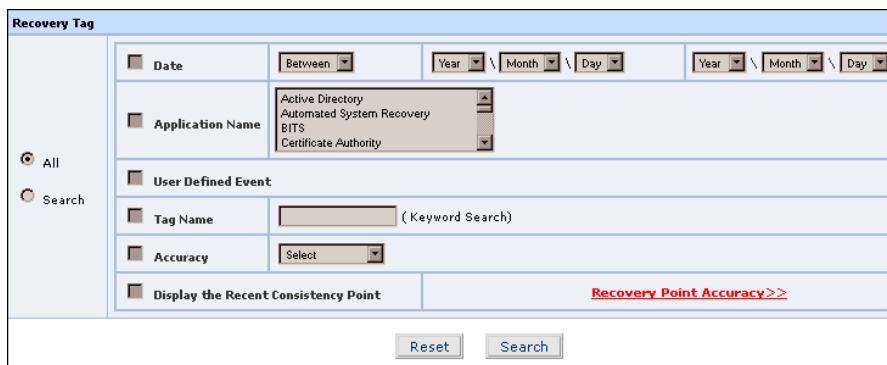
**Step 64.** This takes you to the recovery snapshot options. Under “**Recovery Options**” select “Using Application consistency & Event based”.



Recovery Options		
Logs Available	From 2009/08/12 06:50:59:988 (GMT)	To 2009/08/12 08:28:29:537 (GMT)
Recovery Based On		
<input type="radio"/> Using Time	<input checked="" type="radio"/> Using Application consistency & Event based	

Figure 142

**Step 65.** Then under “**Recovery Tag**” search for an event tag. This displays a list of matching tags, select one of them. “**Recovery Point accuracy**” shows the accuracy of the tag (exact, approximate and not guaranteed). This can be seen in the “**Accuracy**” of the search result.














Recovery Tag	
<input type="checkbox"/> Date	Between <input type="text"/> Year <input type="text"/> \ Month <input type="text"/> \ Day <input type="text"/> Year <input type="text"/> \ Month <input type="text"/> \ Day <input type="text"/>
<input type="checkbox"/> Application Name	Active Directory Automated System Recovery BITS Certificate Authority
<input type="checkbox"/> User Defined Event	
<input type="checkbox"/> Tag Name	( Keyword Search )
<input type="checkbox"/> Accuracy	Select
<input type="checkbox"/> Display the Recent Consistency Point	<a href="#">Recovery Point Accuracy &gt;&gt;</a>

[Reset](#) [Search](#)

Figure 143



**Step 66.** Select an event to recover to and then scroll down to select the “Drive Type”.

Search Result				
	Accuracy	Timestamp	Application	Tag Name
		2009/8/12 7:15:25:109	File System	FileSystem4a826c08
		2009/8/12 7:15:25:109	User Defined	Theta
		2009/8/12 7:11:34:291	File System	FileSystem4a826b21
		2009/8/12 7:11:34:291	User Defined	Beta
<< < 1 > >>				
Recovery Points Accuracy:  - Exact  - Approximate  - Not guaranteed				

**Figure 144**






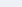

**Step 67.** In the “Drive Type” section click on “Virtual” and the sections beneath it changes displaying

**Step 68.** While performing multiple read write virtual snapshots, you will need to fill in the bulk vsnap options.

Bulk Vsnap Options				
Read/Write	Parent Directory for Vsnap Logs	Parent Directory for Mountpoints		
<input checked="" type="checkbox"/>	h:\VRW_logs_Event			
Virtual Drive Options				
Target Volume	Read/Write	VSnap Write Log Path	Mount Point	Virtual Device
DR-SERVER:E (Logistics)	<input checked="" type="checkbox"/>	h:\VRW_logs_EventVog1 J,H are drives suggested for storing Data log files.		Z
DR-SERVER:F	<input checked="" type="checkbox"/>	h:\VRW_logs_EventVog2 J,H are drives suggested for storing Data log files.		Y
DR-SERVER:G	<input checked="" type="checkbox"/>	h:\VRW_logs_EventVog3 J,H are drives suggested for storing Data log files.		X
<div>Next&gt;&gt; Cancel</div>				

**Figure 145: Linux and Windows Event-based Snapshots**

**Step 69.** Observe the status of the snapshots through the “Recovery” screen.

Recovery Pair Status												
	Host	Host Drive	Recovery Drive	Drive Type	Status	Progress	Expected Recovery Point	Actual Recovery Point	Recovery based on	Info Message	Export Message	View
	DR-SERVER	G	X	Virtual	Ready	0%	2009/8/12 7:11:34:291	-	Tag Based Tag Beta Accuracy 	-		
	DR-SERVER	F	Y	Virtual	Ready	0%	2009/8/12 7:11:34:291	-	Tag Based Tag Beta Accuracy 	-		
	DR-SERVER	E ( Logistics )	Z	Virtual	Ready	0%	2009/8/12 7:11:34:291	-	Tag Based Tag Beta Accuracy 	-		
<div><div>Release Drive</div><div>Force Delete</div></div>												

**Figure 146**



## 5.2.4 Scheduled Snapshots

Scheduled snapshots can be configured only through the CX UI. Scheduled snapshots can

- Execute at a later point of time.
- Repeat execution at regular intervals
- Wait for an event or time before executing the snapshot. (used for pairs without CDP retention option)

### 5.2.4.1 Event-based Scheduled Snapshot through CX UI for Linux Target

**Step 70.** Access the CX UI, click on **“Recovery”**, and then on **“Scheduled Snapshots”**. Select the replication pair and then click on **“Create Snapshot”**

Volume Recovery: Scheduled Snapshots						
Recovery Snapshots		Scheduled Snapshots				
Local hosts replicating to remote						
	Server	Pri Volume	Remote Server	Volume	Last Vx Outpost Agent Heartbeat Time	Resync
	PROD-SERVER	E ( Logistics )	DR-SERVER	E ( Logistics )	2009-08-12 14:03:16	
	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	2009-08-12 14:03:16	
	PROD-SERVER	G ( Engineering )	DR-SERVER	G	2009-08-12 14:03:16	

[Create Snapshot](#) [Reset](#)

Figure 147:

**Step 71.** Under **“options”** select **“Event based”**. This enables the text area next to it, enter the name of the bookmark (consistency tag), and scroll down to select **“Drive Type”** as **“physical”** or **“Virtual”**.

**Volume Replication: Scheduled Snapshot Options**

1 Create Snapshot 2 Schedule Snapshot 3 Review

Pair Details			
Server	Primary Volume	Remote Server	Target Volume
PROD-SERVER	E ( Logistics )	DR-SERVER	E ( Logistics )

**Options**

**Type of Snapshot**

☐ time based

☒ Event based

Bookmark prefix

[Standard bookmark prefixes](#)

**Miscellaneous Options**

Pre execution script pathname

Post execution script pathname

(eg: "C:\scripts\testscript.bat" "arg1" "arg2" ....)

**Figure 148:**

**Step 72.** For this example, select the drive type as “**Virtual**”. Select the desired virtual drive and click on “**Next**”.

Drive Type	
<input type="radio"/> Physical	<input checked="" type="radio"/> Virtual
Virtual Drive Options	
<input type="checkbox"/>	Read/Write
Data log Path	J:,Hdrives suggested for storing Data log files.
Virtual Drives	
	Host
<input checked="" type="radio"/>	DR-SERVER
<input type="radio"/>	DR-SERVER
<input type="radio"/>	DR-SERVER
Mount Point	
<input type="radio"/>	
<input type="button" value="Cancel"/> <input type="button" value="Next&gt;&gt;"/>	

**Figure 149:**

**Step 73.** After completing the configuration, you will ne taken to the review screen. Click on “**Finish**” to create the snapshots

Volume Replication:Scheduled Snapshot Options							
<div> <div>1 Create Snapshot</div> <div>2 Schedule Snapshot</div> <div>3 Review</div> </div>							
Recovery Details							
Server	Primary Volume	Remote Server	Snapshot Volume	Schedule Mode	Schedule Type	Drive Type	Export Option
DR-SERVER	E	DR-SERVER	B	N/A	Event based	Virtual	No
<input type="button" value=" &lt;&lt;Back"/> <input type="button" value=" Finish"/> <input type="button" value=" Cancel"/>							

**Figure 150**

**Step 74.** You will be returned back to the “**Scheduled snapshots**” under “**Recovery**” page. You will find the newly configured snapshot listed under the “**Scheduled Snapshots**”

Scheduled Snapshots												
<input type="checkbox"/>	Host	Host Drive	Snapshot Drive	Drive Type	Type Of Snapshot	Bookmark	Scheduling Mode	Next Scheduled At	Configured At	Export Message	Action	View Export
<input checked="" type="checkbox"/>	DR-SERVER	E ( Logistics )	B	Virtual	Event based	Mars	-	-	-		<a href="#">Edit</a>	
<input type="button" value="Delete Scheduled Jobs"/>												

**Figure 151**

You may edit a physical scheduled snapshot by clicking on the “**Edit**” link under “**Action**”

**Step 75.** To simulate an event we create a bookmark named “Mars” on the source host.

```
C:\Program Files\InMage Systems>vacp -v e: -t "Mars"
Parsing command line arguments ....

Issuing Tags to Volume Set Number : 1
Validating command line arguments ...
User Defined Tag: Mars
Generating Tag: FileSystem4a828f32
Generating "Revocation" tag ...
Preparing the applications for consistency ...
Preparing Files E:\* (recursively)
Starting snapshot set

Using MS Software Shadow Copy provider: b5946137-7b9f-4925-af80-51abd60b20d5
Freezing the applications for consistency ...
Committing shadow copy for the set...
E:\ is mapped to Unique volume \\?\Volume{8ede41b2-5553-4ae0-b5f1-2c93ec4aa70a}\

Checking driver write order state for the given volumes
For volume E:\ driver write order state is: Data
Sending tags to the driver ...
```

**Figure 152: Generating a Consistency Tag on Linux**

**Step 76.** As soon as the consistency tag reaches the target, a snapshot is created, and in addition to that it further waits for the consistency tag (with the same name) to occur the next time. Every time the consistency tag reaches the CX a snapshot is created on the DR server

Snapshot Drives Status											
	Host	Host Drive	Snapshot Drive	Type Of Snapshot	Book Mark Name	Start Time	End Time	status	Percentage	Message	History
<input type="checkbox"/>	DR-SERVER	E ( Logistics )	B	Event based	Mars			waiting on event	0 %	-	
<input type="checkbox"/>	DR-SERVER	E ( Logistics )	B	Event based	Mars	2009-08-12 15:04:04	2009-08-12 15:04:04	Complete	100 %	-	<a href="#">view</a>
<div> <div>Release Drive</div> <div>Forced Delete</div> </div>											

**Figure 153:**

**Step 77.** To release the snapshot drive, select the snapshot drive, and click on “Release Drive”.

**Step 78.** To remove any stale entries, select the snapshot drive, and click on “Forced Delete”. The forced delete removes the entry from the CX UI.

**Step 79.** Click “View”, under “History” to view all time based past executions of the same snapshot.

Snapshot Drives Status										
Host	Host Drive	Snapshot Drive	Type Of Snapshot	Book Mark Name	Start Time	End Time	status	Percentage	Message	
DR-SERVER	E ( Logistics )	B	Event based	Mars	2009-08-12 15:04:04	2009-08-12 15:04:04	Complete	100 %	-	
<div> <div>1 &gt;&gt;</div> <div>Back</div> </div>										

**Figure 154**

### 5.2.4.2 Time-based Scheduled Snapshot through CX UI for Windows Target:

**Step 80.** Access the CX UI, click on “**Recovery**”, and then on “**Scheduled Snapshots**”. Select the replication pair and then click on “**Create Snapshot**”.







Volume Recovery: Scheduled Snapshots						Logg
<a href="#">Recovery Snapshots</a>   <a href="#">Scheduled Snapshots</a>						
Local hosts replicating to remote						
	Server	Pri Volume	Remote Server	Volume	Last VX Outpost Agent Heartbeat Time	Resync
	PROD-SERVER	E ( Logistics )	DR-SERVER	E ( Logistics )	2009-08-12 15:26:46	
	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	2009-08-12 15:26:46	
	PROD-SERVER	G ( Engineering )	DR-SERVER	G	2009-08-12 15:26:46	
<a href="#">Create Snapshot</a> <a href="#">Reset</a>						

Figure 155:

**Step 81.** In the “**Step One**” screen, under “**options**” select “**Time based**”, and scroll down to select “**Drive Type**” as “**physical**” or “**Virtual**”.

Volume Replication: Scheduled Snapshot Options			
<div> <div>1 Create Snapshot</div> <div>2 Schedule Snapshot</div> <div>3 Review</div> </div>			
<b>Pair Details</b>			
Server	Primary Volume	Remote Server	Target Volume
PROD-SERVER	F ( Human Resources )	DR-SERVER	F
<b>Options</b>			
Type of Snapshot			
<input checked="" type="checkbox"/>	time based		
<input type="checkbox"/>	Event based	Bookmark prefix <input type="text"/> <a href="#">Standard bookmark prefixes</a>	
Miscellaneous Options			
	Pre execution script pathname <input type="text"/>		
	Post execution script pathname <input type="text"/>		
(eg: "C:\scripts\testscript.bat" "arg1" "arg2" ....)			

Figure 156:



#### Caution:

Using a cdpcli command directly in the post or pre script is not supported, To overcome this cdpcli feature can be used in a batch file that can be called from the pre or post scripts

**Step 82.** Now, select the appropriate volume for snapshot, and click on “Next”.

Drive Type				
<input checked="" type="radio"/> Physical		<input type="radio"/> Virtual		
Unused Snapshot Drives				
	Host	Drive	Capacity	Filesystem
<input checked="" type="checkbox"/>	DR-SERVER	K ( New Volume )	1073740800	NTFS
<div style="text-align: right;"> <input type="button" value="Cancel"/> <input type="button" value="Next&gt;&gt;"/> </div>				

**Figure 157:**

**Step 83.** This opens up the “Scheduling Mode” page. This is similar to the [Group scheduling modes](#) on page 105 with two differences. FX group scheduling modes end at scheduling a job on “weekly” basis. In the “snapshot schedule” it is extended to “monthly” and even “yearly”. The other difference is the option “Run on Demand” which is unique only to the FX scheduler. Select appropriately and click on “Next” to continue.

Snapshot Schedule	
Scheduling Mode	
<input checked="" type="radio"/> Run Once	
<input checked="" type="radio"/> Run Now	
<input type="radio"/> Run At:	
On	2009 \ 08 \ 12 (yyyy\mm\dd) At 00 : 00 (hrs:mins)
<input type="radio"/> Scheduled	
<input type="radio"/> Run Every:	
	0 Days, 0 Hours, 0 Minutes
<input type="radio"/> Daily At:	
	00 : 00 (hrs:mins)
<input type="radio"/> Weekly On:	
	Sunday At 00 : 00 (hrs:mins)
<input type="radio"/> Monthly On:	
	The first day of every month(s) At 00 : 00 (hrs:mins)
<input type="radio"/> Yearly On:	
	The first day of January At 00 : 00 (hrs:mins)
<div style="text-align: right;"> <input type="button" value("&lt;&lt;back"=""/> <input type="button" value("next&gt;&gt;"=""/> </div>	

**Figure 158:**



**Notes:**

To take a snapshot for a pair without CDP retention option. Perform a scheduled time based snapshot and select “Run Now” to create a snapshot.



**Step 84.** You should see the Review screen now, click on **“Finish”** to create the snapshot configuration.

### Volume Replication: Scheduled Snapshot Options

1

2

3

Create Snapshot

Schedule Snapshot

Review

Server	Primary Volume	Remote Server	Snapshot Volume	Schedule Mode	Schedule Type	Drive Type	Export Option
DR-SERVER	F	DR-SERVER	K	Run once	Run Now	Physical	No

<<Back

Finish

Cancel

**Step 85.** You will return back to **“Scheduled snapshots”** under **“Recovery”** page. The new job can be seen under **“Scheduled Snapshots”**.

Scheduled Snapshots												
<input type="checkbox"/>	Host	Host Drive	Snapshot Drive	Drive Type	Type Of Snapshot	Bookmark	Scheduling Mode	Next Scheduled At	Configured At	Export Message	Action	View Export
<input type="checkbox"/>	DR-SERVER	E ( Logistics )	B ( Logistics )	Virtual	Event based	Mars	-	-	-		<a href="#">Edit</a>	
<input type="checkbox"/>	DR-SERVER	F	K ( New Volume )	Physical	Time based	-	Run Now	-	-		<a href="#">Edit</a>	

Delete Scheduled Jobs

**Figure 159:**

**Step 86.** Under **“Snapshot Drives Status”** you can see the progress of the new job (time based). Click **“View”**, under **“History”**, to see all past time based executions of the same snapshot.

Snapshot Drives Status											
<input type="checkbox"/>	Host	Host Drive	Snapshot Drive	Type Of Snapshot	Book Mark Name	Start Time	End Time	status	Percentage	Message	History
<input type="checkbox"/>	DR-SERVER	E ( Logistics )	B	Event based	Mars			waiting on event	0 %	-	
<input type="checkbox"/>	DR-SERVER	E ( Logistics )	B	Event based	Mars	2009-08-12 15:04:04	2009-08-12 15:04:04	Complete	100 %	-	<a href="#">view</a>
<input type="checkbox"/>	DR-SERVER	F	K	Time based		2009-08-12 15:35:04	2009-08-12 15:35:49	Complete	100 %	-	<a href="#">view</a>

Release Drive

Forced Delete

1

**Figure 160:**

**Step 87.** Click the “**View**” link to see the previous executions of the snapshot. Once the “**Status**” is “**Complete**” then the snapshot is ready for access. Specifying Pre and Post Scripts are optional.

Snapshot Drives Status									
Host	Host Drive	Snapshot Drive	Type Of Snapshot	Book Mark Name	Start Time	End Time	status	Percentage	Message
DR-SERVER	F	K	Time based		2009-08-12 15:35:04	2009-08-12 15:35:49	Complete	100 %	-
1 >>									
Back									

**Figure 161**

For custom event consistency tags (or bookmarks) it is recommended not to use [Standard bookmark prefixes](#)

Scout provides an ability to integrate snapshot process with external applications like tape back software, databases, failover tools, data verification tools, exception handling etc.

The integration ability is provided by means executing pre/post scripts at the target host before and after the snapshot process. You can also specify parameters to the pre/post script and also further control the process using the return codes.

Example: Pre Script Configuration: "c:\preTest.bat" "<arg1>" "<arg2>"

Example: Post Script Configuration: "c:\postTest.bat" "<arg1>" "<arg2>"



**Caution:**

You will need to pass all the arguments to your scripts including the optional switches. Refrain from using the --S , --T and --D switches within the scripts since these are reserved for internal usage

The following table shows the memory usage by virtual snapshots. A maximum of 2048 virtual snapshots can be created on a single machine.

**Table 13**

Number of vsnaps	pagedpool usage	nonpagedpool usage	paged pool overhead(in kb)	Nonpaged pool overhead(in kb)
0	128	4064	NA	NA
1	272	17952	84	8
10	1568	140064	1552	396
100	14528	1361184	14956	6456
250	36128	3399968	37220	16360
500	72128	6795968	75380	32984
1000	144128	13587968	151700	66140
2000	288128	27171968	304460	132056
2048	295040	27986304	310388	135716

## 5.2.5 Performing Target Volume Rollback

Target volume roll back is used to restore the volume back in time. Target Volume Rollback may be initiated from the Scout CX UI or command line.



### Caution:


- This operation stops the replication pair automatically
- This feature is available if retention is enabled for the specific replication pair.
- Resync needs to be completed and replication needs to be in diff sync mode.
- Unmount all virtual snapshots corresponding to this replication pair
- Requires retention logs

**Step 88.** Navigate to “Recovery snapshots” under “Recovery”.

Volume Recovery : Recovery Snapshots					
<a href="#">Recovery Snapshots</a>   <a href="#">Scheduled Snapshots</a>					
Replication Pair Details					
<input type="checkbox"/>	Server	Pri Volume	Remote Server	Volume	Replication Pool
<input type="checkbox"/>	PROD-SERVER	E ( Logistics )	DR-SERVER	E ( Logistics )	2
<input checked="" type="checkbox"/>	PROD-SERVER	F ( Human Resources )	DR-SERVER	F	3
<input checked="" type="checkbox"/>	PROD-SERVER	G ( Engineering )	DR-SERVER	G	4
<div><button>Recover</button> <button>Rollback</button> <button>View Recovery Range</button> <button>Reset</button></div>					

**Figure 162**

**Step 89.** Select the replication pair and then click on “Rollback” a message box appears.



**Figure 163:**

**Step 90.** Click on “OK” to continue. You will also be prompted to delete the retention logs for the selected replication pairs.

**Step 91.** You may choose to roll back the selected target volumes based on time or an event.

Volume Recovery: Rollback Options (Multiple Pairs)					
<b>Pair Details</b>					
Server	Primary Volume	Remote Server	Target Volume	Replication Pool	Mount Point
PROD-SERVER	F ( Human Resources )	DR-SERVER	F	3	
PROD-SERVER	G ( Engineering )	DR-SERVER	G	4	
<b>Recovery Options</b>					
Logs Available	From 2009/08/12 06:50:59:988 (GMT)		To 2009/08/12 10:40:40:325 (GMT)		
Recovery Based On					
<input type="radio"/> Using Time		<input checked="" type="radio"/> Using Application consistency & Event based			

**Figure 164: Time-based Rollback**

**Step 92.** Select the desired consistency tag and click “Submit”

Search Result				
	Accuracy	Timestamp	Application	Tag Name
<input type="radio"/>		2009/8/12 7:15:25:109	File System	FileSystem4a826c08
<input type="radio"/>		2009/8/12 7:15:25:109	User Defined	Theta
<input type="radio"/>		2009/8/12 7:11:34:291	File System	FileSystem4a826b21
<input checked="" type="radio"/>		2009/8/12 7:11:34:291	User Defined	Beta

<< < 1 > >>

Recovery Points Accuracy: - Exact - Approximate - Not guaranteed

**Figure 165: Event-based Rollback**

**Step 93.** Roll back starts when you “Submit”. When rollback is in progress the volume is in locked state

- Existing snapshots and recovery on the volume will be lost.
- Rollback volume cannot be subject to new snapshot, recovery and rollback
- Volume is not appeared in the list
- Replication pair is in disabled mode indicating that replication has stopped.
- While performing rollback for Linux targets, enter a “**mount point**” under “**Pair Details**”

**Step 94.** Once rollback is complete, the replication pair is broken, and rollback volume is enabled under target rollback table.

**Step 95.** To delete record of a rollback, select the host, and click on “**Release Drive**”.

Target Drive Rollback Status								
	Host	Rollback Drive	Status	Progress	Expected Recovery Point	Actual Recovery Point	Recovery based on	Info Message
	DR-SERVER	G ( Engineering )	Complete	100%	2009/8/12 7:11:34:291	-	Tag Based Tag Beta Accuracy	-
	DR-SERVER	F ( Human Resources )	Complete	100%	2009/8/12 7:11:34:291	-	Tag Based Tag Beta Accuracy	-
<div>Release Drive</div>								

**Figure 166:**

**Step 96.** The image below shows two rolled back target volumes, one being a time-based rollback while other being an event based rollback.

Target Drive Rollback Status								
	Host	Rollback Drive	Status	Progress	Expected Recovery Point	Actual Recovery Point	Recovery based on	Info Message
	LIN-DR	/dev/mapper/volume- v1	Complete	100%	2008/10/27 23:27:40:848	-	Tag Based Tag Tag_1 Accuracy	-
	BAKP-SERV	E ( Accounting )	Complete	100%	2008/10/30 9:0:0:0	-	Time Based	-
<div>Release Drive</div>								

**Figure 167:**

## 5.2.6 Recovery Pair Status

Recovery pair status shows the status of any recovery operation:




Recovery Pair Status												
	Host	Host Drive	Recovery Drive	Drive Type	Status	Progress	Expected Recovery Point	Actual Recovery Point	Recovery based on	Info Message	Export Message	View
	DR-SERVER	E ( Logistics )	Z	Virtual	Complete	100%	2009/8/12 10:9:13:800	-	Tag Based Tag Pluto Accuracy 	-		
<div><a href="#">Release Drive</a><a href="#">Force Delete</a></div>												

Figure 168:

45. **"Host"**: Target host name
46. **"Host Drive"**: Drive/ volume of which the snapshot/ recovery is taken
47. **"Recovery Drive"**: Snapshot volume
48. **"Drive Type"**: Type of drive used to contain a snapshot, virtual, WAN drive or virtual mount point
49. **"Status"**: will display the result as Complete or failed
50. **"Progress"**: Percentage in progress: if snapshot then 0 to 100 %, if recovery then 2 stage process 0- 100%
51. **"Expected Recovery Point"**: Desired Recovery (shows time to which recovery should be done).
52. **"Actual Recovery Point"**: Recovery achieved by the agent (shows time to which recovery is achieved).
53. **"Recovery based on"**: Time, event etc.
54. **"Info message"**: If any error during the operation of snapshot/recovery than the error message is displayed here.

## 5.2.7 Target Drive Rollback Status

When you perform a target volume rollback the status can be seen here (and in the “**Protection Status**”).





Target Drive Rollback Status								
	Host	Rollback Drive	Status	Progress	Expected Recovery Point	Actual Recovery Point	Recovery based on	Info Message
	DR-SERVER	G ( Engineering )	Complete	100%	2009/8/12 7:11:34:291	-	Tag Based Tag Beta Accuracy 	-
	DR-SERVER	F ( Human Resources )	Complete	100%	2009/8/12 7:11:34:291	-	Tag Based Tag Beta Accuracy 	-
<div>Release Drive</div>								

Figure 169:

- 55. “**Host**”: Target Host where the volume is being rolled back
- 56. “**Rollback Drive**”: Volume being rolled back.
- 57. “**Status**”: Status of the operation, complete, failed or in progress
- 58. “**Progress**”: % completed
- 59. “**Expected Recovery Point**”: Point in Time where it the recovery has to be made
- 60. “**Actual Recovery point**”: Recovery actually done
- 61. “**Recovery based on**”: Time based or Tag based with the name of the tag is displayed here
- 62. “**Info Message**”: error messages etc

## Recovering Solaris 10 (with zones)

Follow the below steps to recovery zones on Solaris 10.

**Step 97.** You may either choose to rollback the target volumes or perform a virtual snapshot for each replication pair (associated with protecting zones) to a common consistency tag or time. Refer to the section [Recovery](#) section on page 115

**Step 98.** Mount the virtual snapshots (or rolled back target volumes) on the target with their respective source mount point names.

**Step 99.** Navigate to the directory where the zone configuration files are located and apply the configuration file on the target machine. Before applying this file into the target system, edit the file and change the IP address as required and other disk, mount point information if it is different in target from the source system. Also if your zone configuration file contains capped memory then you have problem while applying the configuration file in target it may throw an error “The capped-memory resource or a related resource control already exists and cannot set a resource-specific property from the global scope.” The problem can be resolved and create a zone with the requested settings by moving the add capped-memory section in front of the zone.max-swap setting in configuration file.

```
bash-3.00# zonecfg -z TestZone -f TestZone.config
```

**Step 100.** Although the zone is now created, it is unattached. You will now need to attach the zone.

```
bash-3.00# zoneadm -z TestZone attach -u
```

(-u option will help, in case Source and Target OS update version exactly not matching)  
Zone should now be attached without any issue.

**Step 101.** Boot the zone  

```
-bash-3.00# zoneadm -z TestZone boot
```

**Step 102.** Once the zone is up, you can access the zone using ssh/telnet based.  
For console login  

```
-bash-3.00# zlogin -C TestZone
```

**Step 103.** The target zone will come up with the same IP address as of the source zone, We recommend you halt the source zone before booting the target zone, this enables the new zone to appear in the network without any IP conflict



### Notes:

To Halt a zone use the “-bash-3.00# zoneadm -z TestZone halt” command



# Part 4: Monitoring

This part contains chapters 6, 7, 8 and 9

## **Chapter 6**

**“Protection Status”** on the CX UI is explained in this chapter

## **Chapter 7**

Traps, Email alerts and agent settings are explained in this chapter

## **Chapter 8**

Log files are generated at the CX server and at the agent side. Each of these log files record a specific set of operations. This chapter covers both CX and agent log files.

## **Chapter 9**

Analyzer and trending on the CX user interface are covered in this chapter.

## 6 Checking Status

In this chapter, we will learn to monitor active operations through the “**Protection Status**” on the CX UI. This screen is a read only screen and does not accept any user input anywhere. The sole purpose of this screen is to show the status of any active operations such as replications (both FX and VX), recovery operations etc.

After this chapter you will be able to

Monitor VX replication

- Understand each of the fields under “**Volume protection Status**” and interpret their values
- Differentiate when a target volume is visible
- Determine when a pair is throttled and how to resolve it.

Monitor recovery operations

Monitor FX replication

- FX executions and their logs
- FX trending

## 6.1 Protection Status

### 6.1.1 Protection Status -Volumes

“Protection Status” is used to monitor

- Active Volume replication
- File replication history and
- Any active recovery operations

“Protection Status” is the third icon on the management panel (to the left). A fresh installation of CX will not have any replication pairs and this screen will be blank. There are three tabs in this screen:

**Volumes:** All VX replication pairs are listed under the “Volume Protection Status”. You can limit the number of replication pairs shown per screen through the “List # Records/Page”. You can also use the search feature to list the desired replication pairs based on source host, target host or volume (source or target)

Protection Status - Volumes

Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Sep-2-2009 11:29:32

VolumesFilesSnapshots

Source Host

Target Host

Volume

Search

1-4 of 4 Records

List4Records/Page

Page1of 1

Volume Protection Status										
Server	Volume	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)		Resync progress	RPD	Status	Resync Required	View Details
				On CX-PS	On Target					
W2K3E321145PAS1->PROD-SERVER	I ( New Volume ) -> I	0	0	0	0	0 %	0.58 minutes	Resyncing (Step I)	YES	<a href="#">+</a>
PROD-SERVER->W2K3E321145PAS1	E ( Logistics ) -> E	0	0	0	0	N/A	0.45 minutes	Differential Sync	NO	<a href="#">+</a>
PROD-SERVER->W2K3E321145PAS1	F ( Human Resources ) -> F	0	0	0	0	N/A	0.88 minutes	Differential Sync	NO	<a href="#">+</a>
PROD-SERVER->W2K3E321145PAS1	G ( Engineering ) -> G	0	0	0	0	N/A	0.4 minutes	Differential Sync	NO	<a href="#">+</a>

Figure 170:

## Volume Protection Status

The first section under “Protection Status” is “Volume Protection Status”

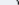
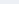
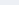
Volume Protection Status													
Server		Volume	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)		Resync progress	RPO		Status	Resync Required	View Details	
					On CX-PS	On Target							
Capacity (MB)	Resync Start Time (Step I)	Resync End Time (Step I)	Resync Start Time (Step II)	Resync End Time (Step II)	Differential Start Time	Last Update Time	Agent Log	Secure CX-PS to Destination	Use Compression	CDP Retention	Resync Mode	Visible	
512.00	2009-08-12 11:57:03	2009-08-12 11:59:18	2009-08-12 11:59:18	2009-08-12 12:02:22	2009-08-12 12:02:22	2009-08-12 16:11:00	Details	No		At Appliance	Configured (Enabled)	Fast	No
PROD-SERVER->DR-SERVER		E ( Logistics ) -> E ( Logistics )		0	0	0	0	N/A	0.35 minutes	Differential Sync		NO	
PROD-SERVER->DR-SERVER		G ( Engineering ) -> G		0	0	0	0	N/A	0.3 minutes	Resyncing (Step II)		YES	
PROD-SERVER->DR-SERVER		F ( Human Resources ) -> F		0	0	0	0	N/A	0.23 minutes	Resyncing (Step II)		YES	

Figure 171:

The first screen has a total of eleven fields each of them denoting information about a specific replication pair.

63. “Server”: Source -> target
64. “Volume”: Source to volume drive mapping
65. “Group”: This displays the group to which the replication belongs
66. “Resync in Transit Step 1(MB)”: Resync data in transition while in step 1 of resync
67. “Resync in Transit Step 2(MB)”: Resync data in transition while in step 2 of resync
68. “Differentials left (MB)”: Remaining differential data to be transferred in MB are shown separately on CX and on Target. Differentials on CX indicate that the differentials have not yet moved to the target host. Differentials on the target indicate that dataprotection.exe has not yet applied these differentials to the target volume. Please refer to the section [Volume Replication –VX Agent](#) on page 14 for more information.
69. “Resync progress”: Percentage of resync progress
70. “RPO”: RPO value in minutes (displays up to 2 hours in minutes and then switches to hours )
71. “Status”: Status of the replication pair (Resyncing (Step I), Resyncing (Step II) and Differential Sync)
72. “Resync required”: If the target is required to be synced with source then “YES” otherwise “NO”.
73. “View details”: Click on “+” to view more details regarding the replication pair. By clicking “+” under “view details” reveals eleven fields:

Capacity (MB)	Resync Start Time (Step I)	Resync End Time (Step I)	Resync Start Time (Step II)	Resync End Time (Step II)	Differential Start Time
512.00	2009-08-12 11:57:03	2009-08-12 11:59:18	2009-08-12 11:59:18	2009-08-12 12:02:22	2009-08-12 12:02:22

Figure 172: Protection Status -> Volume protection Status -> View details

Last Update Time	Agent Log	Secure CX-PS to Destination	Use Compression	CDP Retention	Resync Mode	Visible
2009-08-12 16:12:00	Details	No	At Appliance	Configured (Enabled)	Fast	No

Figure 173: Protection Status -> Volume protection Status -> View details

- **“Capacity (MB)”**: Source Volume Capacity
- **“Resync Start Time(Step 1)”**: Time when Initial resync started
- **“Resync End Time (Step 1)”**: Time when Initial resync completed
- **“Resync Start Time (Step2)”** : Start time of resync step2 (data writes occurred during resync step1)
- **“Resync End Time (Step2)”** : End time of resync step2
- **“Differential start time (Step 2)”**: Time when differential sync started
- **“Last update time”**: Last updated time from CX
- **“Agent log”**: Enabled if logs are present
- **“Secure CX to destination”**: Secure transport is enabled or not / Encryption ( yes / no )
- **“Use compression”**: Compression is enabled or not (yes / no). If yes, then either At CX server or at the host.
- **“CDP retention option”**: CDP retention option is configured (than enabled /disabled) or un-configured
- **“Resync mode”**: Type of resync used for this replication pair e.g. Offload, fast or Direct.
- **“Visible”**: Yes if volume is visible No if volume is not visible

## 6.1.2 Protection Status - Files

“File protection Status” shows status of a file replication

**Files:** All FX replication pairs are listed under the “File Protection Status”. You can limit the number of replication pairs shown per screen through the “List # Records/Page”. You can also use the search feature to list the desired replication pairs based on job description, applications, status, group ID, job ID and exist code.

Protection Status - Files												
Logged in as 'admin@10.0.1.30' - <a href="#">Logout</a> Server Time: Sep-2-2009 11:36:05												
<a href="#">Volumes</a> <b><a href="#">Files</a></b> <a href="#">Snapshots</a>												
Job Description: <input type="text"/> Application: <input type="text"/> Status: <input type="text"/> Group ID: <input type="text"/> Job ID: <input type="text"/> Exit Code: <input type="text"/> <input type="button" value="Search"/>												
1-9 of 9 Records List <input type="text" value="9"/> Records/Page Page <input type="text" value="1"/> of 1												
File Protection Status												
View Details	Job Description	Application	Status	Source Host	Source Directory	Target Host	Target Directory	Scheduled Type	GID	JID	Job Instance	Exit Code
<input type="checkbox"/>	complete backup...	Month End jobs	Completed	PROD-SERVER	c:\data	DR-SERVER	h:\data	Run Every	1	1	1	0
<div> <div>More Details</div> <div>Start Time</div> <div>End Time</div> <div>Last Update Time</div> <div>Data Compression</div> <div>Sync Compression</div> <div>Bytes Changed</div> </div> <div> <div>Log</div> <div>Trending</div> </div>												
	2009-09-02 10:59:20	2009-09-02 11:00:36	2009-09-02 11:00:36	1.73	0.00	1494896						
<input type="checkbox"/>	complete backup...	Month End jobs	Completed	PROD-SERVER	c:\data	DR-SERVER	h:\data	Run Every	1	1	2	0
<input type="checkbox"/>	auto backup...	Weekly jobs	Completed	DR-SERVER	h:\data	PROD-SERVER	g:\data1	Run Every	2	2	3	0
<input type="checkbox"/>	auto backup...	Weekly jobs	Completed	DR-SERVER	h:\data	PROD-SERVER	g:\data1	Run Every	2	2	4	0
<input type="checkbox"/>	auto backup...	Weekly jobs	Completed	DR-SERVER	h:\data	PROD-SERVER	g:\data1	Run Every	2	2	5	0
<input type="checkbox"/>	complete backup...	Month End jobs	Completed	PROD-SERVER	c:\data	DR-SERVER	h:\data	Run Every	1	1	6	0
<input type="checkbox"/>	auto backup...	Weekly jobs	Completed	DR-SERVER	h:\data	PROD-SERVER	g:\data1	Run Every	2	2	7	0
<input type="checkbox"/>	auto backup...	Weekly jobs	Completed	DR-SERVER	h:\data	PROD-SERVER	g:\data1	Run Every	2	2	8	0
<input type="checkbox"/>	auto backup...	Weekly jobs	Completed	DR-SERVER	h:\data	PROD-SERVER	g:\data1	Run Every	2	2	9	0
<input type="button" value="Delete all job history"/> <input type="button" value="Clear logs for selected jobs"/>												

Figure 174:


74. “**Job Description**”: A job description is given while setting up a FX job; this can be altered at a later time through “File Protection”. In the “Protection Status” you will only be able to see the status. This text box is used to set a filter for filtering out required jobs and shows the instance of executing job corresponding to the group and job id. For example start a job in Group 1 and when you run the job 5 times there will be 5 instances with the same group id.
75. “**Application**” is a drop down box with the list of all the applications (set at the time of setting up an FX job). Again this can also be altered through “File Protection”
76. “**Status**” field has a drop down box with a list of status all of the jobs belong to. Displays as starting, target starting and running and at last completed
77. “**Source Host**” contains the name of the host acting as a source for this FX job.
78. “**Source Directory**” is the folder being replicated to the “Target Directory” on the “Target host”

79. **“Scheduled Type”** shows when the job is scheduled to execute.
80. **“GID”** Group ID explains about the scheduling for Groups IDs like Scheduling 1 or 2, 3 jobs in one group
81. **JID:** Is meant for each individual Job id when it starts fresh job Replication
82. **“Job Instance”:** Every time a job runs it is given a unique number called job instance. It will have its own logs etc. If the same job runs again, it will be called as a separate instance. For e.g. if a job runs for ten times there will be ten instances of the same job
83. **Exit code:** Displays a number 0 if job success else it displays the corresponding number (click on exit code to display the list of possible error codes)
84. The check box to the end is used for the command buttons at the bottom of the page.
85. **“Clear logs for selected jobs”** deletes all the log files for the selected jobs
86. **“Delete all job history”** removes all traces of the job from the UI
87. The **“+”** View Details shows more details:
  - **“Start time”:** Time when the replication started (time)
  - **“End time”:** Time when the replication ended (time)
  - **“Last update time”:** Displays last activity of replication (time)
  - **“Data compression”:** Data compressed from source to target (number). A greater number indicates better compression. This is available once the job execution is complete. This requires “compress files” to be enabled under “File/Directory Options”
  - **“Sync compression”:** This displays the percentage saving achieved by transferring differentials (i.e. transferring the changed bytes only). The formula is  $1 - (\text{Actual Bytes transferred} / \text{Total Replication size}) \times 100$ . A greater number indicates better performance.

More Details	Start Time	End Time	Last Update Time	Data Compression	Sync Compression	Bytes Changed
Log Trending	2009-09-02 10:59:20	2009-09-02 11:00:36	2009-09-02 11:00:36	1.73	0.00	1494896

**Figure 175: Sync compression Example**

- **“Bytes changed”:** Total number of bytes transferred from source to target during that particular schedule.
- **“Log and Trending”:** Under more details you can see “log” and “Trending”. “Log” displays all the activity of the replication and “trending” displays corresponding graphs of total bytes sent, and total bytes sent (changed bytes). It also displays data compression and sync compression percentages.



**Figure 176: FX graph per execution**

If the FX “**log**” option in protection status is beyond 1 GB then Internet explorer will not be able to handle the file, although Firefox will support up to 2 GB, the recommended workaround is to use “**low verbosity**” in the job options.

The fields Data compression, Sync compression and Trending will be active once the “**Status**” of the job is “**Completed**”



## 6.1.3 Protection Status - Snapshots

### (Snapshot / Recovery / Rollback) Pair Status

All recovery operations are displayed under this section. Once the process is complete they are removed from "Pair status".

Protection Status - Snapshots											
Volumes Files Snapshots											
(Snapshot / Recovery / Rollback) Pair Status											
Host	Host Drive	Snapshot / Recovery / Rollback Drive	Drive Type	Progress	Start Time	End Time	Expected Recovery Point	Actual Recovery Point	Status	Info Message	
DR-SERVER (Recovery)	E	Y	Virtual Drive	0%	N/A	N/A	2009/9/2 5:27:30:520	-	Ready	-	
DR-SERVER (Recovery)	F	X	Virtual Drive	0%	N/A	N/A	2009/9/2 5:27:30:520	-	Ready	-	

Figure 177:

88. **"Host":** Target host name
89. **"Host Drive":** Drive/ volume of which the snapshot/ recovery is taken
90. **"Snapshot /Recovery / Rollback Drive"**
91. **"Drive Type":** Type of drive used to contain a snapshot, virtual, WAN drive or virtual mount point
92. **"Progress":** Percentage in progress: if snapshot then 0 to 100 %, if recovery then 2 stage process 0-100%
93. **"Start time":** start time of the snapshot.
94. **"End time":** end time of the snapshot.
95. **"Expected Recovery Point":** Desired Recovery (shows time to which recovery should be done).
96. **"Actual Recovery Point":** Recovery achieved by the agent (shows time to which recovery is achieved).
97. **"Status":** Status of snapshot (queued, ready, In progress, completed) /recovery (queued, ready, snapshot In progress, snapshot completed, recovery in progress, completed)
98. **"Info message":** If any error during the operation of snapshot/recovery than the error message is displayed here.

## 6.2 Possible situations while replication

### 6.2.1 Resync required set to Yes

Occasionally the “**resync required**” field is set to “**Yes**” indicating data inconsistencies between the source and target volumes. Given below is the list of conditions and corrective action to be taken to ensure data consistency. Each of the condition will again generate an email alert and a trap.

**Table 14**

Action	Reaction
Target Volume visible in Read/Write mode	Perform a Resync or set Auto resync options
Source or target host down for a long time	Perform a Resync or set Auto resync options
CX failover	Perform a Resync or set Auto resync options
CX backup restored	Perform a Resync or set Auto resync options
Source Volume Resized	Resize the target volume, delete bitmap files on the source host and resume replication from CX UI.

## 6.2.2 Resizing source volume

At times you may require resizing the source volume capacity. You may follow any of the two recommended procedures

### Resize source volume

**Recommended Procedure 1:** Resize the target volume before resizing the source volume

This is achieved in nine sequential steps

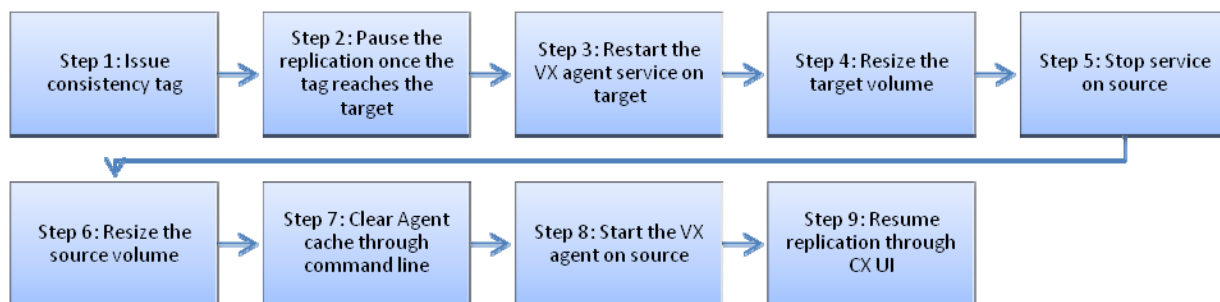


Figure 178

**Step 104.** Issue consistency tag: A vacp consistency tag is issued on the source volume as a precaution. This will enable you to recover back to the consistency tag when required.

**Step 105.** Pause replication: Once the consistency tag reaches the target, [pause the replication pair](#) from the CX UI

**Step 106.** Access the target host and restart the VX agent

**Step 107.** Resize the target volume such that its either equal or larger than the source volume after resize

**Step 108.** Stop the VX agent service on the source host

**Step 109.** Resize the source volume. Ensure that the target volume is not smaller than the source volume

**Step 110.** Navigate to the VX agent install path and issue the following command to clear the agent cache

Windows: `drvutil --stopfiltering <drive letter or mount point> -deletebitmap`

linux: `inmstkops-bin r <protected volume name>`

**Step 111.** Start the source VX agent

**Step 112.** Resume the replication pair through CX UI and the replication pair should progress normally.

## Alternative Procedure 2: Resizing the target volume after resizing the source volume

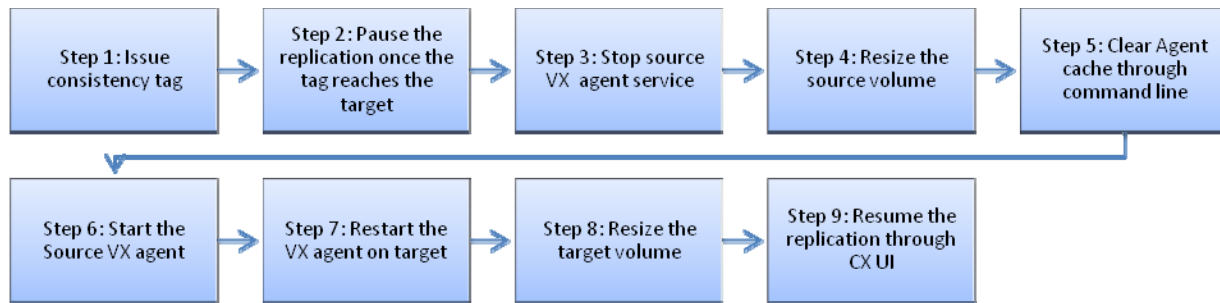


Figure 179

**Step 113.** Issue consistency tag

**Step 114.** Pause the replication pair once the tag reaches the target host

**Step 115.** Stop the VX agent service on the source

**Step 116.** Resize the source volume as desired

**Step 117.** Navigate to the VX agent install path and issue the following command to clear the agent cache

Windows: **drvutil --stopfiltering <drive letter or mount point> -deletebitmap**

linux: **inmstkops-bin r <protected volume name>**

**Step 118.** Start source VX agent

**Step 119.** Restart target VX agent

**Step 120.** Resize the target volume such that its equal or larger than the source volume

**Step 121.** Resume the replication pair through CX UI and the replication pair should progress normally.

## Resuming from a source volume resize

When the source volume has been resized without following the recommended procedures, there is a six step procedure to continue with the replication pair.

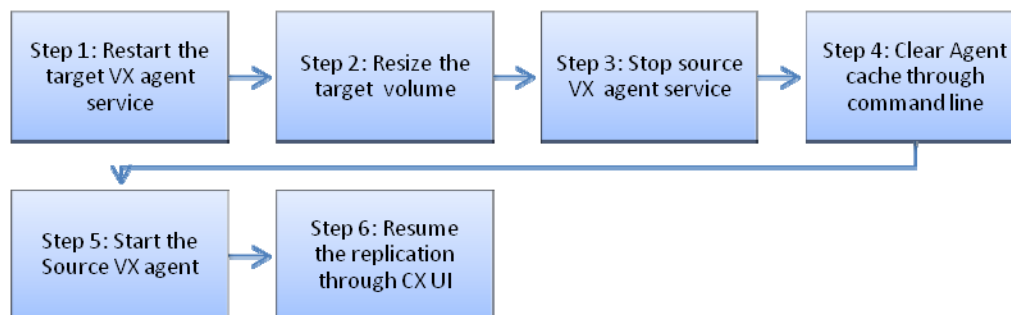


Figure 180

**Step 122.** You should see that the replication pair is paused. You will also be intimated about the source volume resize through email and trap alerts if they are enabled.

**Step 123.** Restart the VX agent service on the target host

**Step 124.** Resize the target volume such that its either equal or larger than the source volume

**Step 125.** Stop the VX agent service on the source host

**Step 126.** Navigate to the VX agent install path and issue the following command to clear the agent cache

Windows: **drvutil --stopfiltering <drive letter or mount point> -deletebitmap**

linux: **inmstkops-bin r <protected volume name>**

**Step 127.** Start the Service on the Source

**Step 128.** Resume the replication through the CX-CS UI



### Notes:

Source volume resize does not support Volpack or Virtual Volume.

Resize does not support shrinking of volume size.

Do not perform a source volume resize while the replication pair is in Resync step 1

It is recommended to perform a resync immediately after following any of the above three procedures

By following these procedures you will be able to maintain the older retention logs and also perform recover operations to events or time before the volume resize

## 6.2.3 Throttling

### What is throttling?

Throttling is a mechanism where the source agent stops sending differential data to the CX server, this allows draining of all the differentials from the CX to the target, once all the differentials are drained then the source agent resumes sending the data.

Volume Protection Status				
Server		Volume	Group	
LIN-PROD->LIN-DR		/dev/mapper/volume-v1 - > /dev/mapper/volume-v1 [Throttled]	Volume /dev/mapper/volume-v1	
Capacity (MB)	Resync Start Time (Step I)	Resync End Time (Step I)	Resync Start Time (Step II)	Resync End (Step II)
1024.00	2008-10-31 04:28:29	2008-10-31 04:34:20	2008-10-31 04:34:20	2008-10-31 04:36:10

Figure 181:

### When does it happen?

This happens primarily when the target host is lagging behind. Reasons include

- Low bandwidth,
- Large amounts of data changes on the source volume,
- If the cache folder for the replication pair is filled up on the CX server.
- WAN outage

### What should be done?

After a replication pair is set, there are two values under the “[Pair Settings](#)” (page 63), one is “**Resync Files Threshold (MB)**”, and the other is “**Differential Files Threshold (MB)**”.

When the specified amount of differential data is accumulated on the CX server, the pair is throttled. This can happen when the pair is in “**Resync**” or “**Differential sync**”. The value can be set under “**Replication Pair Details -> pair settings**”

Resync Files Threshold (MB)	Differential Files Threshold (MB)
2048	8192

Figure 182:

The default value for differential threshold is 8GB and resync threshold is 2 GB. You may choose to tweak these values to enhance performance. You may also consider increasing bandwidth or tweaks bandwidth policies

You may identify and move (or delete) large files on the CX-PS cache through the command “**du / --max-depth=2 | sort -n**”. This should solve the issue where CX-PS cache is full.

## 6.2.4 Disk Space Warning Threshold



Figure 183

If “/”, “/var/log”, “/home/svsystems” on the CX exceeds 80% disk usage then throttling begins. You may change this % through the CX UI by clicking on “**System**” -> “**CX Settings**”-> “**Disk Space Warning Threshold**”

During throttling the source host stops sending data to the CX-PS. The CX-PS server on the other hand sends data to the target host and the target is updated, once the differentials are drained on the CX-PS, the replication pair is restored to its original state.

### FTP Mode

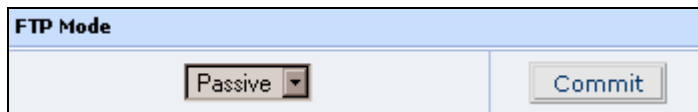


Figure 184

FTP is used during data transfer from CX-PS to remote site, by default this is “**Active**” FTP, to change that select “**Passive**” from the drop down and then click on “**Commit**”. This is particularly used when a firewall is in place. Refer to section “**Firewall considerations**” in the Installation guide.



#### Notes:

While using “**Active**” FTP on windows 2008, disable firewall else the firewall will block all incoming connections.



## 6.3 Agent Heartbeat

Agent heartbeat shows all the agents and process servers pointed to the CX server.

You will find Hostnames, IP address, VX source (sentinel), VX target (Outpost) File replication agent and process Server fields here. This is used mainly to check the last time when the Agent sent an “alive” signal to the CX-CS

System: Agent Heartbeat

v5.10.1.BETA0.1679.1 (RELEASE\_5-10-1\_BETA0\_1679\_Aug\_07\_2009\_InMage).

Logged in as 'admin@10.0.1.30' - [Logout](#)

Server Time: Aug-12-2009 16:15:23

[Monitoring](#)

[Bandwidth Shaping](#)

[CX Settings](#)

[Agent Settings](#)

[License Management](#)

[Agent Heartbeat](#)

[Process Server Traffic Load Balancing](#)

[Process Server Failover](#)

[Remote](#)

[CX](#)

[Versions and Patches](#)

[User Documents](#)

[Logs](#)

[Installers](#)

[RX Settings](#)

Agent Heartbeat

Hostname	IP Address	Sentinel	Outpost Agent	File Replication	Process Server
DR-SERVER	10.0.145.46	2009-08-12 16:15:19	2009-08-12 16:15:19	2009-08-12 16:15:02	
imits030.qa-domain.net	10.0.1.30				2009-08-12 16:15:00
PROD-SERVER	10.0.145.45	2009-08-12 16:15:04	2009-08-12 16:15:04	2009-08-12 16:14:58	
IMITS088	10.0.1.88	2009-08-12 16:13:46	2009-08-12 16:13:46	2009-08-12 16:14:57	
imits145.dev-domain.net	10.0.1.145				2009-08-12 16:14:53

Figure 185

## 7 Configuring Alerts

This chapter explains:-

- Configuring SNMP traps
- Configuring email alerts and
- Modifying agent settings

### 7.1 Trap Listeners

Traps are used to transmit messages from CX-CS to any other host that is configured in the “**Add Trap Listener**”. This is done through SNMP protocol and port 162 is used for this. Traps are used to transmit messages, events etc specific to VX agent. A trap receiving software needs to be installed on the target system (or any other system configured through the UI) which will interpret alerts sent by the CX-CS. The target system can be anywhere as long as it is accessible

A host can be added to interpret traps through the UI. Navigate to “**System→Monitoring→**” and click “**Edit**” button of “**Configured System Administrators**” then enter the host name or the IP address under “**Add trap listener**” and click on “**Add**”.

Add Trap Listener		
Trap Listener (IP/Hostname)	Trap Port	
<input type="text" value="10.0.27.35"/>	<input type="text" value="162"/>	<input type="button" value="Add"/>

Figure 186

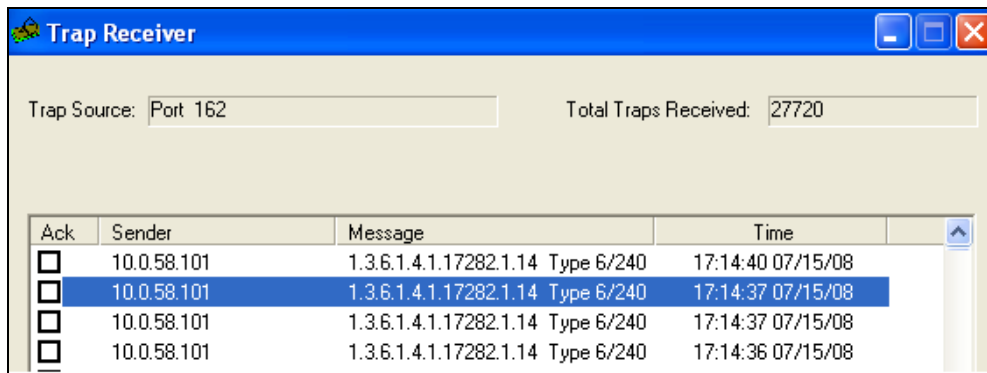
If a host name is to be used here, ensure that a fully qualified host name is used (e.g., Neptune.drsite.in, where Neptune is the host name and drsite.in is the domain name).

After adding a host, it will appear in under the configured trap listeners. This is also the place where a trap listener can be deleted.

Configured Trap Listeners		
	Trap Listener	Trap Port
<input checked="" type="radio"/>	10.0.27.35	162

Figure 187

You may choose to use any of the trap listeners and listen to the trap port (default being 162). The CX server sends out traps to the email address that was registered. A trap message will contain a number as shown in the picture below.



Ack	Sender	Message	Time
<input type="checkbox"/>	10.0.58.101	1.3.6.1.4.1.17282.1.14 Type 6/240	17:14:40 07/15/08
<input type="checkbox"/>	10.0.58.101	1.3.6.1.4.1.17282.1.14 Type 6/240	17:14:37 07/15/08
<input type="checkbox"/>	10.0.58.101	1.3.6.1.4.1.17282.1.14 Type 6/240	17:14:37 07/15/08
<input type="checkbox"/>	10.0.58.101	1.3.6.1.4.1.17282.1.14 Type 6/240	17:14:36 07/15/08

**Figure 188**

The table below explains the reasons for the alert and corrective action to be taken.

The table here gives an overview of traps their conditions and corrective measures to be taken

**Table 15: Types of Traps**

Trap Code	Reasons / conditions that can generate the Alert / Trap	Corrective action to be taken
240	When any host installed with either VX or FX or both Agents do not have the corresponding Scout FX or VX service running or the host itself was powered-down or the concerned host was not reachable through the network.	Start agent (VX,FX) service if stopped. Check network connectivity. Contact Hitachi Data Systems or the concerned Representative if issue persists
260	CX server's / or /var/log or /home partitions may have less space than the default or set size limit (default=80%)	Increase the disk space on the CX server / or <b>"/var/log"</b> or <b>"/home"</b> partitions
280	RPO value for that replication pair was more than the set value or File Replication has not made progress within the last one minute	Make sure that the replication for that corresponding pair is not stopped or stuck. VX replication pair will progress if the target is made visible
265	There is a problem with storage mounted.	Address the storage attached and then reboot the system
215	When the Target volume of a replication pair is not in sync with the source [due to sudden power failure during replication or after a Visible-RW of the Target volume resync was not forced]	Perform a Resync on that replication pair
315	This code starts one day before the VX license will expire and continue to appear even after the license is expired	Contact Hitachi Data Systems or the concerned Representative for a new license
310	This code starts one day before the FX license will expire and continue to appear even after the license is expired	
320	This code starts one day before the CX license will expire and continue to appear even after the license is expired	
330	VX license will expire in seven days	
325	FX license will expire in seven days	
335	CX license will expire in seven days	
250	Source Volume Resized	Please refer to the section <a href="#">Preparing for source volume resize</a> on page 155

## 7.2 Adding Users and Email Alerts

### 7.2.1 Types of Users

There are two types of users that can access the CX-CS UI, the administrators and the users.

Administrators have full access to the CX-CS UI while the user's access is limited to monitoring active backups and managing log files. The following table lists out operations accessible for administrators and users. You can view, edit, add or delete administrators, users and configure e-mail alerts specific to each user account. User accounts on the other hand can view, edit and delete their own accounts and cannot do the same on any administrator account.

**Table 16:**

Operation	Administrator	User
Volume Protection	Full access	No access
File Protection	Full access	No access
Protection Status	Full access	Full access
Recovery	Full access	No access
System	Full access	Limited access
Trending	Full access	Full access

As an administrator you can view, edit, add or delete administrators, users and configure e-mail alerts specific to each user account. Only administrators can add/ remove agents to the CX-CS UI, Set, edit or stop replication pairs (FX and VX) while users can monitor replication pair progress, agent heart beat, manage log files etc.

### 7.2.2 To add an Administrator or User

99. On the “**System**” menu click on the “**Monitoring**” link. The “**Configured System Administrators**” window appears.

100. You can view, edit or delete administrators through this screen. Also email alerts can be configured through the same screen.

101. To add a new user click on “**Add Administrator**”.

Configured System Administrators					
UID	Username	Full Name	View	Edit	Delete
1000000000	admin	CX Administrator	<input type="button" value="View"/>	<input type="button" value="Edit"/>	<input type="button" value="Delete"/>
<input type="button" value="Add Administrator"/>					

**Figure 189 Adding a System Administration**

The “Add Administrator” screen opens up

Add Administrator	
Administrator Details	
Full Name:	backup
UID:	1945646068
User Name:	backupadmin
Admin Access:	<input checked="" type="checkbox"/>
Password	
Enter Password:	*****
Re-enter Password:	*****
E-mail Notification	
E-mail Address:	backup@inmage.net
<input type="button" value="Accept"/> <input type="button" value="Cancel"/>	

Figure 190 Add Administrator

- Step 129.** Enter the full name of the administrator then enter the user name
- Step 130.** To create an administrator account, check the “Admin Access” option. When the option is not checked a user account is created.
- Step 131.** Enter a password for the user and re-enter it to confirm it.
- Step 132.** Enter the e-mail address to be used for CX-CS generated email alerts.
- Step 133.** Click on “Accept” to create the user, a dialog box appears indicating that the user is created successfully, click on “Ok”.
- Step 134.** The “Configured System Administrators” screen appears with the list of added users.

Configured System Administrators					
UID	Username	Full Name	View	Edit	Delete
673325370	backupadmin	backup	<input type="button" value="View"/>	<input type="button" value="Edit"/>	<input type="button" value="Delete"/>
1000000000	admin	CX Administrator	<input type="button" value="View"/>	<input type="button" value="Edit"/>	<input type="button" value="Delete"/>
<input type="button" value="Add Administrator"/>					

Figure 191 Configuration Screen with new administrator details

You can see user details by clicking on the “View” button. You can delete the user account by clicking on “Delete” button.

To modify the user details or to tweak email alerts Click on the “Edit” button for the corresponding user. The “Administrator Configuration” screen appears, there are three sections in this screen

102. **“Update Administrator Settings”** used to change the user name, user password, email id used for CX server’s email alerts and the time interval of email alerts. The default time between email alerts is 30 minutes. All alerts are consolidated into one single email and sent once every 30 minutes.

Update Administrator Settings		
Administrator		New Values
Full Name:	backup	<input type="text"/>
User Name:	backupadmin	<input type="text"/>
Admin Access:	Enabled	<input checked="" type="checkbox"/>
Password:	Old Password:	<input type="password"/>
	New Password:	<input type="password"/>
	Re-type New Password:	<input type="password"/>
E-mail Address:	backup@inmage.net	<input type="text"/>
E-mail Dispatch Interval:	30 minutes	<input type="text"/>
<div>Return Accept</div>		

**Figure 192:**

103. “**Alert Notification**” is used to select only the required email alerts to trap, while the others are ignored. An alert is sent out at the specified interval if any of the following conditions are met. You can check for e-mail alerts or trap listener or both to get alerts.

Alert Notification		
Alert Category	E-mail	Trap
	<input type="checkbox"/>	<input type="checkbox"/>
RPO SLA threshold exceeded	<input type="checkbox"/>	<input type="checkbox"/>
Resync required for VX replication pair	<input type="checkbox"/>	<input type="checkbox"/>
VX/FX/Switch agent not responding	<input type="checkbox"/>	<input type="checkbox"/>
CX secondary storage warnings and alerts	<input type="checkbox"/>	<input type="checkbox"/>
Issues with FX jobs and pre/post scripts	<input type="checkbox"/>	<input type="checkbox"/>
Agent logged alert message	<input type="checkbox"/>	<input type="checkbox"/>
CX, VX, FX license expiry and related issues	<input type="checkbox"/>	<input type="checkbox"/>
Bandwidth shaping alerts	<input type="checkbox"/>	<input type="checkbox"/>
Debug information for CX services	<input type="checkbox"/>	<input type="checkbox"/>
Capacity threshold exceeded	<input type="checkbox"/>	<input type="checkbox"/>
Capacity Utilization Reached Limit	<input type="checkbox"/>	<input type="checkbox"/>
Health Report every <input type="text" value="0"/> day(s)	<input type="checkbox"/>	<input type="checkbox"/>
Move Retention Log	<input type="checkbox"/>	<input type="checkbox"/>
Insufficient Retention Space	<input type="checkbox"/>	<input type="checkbox"/>
Source volume resize	<input type="checkbox"/>	<input type="checkbox"/>
Process server un-install	<input type="checkbox"/>	<input type="checkbox"/>
CX Node Failover	<input type="checkbox"/>	<input type="checkbox"/>
Process Server Failover	<input type="checkbox"/>	<input type="checkbox"/>

**Figure 193:**



There are a total of eighteen types of email alerts. The following table describes the conditions that raise a alert and corrective action for the same.

**Table 17**

Event Description	Reasons / conditions that can generate	Corrective action to be taken
RPO SLA threshold exceeded	a) The target can drain slowly.	Monitor the services at both source and target and CX and ensure the services are running.
Resync required for VX replication pair	a) The re-sync required might be set due to data inconsistency. b) On source volume resize c) If the target drive is exposed in read/write mode. d) Process server fail over for the replication pairs e) Configuration server fail over in case of the High availability scenario	Restart a re-sync from UI or configure a auto re-sync policy to get the pair to do restart
VX/FX/Switch agent not responding	This email is sent when the agent is unable to communicate with the CX server within 900 seconds. Either the 1. Agent service may not be started 2. A fire wall may be blocking the agent, 3. Network failure or 4. The host itself may be down  The timeout may be configured under <b>"System-&gt; Agent settings"</b>	Disable any firewalls (or allow FX agent through the firewall). Check and ensure that FX service is running Ping the target and CX servers to see if the host can communicate with them and vice versa. Check if agent is pointed to the right CX server
CX secondary storage warnings and alerts	If the disk space at CX exceeds than disk space warning threshold, which is configured from UI.	An email alert is sent when the disk usage has reached 80% (as configured under <b>"System-&gt; CX settings"</b> ) for the following four volumes on the CX server. <ul style="list-style-type: none"> <li>• /</li> <li>• /home</li> <li>• /tmp and</li> <li>• /var</li> </ul>
Agent logged alert message	An email alert is sent out when the VX or FX agent has encountered an error.	Check the Agent log
CX, VX, FX license expiry and related issues	Email alert is sent seven days before license expiry and continues until the new license is uploaded.	Contact Hitachi Data Systems or the concerned Representative for obtaining new Licenses

Bandwidth shaping alerts	An email alert is sent out when a policy is created, deleted, or is invalid (Trap listener is not available for this alert.)	Monitor the bandwidth as per the configured policy
Capacity threshold exceeded	This is applicable for capacity based license. When the threshold is exceeded as set under “ <b>System-&gt; license management - Set Capacity Utilization Threshold</b> ” an email alert is sent out. (You can get only e-mail alert. Trap listener is not available for this alert.)	Upgrade license to higher capacity
Capacity utilization reached limit	Again this option is applicable for capacity based licensing, when the full capacity of the license is consumed by replication pairs, an email alert is sent out until a higher capacity license is uploaded or a replication pair is removed in order to adjust within the defined license capacity. (You can get only e-mail alert. Trap listener is not available for this alert.)	Upgrade license to higher capacity
Health report every	Enable this option to receive a <a href="#">protection report</a> . By default the interval is set to one day. Protection report will always. (You can get only e-mail alert. Trap listener is not available for this alert.)	if the protection coverage is lesser as indicated by the report, monitor for the following on your system: a) RPO characteristics for the volume b) Has the pair been throttled? c) Has resync required been set to yes
Source volume resize	If the source volume capacity is resized greater than the current configured source volume pair	As per the documentation.
CX node Failover	CX HA failover. Primary CX server has gone down (This alert is only available in HA set-up)	Bring the primary CX online and perform a failback
Debug information for CX services	An email is sent out if there are any errors in the host logs (This alert is not displayed on Dashboard) (You can get only e-mail alert. Trap listener is not available for this alert.)	Check the host logs
Issues with FX jobs and pre/post scripts	FX job failed while replicating or script path error (non-existent script path or name) (You can get only e-mail alert. Trap listener is not available for this	Rectify the script path Check service privileges Disable firewall for FX agent

	alert.)	
Move Retention Log	When retention logs are moved	This message is for information only.
Insufficient Retention Space	Retention directory exceeding configured space	Either edit the retention policy and increase the retention space to accommodate more logs. OR move the retention directory to a larger volume with enough space and increase the size of the retention directory.
Process server un-install	When a process server pointed to this CX-CS is uninstalled (You can get only e-mail alert. Trap listener is not available for this alert.)	This message is for information only.
CX Node Failover	CX HA failover	Bring the primary node online and perform a fallback
Process Server Failover	This occurs when a replication pair is moved to a different process server.	Perform a manual resync.

You may also choose to set a custom subject to the CX alert emails through the “E-Mail Subject”

The screenshot shows a window titled "E-Mail Subject". Inside, there is a "Subject" label followed by a text input field containing "CX.Alert". At the bottom of the window, there are two buttons: "Return" and "Accept".

Figure 194

#### 104. “CX Auto Timeout”

A user will be logged off automatically after a period of inactivity. By default the value is set to three hours (180 minutes). Enter a time in minutes starting from 1 to 9999 minutes. For e.g. when a user has logged into the CX UI and has been idle for three hours (or the configured amount of time), then he will be logged off automatically.

The screenshot shows a window titled "CX Auto Timeout". Inside, there is a label "Auto Timeout for CX:" followed by a text input field containing "180" and a label "(1-9999) min". At the bottom of the window, there are two buttons: "Return" and "Accept".

Figure 195:



Notes:

Click “Return” to cancel all the changes. You can also configure the email alerts generated by the CX server to be sent to the administrator

### 7.2.3 Mail server settings

The “**Mail Settings**” (under “**System-> Monitoring**”) provides an interface to configure the email server. By default the mail server host name is set as local host. When you change any of the mail server settings, enter the mail server name in this field.

CX Server Name/ID: The name entered here will reflect in all email alerts sent by this CX server. This is helpful while working with multiple CX servers.

Mail Settings	
Mail Server Hostname	CX Server Name/ID
<input type="text" value="Bunker"/>	<input type="text" value="Afa-CX"/>
<input type="button" value="Update"/>	

Figure 196

## 7.3 Agent Settings

### 7.3.1 Agent settings

You can specify the timeout duration for the agents with respect to the CX server. The CX server waits for a response from the agents for a specified time. If there is no response from the agent after the specified time, the connection is automatically timed out.

Agent Settings					
	Server	Agent Type	Agent Timeout (Seconds)	CX NAT IP	Alias
<input checked="" type="radio"/>	InMageProfiler	VX	<input type="text" value="900"/>	<input type="text"/>	<input type="text"/>
<input checked="" type="radio"/>	PROD-SERVER	VX	<input type="text" value="900"/>	<input type="text"/>	<input type="text"/>
<input checked="" type="radio"/>	DR-SERVER	VX	<input type="text" value="900"/>	<input type="text"/>	<input type="text"/>

Figure 197 Change Agent Settings

105. On the “**System**” menu click on the “**Agent Settings**” tab. The Agent Settings screen appears

106. Select the check box for the agent for which you want to specify the timeout.

107. Enter the “**timeout**” for the agent and click “**Change Settings**”

108. “**CX NAT IP**”: enter the NAT IP address for the CX server. (This is the place to set NAT IP for CX server)

109. **Alias**: You may choose to give an alias name for any host on the CX UI. This will not change the actual host name but show the alias name within brackets on the CX UI.



#### Notes:

The default agent timeout value is 900 seconds.

### 7.3.2 Process Server

You may assign a NAT IP to a process server under the “**System -> Agent settings**” screen. Select the desired process server then enter a NAT IP and click on “**Change Settings**”

Process Server		
	IP Address	NAT IP Address
<input checked="" type="radio"/>	PROCESSSERVER2.INIMAGE.IN (10.0.243.2)	<input type="text" value="221.134.125.11"/>
<input type="button" value="Change Settings"/>		

Figure 198

## 8 Log Management

This chapter deals with log management. After this chapter you will be able to:

- Set space utilization policies for logs
- Getting CX logs to send to support personnel

### 8.1 Host – Log details

For better support, you might need to download Scout system logs files and pass it across the Scout support personnel. CX Server administrator console provides a feature to just do that.

**Step 135.** Click on “**System**” menu and navigate to “**Logs**” tab, you can further find navigate to “**Host Logs**” “**CX Logs**” and “**Download Logs**”



Notes:

Host logs are generated by the source and target hosts, where VX and FX agents are installed. VX agents by default shall log all severe and critical errors to host logs on the CX Server.

**Step 136.** Navigate to the desired host and click on it to expand the list of log files for the selected hosts

Host - Log Details
<u>DR-SERVER</u>
imits030.qa-domain.net
<u>PROD-SERVER</u>
<u>IMITS088</u>
imits145.dev-domain.net

Figure 199 Host Log files browse by host

Host Log files nomenclature

Table 18: Host logs

Log name	Description
driveletter_rpo	View or Download the RPO trends observed for the drive (e.g., H_rpo)
driveletter_sentinel	View or Download VX Source (sentinel) logs for that particular drive (e.g., H_sentinel)
hosts	View or Download all the severe and critical errors for that hosts reported by the agents
Driveletter_perf	View or download all log info about differentials

## 8.2 CX - Log Details

Click on “CX Logs” to download all logs pertaining to CX logs

<b>CX - Log Details</b>
<a href="#">tman_monitor_ps</a>
<a href="#">audit</a>
<a href="#">tman_volsync</a>
<a href="#">Message</a>
<a href="#">tmanager</a>
<a href="#">tman_monitor</a>
<a href="#">bpmtrace</a>
<a href="#">perf</a>
<a href="#">tman_monitor_disks</a>
<a href="#">mrtgtrace</a>
<a href="#">tmanager_ps</a>
<a href="#">TrapLog</a>
<a href="#">xferlog</a>
<a href="#">rsyncd</a>

Figure 200 CX log files

CX Log files nomenclature

Table 19: CX Logs

Log name	Description
Audit	View or Download log of User actions
tman_monitor_disks	View CX internal database connectivity, logs offline disks
tmanager_ps	View or Download logged messages for volume and file monitoring process for process server
Tmanager	View or Download logged CX messages
tman_monitor	View or Download logged messages for volume and file monitoring process
Message	View or Download logged Scheduler messages.
tman_volsync	View or Download logged tman_volsync messages
Bpmtrace	View or Download logged VX Bandwidth Shaping messages
Mrtgtrace	
tman_monitor_ps	
TrapLog	View or Download CX logged trap messages
Xferlog	
rsyncd	



### 8.3 Download Logs

To download all logs in a compressed achieve, select the required hosts and/or RPO and/or Perf logs, then select the type of achieve either a zip or tar and then click on “**Achieve Logs**”, this creates the required archive, then click on the “**Download Logs**” which appears after you click on “**Archive logs**”. The same applies to “**CX-Download Logs**”

Figure 201 Download Logs

### 8.4 Audit log:

The audit log stores all user activity on the CX UI. You may click on the “**Audit Log**” to view it as shown below.

Audit Logs			
User	Date/Time	IP Address	Details
	2008-12-26 08:00:03	10.0.1.88	FX agent registered to CX with details (hostName::PROD-SERV, ipaddress:10.0.1.88,operatingSystem::WINDOWS, hostid::EE3FF788-F450-0943-86FAA72385FD7F10, sentinelEnabled::0, outpostAgentEnabled::0, filereplicationAgentEnabled::1)
	2008-12-26 08:00:03	10.0.1.30	FX agent registered to CX with details (hostName::BAKP-SERV, ipaddress:10.0.1.30,operatingSystem::WINDOWS, hostid::BE982596-FC32-3F42-A94E320EB63ACA81, sentinelEnabled::0, outpostAgentEnabled::0, filereplicationAgentEnabled::1)
	2008-12-26 08:00:19	10.0.1.30	VX agent registered to CX with details (hostName::BAKP-SERV, ipaddress:10.0.1.30, operatingSystem::Windows NT 5.2 ver = Build 3790, hostid::BE982596-FC32-3F42-A94E320EB63ACA81, sentinelEnabled::1, outpostAgentEnabled::1, HostUpdateTime::1230258619, Version::430000, vxAgentPath::C:\Program Files\InMage Systems, patchDetails(, agentTimeStamp = '2008-12-26 17:10:53', agentTimeZone = '+0530', InVolCapacity = 20974428160, InVolFreeSpace = 14741786624, SysVolPath = 'C:\WINDOWS', SysVolCap = 20974428160, SysVolFreeSpace = 14741786624))
admin	2008-12-26 08:00:25	10.0.0.81	Login success

Figure 202

The following are the types of actions that are logged into the audit log

- New administrator add
- Trap listener add
- Trap listener delete
- Disk space warning Threshold configuration
- Agent settings configuration
- Replication pair creation
- Replication pair updated
- Snapshot creation(physical)
- Snapshot creation(virtual)
- Snapshot update
- Snapshot delete
- Recovery snapshot creation(physical)
- Recovery snapshot creation(virtual)
- Replication pair delete
- Rollback recovery snapshot creation
- CX backup
- CX restore
- All FX logs deletion
- FX logs deleted for specific Jobs
- FX logs auto deletion
- FX job create
- FX job edit
- FX job start
- FX job stop
- FX job remove
- FX template creation
- FX template deletion

## 8.5 CX Log Rotation

### Linux

CX rotates the logs if 80% of `“/var”` is used and the log file is equal or greater than 10MB. Currently CX rotates the following logs under `“/var/log”`

- `“/var/log/xferlog”`
- `“/var/log/wtmp”`
- `“/var/log/tls.log”`
- `“/var/log/boot.log”`
- `“/var/log/mysqld.log”`
- `“/var/log/secure”`
- `“var/log/spooler”`
- `“var/log/maillog”`
- `“/var/log/messages”`
- `“/var/log/warn”`
- `“/var/log/httpd/*.log”` (all logs under httpd)
- `“/var/log/rsyncd.log”`
- `“/tmp/phpdebug.txt”`

And `“phpdebug.txt”` in `“/tmp”` will be rotated if its size exceeds 10MB. All the files are zipped and placed under the same folder.

### Windows

Windows based CX server rotates all the log files under `“c:\home\svsystems\var”`

## 8.6 Clear FR Logs by Date

Every file replication job generates a log in the Scout CX server, each time it runs. These logs accumulate over time. You can specify the period after which you would like these logs to be deleted from the server.

110. On the **“System”** menu click on the **“CX Settings”** tab. The **“Clear FR Logs”** by Date screen appears

111. Select the time interval for automatic deletion of logs and click **“Commit”**.

Clear FR Logs by Date	
Auto Delete After	
<div>None ▾</div>	<div>Commit</div>

Figure 203: Clear FR Logs by Date

## 8.7 Retention Reserve Space Settings

Retention logs are stored on the target host. By default a free space of 256 MB is maintained on the retention volume. You may choose to increase or decrease this free space limit from this interface. For example, when the retention volume is of 1024 MB, by default only 768 MB will be used by the retention logs leaving 256 MB as free space. To use volumes containing free space of less than 256 MB, navigate to **“System” -> “Agent Settings”** and then reduce the **“Retention Reserve Space settings”** and click on **“Commit”**

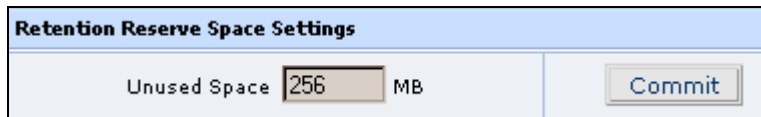
A screenshot of the 'Retention Reserve Space Settings' window. It has a light blue header with the title. Below the header, there is a text input field labeled 'Unused Space' containing the value '256', followed by 'MB'. To the right of this field is a button labeled 'Commit'.

Figure 204:

## 8.8 Installers

The “Install” tab contains “Manage” and “Installers” tabs. This is primarily used for installing agent software on remote clients through the CX UI. Please refer to the Hitachi Dynamic Replicator - Scout Installation Guide for detailed steps in using the CX UI based install.

You may also upload and download agent installers through **“System-> Installers”** screen

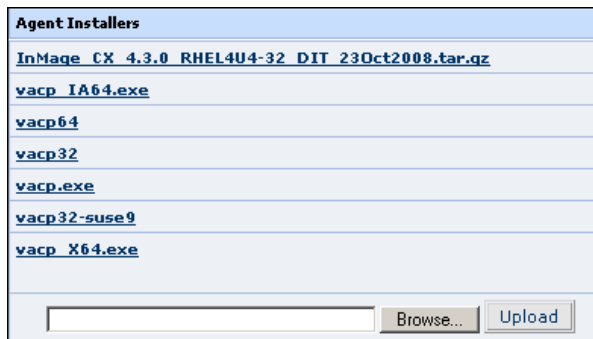
A screenshot of the 'Agent Installers' window. It has a light blue header with the title. Below the header, there is a list of files: 'InMage CX 4.3.0 RHEL4U4-32 DIT 23Oct2008.tar.gz', 'vacp IA64.exe', 'vacp64', 'vacp32', 'vacp.exe', 'vacp32-suse9', and 'vacp X64.exe'. At the bottom, there is a text input field, a 'Browse...' button, and an 'Upload' button.

Figure 205

## 9 Analyzer and Trending

In this chapter we will see trending and analyzer sections. While trending displays self explanatory graphs, analyzer helps acquire vital information about the environment. With this information one can strike the balance between the required RPO and the available bandwidth, considering the data change rates. After this chapter you will be able to

- Use trending graphs
- Use analyzer

### 9.1 Trending

#### 9.1.1 When is Trending Used

Trending is the last button. You should see six links on the “**Trending**” screen. This gives a graphical representation of compression achieved on a monthly basis, it gives in detail graphs of each of the hosts. RPO graphs, data change rates etc.




Figure 206

## 9.1.2 Network Trending

Network trending shows the amount of traffic (both inbound and outbound) on a daily, weekly, monthly basis. Roll the mouse over the “**Network Trending**” screen to view the list of process servers. Click on the desired process server to see trending graphs specific to the process server.




Figure 207

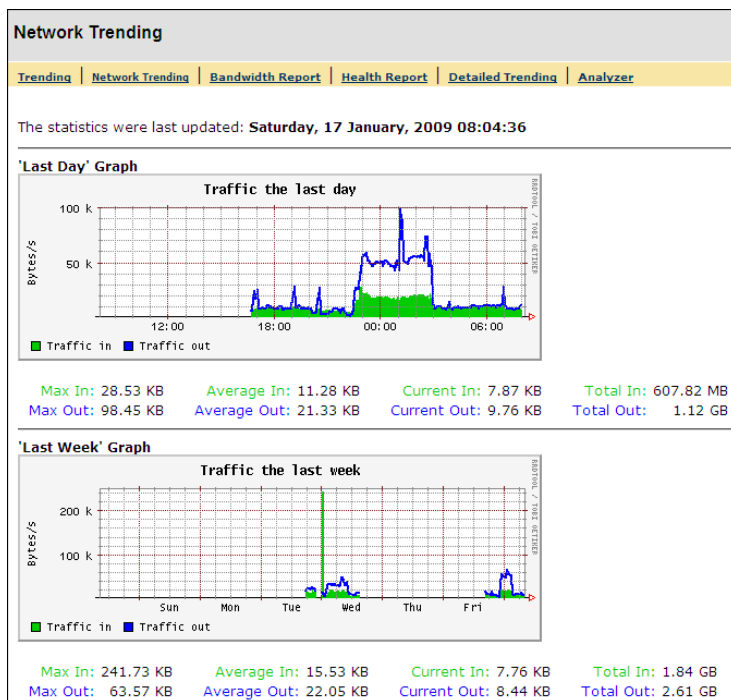


Figure 208

### 9.1.3 Bandwidth Report

All the incoming and outgoing FTP traffic for each Process Server is shown within these reports. The FTP traffic is logged in the “**xferlog**” file maintained by proftpd. This file is scanned to find incoming and outgoing FTP traffic for each host.



#### Caution:

For hosts attached to windows Process Server, Bandwidth report will not be generated

The Bandwidth Report screen contains two tabs, “**Bandwidth report**” and “**Custom Report**”

#### 9.1.3.1 Bandwidth report tab

This contains a drop down menu where you may select the desired host. By default the whole month’s data is displayed as shown below. You may also see “**Last Day**”, “**Last Week**”, “**Last Month**” and “**Last Year’s**” graphs by clicking on the respective links on the right hand side.

Bandwidth Report

Trending

Network Trending

Bandwidth Report

Health Report

Detailed Trending

Analyzer

Bandwidth Report

Custom Report

Bandwidth Report for IMIT14.INMAGE.IN [ 10.0.119.14 ]

Select Host

IMIT14.INMAGE.IN

Last Day

Last Week

Last Month

Last Year

Month: 2009 Jan

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	
In	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	17.93 KB	
Out	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	29.48 KB	
Max	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	29.48 KB	
Sum	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	47.41 KB	

	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Month	
In	43.01 KB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	60.94 KB	
Out	67.97 KB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	97.45 KB	
Max	67.97 KB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	97.45 KB	
Sum	110.98 KB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	158.39 KB	

Logged in as 'admin@10.0.42.10' - Logout

Server Time: Jan-17-2009 08:18:07

Figure 209

### 9.1.3.2 Custom Report

At times you may require a report between specific times. Select the desired host then select the date and time and click on “**Generate Report**” to see a custom report. Optionally you may select the “**Complete Host Report**” option to display “**last day, last week, last month and last year’s**” data as well

**Custom Bandwidth Report**

Trending | Network Trending | **Bandwidth Report** | Health Report | Detailed Trending

Bandwidth Report | **Custom Report**

**Query Form**

Select Host(s) IMI-VM1  
IMIT14.INIMAGE.IN

Start Date 01/01/2009 00:00

End Date 01/17/2009 23:59

Complete Host Report ☐ [Last Day, Last Week, Last Month, Last Year]

Generate Report Reset

Figure 210

An example of “**custom report**” is shown in the picture below

Bandwidth Report [ 01/12/2009/00:00 - 01/17/2009/23:59 ]				
Date	In	Out	Max	Sum
<a href="#">January 12, 2009</a>	0 B	0 B	0 B	0 B
<a href="#">January 13, 2009</a>	0 B	0 B	0 B	0 B
<a href="#">January 14, 2009</a>	0 B	0 B	0 B	0 B
<a href="#">January 15, 2009</a>	0 B	0 B	0 B	0 B
<a href="#">January 16, 2009</a>	17.93 KB	29.48 KB	29.48 KB	47.41 KB
<a href="#">January 17, 2009</a>	43.01 KB	67.97 KB	67.97 KB	110.98 KB
Total	60.94 KB	97.45 KB	97.45 KB	158.39 KB

Figure 211



## 9.1.4 Health Report

Reports screen contains two tabs “**Health Report**” and “**Custom Report**”. The “**Health Report**” is generated for every source volume from the time the replication is set. This report contains details such as data change rates, amount of data compressed, amount of data not compressed, max RPO reached for that day etc. A “**custom report**” can also be generated for the desired volume (or group of volumes) by specifying the “**Start date**” and “**End Date**”.

Health Report

Logged in as 'admin@10.0.42.10' - [Logout](#)

Server Time: Jan-17-2009 08:52:22

[Trending](#) | [Network Trending](#) | [Bandwidth Report](#) | [Health Report](#) | [Detailed Trending](#) | [Analyzer](#)

[Health Report](#) | Custom Report

Last updated: **Saturday, 17 January, 2009 08:32:55**

Health Report [ Dec 17, 2008 08:30 - Jan 17, 2009 08:30 ]

IMIT14.INMAGE.IN

IMIT14.INMAGE.IN (/dev/mapper/raj-v1) - PROTECTED

[Change Rate](#) | [RPO](#) | [Retention](#) | [Health](#)

Date	Data changes (in MBytes)		Max RPO	Retention Window (Days)	No. of hours RPO not met	Retention log reset?	Throttled Duration (Hours)	Protection Coverage
	With Compression	Without Compression						
Jan 16, 2009	0	0	3.08 min	0	23.9	N	0	1%
Jan 17, 2009	0	0	0.18 min	0	5.57	N	0	67%
Total:	0	0			29.47		0	34%

Figure 212



### Notes:

A blank health report is shown for the first hour after the CX configuration server is installed. This is the time taken for generating the report

### 9.1.4.1 Health Report

Step 137. To view the reports click on “**Trending**” and then on “**Reports**”.

Health Report [ Aug 17, 2008 12:30 - Sep 17, 2008 12:30 ]								
SOURCECX								
<b>SOURCECX (G) - PROTECTED</b>								
Date	Data changes (in MBytes)		Max RPO	Retention Window (Days)	No. of hours RPO not met	Retention log reset?	Throttled Duration (Hours)	Protection Coverage
	With Compression	Without Compression						
Sep 16, 2008	0.19	0.66	0.99 min	0	1.07	Y	0	1%
Sep 17, 2008	0.15	0.48	0.99 min	0	12.2	N	0	2%
<b>Total:</b>	0.34	1.14			13.27		0	1%
<b>SOURCECX (J) - PROTECTED</b>								
Date	Data changes (in MBytes)		Max RPO	Retention Window (Days)	No. of hours RPO not met	Retention log reset?	Throttled Duration (Hours)	Protection Coverage
	With Compression	Without Compression						
Sep 15, 2008	0	0	0 min	0	1.15	Y	0	0%
Sep 16, 2008	0.19	0.62	0.95 min	1.02	3.65	Y	0.38	100%
Sep 17, 2008	0.15	0.49	1 min	1.02	0	N	0	100%
<b>Total:</b>	0.34	1.11			4.8		0.38	67%

Figure 213:

Fields in the above figure are explained below:

- “**Date**”: Date of the generated Report
- “**Data changes -> With compression**”: All the compressed data during transit is displayed in MB.
- “**Data changes Without compression**”: All the transmitted data (specific to the source volume selected) without compression is shown in MB.
- “**Max RPO**”: Maximum RPO achieved for the given date.
- “**Retention windows (days)**”: Displays current retention day.
- “**No. of hours RPO not met**”: Total time that the RPO didn’t meet the required SLA.
- “**Retention log reset**”: If the replication pair is set to Resync then the retention logs are reset when a resync is performed on that replication pair. The value “N” indicates that retention logs are intact for the given date.
- “**Throttled Duration (Hours)**”: Total number of hours the replication pair (with the selected source volume) has throttled
- “**Protection Coverage**”: The value usually shows 100% unless
  - RPO has exceeded the threshold value
  - A resync is performed on the replication pair
  - A replication pair has throttled
  - If there is a retention policy violation

### 9.1.4.2 Custom Report

This report is generated per volume or a set of volumes on a daily basis. It shows data change rate with or without compression, Maximum RPO for the day etc.

**Step 138.** To view the report, click on **“Trending”**, then on **“Reports”**, and select **“Custom Report”** tab

**Step 139.** Select the desired volume (or group of volumes), enter **“Start date”** and **“End Date”** in **“mmddyyyy”** format along with the time, and then click on **“Generate Report”**

Figure 214:

**Step 140.** This opens **“Health Report”** page and each of the fields are explained after the figure.

Health Report [ Sep 16, 2008 00:00 - Sep 16, 2008 23:59 ]								
SOURCECX								
SOURCECX (G) - PROTECTED								
Date	Data changes (in MBytes)		Max RPO	Retention Window (Days)	No. of hours RPO not met	Retention log reset?	Throttled Duration (Hours)	Protection Coverage
	With Compression	Without Compression						
Sep 16, 2008	0.19	0.64	0.99 min	0	14.92	Y	0	1%
<b>Total:</b>	0.19	0.64			14.92		0	1%

Figure 215:

Fields for the above figure are explained below:

- **“Date”**: Date of the generated Report
- **“Data changes -> With compression”**: All the compressed data during transit is displayed in MB
- **“Data changes Without compression”**: All the transmitted data (specific to the source volume selected) without compression is shown in MB
- **“Max RPO”**: Maximum RPO achieved for the given date.
- **“Retention windows (days)”**: Displays current retention day
- **“No. of hours RPO not met”**: Total time that the RPO didn’t meet the required SLA.

- **“Retention log reset”**: If the replication pair is set to Resync then the retention logs are reset when a resync is performed on that replication pair. The value “N” indicates that retention logs are intact for the given date.
- **“Throttled Duration (Hours)”**: Total number of hours the replication pair (with the selected source volume) has throttled
- **“Protection Coverage”**: The value usually shows 100% unless
  - RPO has exceeded the threshold value
  - A resync is performed on the replication pair
  - A replication pair has throttled
  - If there is a retention policy violation

From the **“Health Report”** click **“Health”** as shown in the figure (marked in rectangular shape)

Health Report [ Aug 18, 2008 09:30 - Sep 18, 2008 09:30 ]								
SOURCECX								
SOURCECX (G) - PROTECTED						Change Rate   RPO   Retention   Health		
Date	Data changes (in MBytes)		Max RPO	Retention Window (Days)	No. of hours RPO not met	Retention log reset?	Throttled Duration (Hours)	Protection Coverage
	With Compression	Without Compression						
Sep 16, 2008	0.19	0.66	0.99 min	0	1.07	Y	0	1%
Sep 17, 2008	0.28	0.93	0.99 min	0	0.63	N	0	2%
Sep 18, 2008	0.1	0.36	26.3 min	0	8.98	N	0	1%
<b>Total:</b>	0.57	1.95			10.68		0	1%

Figure 216:



#### Notes:

An email is sent when Email alerts are configured

### 9.1.4.3 Health Report Graphs

#### Change Rate Graph

Click on “**Change Rate**”, this opens up the graph page which contains last day, week, month, and year graphs that shows compressed and uncompressed data. X-axis indicates “**Time**” and Y-axis is “**changes in bytes**”. Graphs for last day, last week, last month and last year are shown on this screen.

The graph below shows that at 18:00 hrs, approximately 0.5KB of data is compressed. Similarly for 00:00 hrs, 06:00 hrs, and 12:00 hrs the same amount of data is compressed. Green color indicates uncompressed data while blue indicates compressed data.




Figure 217:

#### RPO Graph

RPO graph is displayed when RPO link is clicked. Here, x-axis indicates “**Time**” and y-axis indicate “**RPO in minutes**”. Last week graph has “**week days**” in x-axis, last month has “**week number**” in the x-axis, and last year has “**month name**” in the x-axis.

The graph below shows that the “**RPO**” is approximately 0.7 minutes at 18:00hrs, RPO is approximately 1 minute at 00:00 hrs, RPO is approximately 0.6 minute at 06:00 hrs and RPO is approximately 0.6 minute at 12:00 hrs.




Figure 218:

**Retention Graph**

The retention graph’s x-axis indicates “**Time**” and y-axis indicates “**Retention Window in days**”. Weekly graph has “**week days**” in x-axis. Monthly graph has “**week number**” in the x-axis, and yearly graph has “**month name**” in the x-axis. Brown color indicates that retention has not met.

This graph shows that the replication started at 06:00 hrs and maintained “**Retention**” after 1800 hrs.




Figure 219:

Similarly you will find weekly, monthly and yearly graphs.

**9.1.4.4 Health Graph**

Click on “**Health**” link you should see a health report screen wherein you get a list box. This list box contains Protection health, RPO health, throttle, retention health, Resync Health, and Replication Accuracy. By default Protection Health graphs is displayed. When you click on the list box a list is displayed. Each health graphs when selected displays corresponding graphs.




Figure 220:

## Protection Health

By default “**Protection Health**” screen appears. X-axis indicates time, week, week number, and month names. In the figure below, last day graph indicates that at 12.00 hrs. “**Protection Health**” was degraded from 16.00 hrs onwards in “**Last Day**”. “**Protection Health**” status became degraded at the end of the day in “**Last Week**” graph. “**Protection Health**” became degraded at the end of week 51 in “**Last Month**” graph. “**Protection Health**” was not met in “**Last year graph**”.




Figure 221

Similarly you will find weekly, monthly and yearly graphs.

## RPO Health

The graph below shows when the RPO has met as green bars. Periods when the RPO has not met are indicated by red bars.




Figure 222:

## Throttle Health

This graph indicates the “**Throttle Health**” was acceptable. Again green bars indicates health while the red bars indicate degradation. Similarly throttle graphs are available on weekly, monthly and yearly basis.




Figure 223:

## Retention Health

This graph shows that the retention health was acceptable most of the times. The red bar indicates that the retention health dropped just before 06: 00 hrs. “**Retention Health**” graph, like other graphs is available on weekly, monthly and yearly basis also




Figure 224:



**Resync Health**

This graphs shows both healthy and degraded “**Resync Health**” at different points in time. Again this graph is available on daily, weekly,monthly and yearly basis.




Figure 225:

**Replication Accuracy:**

This graphs indicates that the “**Replication Accuracy**” was healthy between 16:00 and 20:00 hrs and approximate just before 16:00 hrs. This graph is also available on weekly,monthly and yearly basis




Figure 226

## 9.1.5 Detailed Trending

Shows the list of all the graphs available, they can be downloaded in PNG format.

Detailed Trending

Logged in as

Server

Trending	Network Trending	Reports	Detailed Trending	Analyzer
File Name			Last Modified	Size (Bytes)
<a href="#">Host Distribution Compressed.png</a>			December 30 2008 03:41:20	3163
<a href="#">Host Distribution Uncompressed.png</a>			December 30 2008 03:41:19	3243
<a href="#">LIN-PROD dev mapper volume-v1 2008 Dec.png</a>			December 30 2008 03:41:19	2531
<a href="#">LIN-PROD dev mapper volume-v1 2008 Dec 29.png</a>			December 30 2008 03:41:19	3313
<a href="#">LIN-PROD dev mapper volume-v1 2008 Dec 29 rpo.png</a>			December 30 2008 03:42:17	3809
<a href="#">LIN-PROD dev mapper volume-v1 2008 Dec 30.png</a>			December 30 2008 03:41:19	2821
<a href="#">LIN-PROD dev mapper volume-v1 2008 Dec 30 rpo.png</a>			December 30 2008 03:42:17	3294
<a href="#">LIN-PROD dev mapper volume-v1 2008 Dec rpo.png</a>			December 30 2008 03:42:18	2703
<a href="#">LIN-PROD dev mapper volume-v2 2008 Dec.png</a>			December 30 2008 03:41:19	2522

Figure 227

## 9.1.6 Analyzer

### 9.1.6.1 When is analyzer used

Analyzer is used to find the bandwidth required for a desired RPO for a single or a group of replication pairs. This is calculated considering last seven days of the replication pair's data change rates, compression achieved, retention storage used etc.

The analyzer can be found under “**Trending**”, The User interface shows two separate sections, the first one “**Protection Options**” and the second one is “**Pairs Configured**”

Protection Options	
Cumulative bandwidth available (Kbits/Sec):	<input type="text" value="256"/>
Desired RPO (Min.):	<input type="text" value="1"/>
Bandwidth Adjustment Factor:	<input type="text" value="0.35"/>
Retention Window (Days):	<input type="text" value="3"/>

Figure 228:

“Protection Options” will accept inputs such as “Cumulative bandwidth available”, “Desired RPO” and “Retention Window”. The “Overhead factor” is used for calculation. Fill the relevant Fields and scroll down to select at least one of the replication pairs and then click on “Analyze”

Pairs Configured													
Include	Source Host: Source Volume	Target Host: Target Volume	Bandwidth Required For RPO ~ = 0 Kbits/Sec (PEAK)		Bandwidth Required For RPO ~ = 0 Kbits/Sec (AVERAGE)		Cumulative data changes(in MBytes)			Average data change rate (MBytes/Sec)		Retention Storage Required (in MBytes)	Target Storage Required (in MBytes)
<input checked="" type="checkbox"/>			With Compression	Without Compression	Average with Compression	Average without Compression	With Compression	Without Compression	Monitoring Interval (Days)	Compression Enabled	Compression Disabled		
<input checked="" type="checkbox"/>	LIN-PROD:/dev/mapper/volume-v1	LIN-DR:/dev/mapper/volume-v1	0.0283	0.0430	0.0199	0.0305	0.6229	0.9522	4	0.0001	0.0001	0.7142	1024.7142
<input checked="" type="checkbox"/>	LIN-PROD:/dev/mapper/volume-v2	BAKP-SERV:I	0.0209	0.0329	0.0152	0.0233	0.4760	0.7276	4	0.0001	0.0001	0.5457	1024.5457
<input checked="" type="checkbox"/>	PROD-SERV:C:\p	BAKP-SERV:C:\p	0.0209	0.0323	0.0161	0.0247	0.5016	0.7731	4	0.0001	0.0001	0.5798	512.5793
<input checked="" type="checkbox"/>	PROD-SERV:F	LIN-DR:/dev/mapper/volume-v3	0.0319	0.2135	0.0190	0.0643	0.5940	2.0088	4	0.0001	0.0001	1.5066	513.5061
<input checked="" type="checkbox"/>	PROD-SERV:F	BAKP-SERV:F	0.0319	0.2135	0.0190	0.0643	0.5940	2.0088	4	0.0001	0.0001	1.5066	513.5061
Total:			0.1339	0.5352	0.0892	0.2071						4.8529	3588.8514
<div>AnalyzeReset</div>													

Figure 229:

This will display the result above the “Protection Options” and shows if the desired RPO is achieved with (or without) compression.

Analyzer				
Logged in as 'admin' - <a href="#">Logout</a>				
<a href="#">Trending</a>	<a href="#">Network Trending eth0</a>	<a href="#">Reports</a>	<a href="#">Detailed Trending</a>	<a href="#">Analyzer</a>
<b>Result:</b> Desired RPO can be Achieved with and without Compression RPO Achieved with compression(in min.): 0.0000 RPO Achieved without compression(in min.): 0.0000  Bandwidth Given for compression (with compression in Kbits/sec) : 256 Bandwidth Given for without compression (without compression in Kbits/sec): 256				

Figure 230:

When the “Retention Window” is increased under the “Protection Options” then the “Retention Storage Required” under “Pairs Configured” will change accordingly.

**“Recommended CX Configuration”:** After the Result is displayed, click on **“View Configuration”** to see recommended CX configuration depending on the data change rates.

Recommended CX Configuration	
<a href="#">View Configuration</a>	Average data change rate: 1.1155 MBytes/day

Figure 231

Click on **“Back”** to return to the previous screen

Recommended Configuration		
No of Recommended Hosts	4	
Number of Volumes/Drives Protected	10	
Source Data Pool Size	2GB	
Target Memory Per Replication Pair	128MB	
Max Disk/Volume Size Supported	2TB	
Number of Volumes per Host	Unlimited	
VSNAP Per Replication Volume	Unlimited	
Hardware/Software Configuration		
Environment	Hardware	Software
Linux CX Server	Processor-Quad Core 2.5 MHZ processor Physical memory:8GB Network:2 * 1GB/s[Teaming smart Load Balancing] Disks:RAID-0[4 x 250 GB SATA 7200 RPM drives]=1Terabyte	RHEL5[64bit]
<a href="#">Back</a>		

Figure 232

# Part 5: Other Administration Tasks

This chapter deals with CX administrative tasks such as backing up CX and License management. After this chapter you will be able to:

- Backup and restore CX settings
- Add/ Remove agent licenses pointed to this CX server
- Upload new license

## 10 Configuration settings

### 10.1 CX Backup/Restore

This feature allows you to backup or restore the current settings on the CX box. The backup and restore pertains to the database and CX configuration files.

To backup the Scout CX settings, click on “System” -> “CX Settings”

Backup/Restore CX Settings	
Filename	Action
<input type="text" value="backup_name"/>	<input type="button" value="Backup"/>
<input type="text"/> <input type="button" value="Browse..."/>	<input type="button" value="Restore"/>

Figure 233: CX backup

**Step 141.** Enter a valid backup name in the input box next to the “Backup” button.



#### Notes:

A valid backup name begins with an alphabet and can have subsequent characters as alphabets, numbers, and underscores.

**Step 142.** Click on the “Backup” button and you will be asked to choose a location to save the backup file as shown in the figure below.




Figure 234: Saving the CX Backup File

**Step 143.** Click “Save to disk” and you will be asked to choose a directory to save the file. Select a directory and click “Save”. The CX settings are backed up to your computer as a file that bears the same name as specified in the backup name input box with the “.amc” extension as suffix. It is important to backup CX setting periodically to restore the settings at a later time (usually if the CX server is destroyed, or it has to be replaced with old settings)

## 10.2 License Management

Refer to the [License Management](#) section on page 29.

### 10.2.1 License Validity and Expiry

A license file is bound to a single CX box. The licensing module validates this by checking the MAC address of one of the interfaces on the host as specified by the customer, with the information in the license file. If there is a mismatch, none of the VX or FX jobs will work. VX agents at source will be throttled, FX agents will not receive their job configuration information and unlicensed hosts will not be allowed to receive a license. A MAC mismatch is alerted to the user under license details and so is a license expiry.



Figure 235: License expired for Capacity-based License

The licensing scheme also supports time limited evaluation licenses. Each license has its own expiration date, after which, the license is unusable. If the license of an agent participating in a replication expires, the replication will come to a halt. License expiration for licensed agents is alerted to the user through red highlighted errors in the license management UI as shown.

You will be alerted about the expiry of the license for a particular agent by a red highlighted error in the Licensed Hosts window as shown below.

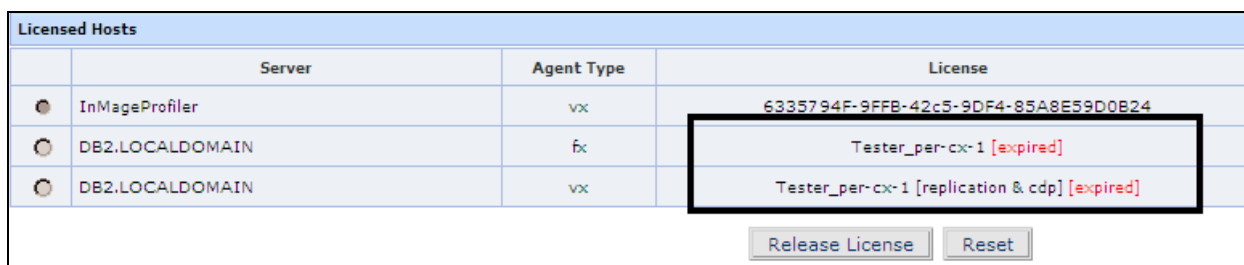


Figure 236 License Expiration Alerts

## 10.2.2Capacity Calculator

A capacity calculator is shown under “**License Details**”. It is used to add up the desired volume capacity. To move to a higher capacity license, identify the list of volumes that are to be protected through the capacity calculator

License Details	
License Upload Date	Jul 09, 2008
License Version	1.2
Customer Name	Perform
Partner Name	InMage
Host Mac Address	00:0C:29:53:2E:28
License Mac Address	00:0c:29:53:2e:28 00:98:87:67:56:45 98:89:89:87:98:87 00:00:00:00:00:00 01:01:57:78:76:67
Expiration Date	Permanent
License Count	6 (3 vx, 3 fx)
License Capacity (GB)	3
<a href="#">Capacity Calculator</a>	

Figure 237

Select the list of volumes and the “**Expected capacity growth**” value as a percentage of the total capacity of selected volumes. The result will be displayed on the screen. You may then make a request for a higher capacity license

Capacity Calculator	
Select the devices you want to protect	<a href="#">Check All</a>   <a href="#">Uncheck All</a>
<input type="checkbox"/> SQL2K8SRC-64	
<input type="checkbox"/> W2K3-TGT-64BIT	
<input type="checkbox"/> PROD-SERV	
<input type="checkbox"/> E:\Mount_prod	
<input checked="" type="checkbox"/> E	
<input checked="" type="checkbox"/> F	
<input checked="" type="checkbox"/> G	
Expected capacity growth	<input type="text" value="5"/> %
<input type="button" value="Calculate"/>	

Figure 238



### 10.2.3 Upgrading a License

If a capacity based license is uploaded over an existing host based license, you can upgrade to the new license without disrupting any backup or recovery operations.

On the CX CS UI, click on “**System -> License Management**” then scroll down to “**Licensed Hosts**”. All VX agents with host based licenses will show an “**Assign Capacity License**” link under the Upgrade License column. To upgrade the VX license to capacity based license, just click on the link and the new license is assigned to the agent without disrupting any operation.

Licensed Hosts				
	Server	Agent Type	License	Upgrade License
	InMageProfiler	vx	6335794F-9FFB-42c5-9DF4-85A8E59D0B24	
	PROD-SERV	vx	Perform_per-vx-1 [replication & cdp]	<a href="#">Assign Capacity License</a>
	PROD-SERV	fx	InMageSy-stem-sinc-fxFR-licensetemp8	
	BAKP-SERV	vx	InMageSy-stem-sinc-vxSE-licensetemp2	<a href="#">Assign Capacity License</a>
	BAKP-SERV	fx	InMageSy-stem-sinc-fxFR-licensetemp1	
<div>Release LicenseReset</div>				

Figure 239

## 10.2.4CX Cluster

### 10.2.4.1 Introduction to CX High availability

CX Cluster is a high availability feature which involves a secondary CX server in standby mode ready to take over when the primary CX server is down.

This involves two identical CX servers based on RHEL5 64 bit, or RHEL5 Update 1 or RHEL5 Update 2. While one CX server is active and serves all clients the other is in standby mode. All CX configuration information is maintained on both CX servers through a set of FX jobs scheduled to execute at regular intervals.

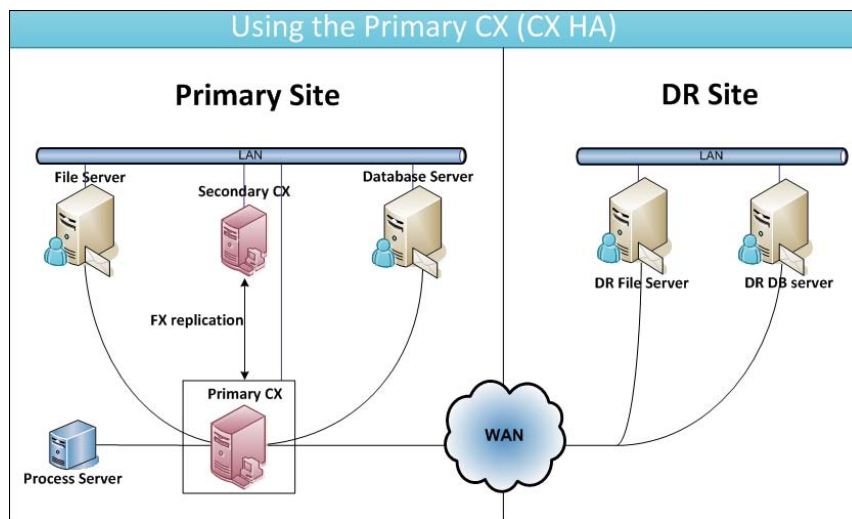


Figure 240

The secondary CX server takes over when the primary CX server is down and ensures all backup activities progress. This failover is transparent; you will however receive an email alert when already configured.

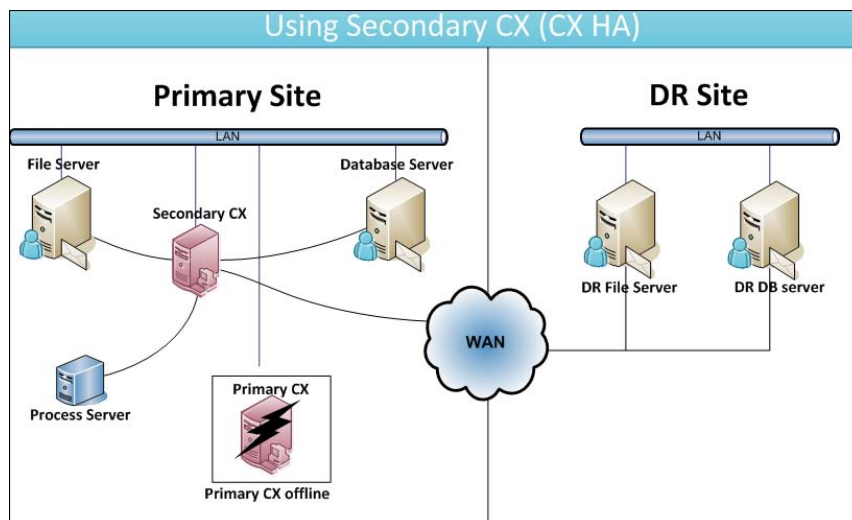


Figure 241: Primary CX down

### 10.2.4.2 Step 1: Install CX HA

Before installing HA ensure that all the configuration servers, VX agents and FX agents are pointed to the cluster IP address rather than the individual CX.

Please refer to the installation guide for detailed steps for installing or upgrading CX HA




Figure 242:

When only CX configuration server is installed on the primary CX server, ensure that you install only CX server on the secondary CX server. Similarly when CX and PS are installed on the primary CX machine, ensure that you install CX and PS on the secondary machine as well.

### 10.2.4.3 Step 2: Checking HA Services

112.To check the HA service use the command “`/etc/init.d/heartbeat status`”. To check the CX services use the command “`/etc/init.d/tmanagerd status`”. Both the commands will give the list of services and their respective states. If any service is stopped, start them and verify if all of them are up and running.

```
[root@rhel5ha2 tmp]# service heartbeat status
heartbeat OK [pid 21730 et al] is running on rhel5ha2 [rhel5ha2]...
[root@rhel5ha2 tmp]# service tmanagerd status
tmanager.pl (pid 21735 21216 21204 21193 21184 21172 21161 21147 21135 21116 21103 21090 21081 21069 21058 21045 21029 21018 21003 20988 20978 20962 20949 20907 20866 20788 20785) is running...
bpm.pl (pid 20816) is running...
gentrends.pl (pid 20847) is running...
scheduler (pid 20799) is running...
httpd (pid 20412 20408 20407 20406 20404 20403 20401 20400 20391) is running...
mysqld (pid 20085) is running...
inmsync (pid 20793) is running...
proftpd (pid 20914) is running...
mrtg (pid 22535) is running...
[root@rhel5ha2 tmp]#
```

Figure 243



#### Notes:

Ensure that both the CX servers have the same time

### 10.2.4.4 Step 3: Configure CX HA from UI

When CX HA is installed and all services are up and running, you should see the “**CX Cluster Member Information**” as shown below.

“**CX Cluster Member Information**” displays five fields as explained below

- “**Nodes**”: Displays node numbers
- “**Host Name**”: Shows the hosts names for active and passive nodes
- “**IP Address**”: IP address of each of the node
- “**Role**”: This field shows which CX server is active. One will always be passive while the other is active.
- “**Last DB Sync TimeStamp**”: Shows the last time when the “**Passive**” node has checked the cluster status. i.e. Checked if the active node is still active. For the “**Active**” node this field shows the last time when the CX database was successfully backed up to the passive node. By default these fields show “0” when HA is not configured. Click on “**Enable DB Sync**” to enable CX HA

<a href="#">Monitoring</a>	<a href="#">Bandwidth Shaping</a>	<a href="#">CX Cluster</a>	<a href="#">CX Settings</a>	<a href="#">Agent Settings</a>	<a href="#">License Documents</a>	<a href="#">Logs</a>	<a href="#">Agent Installers</a>	<a href="#">Fabric Configuration</a>
CX Cluster Member Information								
Nodes	Host Name	IP Address	Role	Last DB Sync TimeStamp				
Node. 1	IMIT2S006	10.0.2.6	PASSIVE	0000-00-00 00:00:00				
Node. 2	IMIT2S08	10.0.2.8	ACTIVE	0000-00-00 00:00:00				
					<a href="#">Enable DB Sync</a> <a href="#">Disable DB Sync</a>			

Figure 244



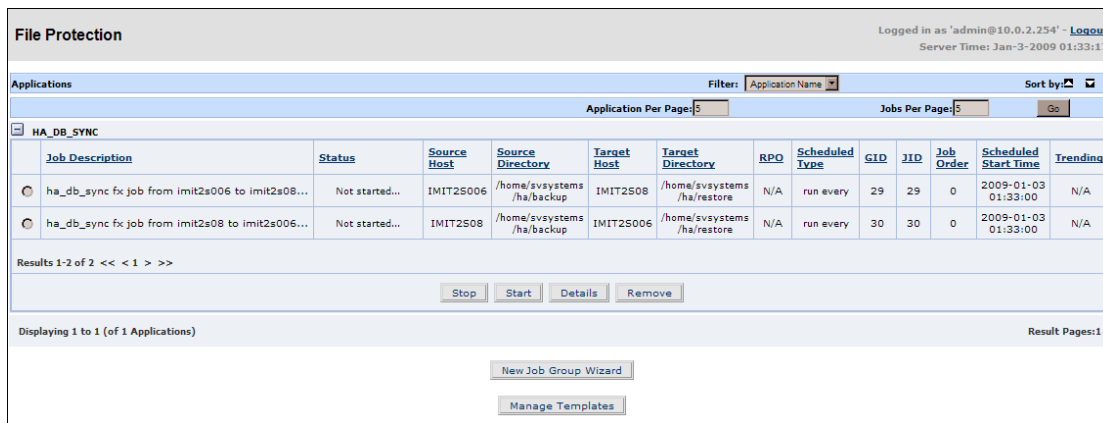
#### Caution:

Ensure that all the VX, FX, PS are pointed to the CX clustered IP and not the individual CX IP address

When the FX jobs for CX HA fail with -255 error code, you may use the Use the “**mysqladmin flush-hosts**” command on the identified node to flush the hosts on all the CX servers. This error is usually caused when the “**mysqld**” service blocks a host after 10 connection errors

#### 10.2.4.5 Step 4: Monitor CX HA

Once CX HA is enabled, you should be able to see two FX jobs set. These jobs perform CX data base backup and check for active node status. The CX DB sync always happens from active to passive nodes. When the active node is down, a failover is performed without any manual intervention.

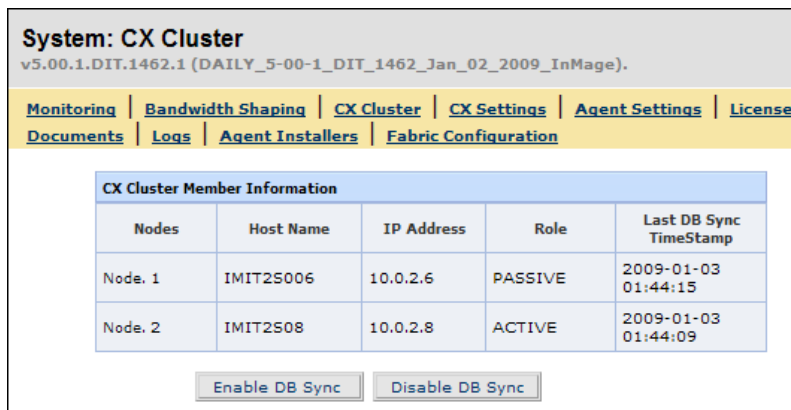


The screenshot shows the 'File Protection' interface. At the top, it says 'Logged in as "admin@10.0.2.254" - Logout' and 'Server Time: Jan-3-2009 01:33:17'. Below this is a section for 'Applications'. There is a filter dropdown set to 'Application Name' and a 'Sort by' dropdown. Below the filter, there are input fields for 'Application Per Page' (set to 5) and 'Jobs Per Page' (set to 5), followed by a 'Go' button. A checkbox labeled 'HA\_DB\_SYNC' is selected. Below this is a table with the following columns: Job Description, Status, Source Host, Source Directory, Target Host, Target Directory, RPO, Scheduled Type, GID, JID, Job Order, Scheduled Start Time, and Trending. There are two rows of data. The first row shows a job from imit2s006 to imit2s08 with status 'Not started...'. The second row shows a job from imit2s08 to imit2s006 with status 'Not started...'. Below the table, it says 'Results 1-2 of 2 << < 1 > >>'. There are buttons for 'Stop', 'Start', 'Details', and 'Remove'. At the bottom, it says 'Displaying 1 to 1 (of 1 Applications)' and 'Result Pages:1'. There are also buttons for 'New Job Group Wizard' and 'Manage Templates'.

Job Description	Status	Source Host	Source Directory	Target Host	Target Directory	RPO	Scheduled Type	GID	JID	Job Order	Scheduled Start Time	Trending
ha_db_sync fx job from imit2s006 to imit2s08...	Not started...	IMIT2S006	/home/svsystems/ha/backup	IMIT2S08	/home/svsystems/ha/restore	N/A	run every	29	29	0	2009-01-03 01:33:00	N/A
ha_db_sync fx job from imit2s08 to imit2s006...	Not started...	IMIT2S08	/home/svsystems/ha/backup	IMIT2S006	/home/svsystems/ha/restore	N/A	run every	30	30	0	2009-01-03 01:33:00	N/A

Figure 245: FX jobs set as part of the CX HA

Observe that the “Last DB Sync TimeStamp” now shows successful DB sync for the active node and successful CX HA status check by the passive node



The screenshot shows the 'System: CX Cluster' interface. At the top, it says 'v5.00.1.DIT.1462.1 (DAILY\_5-00-1\_DIT\_1462\_Jan\_02\_2009\_InMage)'. Below this is a navigation bar with links: Monitoring, Bandwidth Shaping, CX Cluster, CX Settings, Agent Settings, License, Documents, Logs, Agent Installers, and Fabric Configuration. Below the navigation bar is a section titled 'CX Cluster Member Information'. It contains a table with the following columns: Nodes, Host Name, IP Address, Role, and Last DB Sync TimeStamp. There are two rows of data. The first row shows 'Node. 1' with Host Name 'IMIT2S006', IP Address '10.0.2.6', Role 'PASSIVE', and Last DB Sync TimeStamp '2009-01-03 01:44:15'. The second row shows 'Node. 2' with Host Name 'IMIT2S08', IP Address '10.0.2.8', Role 'ACTIVE', and Last DB Sync TimeStamp '2009-01-03 01:44:09'. Below the table, there are buttons for 'Enable DB Sync' and 'Disable DB Sync'.

Nodes	Host Name	IP Address	Role	Last DB Sync TimeStamp
Node. 1	IMIT2S006	10.0.2.6	PASSIVE	2009-01-03 01:44:15
Node. 2	IMIT2S08	10.0.2.8	ACTIVE	2009-01-03 01:44:09

Figure 246

#### 10.2.4.6 Step 5: CX Failover

The HA service monitors MySQL and http services, if any one of them is down then a failover is performed to the secondary CX server (provided the MySQL, http and heartbeat services are up and running on the secondary CX server).

#### 10.2.4.7 CX Failback

The secondary CX server will continue to be as an active node even after the primary CX server is back online. To failback to the primary CX server, manually start the MySQLd, httpd, heart beat and nodentwd services on the primary CX server, then stop any of those services on the secondary CX server.

## 10.2.5 Configure Remote CX

### 10.2.5.1 Introduction to remote CX recovery

While the “CX cluster” feature is used for high availability, the “**Remote CX**” feature is used for remote recovery. This solution describes using another standby CX server to counter loss of the primary CX server. The standby CX server is placed over a WAN link within the DR site. Then both the CX servers are linked through the primary CX server’s interface. This requires FX agents to be installed on both of them. You may continue with your normal backup and recovery operations through the primary CX server.

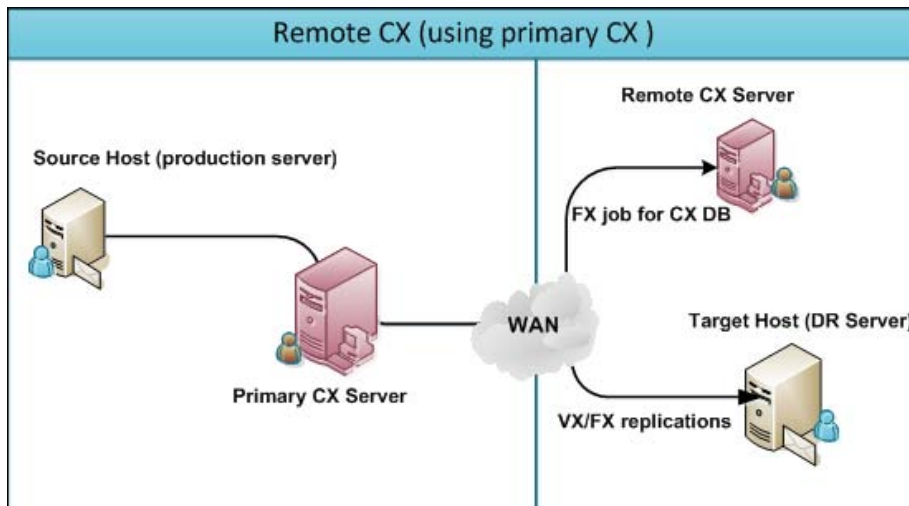


Figure 247

When the primary CX server is lost, all the replications are automatically diverted to the standby CX server. To start using the standby CX server UI, the link created between both the CX servers needs to be removed.

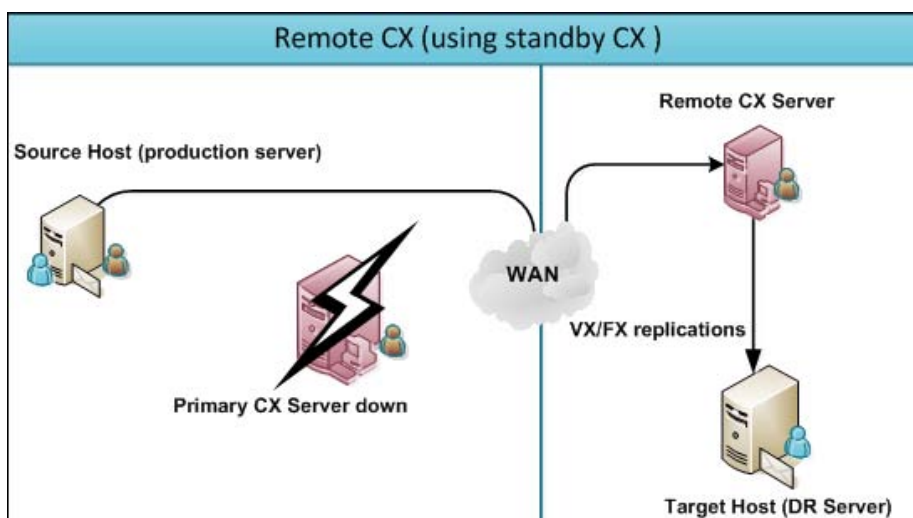


Figure 248

This is broadly classified into four major sections as shown below. Each section is a collection of steps.

Deployment	Protecting primary CX	Using primary CX	Failover to Standby CX
<ul style="list-style-type: none"><li>•Install primary CX</li><li>•Install FX on Primary CX</li><li>•Install Secondary CX</li><li>•Install FX on secondary FX</li><li>•Point Both FX agents to primary CX server</li><li>•Upload appropriate licences to both CX servers</li></ul>	<ul style="list-style-type: none"><li>•Link Both CX servers</li><li>•Monitor CX database replication.</li></ul>	<ul style="list-style-type: none"><li>•Set Vx, FX replications etc through the primary CX server's UI</li></ul>	<ul style="list-style-type: none"><li>•Make the standby CX server as active</li><li>•Break the link between both CX servers</li><li>•Start the tmanagerd service on standby CX server</li></ul>

**Figure 249**

#### **10.2.5.2**

CX HA is not similar to Remote CX feature, for a list of differences between them refer to the section [Differences between CX HA and Remote CX features](#) on page 211.

### 10.2.5.3 Protecting Primary CX Server

**Step 144.** Login to the primary CX server then click on “**System-> Configure Remote CX**”, to enter the following values

- “**IP Address**”: Enter the Standby CX server’s IP address
- “**NAT IP Address**”: When the CX server is using a NAT IP address, enter the NAT IP.
- “**Port Number**”: Enter the HTTP Port of the standby CX server; this is set to 80 by default
- “**Time Out (In Minutes)**”: The agent tries for the specified time before switching to the standby CX server. This is set to 10 minutes by default.

**Step 145.** Click on “**Configure**” to continue.

The screenshot shows the 'System: Domain Pairing' interface. At the top, it says 'v4.3.0.DIT.1390.1 (DAILY\_4-3-0\_1390\_Oct\_22\_2008\_InMage)' and 'Logged in as 'admin' - Logout'. Below this is a navigation bar with links: Monitoring, Bandwidth Shaping, CX Settings, Agent Settings, License Management, Agent Heartbeat, and Configure Remote CX (highlighted). Underneath the navigation bar are links for Documents, Logs, and Agent Installers. The main content area is titled 'Configure Standby CX' and contains a form with the following fields: IP Address (10.0.160.201), NAT IP Address (empty), Port Number (80), and Time Out (in minutes) (10). At the bottom of the form are 'Configure' and 'Reset' buttons.

Figure 250

**Step 146.** The link is set between the primary and standby CX server. This sets the FX job between the primary and standby CX servers. The FX job here is set automatically without any manual intervention. This job is set to run every five minutes.

The screenshot shows the 'System: Domain Pairing' interface. At the top, it says 'v4.3.0.DIT.1390.1 (DAILY\_4-3-0\_1390\_Oct\_22\_2008\_InMage)' and 'Logged in as 'admin' - Logout'. Below this is a navigation bar with links: Monitoring, Bandwidth Shaping, CX Settings, Agent Settings, License Management, Agent Heartbeat, and Configure Remote CX (highlighted). Underneath the navigation bar are links for Documents, Logs, and Agent Installers. The main content area is titled 'Stand by CX Details' and contains a table with the following data:

	IP Address	NAT IP Address	Port Number	Pairing Type	Timeout (in minutes)
C	10.0.160.201		80	passive	10

At the bottom of the table are 'Release Pair' and 'Reset' buttons.

Figure 251



**Step 147.** The FX job appears automatically under the primary CX server and then on the secondary CX.

File Protection Status												
Filter	Job Description	Application	Status	Source Host	Source Directory	Target Host	Target Directory	Scheduled Type	CID	JID	Job Instance	Exit Code
Set Clear		Select	Select						Select	Select		Select
	Database Sync fo...	DB SYNC	Completed	PRIM-RHEL5U1	/home/svsystems/bin/db	STAND-RHEL5U1	/home/svsystems/bin/db	Run Every	1	1	1	0

**Figure 252**

**Step 148.** Login to the standby CX server's UI and navigate to "System-> Configure Remote CX". You should be able to see the primary and standby CX servers.

System: Domain Pairing

v4.3.0.DIT.1390.1 (DAILY\_4-3-0\_1390\_Oct\_22\_2008\_InMage).

Logged in as 'admin' - [Logout](#)

[Monitoring](#)

[Bandwidth Shaping](#)

[CX Settings](#)

[Agent Settings](#)

[License Management](#)

[Agent Heartbeat](#)

[Configure Remote CX](#)

[Version & Patches](#)

Installers

Configured CX Details				
IP Address	NAT IP Address	Port Number	Pairing Type	Timeout (in minutes)
10.0.160.200		80	active	10
10.0.160.201		80	<div>Set Active</div>	10

**Figure 253**



**Caution:**

It is important that both the primary and standby CX servers have the same time; else the FX job may not function as intended.

### 10.2.5.4 Using the Primary CX Server

You may set VX replication pairs, FX jobs, etc through the primary CX UI.

Protection Status

Logged in as 'admin' - [Logout](#)

Server Time: Nov-1-2008 08:42:36

Volume Protection Status

Server	Volume	Group	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)	Resync progress	RPO	Status	Resync Required	View Details
WIN2K8-SRC->WIN2K8-TRG	E ( vol1 ) -> J	Volume J	0	0	0	N/A	0.78 minutes	Differential Sync	NO	<a href="#">+</a>

Figure 254

All the FX, jobs, VX replications etc should appear on the standby CX-CS UI. However you will not be able to alter anything here.

The primary CX server should display command buttons etc enabling you to set or remove FX jobs or VX replications etc.

File Protection

Logged in as 'admin' - [Logout](#)

Server Time: Nov-1-2008 08:17:33

Applications

Filter: 

Application Name

Sort by:

Application Per Page: 

5

Jobs Per Page: 

5

Go

DB SYNC

Job Description	Status	Source Host	Source Directory	Target Host	Target Directory	RPO	Scheduled Type	GID	JID	Job Order	Scheduled Start Time	Trending
<div></div> database sync for linux...	Not started...	PRIM-RHEL5U1	/home/svsystems/bin/db	STAND-RHEL5U1	/home/svsystems/bin/db	2 m 55 s	run every	1	1	0	2008-11-01 08:19:00	<a href="#">View</a>

Results 1-1 of 1 << < 1 > >>

Stop

Start

Details

Remove

Displaying 1 to 1 (of 1 Applications)

Result Pages:1

Figure 255

The standby CX server will not show any command buttons and should remain as a read-only interface

File Protection

Logged in as 'admin' - [Logout](#)

Server Time: Nov-1-2008 08:18:15

Applications

Filter:

Sort by

Application Per Page:

Jobs Per Page:

Go

☐ DB SYNC

Job Description	Status	Source Host	Source Directory	Target Host	Target Directory	RPO	Scheduled Type	GID	JID	Job Order	Scheduled Start Time	Trending
database sync for linux...	Starting...	PRIM-RHEL5U1	/home/svsystems/bin/db	STAND-RHEL5U1	/home/svsystems/bin/db	7 m 52 s	run every	1	1	0	2008-11-01 08:14:01	N/A

Results 1-1 of 1 << < 1 > >>

Displaying 1 to 1 (of 1 Applications)

Result Pages: 1

Figure 256

### 10.2.5.5 Failing over to the Standby CX Server

**Step 149.** Switch to the secondary CX UI and navigate to “**System-> Configure Remote CX**”, then click on “**Set Active**” to start using the standby CX server’s UI.

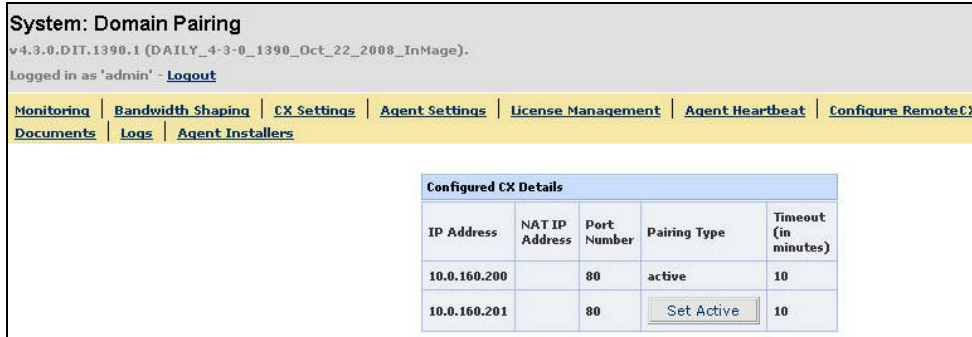


Figure 257

**Step 150.** This should make the standby CX server as active, select the stand by CX server entry and click on “**Release Pair**” to break the link between both the CX servers. This will also delete the FX job.

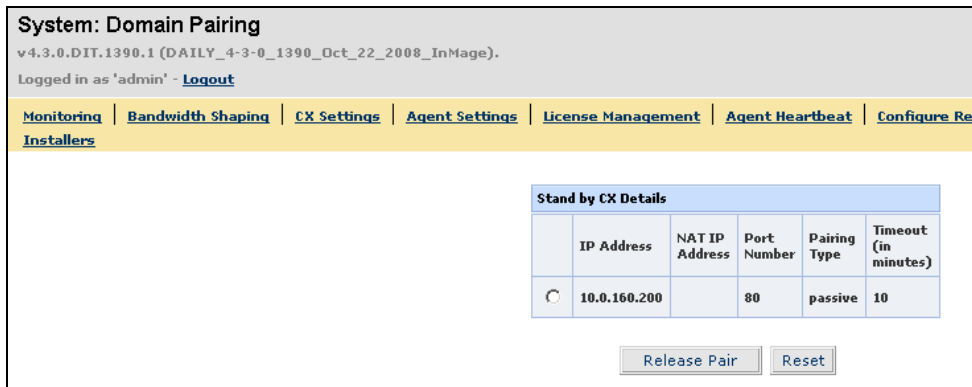


Figure 258

### 10.2.5.6 Post Failover

All the VX replication pairs are set for a “**Resync required**” as “**YES**”

All the FX jobs are scheduled to “**Run On Demand**”

Start the tmanagerd service on the standby CX server to start using it

Protection Status

Logged in as 'admin' - [Logout](#)

Server Time: Nov-1-2008 08:55:38

Volume Protection Status

Server	Volume	Group	Resyncs In Transit Step1 (MB)	Resync In Transit Step2 (MB)	Differentials Left (MB)	Resync progress	RPO	Status	Resync Required	View Details
WIN2K8-SRC->WIN2K8-TRG	E ( vol1 ) -> J	Volume J	0	0	0	N/A	0.12 minutes	Differential Sync	YES	

Figure 259

## 10.2.6 Differences between CX HA and Remote CX features

CX HA and Remote CX features are not to be mistaken. Given below are the differences between them.

**Table 20**

	CX HA or CX Cluster	Remote CX
Why is it used	To overcome loss of the primary CX server by instantly failing over to the secondary CX	Intended for remote recovery when the primary CX server (or the cluster CX server) on the primary site is down.
Supported platforms	RHEL 5 64 bit LINUX and above	LINUX and WINDOWS
Topology	Secondary CX is placed locally and is on a LINUX HA configuration	Secondary CX is placed remotely and is not on a HA configuration.
Why is it used	High Availability	Remote Recovery
Failover	Automatic failover	Manual failover by choosing the secondary node as active
Failback	Requires manual intervention	User should configure a remote CX from the new failed over node to do a failback
Miscellaneous	Works only on LINUX	Works on windows and Linux when:  Primary -> Remote (both windows)  Primary -> Remote (both Linux)

## 10.2.7 Process Server Failover

At times you may need to move VX replication pairs handled by one process server to another. This may be due to

- Over load on the PS
- Low resources on the PS
- PS that went down or lost

The picture below shows a CX configuration server communicating with two process servers, a production server and a DR server. The primary PS is handling the replication pair while the backup PS is idle.

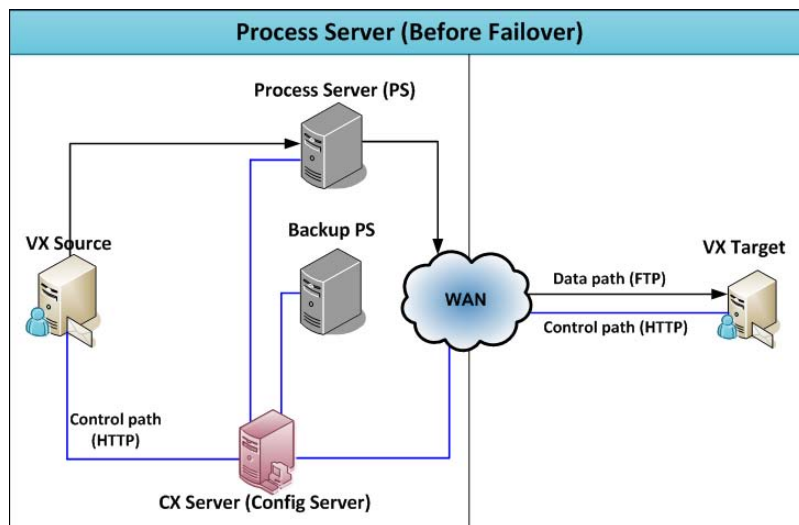


Figure 260

When the process server is down, you will need to perform a “**Process Server Failover**” through the CX UI.

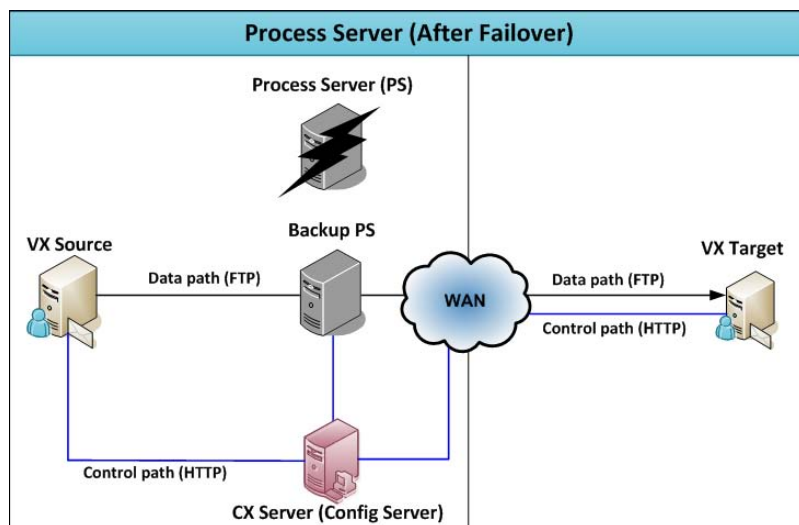


Figure 261

Navigate to “**System** “-> “**Process Server Failover**”, you should now see the screen given below.

Select the old process server with the VX replication pairs under the “**Current Process Server**” then select the new process server under “**Failover Process Server**” and click on “**Failover**”. This will move all the replication pairs from the old process server to the new process server. The corresponding VX replication pairs will require a Resync.

Process Server Failover			
Current Process Server		Failover Process Server	
<div>PS (10.0.42.10) PS_56 (10.0.56.1) ProcessServer (10.0.42.15)</div>		<div>PS (10.0.42.10) PS_56 (10.0.56.1) ProcessServer (10.0.42.15)</div>	
<b>Process Server Details</b>		<b>Process Server Details</b>	
Host Name	PS_56	Host Name	PS
IP Address	10.0.56.1	IP Address	10.0.42.10
Heartbeat	2009-01-17 06:59:11	Heartbeat	2009-01-17 20:46:02
No. of Replication Pairs	5	No. of Replication Pairs	0
<b>Pair Details</b>		<b>Pair Details</b>	
Server	Pri Volume	Remote Server	Volume
IMI-VM1	E	IMI-VM1	F
IMI-VM1	E	IMI-VM1	G
IMIT14.INMAGE.IN	/dev/mapper/raj-v2	IMIT14.INMAGE.IN	/dev/mapper/raj-v3
IMIT14.INMAGE.IN	/dev/mapper/raj-v1	IMIT14.INMAGE.IN	/dev/mapper/raj-v4
IMIT14.INMAGE.IN	/dev/mapper/raj-v6	IMIT14.INMAGE.IN	/dev/mapper/raj-v5
<div>FailoverReset</div>			

Figure 262

### 10.2.8Versions & Patches

This shows the “**CX Patch History**” and “**Agent Versions and Patch history**” sections. The CX patch history contains all the patches and updates performed on the CX server. The Agent versions and patch history contains all patches on the agent (both FX and VX)

System: Agent Versions				
v5.00.1.GA-TRIAL.1473.1 (RELEASE_5-00-1_GA-TRIAL_1473_Jan_13_2009_InMage).				
Logged in as 'admin@10.0.42.10' - Logout Server Time: Jan-17-2009 19:04:32				
<div>MonitoringBandwidth ShapingCX SettingsAgent SettingsLicense ManagementAgent HeartbeatProcess Server FailoverRemote CXVersions and PatchesUser Document</div> <div>LogsAgent Installers</div>				
<b>CX Patch History</b>				
+ localhost.localdomain				
<b>Agent Versions and Patch History</b>				
Patch History	Hostname	Volume Replication	File Replication	Sentinel Driver
+	PS			
+	IMI-VM1	RELEASE_5-00-1_GA-TRIAL_1469_9_Jan_2009_InMage		RELEASE_5-00-1_GA-TRIAL_1469_9_Jan_2009_InMage
+	IMIT14.INMAGE.IN [LinuxVX]	RELEASE_5-00-1_GA-TRIAL_1471_Jan_12_2009_INMAGE		RELEASE_5-00-1_GA-TRIAL_1471_Jan_12_2009_INMAGE
+	PS_56			
+	ProcessServer			

Figure 263

## 10.3 Process Server Traffic Load Balancing

A single process server may be used by many hosts leading to degraded performance. You may choose to add additional NICs to the process server and assign hosts to desired NICs. The picture below shows a process server is handling multiple hosts through different NICs at the same time.




Figure 264

After installing NICs on the process server, switch to the CX UI and navigate to “**System- Process Server Traffic Load Balancing**”. You should see the interface as shown below.

System: Process Failover

Logged in as 'admin@10.0.1.30' - [Logout](#)  
Server Time: Aug-12-2009 17:57:15

[Monitoring](#) | [Bandwidth Shaping](#) | [CX Settings](#) | [Agent Settings](#) | [License Management](#) | [Agent Heartbeat](#) | [Process Server Traffic Load Balancing](#) | [Process Server Failover](#) | [Remote](#)

[CX](#) | [Versions and Patches](#) | [User Documents](#) | [Logs](#) | [Installers](#) | [RX Settings](#)

Agent - Process Server NIC Mapping

Select Volume Replication Agent

DR-SERVER (10.0.145.46)  
PROD-SERVER (10.0.145.45)  
IMITS088 (10.0.1.88)

Select Process Server

imits030.qa-domain.net (10.0.1.30)  
imits145.dev-domain.net (10.0.1.145)

Select NIC to Map

Figure 265

All the VX hosts pointed to the CX server appear under the “**Select Volume Replication Agent**” list. All Process servers connected to the CX server appear under the “**Select Process Server**” list. Select the desired VX agent host that is to be assigned to a NIC on the process server, you should see the host name, IP address and heartbeat of the VX agent under the “**Volume Replication Agent Details**”.

Then select the process server, you should see similar details under the “**Process Server Details**”. This should display the list of NICs available on the process server under “**Select NIC to Map**”. Select the desired NIC and click on “**Submit**”

Agent - Process Server NIC Mapping

Select Volume Replication Agent

DR-SERVER (10.0.145.46)  
PROD-SERVER (10.0.145.45)  
IMITS088 (10.0.1.88)

Select Process Server

imits030.qa-domain.net (10.0.1.30)  
imits145.dev-domain.net (10.0.1.145)

Select NIC to Map

imits145.dev-domain.net (10.0.1.145)  
imits145.dev-domain.net (192.168.95.1)  
imits145.dev-domain.net (192.168.205.1)

Volume Replication Agent Details

Host Name	DR-SERVER
IP Address	10.0.145.46
Heartbeat	2009-08-12 17:56:48

Process Server Details

Host Name	imits145.dev-domain.net
IP Address	10.0.1.145
Heartbeat	2009-08-12 17:56:53

NIC Details

Device Name	eth0
IP Address	10.0.1.145
NAT IP Address	N/A

Figure 266

A prompt appears, click on “Ok” to proceed. This will assign the VX agent to the selected NIC on the process server, a success dialogue is displayed, click on “Ok” to return to the same screen. You should see this entry under the “Already configured Agent-process Server NIC mapping”. This is also the place where you can delete this mapping.


Already configured Agent - Process Server NIC Mapping			
	Volume Replication Agent	Process Server	NIC
	DR-SERVER [10.0.145.46]	imits145.dev-domain.net [10.0.1.145]	eth0 [10.0.1.145]
<div> Delete Reset </div>			

Figure 267



#### Notes:

When there is no Agent-PS NIC mapping , by default all the load is handled by eth0  
Please refer to the Firewall section in the install guide for a detailed list of ports to be enabled

## 10.4 User Documents

A list of user documents is displayed. Some of them are Installation guide, Quick start guide, Administration guide, and Troubleshooting guide.

## 10.5 RX Settings

You may point the CX-CS to the RX server through this screen. Enter the RX server’s IP address, the HTTP port on which it’s running. The “Synchronization Interval” is defaulted to five minutes; this indicates that the CX server updates the RX server every five minutes. You may choose the “PUSH Method” when the CX server is behind a firewall. Click on “Change Settings” to save the settings and tie the CX server to the RX server.

System: RX Settings  
v5.10.1.DIT.1525.1 (DAILY\_5-10-1\_DIT\_1525\_Mar\_06\_2009\_InMage).

Monitoring | Bandwidth Shaping | CX Settings | Agent Settings | License Management  
CX | Versions and Patches | User Documents | Logs | Agent Installers | RX Settings

RX Settings

RX IP Address	10.0.42.15
HTTP Port	80
Synchronisation Interval (minutes)	5
Method for Data Synchronisation	<input checked="" type="radio"/> PULL Method (!) <input type="radio"/> PUSH Method (!)

Unregister RX
Change Settings

Figure 268



#### Notes:

Click on the “Unregister RX” to stop the CX server from communicating with the RX server. You may however perform a similar action from the RX UI as well. Please refer to the RX usage document for detailed steps



# Part 6: Protecting Applications

This chapter briefly covers the application support offered by Scout. Popular enterprise applications like Exchange and SQL servers are considered as examples. The process of protecting these applications and the scripts that are involved are briefly discussed. After this chapter you will be able to:

Understand Scout's application support for Exchange and SQL

## 11 Application Support

### 11.1 Application support

Scout exclusively supports enterprise applications like

- Exchange server
- SQL Server
- Oracle database
- Blackberry Exchange server etc

Volumes holding these applications are detected automatically by the VX agent and are shown on the CX UI. Depending upon the resources and the DR policy in place, you may choose to have a local and a remote backup at the same time (or just one of them).




Figure 269:

The CX server is capable of discovering application servers, issuing consistency tags (or bookmarks) while backing them up. This is done through "[built in FX job templates](#)" on page 80. FX templates can also perform failover and failback on application servers.



#### Notes:

FX agent needs to be installed on the source and target hosts  
VX and FX agents should be installed in their default paths  
FX agent service should be up and running with admin user privileges  
FX agent should be licensed accordingly through corresponding CX server  
For more details about failover and failback of application servers, refer to the solution documents

Enterprise applications can be failed over to the DR server through FX jobs. When the production server is down, all uses connected to the production server are automatically diverted to the remote server (DR server) and business operations are continued. This process is transparent to the user.




Figure 270: Failover to target system

When the production server is back online, a failback is performed to the production server thus reconnecting all users earlier connected to it.

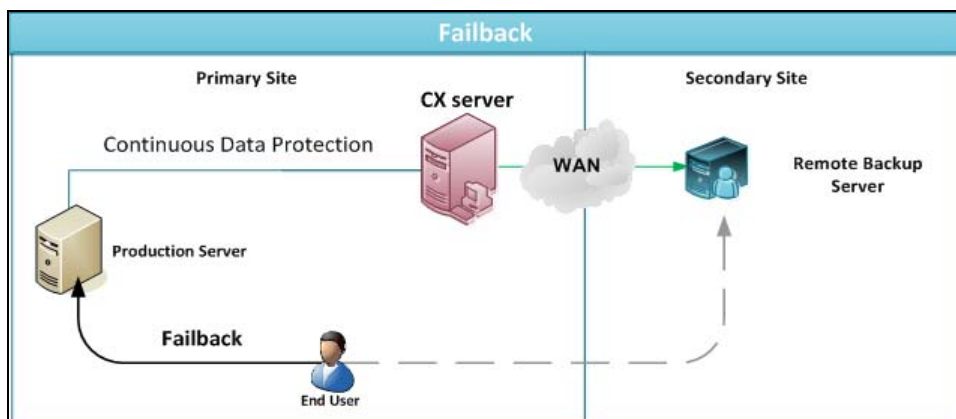


Figure 271: Failback to source system

## 11.2 Exchange Server Support

Exchange server plays a critical role in a fast paced business environment where timely communication is critical. Protecting and recovering an exchange server is a challenge faced by many administrators. Backing up exchange has some challenges such as

- Large number of users
- Exchange spanning over multiple volumes
- Maintaining consistency on the backup server

While recovering it is imperative to recover to a point with negligible or no data loss while ensuring data integrity.

### 11.2.1 Protection

After the VX and FX agents are installed on the source and target hosts, proceed to backing up the exchange server. The CX-CS UI detects volumes containing exchange server (shown under “**Application**” column under “**Volume protection**” -> “**Primary drives**”



FS	Application	Capacity (
NTFS	Unknown	107747
NTFS	Microsoft Exchange 2007	215502
NTFS	Unknown	107747
NTFS	Unknown	107751

Figure 272:

Replicate all the volumes of the same exchange server in a volume group and once the pairs reach “**Differential Sync**”, the backup is ready with only differentials being applied. You may choose to define bandwidth policies to take full advantage of the Scout server

## 11.2.2 Recovery Types

This section gives an overview of simple exchange recovery. Exchange recovery is performed if there is an outage of source exchange server, or as part of a mock drill.

**The types of recovery are as below**

- Rollback the target exchange server and continue business operations
- Take a read write virtual snapshot of the target exchange, thereby keeping the actual target exchange volume(s) unchanged and operate on the R/W Virtual snapshot
- Failover to the target exchange server until the source is back online and then failback to the source exchange server
- Access the target exchange volumes and manually select the required files

### Site wide outage

If there is a site wide outage then the CX server will be offline as well, under this condition, you may perform a recovery through the console of the target exchange system. The following files will be used to perform a failover and a failback

#### Application.exe

Application.exe performs

- Planned failover of exchange and SQL server
- Unplanned failover of exchange and SQL server
- Failback of exchange and SQL server

#### Recovery process through CX UI

An FX job is created with the source exchange and target exchange selected as source and target respectively, the FX template used should be Exchange planned failover. Once the job completes exchange failover will be complete.



#### Notes:

**Failback will also use the same FX template (Exchange planned failover)**

**Planned failover issues a consistency tag on the source and the target is recovered to the same consistency tag. This ensures that there is minimal or no data loss.**

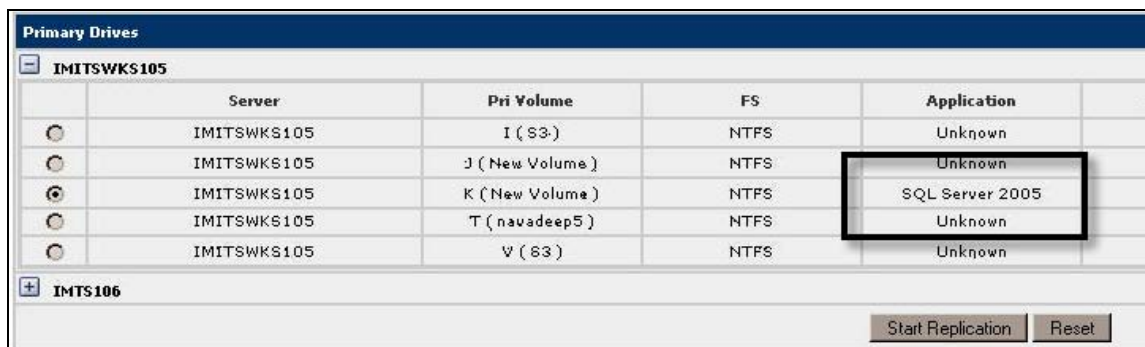
In case of partial recovery, the target is made visible or a snapshot (physical or virtual) is taken and then copied it back to the source host. Refer to the exchange failover guide for detailed explanation and process.

## 11.3 SQL Server

Scout supports SQL server 2000 and 2005. SQL server usually spans over multiple volumes. The challenges faced by backup administrators will mostly remain same as in case of exchange backup.

### 11.3.1 Protection

The VX agent will detect installations of SQL server 2000 (or 2005) and will display under the “Application” field as shown below. If there is more than one volume, then replicate them in a volume group. For a complete process, refer to the solution document for SQL server.



Primary Drives				
[-] IMITSWKS105				
	Server	Pri Volume	FS	Application
	IMITSWKS105	I (S3)	NTFS	Unknown
	IMITSWKS105	J (New Volume)	NTFS	Unknown
	IMITSWKS105	K (New Volume)	NTFS	SQL Server 2005
	IMITSWKS105	T (navadeep5)	NTFS	Unknown
	IMITSWKS105	V (S3)	NTFS	Unknown

[+] IMTS106

Start Replication Reset

Figure 273:

### 11.3.2 Recovery types

Type of recovery depends on the kind of outage that occurs. If there is a site wide outage then recovery will be performed through target host's console.

If there is a partial outage such as data corruption or any other reason where the source host is functional but unable to start off the application server then recovery can be performed either through CX UI or through console.

**Types of recovery are as below**

- Rollback the target exchange server and continue business operations
- Take a read write virtual snapshot of the target SQL server, thereby keeping the actual target SQL server volume(s) unchanged and operate on the R/W Virtual snapshot
- Failover to the target SQL server until the source is back online and then failback to the source SQL server
- Access the target SQL server volumes and manually select the required files



Notes:

Refer to the respective solution documents for detailed steps.

# Part 7: Command Line Tools

In earlier chapters we have seen how to perform operations through the CX user interface, in this chapter we focus on console based utilities. These commands perform the same tasks as of the CX interface. This chapter is broadly divided into two major sections, namely vacp.exe and cdpcli.exe. vacp is executed on the Production server's protected volumes to issue consistency markers, while cdpcli.exe is executed on the DR-Server's console to perform the same tasks as that of the CX interface. This is useful when the CX server is destroyed or if there is a site level outage. Linux and Windows based screenshots are differentiated by white and black backgrounds respectively. After this chapter, you will be able to

- Issue consistency markers through the production server's console
- Perform console based recovery operations through cdpcli utility on both windows and Linux platforms
- Differentiate cdpcli utility with that of the CX interface recovery operations.
- Other Application support console tools and tweaks

## 12 Console Tools

Scout comes with a variety of command line utilities, most of them however are used internally by the solution while others may be used when required.

**Table 21**

Utility	Purpose	Used on	Platforms
Vacp	This is used on the production servers to generate consistency events.	Production servers	Windows ,Linux and Solaris
Winop	AD replication to another AD DNS update is supported when Add / Remove NETBIOS name Host SPN, exchange specific SPN entries can be added/removed	Production servers and DR Servers	Windows only
Exfailover	Used to perform Exchange failover	DR Servers	Windows only
Dns	To perform DNS failover or failback	Production servers and DR Servers	Windows only
Clusutil	To convert cluster machines to standalone machines	DR Servers	Windows only
Mapdriveletter	To assign drive letters to volumes	DR Servers	Windows only
Cdpci	To perform CLI based recovery	DR Servers	Windows , Linux and Solaris
bwreport.pl	To generate bandwidth usage statistics from command line	CX configuration server	Windows and Linux
ICAT	To archive content to multiple archival repositories at the same time	DR Servers	Windows and Linux



**Notes:**

We recommend that you navigate to the InMage agent directory before executing any of the command line tools.



## 12.1 Issuing Consistency Tags

Consistency tags are like check points to which a consistent recovery can be made in the event of a disaster. They are issued on the source volume of an active replication pair in “**Differential Sync**” mode (Requires CDP retention option to be enabled on the replication pair). Consistency tags should be issued at regular intervals or after considerable amount of changes on the source volume. Tags will be issued through “**vacp.exe**” found under the installation folder.

### 12.1.1 VACP Command Line Options (Windows)

This section describes various command lines options of VACP utility.

The VACP tool can be used for creating user defined tags, file system consistency tags, application consistency tags and components per application tags. If the source host is on windows XP use the command “**vacp.exe**”. If vacp is failing use the command **vssadmin list writers** to display the list of available writers and check their state (should be stable). Microsoft updates VSS from time to time, it is highly recommended that all the [updates and hot fixes](#) be installed and incase of timeout errors check the [Microsoft support](#) site to troubleshoot.

#### Syntax:

```
Vacp.exe -a <name of the applications> -v <volume 1>, < volume 2> -f -s -x -t  
<tag1>, <tag2> [-remote -serverdevice < device1>, <device2> -serverip <server ip  
address> | [-serverport <portnumber>]] -h
```

Table 22: VACP options

Flags	Description
-h	Displays the usage information on the console. This option cannot be used with any other option.
-p	Enumerates all the supported applications in the system
-a <List of applications>	Specifies one or more application names. Application names are case-insensitive. “All” keyword can be used to create tags on all applications in the system. Multiple application names are separated by semi colons.
-a<Application name>/ <component name >	Similar to -a and additionally supports one or more component names
-v <List of volumes>	Specifies one or more volume names. Volume names are case-insensitive. “All” keyword can be used to create tags on all volumes in the system. Example: E: G: VACP terminates with an error message when an invalid volume letter is specified.
-t <List of user tags>	Specifies one or more user tags. The maximum length of the user tag should not exceed 250 ASCII characters. Duplicate user tags are allowed.
-x	Insert tags without any consistency mechanism. This option must be

	specified along with -v and -t options. “-a” and “-x” are mutually exclusive.
- f	Performs full backup and truncates application/database specific log files. This is applicable only on windows 2003 and is not valid for windows XP
- s	Issues a tag forcibly irrespective of the driver mode (specific to windows)
- w	Exchange 2007 consistency tag (Specific only to Windows 2003 sp1 and above)
-remote	This is used in combination with -serverdevice, -serverip and optionally -serverport. These switches are used only on client server architecture such as Citrix XenServer, Hyper-V.
-serverdevice	Corresponding volume on the base operating system where a tag is to be issued
-serverip	IP address of the server where the vacp server is running.
-serverport	Optional switch is used to specify the port number
-sync	By default all consistency tags issued are asynchronous in nature, at times these tags may be dropped due to heavy I/O. You may choose to use the “-sync” switch to issue a synchronous consistency tag. Vacp waits until the tags are drained by the s2.exe. By default the wait time is infinite, however you may use the “-tagtimeout” to restrict it to the desired time in seconds. The Sync option ensures that correct information is communicated to the respective application
-tagtimeout	This switch is used in combination with the -sync option. This determines how many seconds should the vacp wait it receives a success or failure notification from the driver.
-verify	To check the health of VSS on the system

Apart from file system level data consistency, “**vacp.exe**” also provides application support. Currently vacp supports

- Exchange 2003 and 2007
- SQL server 2000, 2005 and 2008
- System registry etc.



#### Notes:

CDPCLI utility is used on the target console for performing recovery operations.

One or more user defined tags can be specified as ASCII strings with "-t" command line option. Duplicate user defined tags are also allowed.

VACP consistency tags or consistency events are issued on the source host

File system consistency tags are unique and have the following format "FileSystemXXXXXXXX" where XXXXXXXX is sequence number

File system consistency tags are always generated unless otherwise specified by "-x" option

Application consistency tags are unique and have the following format "APPXXXXXXXX" where APP is application tag name and XXXXXXXX is sequence number.

One can ensure consistency for multiple application using "-a" option.

vacp cannot be executed on locked drives due to limitation of windows shadow copy service.

While working with vacp cli it is mandatory that component names are to be enclosed in double quotes.

Tags can be issued on multiple volumes at the same time. You may use a semi colon or a comma to separate volumes.

Tags issued on windows XP (using vacpxp) will always and only perform full backup unlike windows 2003 (2k3 performs full backup and copy backup)

While working with exchange 2007 on windows 200 Sp1, ensure you use -w option.

While working with SQL server 2005 and 2008 ensure that the service "SQL server VSS writer" is set to start automatically else VACP tags will fail.

While using Symantec antivirus Ver 10.0.100, always install SAV 10.0 Maintenance Patch 1, else vacp fails

Vacp can handle multiple operations at the same time; i.e. more than one option can be used. A semicolon is used to separate multiple parameters for individual options

E.g.: vacp -a exchange/"first store group" -v c;d: -t "tag1";"tag2"

In the above example we use two keys "-v" and "-t", both are separate by a semicolon.

### 12.1.1.1 Issuing vacp consistency tags

There are two ways to issue vacp consistency tags:




Figure 274

- **Asynchronous:** The asynchronous tag is used by default. Rarely the consistency tag may be dropped when there are heavy I/Os on the corresponding disks. Although the result of the tag is a success, it may not appear on the CX UI. Therefore when there are missing tags, it is recommended to switch to Synchronous tags.

Flow of Asynchronous tags



Figure 275

- **Synchronous:** Synchronous tags are used to ensure a confirmation right after issuing a consistency tag. The success or failure of the tag is seen after issuing the tag.

Flow of Synchronous tags




Figure 276



#### Notes:

For Synchronous tags, failure is intimated when a timeout is reached  
For Synchronous tags when a tag failed on any of the volumes, the end result indicates that tag has failed, this is contrary to Asynchronous tags.

### 12.1.1.2 Examples

Issuing consistency tags on windows machine can be done through “**vacp.exe**” under the installation folder. (For windows XP the command is “**vacpxp.exe**”)

**vacp -h**

```
C:\Program Files\InMage Systems>vacp -h
Parsing command line arguments ....
Usage:
vacp.exe [-a <app1[</comp1>..];..]<vol1;..>[-t <tag1;..>][-f][-x][-p<app1;..>][-s][-h]
-a <appName1[</comp1></compt2>..];appName2[</comp1></compt2>..];..>
    specifies one or more applications with zero or more components.
    When component names<optional> are mentioned, consistency would be
    ensured only for those components. Otherwise consistency of entire
    application would be ensured.Specify "all" to generate
    consistency tags for all applications.
    Eg: sql/masterdb/testdb;exchange/"First Storage Group";systemfiles etc
-v <vol1;vol2;..>
    specifies one or more volumes. Specify "all" to generate
```

Figure 277:

**vacp -p**

Lists all the applications and their respective components as shown below

```
C:\Program Files\InMage Systems>vacp -p
Parsing command line arguments ....
1:> Application Name = SystemFiles
    Affected Volumes = C:\
    Component[1]: System Files
2:> Application Name = SQL
    Affected Volumes = H:\ n:\
    Component[1]: master
    Component[2]: model
    Component[3]: msdb
    Component[4]: testdb_n
3:> Application Name = Exchange
    Affected Volumes = G:\
    Component[1]: d95ebafc-e9a6-47d3-a5f1-db00c857900f <PSG>
4:> Application Name = IISMETABASE
    Affected Volumes = C:\
    Component[1]: IISMETABASE
5:> Application Name = EventLog
    Affected Volumes = C:\
    Component[1]: Event Logs
6:> Application Name = Registry
    Affected Volumes = C:\
    Component[1]: Registry
7:> Application Name = COM+REGDB
    Affected Volumes = C:\
    Component[1]: COM+ REGDB
8:> Application Name = WMI
    Affected Volumes = C:\
    Component[1]: WMI <Windows Managment Instrumentation>
C:\Program Files\InMage Systems>_
```

Figure 278:

**vacp -a**

**vacp -a <application\_name>**

Where application\_name is the name of the application (using the -p option will give you a complete list of supported applications which are running)

```
C:\Program Files\InMage Systems>vacp -a "Systemfiles"
Parsing command line arguments ....
Validating command line arguments ...
Generating Tag: Systemfiles46cc2193
Generating Tag: FileSystem46cc2193
Preparing the applications for consistency ...
Preparing Files C:\* (recursively)
```

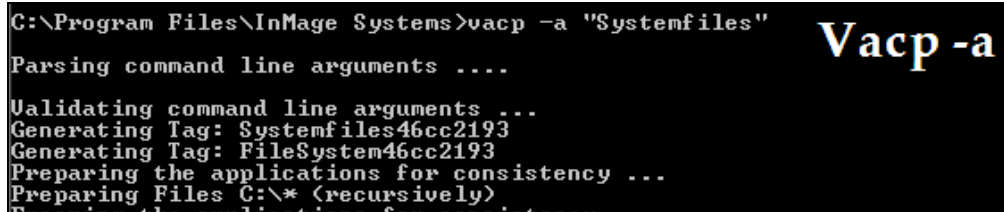


Figure 279:

**vacp -v**

**vacp -v ["Drive\_letter"] -t ["Event\_name"]**. The example is shown in the figure below

```
C:\Program Files\InMage Systems>vacp -v p: -t "Event_Neptune"
Parsing command line arguments ....
Validating command line arguments ...
ENTERED: InMageUssRequestor::GatherUssAppsInfo
ENTERED: InMageUssRequestor::Initialize
EXITED: InMageUssRequestor::Initialize
```

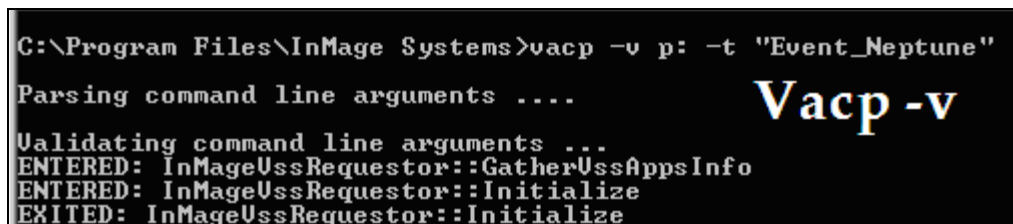


Figure 280:

To issue the same tag to more than one volume you can use one of the following syntax

**Vacp -v ["First\_Drive\_letter"]; ["Second\_Drive\_Letter"] -t ["Tag\_name"]**

```
C:\Program Files (x86)\InMage Systems>vacp -v g:;h: -t "Combo"
Parsing command line arguments ....
Entered IsVolumeLocked
Entered IsThisNtfsHiddenVolume
Opening volume: g:
Entered OpenVolume
Entered OpenVolumeExtended
Opening \\.\g: volume
```

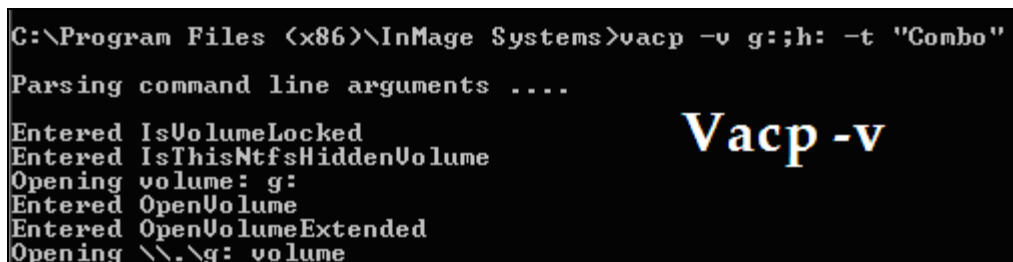


Figure 281:

## **vacp -t**

vacp -t is used to specify the name of the consistency tag. It is used along with other switches (options) of vacp.exe.

```
C:\Program Files\InMage Systems>vacp -v g: -t "Quarterly_audit"
Parsing command line arguments ....
Validating command line arguments ...
Generating Tag: Quarterly_audit
Generating Tag: FileSystem46ce9f5c
Preparing the applications for consistency ...
Preparing Files G:\* <recursively>
Freezing the applications for consistency ...
g: is mapped to Unique volume \\?\Volume{591eecda-4418-11dc-b9e1-005056c00008}\
Checking driver mode for given volumes
Volume \\?\Volume{591eecda-4418-11dc-b9e1-005056c00008}\ is not
Sending tags to the driver ...
Successfully sent tags to the driver ...
Marked all applications as successfully backed up
Resuming the applications after point-in-time consistency check-p
Exiting gracefully ...
```

## **Vacp -t**

Figure 282:

## **vacp -x**

vacp -x key is used to issue a blind tag without any consistency.

```
C:\Program Files\InMage Systems>vacp -v g: -x -t "Quarterly_audit_x_consistant"
Parsing command line arguments ....
Validating command line arguments ...
Generating Tag: Quarterly_audit_x_consistant
g: is mapped to Unique volume \\?\Volume{591eecda-4418-11dc-b9e1-005056c00008}\
Checking driver mode for given volumes
Volume \\?\Volume{591eecda-4418-11dc-b9e1-005056c00008}\ is not part of replica
Sending tags to the driver ...
Successfully sent tags to the driver ...
```

## **Vacp -x**

Figure 283:

## **vacp -f**

vacp -f is used to perform a full backup (only for windows 2003).

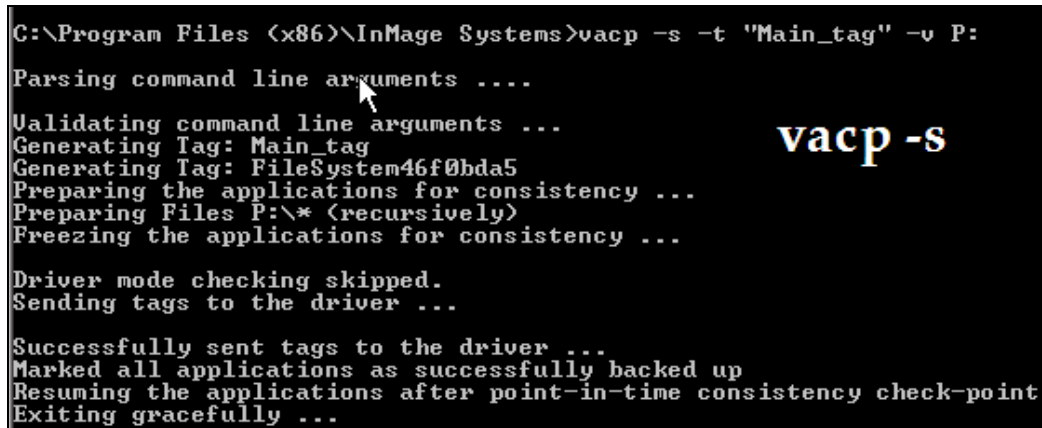
```
C:\Program Files\InMage Systems>vacp -v g: -x -f -t "Quarterly_audit_fulbkp"
Parsing command line arguments ....
Validating command line arguments ...
Generating Tag: Quarterly_audit_fulbkp
g: is mapped to Unique volume \\?\Volume{591eecda-4418-11dc-b9e1-005056c00008}\
Checking driver mode for given volumes
Volume \\?\Volume{591eecda-4418-11dc-b9e1-005056c00008}\ is not part of repl
Sending tags to the driver ...
Successfully sent tags to the driver ...
```

## **Vacp -f**

Figure 284:

## **vacp -s**

vacp -s is used to issue a consistency tag irrespective of the driver mode. The tag is forced in this case.

A screenshot of a Windows command prompt window. The title bar reads 'C:\Program Files (x86)\InMage Systems'. The command entered is 'vacp -s -t "Main\_tag" -v P:'. The output shows the process of parsing, validating, and generating a consistency tag for the file system '46f0bda5'. It also shows that driver mode checking is skipped and tags are sent to the driver. The text 'vacp -s' is overlaid in large, stylized yellow and white font on the right side of the screenshot.

```
C:\Program Files (x86)\InMage Systems>vacp -s -t "Main_tag" -v P:
Parsing command line arguments ....
Validating command line arguments ...
Generating Tag: Main_tag
Generating Tag: FileSystem46f0bda5
Preparing the applications for consistency ...
Preparing Files P:\* (recursively)
Freezing the applications for consistency ...

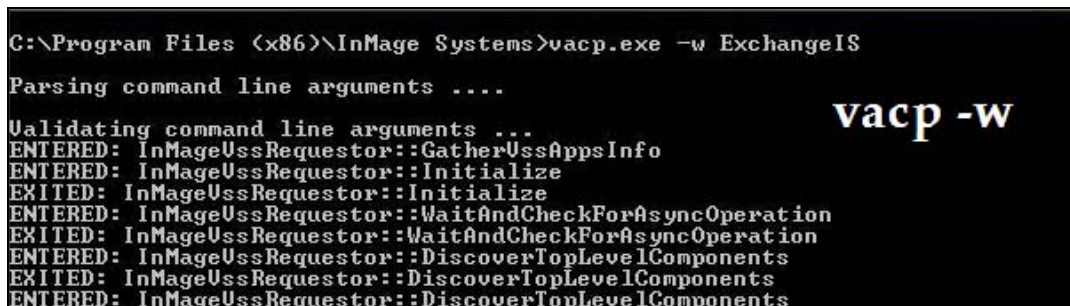
Driver mode checking skipped.
Sending tags to the driver ...

Successfully sent tags to the driver ...
Marked all applications as successfully backed up
Resuming the applications after point-in-time consistency check-point
Exiting gracefully ...
```

Figure 285:

## **vacp -w**

vacp -w is used for exchange 2007 and ensures consistency and supports ExchangeIS (exchange information store) and ExchangeREPL (exchange replication) switches.

A screenshot of a Windows command prompt window. The title bar reads 'C:\Program Files (x86)\InMage Systems'. The command entered is 'vacp.exe -w ExchangeIS'. The output shows the process of parsing, validating, and discovering top-level components for ExchangeIS. The text 'vacp -w' is overlaid in large, stylized yellow and white font on the right side of the screenshot.

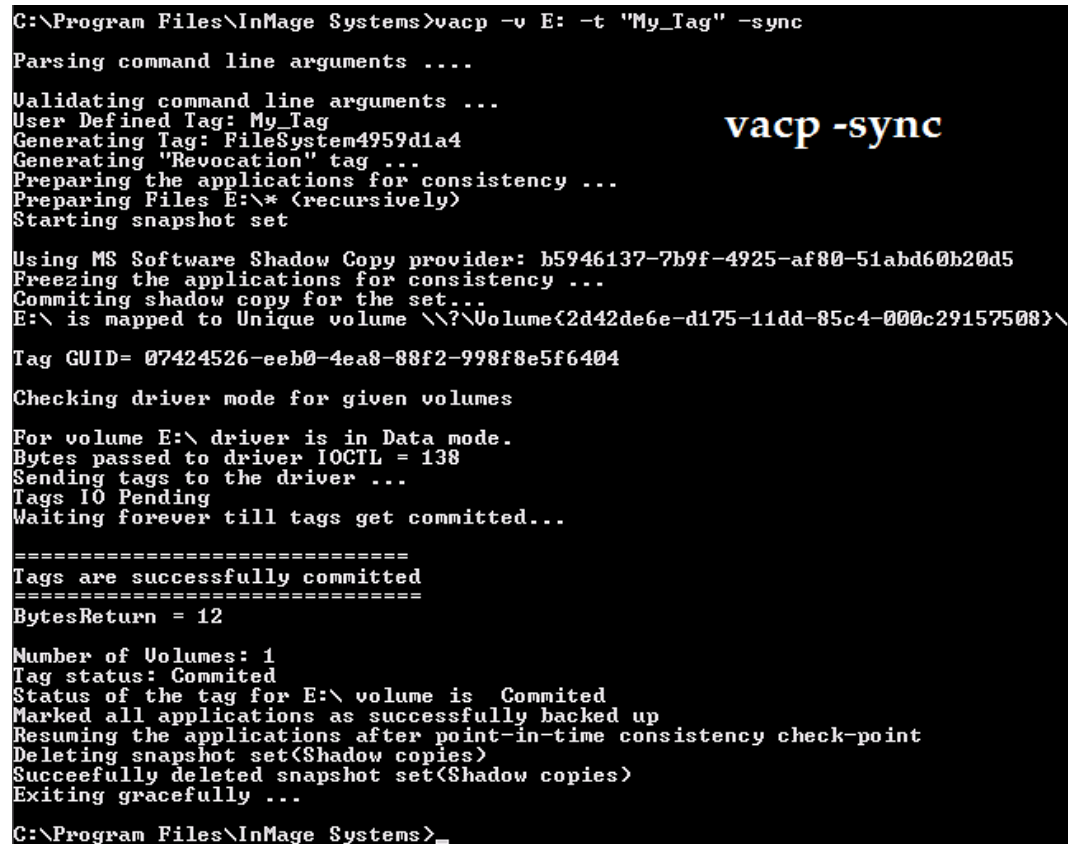
```
C:\Program Files (x86)\InMage Systems>vacp.exe -w ExchangeIS
Parsing command line arguments ....
Validating command line arguments ...
ENTERED: InMageUssRequestor::GatherUssAppsInfo
ENTERED: InMageUssRequestor::Initialize
EXITED: InMageUssRequestor::Initialize
ENTERED: InMageUssRequestor::WaitAndCheckForAsyncOperation
EXITED: InMageUssRequestor::WaitAndCheckForAsyncOperation
ENTERED: InMageUssRequestor::DiscoverTopLevelComponents
EXITED: InMageUssRequestor::DiscoverTopLevelComponents
ENTERED: InMageUssRequestor::DiscoverTonLevelComponents
```

Figure 286:



## Vacp -sync

By default all consistency tags issued are asynchronous in nature, at times these tags may be dropped due to heavy I/O. You may choose to use the “-sync” switch to issue a synchronous consistency tag. Vacp waits until the tags are drained by the s2.exe. By default the wait time is infinite, however you may use the “-tagtimeout” to restrict it to the desired time in seconds.

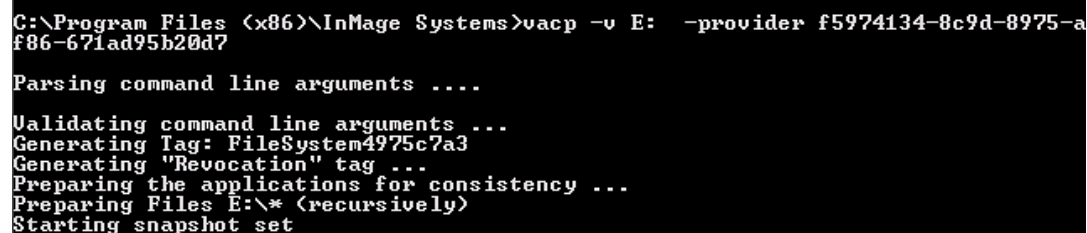


```
C:\Program Files\InMage Systems>vacp -v E: -t "My_Tag" -sync
Parsing command line arguments ....
Validating command line arguments ...
User Defined Tag: My_Tag
Generating Tag: FileSystem4959d1a4
Generating "Revocation" tag ...
Preparing the applications for consistency ...
Preparing Files E:\* (recursively)
Starting snapshot set
Using MS Software Shadow Copy provider: b5946137-7b9f-4925-af80-51abd60b20d5
Freezing the applications for consistency ...
Committing shadow copy for the set...
E:\ is mapped to Unique volume \\?\Volume{2d42de6e-d175-11dd-85c4-000c29157508}\
Tag GUID= 07424526-eeb0-4ea8-88f2-998f8e5f6404
Checking driver mode for given volumes
For volume E:\ driver is in Data mode.
Bytes passed to driver IOCTL = 138
Sending tags to the driver ...
Tags IO Pending
Waiting forever till tags get committed...
=====
Tags are successfully committed
=====
BytesReturn = 12
Number of Volumes: 1
Tag status: Committed
Status of the tag for E:\ volume is Committed
Marked all applications as successfully backed up
Resuming the applications after point-in-time consistency check-point
Deleting snapshot set(Shadow copies)
Succesfully deleted snapshot set(Shadow copies)
Exiting gracefully ...
C:\Program Files\InMage Systems>_
```

Figure 287

## Vacp -provider

You may also make use of any registered third party providers by using the “-provider” switch as shown in the picture below

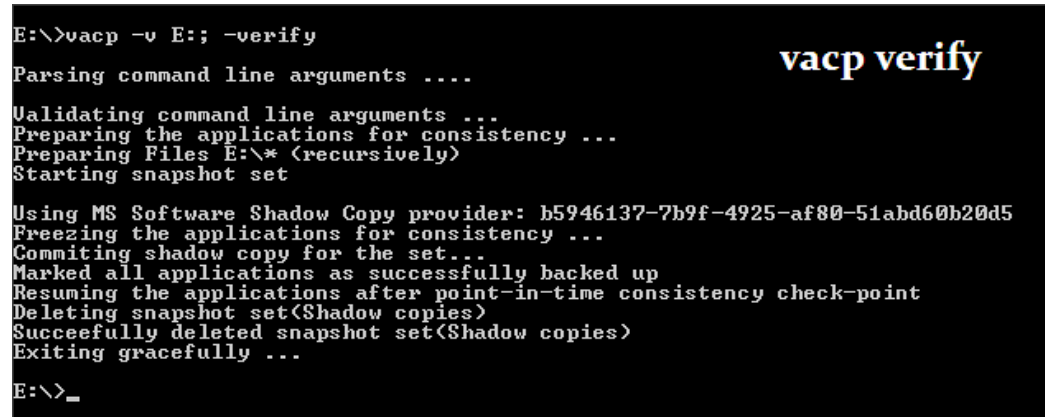


```
C:\Program Files (x86)\InMage Systems>vacp -v E: -provider f5974134-8c9d-8975-a
f86-671ad95b20d7
Parsing command line arguments ....
Validating command line arguments ...
Generating Tag: FileSystem4975c7a3
Generating "Revocation" tag ...
Preparing the applications for consistency ...
Preparing Files E:\* (recursively)
Starting snapshot set
```

Figure 288

## Vacp -verify

The VSS service should be up and running at all times on the production server since vacp internally uses it. To check the health of a VSS, you can use the “-verify” switch



```
E:\>vacp -v E:; -verify
Parsing command line arguments ....
Validating command line arguments ...
Preparing the applications for consistency ...
Preparing Files E:\* (recursively)
Starting snapshot set
Using MS Software Shadow Copy provider: b5946137-7b9f-4925-af80-51abd60b20d5
Freezing the applications for consistency ...
Committing shadow copy for the set...
Marked all applications as successfully backed up
Resuming the applications after point-in-time consistency check-point
Deleting snapshot set(Shadow copies)
Succceefully deleted snapshot set(Shadow copies)
Exiting gracefully ...
E:\>_
```

Figure 289

## 12.1.2 VACP Command Line Options (Linux/ Solaris)

### Syntax:

```
vacp.exe [-a <app1[</comp1>..];..>][-v <vol1;..>][-guid <volguid1;..>][-t <tag1;..>][-w <writer instance name;..>][-f][-x][-p<app1;..>][-s][-provider <Provider ID>][-notag][-sync][-tagtimeout <Timeout value in second>][-remote -serverdevice<device1,device2..>-serverip <server IP> -serverport <server port>][-h]
```

Table 23: VACP options (Linux)

Flags	Description
-h	Displays the usage information on the console. This option cannot be used with any other option.
-v <List of volumes>	Specifies one or more volume names. Volume names are case-insensitive. "All" keyword can be used to create tags on all volumes in the system. Example: E: G: VACP terminates with an error message when an invalid volume letter is specified.
-t <List of user tags>	Specifies one or more user tags. The maximum length of the user tag should not exceed 250 ASCII characters. Duplicate user tags are allowed.
-x	Insert tags without any consistency mechanism. This option must be specified along with -v and -t options. "-a" and "-x" are mutually exclusive.
-remote	This is used in combination with -serverdevice, -serverip and optionally -serverport. These switches are used only on client server architecture such as Citrix XenServer, Hyper-V.
-serverdevice	Corresponding volume on the base operating system where a tag is to be issued
-serverip	IP address of the server where the vacp server is running.
-serverport	Optional switch is used to specify the port number
-guid : <volguid1;volguid2;..>	Specifies one or more volume GUIDs or volume mount point GUIDs. (e.g., C4E7FA17-C0D0-4139-880F-E0874B907FA0;18EF281B-2EE8-4095-84DB-C938FE82987D) Here in the above example two volume GUIDs are separated by a semi colon.
-provider: <Provider ID>	Specifies VSS provider ID that VACP should use. By default VACP uses Microsoft Software Shadow copy provider. Example -provider f5974134-8c9d-8975-af86-671ad95b20d7
-notag	When specify, VACP will not create and issue the tag.

It simply creates snapshot for the consistency and delete it.



#### Caution:

While issuing consistency tags on the Solaris platform, always execute the vacp command while you are under the inmage VX agent installation folder. When you execute the vacp command from any other directory you may come across the “**ld.so.1: vacp: fatal: libstdc++.so.6: open failed: No such file or directory**” error

### 12.1.2.1 Examples

While using vacp on Linux, use the command “**./vacp**” (the dot followed by forward slash should precede the vacp)

#### vacp or vacp -h

```
[root@Lin-DR bin]# pwd
/usr/local/InMage/Vx/bin
[root@Lin-DR bin]# ./vacp
Usage:
./vacp  [-v <vol1,...>]  [-remote -serverdevice <device1,device2,...>]  [-t <tag1,...>] [-serverip <serve IPAddress>] [-serverport <ServerPort>] [-x] [-h]

-v <vol1,vol2,...>
    Specify volumes on which tag has to be generated.
    Specify "all" to generate tags on all volumes in the system
```

**VACP help**

Figure 290:

#### vacp – t, vacp –v

While –t is used to specify the name of the consistency tag –v is used to specify the volume.

```
[root@R4U5_Source bin]# ./vacp -v /dev/Volumegroup/V3 -t "Tag_One"

Generating tag names ...
Tag: FileSystem46ce89d0
Tag: Tag_One
Sending Following Tag Request ...
Flags = 3
Num. Volumes = 1
Volume: 1 Name: /dev/mapper/Volumegroup-V3 Length:26
Num. Tags = 2
Tag: 1 Length:12
Tag: 2 Length:23
tags successfully issued
```

**vacp -t**

Figure 291:

## **vacp -x**

Using **-x** will issue a blind tag and will not have any consistency mechanism.

```
[root@R4U5_Source bin]# ./vacp -v /dev/Volumegroup/V3 -t "Tag_two" -x

Generating tag names ...
Tag: Tag_two

Sending Following Tag Request ...
Flags = 1
Num. Volumes = 1
Volume: 1 Name: /dev/mapper/Volumegroup-V3 Length:26
Num. Tags = 1
Tag: 1 Length:12
tags successfully issued
```

## **Vacp -x**

**Figure 292:**

## **vacp -remote**

The **-remote** switch is used in combination with **-serverdevice**, **-serverip** and **-serverport**. This is usually used when a tag is to be issued on a client system. This requires vacp client and vacp server to interact with each other to successfully issue a consistency tag. For example, Citrix Xenserver vacp server is running on the base operating system and guest virtual machines will have vacp client (downloaded from the CX UI). A tag is initiated from the guest virtual machine and passes over to the vacp server.

```
[root@linuxg1 local]# ./vacp -remote -v /dev/mapper/VolGroup00-LogVol00 -t "my
tag" -serverdevice /dev/mapper/UG_XenStorage--d3cad621--52f6--3565--5cba--a225
c249e5b9-LV--0ac2d8a6--164d--4c3a--b74a--e58c8d45f188 -serverip 10.0.1.146 -ser
verport 20003

Generating tag names ...
Tag: FileSystem4877465c
Tag: my tag

Sending Following Tag Request ...
Flags = 3
Num. Volumes = 1
Volume: 1 Name: /dev/mapper/UG_XenStorage--d3cad621--52f6--3565--5cba--a225c249
e5b9-LV--0ac2d8a6--164d--4c3a--b74a--e58c8d45f188 Length:112
Num. Tags = 2
Tag: 1 Length:11
Tag: 2 Length:23
```

## **vacp -remote**

**Figure 293:**

## 12.2 CX UI v/s cdpcli

The following are the list of features that the CX UI and cdpcli support. While most of them are supported by both, a few of the features are restricted to either one of them.

**Table 24: CX UI V/s CDPCLI**

Feature	CX UI	Windows cdpcli	Linux cdpcli
Simple Snapshot	✓	✓	✓
Time based snapshot	✓	✓	✓
Event based snapshot	✓	✓	✓
Simple virtual snapshot	✓	✓	✓
Time based virtual snapshot	✓	✓	✓
Event based virtual snapshot	✓	✓	✓
Virtual volume	✗	✓	✓
Validate retention logs	✗	✓	✓
Showssummary of retention logs	✗	✓	✓
Scheduled snapshots	✓	✗	✗
Applytracklog (Read write tracking virtual snapshot feature)	✗	✓	✓
Rollback (Target volume rollback)	✓	✓	✓
Roll back multiple volumes	✓	✓	✓
Snapshot multiple volumes	✓	✓	✓



### Notes:

Before executing cdpcli commands ensure to login as an administrator on windows 2008 platform.

## 12.3 CDPCLI Interface

### 12.3.1 When is cdpcli used

Cdpcli is a console based application used on the target host to perform recovery operations like snapshots, target volume rollback, managing retention logs etc. This tool is mostly used when the source site is unreachable from the target site or when the CX server is down. This is an alternative to the CX UI based recovery.

Cdpcli is used on the target host for volume recovery. Recovery through cdpcli is performed when the CX server is offline. Gain access to the target host and navigate to the agent installation path to find the file “cdpcli.exe”.

### 12.3.2 Cdpcli on windows

Using the command directly without any arguments will display the list of supported operations as shown in the figure below.

```
C:\Program Files\InMage Systems>cdpcli.exe
usage error: please use cdpcli.exe --h [operation] to view usage information
where operation can be one of the following:
validate                - To validate retention database
showsummary             - To view summary information
listevents              - To list consistency events
snapshot               - To take a snapshot
recover                 - To perform recovery snapshot
rollback               - To perform rollback
hide                   - To hide a volume
unhide_ro              - To unhide a volume in read-only mode
unhide_rw              - To unhide a volume in read-write mode
vsnap                  - To perform virtual snapshot
virtualvolume          - To create virtual volume
listcommonpoint         - To list common recovery point
iopattern              - To get the io pattern
fixerrors              - To fix the database issues
export                 - To export retention information from t
he database
displaystatistics      - To display the current replication statistics
```

Figure 294:

#### 12.3.2.1 Validate

To validate retention logs, use the cdpcli command with the argument validate

The syntax is

```
cdpcli --validate --db=<Path of the db file within the retention logs>
```

Or

```
cdpcli --validate --vol=<target_volume>
```

```
C:\Program Files\InMage Systems>cdpcli --validate --vol=i:
Note:
Validation of the database requires exclusive access to t
No other application/service will be able to access the d
while validation is in progress
-----
verifying data file :D:/dl1/970de395d0/03DDEB26-3024-D94B-A
pv1_diffsync_000438e6678081aa.dat
End of File:648 bytes
```

Figure 295:

### 12.3.2.2 Showsummary

Using the showsummary switch will display

- Total number of consistency tags and their type
- Time range of retention logs
- Space occupied on disk
- Location of the retention logs
- total number of retention logs

The syntax is

**cdpcli--showsummary --vol=<target vol>**

```
C:\Program Files\InMage Systems>cdpcli --showsummary --vol=i:

Database:          D:/dl1/970de395d0/03DDEB26-3024-D94B-A974875273456A4F/
I/cdpv1.db
Version:           1
Revision:          1
Log Type:          Roll-Backward
Disk Space (app):  17464 bytes
Total Data Files:  1
Recovery Time Range(GMT): 2007/8/30 7:50:0:640:625:0 to
                        2007/8/30 7:52:47:421:875:0
```

Figure 296:

### 12.3.2.3 List Events

As the name suggests, list events display a complete list of vacp consistency tags issued on the source.

This command displays

- Timestamp of the consistency tag
- Accuracy of the tag
- Application name
- Event name

The syntax is

**cdpcli --listevents --vol=<target volume>**

```
C:\Program Files\InMage Systems>cdpcli --listevents --vol=i:

-----
No.   TimeStamp(GMT)          Accuracy    Application  Event
-----
1     2007/8/30 8:5:16:718:750:0  Exact      FS           FileSystem46d67a34
2     2007/8/30 8:5:16:718:750:0  Exact      USERDEFINED  Tag_Final
3     2007/8/30 8:4:47:203:125:0  Exact      FS           FileSystem46d67a18
4     2007/8/30 8:4:47:203:125:0  Exact      USERDEFINED  Tag_Beta
5     2007/8/30 8:4:20:171:875:0   Exact      FS           FileSystem46d679fc
6     2007/8/30 8:4:20:171:875:0   Exact      USERDEFINED  Tag_alfa
Total Events:6
```

Figure 297:



**Table 25**

Flags	Description
<code>--vol=&lt;volume name&gt;</code>	Target volume name for list event
<code>--app=&lt;application name&gt;</code>	Displays lists of vacp consistency tags issued on the source with respect to the specified application. This gives the information about the timestamp, accuracy, and event name of the tags.
<code>--event=&lt;event name&gt;</code>	Displays lists of vacp consistency tags issued on the source with respect to the specified event name. It also gives information regarding timestamp, accuracy, application name of the tag.
<code>--at=&lt;time&gt;</code>	Displays lists of vacp consistency tags issued on the source at that particular stamp. It also gives information regarding accuracy, application name, and event name of the tag.
<code>--beforetime=&lt;time&gt;</code>	Displays lists of vacp consistency tags issued on the source before the specified timestamp. It also gives information regarding accuracy, application name, and event name of the tags.
<code>--aftertime=&lt;time&gt;</code>	Displays lists of vacp consistency tags issued on the source after the specified timestamp. It also gives information regarding accuracy, application name, and event name of the tags.

Listing events specific to an application

`cdpcli.exe --listevents --vol=<volume name> --app=<tagname>.`

```
C:\Program Files\InMage Systems>cdpcli.exe --listevents --vol=F: --app=USERDEFINED
-----
No.   TimeStamp(GMT)           Accuracy   Application Event
-----
1     2008/12/23 7:11:30:333:625:0 Exact      USERDEFINED tagtwo
2     2008/12/23 6:46:5:130:500:0 Exact      USERDEFINED tagone
Total Events:2
C:\Program Files\InMage Systems>
```

**Figure 298:**

You may use the switch "`--eventnum`" to perform a recovery to the specific event number. This switch is to be used in combination with the "`--app`" switch. When the "`--app`" switch is not specified by default the application is set to "`USERDEFINED`"

To view a consistency tag's timestamp, accuracy etc

**cdpcli.exe --listevents --vol=<volume name> --event=<Event>**

```
C:\Program Files\InMage Systems>cdpcli.exe --listevents --vol=F: --event=tagone
```

No.	TimeStamp(GMT)	Accuracy	Application	Event
1	2008/12/23 6:46:5:130:500:0	Exact	USERDEFINED	tagone

Total Events:1

```
C:\Program Files\InMage Systems>
```

Figure 299:

To view consistency tags at a specified time

**cdpcli.exe--listevents --vol=<volume name> --at= <time to be displayed>**

```
C:\Program Files\InMage Systems>cdpcli.exe --listevents --vol=F: --at=2008/12/23 7:11:30:333:625:0
```

No.	TimeStamp(GMT)	Accuracy	Application	Event
1	2008/12/23 7:11:30:333:625:0	Exact	USERDEFINED	tagtwo
2	2008/12/23 7:11:30:333:625:0	Exact	FS	FileSystem49508f09

Total Events:2

Figure 300:

To view consistency tags before a specified time

**cdpcli.exe --listevents --vol=<Target volume name in the replication pair> --before= <time>**

```
C:\Program Files\InMage Systems>cdpcli.exe --listevents --vol=F: --beforetime=2008/12/23 7:11:30:333:625:0
```

No.	TimeStamp(GMT)	Accuracy	Application	Event
1	2008/12/23 6:46:5:130:500:0	Exact	FS	FileSystem49508912
2	2008/12/23 6:46:5:130:500:0	Exact	USERDEFINED	tagone

Total Events:2

Figure 301:

To view consistency tags after a specified time

**cdpcli.exe --listevents --vol=<Target volume name in the replication pair > --aftertime= <time>**

```
C:\Program Files\InMage Systems>cdpcli.exe --listevents --vol=F: --aftertime=2008/12/23 7:01:30:333:625:0
```

No.	TimeStamp(GMT)	Accuracy	Application	Event
1	2008/12/23 7:11:30:333:625:0	Exact	USERDEFINED	tagtwo
2	2008/12/23 7:11:30:333:625:0	Exact	FS	FileSystem49508f09

Total Events:2

Figure 302:

### 12.3.2.4 Snapshot

The destination volume (which will contain the snapshot) should not be lesser than the source (target volume in the replication pair) volume. The syntax is

**cdpcli --snapshot --source=<Target volume in the replication pair> --dest=<Free volume equal or larger than the source volume>**

```

C:\WINDOWS\system32\cmd.exe - cdpcli --snapshot --source=i: --dest=k:

C:\Program Files\InMage Systems>cdpcli --snapshot --source=i: --dest=k:

Note:
  1. Do not specify a volume which is part of any replication
     as either source or destination volume unless the
     Replication agent service is stopped. It can result in failure
     of the replication process and/or corrupt the volumes
  2. Aborting this process during the operation may leave the
     source/destination volume in dismounted or locked state.
     Please check the troubleshooting reference to restore the
     volume state.

          snapshot being taken from volume I to volume K

Checking for Replication agent service <svagent> status ...
Replication agent service <svagents> is currently stopped
Snapshot Started for destination:K:. Current Time:2006\4\27 10:13:11:639
Snapshot In Progress for destination:K:
Destination: K: progress:1
Destination: K: progress:2
Destination: K: progress:3
  
```

Figure 303:

Table 26

Flags	Description
--dest=<destination volume>	Destination for snapshot.(Required)
--source=<source volume>	Source volume for snapshot.(Required)
--postscript=<script path>	Script to be executed on completion of snapshot.(Optional)
--force=<yes no ask>	Specify this option as “yes” if you want to continue the operation even if replication services are running. Specify “no” if you want to terminate the operation if replication services are running. Specify “ask” if user confirmation is required before proceeding with the operation. The default value is “ask”. (Optional)
--prescript=<script path>	Script to be executed before starting the snapshot copy. Snapshot is taken only If script returns an exit code of zero.(Optional)
--snapshotpair=<snapshotpairdetails>	Allows you to take multiple snapshots in a single command.
--runonce	Runs the prescript and postscript once for the entire process

### 12.3.2.5 Snapshot pair

You may perform a one snapshot per replication pair through the “**snapshotpairs**” switch. For e.g. two snapshots may be performed for two replication pairs through this command.

```
cdpcli.exe --snapshot --snapshotpairs=<target volume, snapshot volume; target volume2, snapshot volume2>
```

```
C:\Program Files\InMage Systems>cdpcli.exe --snapshot --snapshotpairs=F:, E:;
Note:
  1. Specifying a volume which is part of any replication
     as either source or destination volume is not allowed
     unless the Replication agent service is stopped.
  2. All data on destination volume would be overwritten by
     this operation.
  3. Aborting this process during the operation may leave the
     source/destination volume in dismounted or locked state.
     Please check the troubleshooting reference to restore the
     volume state.

Checking for Replication agent service (svagent) status ...
Replication agent service (svagents) is running

Press "y" or "Y" to continue, "n" or "N" to terminate the snapshot ...y
source volume \\.\F: capacity :1070592000 bytes

E:Snapshot started.
Destination volume \\.\E: capacity :1073741824 bytes

E:Copy in progress.
NOTE: this is not an error StopFilter portablehelpers for volume \\?\Volume{85a9
656b-9af1-11dd-b29f-000c29fd822c}
E:10%
E:20%
E:30%
E:40%
E:50%
E:60%
E:70%
E:80%
E:90%
E:100%

E:Status: complete.
C:\Program Files\InMage Systems>_
```

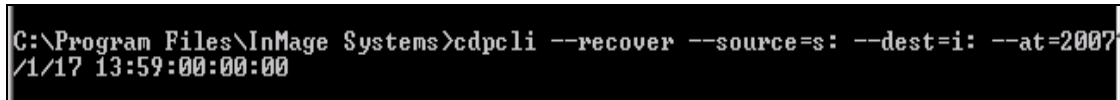
Figure 304:

### 12.3.2.6 Recover

A snapshot may be performed to a valid point in time by using the “**recover**” switch.

The syntax is shown below

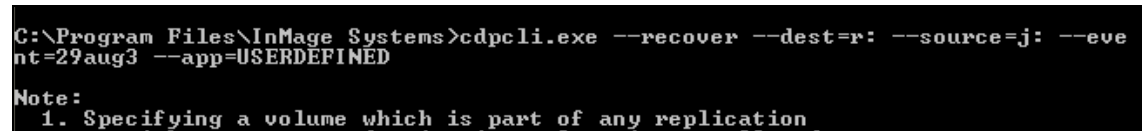
```
cdpcli --recover --source =<target volume in the replication pair> --dest=<volume which will hold the snapshot> --at=<time in yyyy/mm/dd/ hh:mm:ss:ss>
```



```
C:\Program Files\InMage Systems>cdpcli --recover --source=s: --dest=i: --at=2007/1/17 13:59:00:00:00
```

Figure 305: Time based recovery

```
cdpcli --recover --source =<target volume in the replication pair> -- dest=<volume which will hold the snapshot> --event=<event_name> --app=<corresponding application name>
```

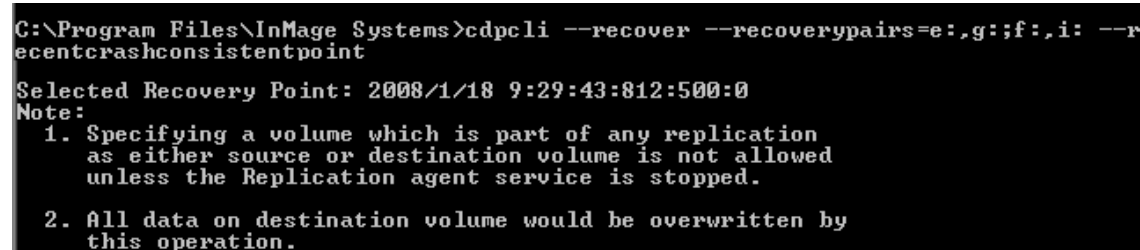


```
C:\Program Files\InMage Systems>cdpcli.exe --recover --dest=r: --source=j: --event=29aug3 --app=USERDEFINED
Note:
1. Specifying a volume which is part of any replication.
```

Figure 306: Event based recovery

You may use the switch “**--eventnum**” to perform a recovery to the specific event number. This switch is to be used in combination with the “**--app**” switch. When the “**--app**” switch is not specified by default the application is set to “**USERDEFINED**”

```
cdpcli -- recover -- recoverypairs =<protected volume, snapshot volume; protected volume, snapshot volume> {--event=<event_name> or -at=<time> or lastcrashconsistentpoint or lastfsconsistentpoint or lastappconsistentpoint=<corresponding application name>}
```



```
C:\Program Files\InMage Systems>cdpcli --recover --recoverypairs=e:,g:f:,i: --recentcrashconsistentpoint
Selected Recovery Point: 2008/1/18 9:29:43:812:500:0
Note:
1. Specifying a volume which is part of any replication as either source or destination volume is not allowed unless the Replication agent service is stopped.
2. All data on destination volume would be overwritten by this operation.
```

Figure 307: Recovering multiple volumes to same point

**Table 27: cdpcli recover options**

Flags	Description
--dest=<destination volume>	destination for recovery snapshot.(Required)
--source=<source volume>	source volume for snapshot.(Required)
--db=<database path>	path for retention database.(optional)
--at=<Time>	recover to the specified timestamp OR
--app=<app name> -- event=<Event>	recover to the specified event.OR
--app=<app name> -- eventnum=<num>	Recover to the specified event number. OR
--app=<app name> -- aftertime=<Time>	Recover to a application consistency event occurring after the specified time. OR
--app=<app name> -- beforetime=<Time>	Recover to a application consistency event occurring before the specified time. specify only one recovery option.(Required) Time format:yr/mm/dd hr:min:sec:millisec:usec:nanosec
--prescript=<script path>	Script to be executed before starting the recovery process. recovery is done only if script returns exit code zero.(Optional)
--postscript=<script path>	script to be executed on completion of recovery process.(Optional)
--force=<yes no ask>	Specify this option as “yes” if you want to continue the operation even if replication services are running. Specify “no” if you want to terminate the operation if replication services are running. specify “ask” if user confirmation is required before proceeding with the operation. The default value is “ask”. (Optional)
--recoverypairs	<p>Recovery pair details format:</p> <p>Source volume 1, target volume 1,mount point 1, retention database 1; Source volume 2, target volume 2, mount point 2, retention database 2; ...</p> <p>Where:</p> <p>Source Volume: Source Volume for recovery.</p> <p>Target Volume: destination Volume for recovery.</p> <p>Mount Point: Directory name to mount the destination volume. Applicable for UNIX only. (Optional)</p> <p>Retention Database: Path for retention database. Normally, this is not required and the settings are fetched from Central Management server. This is required only when communication with central management server is unavailable. (Optional)</p>

<code>--recentcrashconsistentpoint</code>	Perform Recovery operation on the specified volumes to the last common crash consistent point in the specified time range. If the time range is not specified, the most recent common crash consistent point is chosen.
<code>--recentfsconsistentpoint</code>	Perform Recovery operation on the specified volumes to the last common file system consistent point in the specified time range. If the time range is not specified, the most recent common file system consistent point is chosen. This requires at least one common consistency tag across all the volumes that are to be recovered
<code>--recentappconsistentpoint=&lt;app name&gt;</code>	Perform Recovery operation on the specified volumes to the last common application consistent point in the specified time range. If the time range is not specified, the most recent common application consistent point is chosen. This is applicable only to windows OS
<code>--timerange</code>	This option is used to specify search interval to look for the recovery point. This option is used in combination with <code>lastcrashconsistentpoint</code> , <code>lastfstconsistentpoint</code> and <code>lastappconsistentpoint</code> .
<code>--eventnum</code>	Recover to the specified event number. The counting starts with value 1 starting from latest event to oldest one. Note: This option is not available if <code>recoverypairs</code> option is specified.
<code>--aftertime</code>	Recover all the volumes to first common consistency event occurring after the specified time. Time format: yr/mm/dd hr:min:sec:millisec:usec:nanosec
<code>--runonce</code>	Based on this option, prescripts and postscripts are executed either for each recovery pair or only once for the whole process.

### 12.3.2.7 Rollback

Using the “**rollback**” switch will perform a target volume rollback. This can be done based on time and an event. Performing a rollback will break the replication pair.

**Table 28: cdpcli rollback options**

Flags	Description
<code>--rollbackpairs =&lt;rollback pairs details&gt;</code>	Rollback pair details format: target volume 1,mount point 1, retention database 1; target volume 2, mount point 2, retention database 2; Where: Target Volume: destination Volume for rollback. Mount Point: Directory name to mount the destination volume. Applicable for UNIX only. (Optional) Retention Database: Path for retention database. Normally, this is not required and the settings are fetched from Central Management server. This is required only when communication with central management server is unavailable. (Optional)
<code>--rollbackpairs=all</code>	Use this option if you want to rollback all the target volumes. This option is available only when the central management server is reachable. Note: the volumes may have to be manually mounted after the rollback completes.
<code>--dest=&lt;destination volume&gt;</code>	Destination volume for rollback
<code>--mountpoint=&lt;mount point&gt;</code>	Directory name to mount the destination volume. Applicable for UNIX only. (Optional)
<code>--db=&lt;database path&gt;</code>	Path for retention database. Normally, this is not required and the settings are fetched from Central Management server. This is required only when communication with central management server is unavailable. (Optional)
<code>--recentcrashconsistentpoint</code>	Perform rollback operation on the specified volumes to the last common crash consistent point in the specified time range. If the time range is not specified, the most recent common crash consistent point is chosen.
<code>--recentfsconsistentpoint</code>	Perform rollback operation on the specified volumes to the last common file system consistent point in the specified time range. If the time range is not specified, the most recent common file system consistent point is chosen. Requires at least one common consistency tag for all the volumes that are to be rolled back
<code>--recentappconsistentpoint</code>	Perform rollback operation on the specified volumes to the last common application consistent point in the specified time range. If the time range is not specified, the most recent common application consistent point is



	chosen. This is applicable only to Windows OS.
--timerange=<start time, end time>	This option is used to specify search interval to look for the rollback point. This option is used in combination with recentcrashconsistentpoint, recentfstconsistentpoint and recentappconsistentpoint.
--at=<Time>	Perform rollback operation on the volumes to the specified timestamp. The specified time should be within the recovery time range for all the selected volumes
--app=<application name>	This option is used to restrict searching of events to the specified application.
--event=<value>	Rollback all the volumes to the specified event. The specified event should be available for all the specified volumes.
--eventnum=<num>	Rollback to the specified event number. The counting starts with value 1 starting from latest event to oldest one. Note: This option is not available if recoverypairs option is specified.
--aftertime=<Time>	Rollback all the volumes to first common consistency event occurring after the specified time. Time format: yr/mm/dd hr:min:sec:millisec:usec:nanosec
--beforetime=<Time>	Rollback all the volumes to a common latest application Consistency event occurring before the specified time. Time format: yr/mm/dd hr:min:sec:millisec:usec:nanosec
--prescript=<script path>	Script to be executed before starting the rollback process. This script is executed for each rollback pair before starting the rollback if runonce is not specified. Otherwise, it is run once before the rollback process starts. Rollback is done only if script returns exit code zero.(Optional)
--postscript =<script path>	Script to be executed on completion of the rollback process. This script is executed for each rollback pair on completion of the rollback if runonce is not specified. Otherwise, it is run once the whole rollback process completes. (Optional)
--runonce	Based on this option, prescripts and postscripts are executed either for each recovery pair or only once for the whole process.
--force=<yes no ask>	Specify this option as “yes” if you want to continue the operation even if replication services are running. Specify “no” if you want to terminate the operation if replication services are running. Specify “ask” if user confirmation is required before proceeding with the operation. The default value is “ask”. (Optional)
--deleteretentionlog	Deletes retention logs specific after volume is rolled back.

The syntax for time based target volume rollback is

`cdplici -- rollback -- dest=<target volume> or --db=<path of the DB file> -- app= <name of the application > -- event=<name of the event>`

```
C:\Program Files\InMage Systems>cdplici --rollback --dest=i: --event="Tag_alfa"
Note:
1. Specifying a volume which is part of any replication
   as either source or destination volume is not allowed
   unless the Replication agent service is stopped.
```

Figure 308: Event based roll back

```
C:\Program Files\InMage Systems>cdplici --rollback --dest=i: --at="2007/08/30 9:20:00:0"
Note:
1. Specifying a volume which is part of any replication
   as either source or destination volume is not allowed
```

Figure 309: Time based roll back

The name of the application can be seen in the output of the “`listevents`” switch and in the UI it can be seen under the search results for the consistency tag









Search Result				
	Accuracy	Timestamp	Application	Taq Name
		2007/8/28 9:53:55:791	File System	FileSystem46d3f09b
		2007/8/28 9:53:55:791	User Defined	fsystem
		2007/8/28 9:49:12:660	File System	FileSystem46d3ef7f
		2007/8/28 9:49:12:660	User Defined	fsystem

Figure 310:

`cdplici -- recover -- rollbackpairs =<protected volume1; protected volume2> {-- event=<event_name> or -at=<time> or lastcrashconsistentpoint or lastfsconsistentpoint or lastappconsistentpoint=<corresponding application name>}`

```
C:\Program Files\InMage Systems>cdplici --rollback --rollbackpairs=e::f: --recent crashconsistentpoint
Selected Recovery Point: 2008/1/18 9:32:59:78:125:0
Note:
1. Specifying a volume which is part of any replication
   as either source or destination volume is not allowed
   unless the Replication agent service is stopped.
2. All data on destination volume would be overwritten by
   this operation.
3. Aborting this process during the operation may leave the
   source/destination volume in dismounted or locked state.
   Please check the troubleshooting reference to restore the
```

Figure 311: Rolling Back Multiple Target Volumes To The Same Point

## Deleteretentionlog

This command deletes the retention logs after the volume is rolled back. The syntax is

```
cdpcli.exe --rollback --dest =<Target volume> --event= <tag name> --  
deleteretentionlog=<yes or no>
```

```
C:\Program Files\InMage Systems>cdpcli.exe --rollback --dest=F: --event=tagtwo  
--deleteretentionlog=yes  
Selected Recovery Point: 2008/12/23 7:11:30:333:625:0  
Note:  
  1. Specifying a volume which is part of any replication  
     as either source or destination volume is not allowed  
     unless the Replication agent service is stopped.  
  2. All data on destination volume would be overwritten by  
     this operation.  
  3. Aborting this process during the operation may leave the  
     source/destination volume in dismounted or locked state.  
     Please check the troubleshooting reference to restore the  
     volume state.  
Checking for Replication agent service (svagent) status ...  
Replication agent service (svagents) is running  
Press "y" or "Y" to continue, "n" or "N" to terminate the rollback ...y  
F:Rollback Started.  
F:Rollback InProgress.  
Cx is not updated with the deletion of the vsnap N:\  
Deleting USNAP N:\ of Target Volume F ,since it is in RW Mode now...  
F:100%  
F:Status: complete.  
C:\Program Files\InMage Systems>
```

Figure 312:

### 12.3.2.8 Hide

To hide a volume use the switch “**hide**” followed by the volume letter. The syntax is

```
cdpcli --hide <volume to be hidden>
```

```
C:\Program Files\InMage Systems>cdpcli --hide g:  
Note:  
  While Replication status is in resync, hide or unhide operations on  
  target volume should not be performed.  
Checking for Replication agent service (svagent) status ...  
Replication agent service (svagents) is running  
  
Request has been sent successfully to CX server  
Waiting for Cx Info change, Checking for desired state  
G is now hidden
```

Figure 313:

### 12.3.2.9 Unhide\_ro

To expose a protected volume (target volume) in read only mode use the switch “unhide\_ro”

The syntax is

**Cdpccli--unhide\_ro** <Drive letter to be exposed>

```
C:\Program Files\InMage Systems>cdpccli --unhide_ro g:
Note:
  While Replication status is in resync, hide or unhide operations on
  target volume should not be performed.

Checking for Replication agent service (svagent) status ...
Replication agent service (svagents) is running

Request has been sent successfully to CX server
Waiting for Cx Info change, Checking for desired state
GVerification complete, Operation successful
```

Figure 314:

### 12.3.2.10 Unhide\_rw

To expose the protected volume (target volume) in read write mode use the switch “unhide\_rw”

The syntax is

**Cdpccli --unhide\_rw** <drive letter to be exposed>

```
C:\Program Files\InMage Systems>cdpccli --unhide_rw g:
Note:
  While Replication status is in resync, hide or unhide operations on
  target volume should not be performed.

Checking for Replication agent service (svagent) status ...
Replication agent service (svagents) is running

Request has been sent successfully to CX server
Waiting for Cx Info change, Checking for desired state
G is now accessible in read-write mode by all applications
Verification complete, Operation successful
```

Figure 315:

### 12.3.2.11 Virtualvolume

This creates a file on the target system which can be mounted and used as a local disk (which may can also be a target volume). This involves two steps

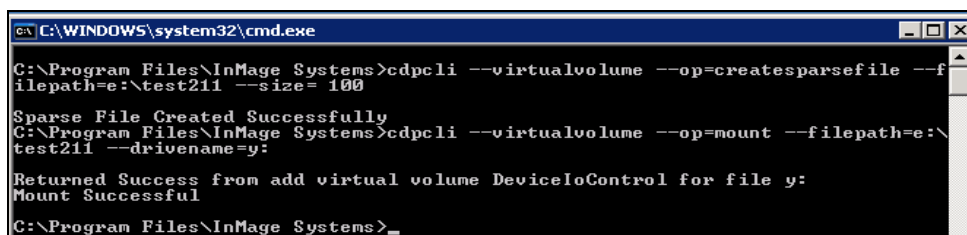
- Creating a sparse file
- Mount the sparse file to make it appear as a volume

To create a sparse file the syntax is

```
cdpcli --virtualvolume --op=createsparsefile --filepath=<complete file path> --size=<enter memory in MB>
```

To mount the sparse file as a volume the syntax is

```
Cdpcli --virtualvolume --op=mount --filepath=<complete file path> --drivename=<available volume letter>
```



```
C:\WINDOWS\system32\cmd.exe
C:\Program Files\InMage Systems>cdpcli --virtualvolume --op=createsparsefile --filepath=e:\test211 --size= 100
Sparse File Created Successfully
C:\Program Files\InMage Systems>cdpcli --virtualvolume --op=mount --filepath=e:\test211 --drivename=y:
Returned Success from add virtual volume DeviceIoControl for file y:
Mount Successful
C:\Program Files\InMage Systems>_
```

Figure 316: Mounting Virtual Volume on Windows

This volume can be formatted and used just like any other volume. Do not use the resident volume (where the sparse file is placed) as a target or for snapshots.

You may list all the virtual volume through the following command

```
cdpcli.exe --virtualvolume --op=list
```

To unmount selective virtual volume

```
cdpcli.exe --virtualvolume --op=unmount -drivename=<drive name>
```

To unmount all virtual volumes use the “- - op=unmountall” switch

Table 29

Flags	Description
--op=<createsparsefile>	Creates a sparse file
--op=<mount >	Mounts the sparse file
--op=<unmount>	To unmount a virtual volume
--op=<unmountall>	To unmount all virtual volumes
--op=<list>	To display list of all virtual volumes
--filepath=<path of the file>	Path of the file to create or mount a sparse file
--drivename=<available drive name>	Drive name for mount or unmount a virtual volume

### 12.3.2.12 Vsnap

Virtual snapshots are of three kinds

- Read only virtual snapshot
- Read write virtual snapshot
- Read write tracking virtual snapshot

The read only virtual snapshots do not support any writes and all writes are discarded

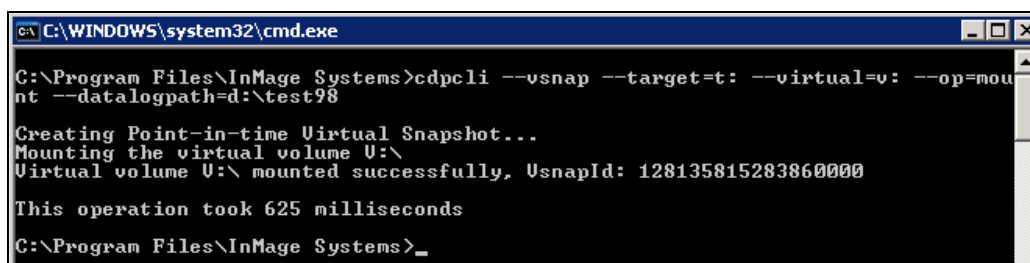
The read write virtual snapshot retains writes onto in the form of log files stored under another volume. Once unloaded the writes too are discarded

The read write tracking virtual snapshot are much advanced than the other two since the writes made to the disk can be retained and applied to another virtual snapshot.

Each of these snapshots can be based on time or an event. So the combinations increase.

#### Plain Virtual Snapshot

This is a plain snapshot that does not require CDP retention option to be enabled on the replication pair



```
GA C:\WINDOWS\system32\cmd.exe
C:\Program Files\InMage Systems>cdpcli --vsnap --target=t: --virtual=v: --op=mount
nt --datalogpath=d:\test98
Creating Point-in-time Virtual Snapshot...
Mounting the virtual volume U:\
Virtual volume U:\ mounted successfully, UsnapId: 128135815283860000
This operation took 625 milliseconds
C:\Program Files\InMage Systems>_
```

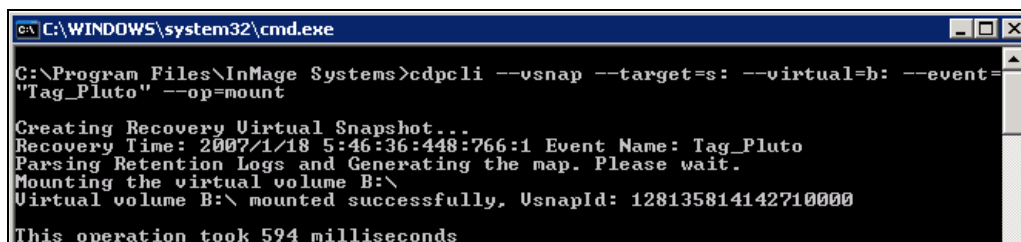
Figure 317: Event based Virtual Snapshot on Windows

#### Event based Read only virtual snapshot.

For event based virtual snapshots the syntax is

**Cdpcli --vsnap --target=<target volume letter with colon> --virtual=<virtual drive letter with colon> --event=<event\_name> --op=mount**

Optionally DB path, pre and post scripts can also be used here.



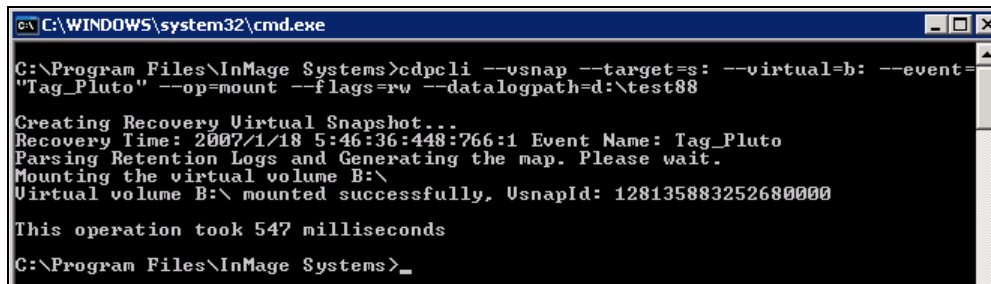
```
GA C:\WINDOWS\system32\cmd.exe
C:\Program Files\InMage Systems>cdpcli --vsnap --target=s: --virtual=b: --event=
"Tag_Pluto" --op=mount
Creating Recovery Virtual Snapshot...
Recovery Time: 2007/1/18 5:46:36:448:766:1 Event Name: Tag_Pluto
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume B:\
Virtual volume B:\ mounted successfully, UsnapId: 128135814142710000
This operation took 594 milliseconds
```

Figure 318: Event based Virtual Snapshot on Windows

## Event based read write virtual snapshot

For event based read/write virtual snapshots the syntax is

**Cdpccli --vsnap --target=<target volume letter with colon> --virtual= <virtual drive letter with colon> --flags=rw -event=<event\_name> --op=mount -datalogpath=<datalogpath>**



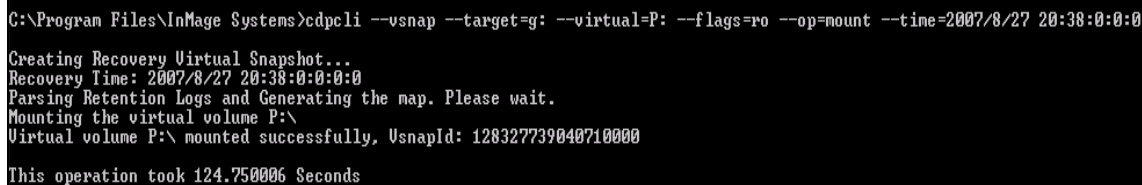
```
C:\WINDOWS\system32\cmd.exe
C:\Program Files\InMage Systems>cdpccli --vsnap --target=s: --virtual=b: --event=
"Tag_Pluto" --op=mount --flags=rw --datalogpath=d:\test88
Creating Recovery Virtual Snapshot...
Recovery Time: 2007/1/18 5:46:36:448:766:1 Event Name: Tag_Pluto
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume B:\
Virtual volume B:\ mounted successfully, UsnapId: 128135883252680000
This operation took 547 milliseconds
C:\Program Files\InMage Systems>_
```

Figure 319: Event based Read Write Virtual Snapshot

## Time based read only virtual snapshot

For time based virtual snapshot use the syntax

**Cdpccli --vsnap -target=<target volume in a replication pair> --virtual<volume where vsnap will be mounted> --flags<ro for read only, rw for read write> --op=<mount to mount, unmount to unmount> --time <time falling between the retention time range>**

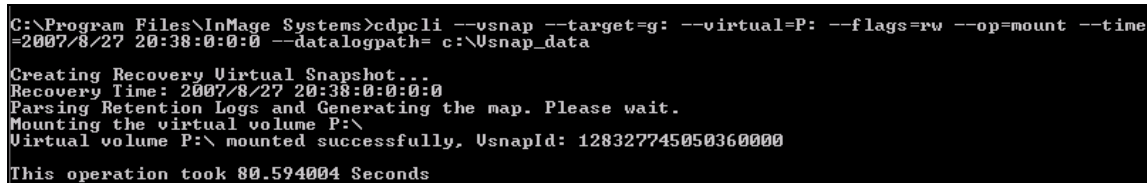


```
C:\Program Files\InMage Systems>cdpccli --vsnap --target=g: --virtual=P: --flags=ro --op=mount --time=2007/8/27 20:38:0:0
Creating Recovery Virtual Snapshot...
Recovery Time: 2007/8/27 20:38:0:0:0
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume P:\
Virtual volume P:\ mounted successfully, UsnapId: 128327739040710000
This operation took 124.750006 Seconds
```

Figure 320:

## Time based read write virtual snapshot

This is similar to a time based read only virtual snapshot except for the “Flags” for read only its “ro” and for read write its “rw”



```
C:\Program Files\InMage Systems>cdpccli --vsnap --target=g: --virtual=P: --flags=rw --op=mount --time=
2007/8/27 20:38:0:0:0 --datalogpath= c:\Usnap_data
Creating Recovery Virtual Snapshot...
Recovery Time: 2007/8/27 20:38:0:0:0
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume P:\
Virtual volume P:\ mounted successfully, UsnapId: 128327745050360000
This operation took 80.594004 Seconds
```

Figure 321:

## Working with Read write tracking virtual snapshots

This section includes creating read write tracking virtual snapshot, unloading a virtual snapshot, unloading a virtual snapshot but keeping the track logs intact and applying track logs to a new virtual snapshot.

For read write tracking virtual snapshot the syntax is

```
cdpcli --vsnap --op=mount --target=<target volume letter> --virtual=[ virtual volume drive letter or mount point] --event=[ event name] --datalogpath= <folder where vsnap writes are stored> --flags=rwt
```

```
C:\Program Files\InMage Systems>cdpcli --vsnap --op=mount --target=h: --virtual=x: --event=exchange451a1484 --datalogpath=f:\vsnap_data --flags=rwt

Creating Recovery Virtual Snapshot...
Recovery Time: 2006/9/27 6:4:57:46:562:3 Event Name: exchange451a1484
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume X:\
Virtual volume X:\ mounted successfully, UsnapId: 128038214112720000
This operation took 562 milliseconds
```

Figure 322: Created a Vsnap (read write tacking) on Windows

Open up the virtual snapshot and create some files on it.

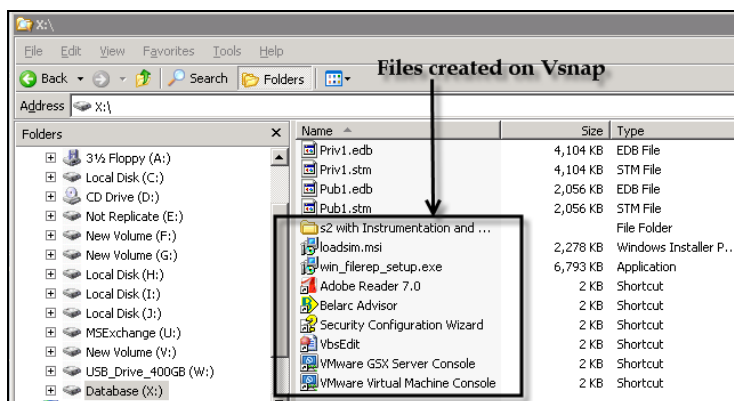


Figure 323:

Unmount the virtual snapshot and keep the track logs intact by using the command

```
Cdpcli --vsnap --op=unmount --virtual=<virtual drive letter> --flags=nodelete
```

```
C:\Program Files\InMage Systems>cdpcli --vsnap --op=unmount --virtual=x: --flags=nodelete

x: Unmounted Successfully
```

Figure 324: Unmount Vsnap on Windows with Logs Intact



## Create another virtual snapshot

```
C:\Program Files\InMage Systems>cdpcli --vsnap --op=mount --target=h: --virtual=T: --event=exchange451a1484 --datalogpath=f:\vsnap_data --flags=rw

Creating Recovery Virtual Snapshot...
Recovery Time: 2006/9/27 6:4:57:46:562:3 Event Name: exchange451a1484
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume T:\
Virtual volume T:\ mounted successfully, VsnapId: 128038224125020000
This operation took 593 milliseconds
```

Figure 325:

Applying old writes of X: to new virtual snapshot, the syntax is

`Cdpcli --vsnap --op=applytracklog --vsnapid=<vsnap id of the old volume on which files were created> --target=<new virtual snapshot> --datalogpath=<datalog path of old virtual snapshot>`

```
C:\Program Files\InMage Systems>cdpcli.exe --vsnap --op=applytracklog --vsnapid=128038214112720000 --target=T: --datalogpath=f:\vsnap_data

Applying Track Logs... Please wait....
Completed the operation Successfully.
```

Figure 326: Applying old Logs to New Vsnap on Windows

To delete track logs the syntax is

`Cdpcli --vsnap --op=deletelogs --vsnapid=<vsnapid to which the tracklogs belong> --datalogpath=<location of datalogs>`

```
C:\Program Files\InMage Systems>cdpcli.exe --vsnap --op=deletelogs --vsnapid=128038214112720000 --datalogpath=f:\vsnap_data

Completed Successfully
```

Figure 327:

To apply tracks of an existing vsnap to another vsnap

The syntax is

`Cdpcli --vsnap --op=applytracklog --virtual=<drive letter of original vsnap with writes> --target=<new vsnap target to which the writes will be applied> --datalogpath=<location of original track logs>`

```
C:\Program Files\InMage Systems>cdpcli.exe --vsnap --op=applytracklog --virtual=X: --target=T: --datalogpath=f:\vsnap_data

Applying Track Logs... Please wait....
Completed the operation Successfully.
```

Figure 328:

## List of virtual snapshots

The syntax for listing all virtual snapshot is

**cdpcli -vsnap --op=list**

```
C:\Program Files\InMage Systems>cdpcli --vsnap --op=list
Following is the list of virtual volumes mounted in the system
    1> M:\
    2> N:\
    3> P:\
    4> Q:\
    5> T:\
    6> X:\
```

Figure 329:

## Unmount

**Cdpcli --vsnap --op=unmount --virtual=<Virtual snapshot drive>**

```
C:\Program Files\InMage Systems>cdpcli --vsnap --op=unmount --virtual=P:
P: Unmounted Successfully
```

Figure 330:

## Recentcrashconsistentpoint

This switch is used to perform recovery operations such as rollback, vsnap, snapshot etc to a common consistency point for all existing target volumes within the target host. For e.g. when the target host contains three target volumes, using this switch will perform a recovery to a recent common consistency point among the three target volumes. When a “--timerange” switch is used in combination with the “recentcrashconsistentpoint” then the search for the common point will be restricted to the time range

The syntax for recentcrashconsistentpoint is

**cdpcli.exe --vsnap --vsnappairs="<target volume1, vsnap 1; target volume 2, vsnap2>" --recentcrashconsistentpoint --op=mount**

```
C:\Program Files\InMage Systems>cdpcli.exe --vsnap --vsnappairs="q:,u:;r:,v:" --recentcrashconsistentpoint --op=mount
Retention DB Path D:\consistencylogs\db4ff983f2\AF102F3E-90D2-8541-80F4FC5934928FD6\Q being used for storing vsnap metadata for the pair Q: -> u:
Retention DB Path D:\consistency2logs\252953d683\AF102F3E-90D2-8541-80F4FC5934928FD6\R being used for storing vsnap metadata for the pair R: -> v:
Selected Recovery Point: 2009/1/9 15:54:22:770:615:2
Acquiring read lock on Q:. attempt:1
```

Figure 331:

## Recentfsconsistentpoint

Take Recovery Vsnap on the specified volumes to the last common file system consistent point in the specified time range. If the time range is not specified, the most recent common file system consistent point is chosen.

The syntax is

```
cdpcli.exe --vsnap --vsnappairs="<target volume1, vsnap 1; target volume 2, vsnap2>" --recentfsconsistentpoint --op=mount
```

```
C:\Program Files\InMage Systems>cdpcli.exe --vsnap --vsnappairs="q:,u:;r:,v:" --recentfsconsistentpoint --op=mount

Retention DB Path D:\consistencylogs\db4ff983f2\AF102F3E-90D2-8541-80F4FC5934928FD6\Q being used for storing vsnap metadata for the pair Q: -> u:
Retention DB Path D:\consistency2logs\252953d683\AF102F3E-90D2-8541-80F4FC5934928FD6\R being used for storing vsnap metadata for the pair R: -> v:
Selected Recovery Point: 2009/1/9 15:46:54:281:773:8
Acquiring read lock on Q:. attempt:1
```

Figure 332

## Recentappconsistentpoint

Take Recovery Vsnap on the specified volumes to the last common application consistent point in the specified time range. If the time range is not specified, the most recent common application consistent point is chosen.

The syntax is

```
cdpcli.exe --vsnap --vsnappairs="<target volume1, vsnap 1; target volume 2, vsnap2>" --recentappconsistentpoint --op=mount
```

```
C:\Program Files\InMage Systems>cdpcli.exe --vsnap --vsnappairs="q:,u:;r:,v:" --recentappconsistentpoint --op=mount

Retention DB Path D:\consistencylogs\db4ff983f2\AF102F3E-90D2-8541-80F4FC5934928FD6\Q being used for storing vsnap metadata for the pair Q: -> u:
Retention DB Path D:\consistency2logs\252953d683\AF102F3E-90D2-8541-80F4FC5934928FD6\R being used for storing vsnap metadata for the pair R: -> v:
Selected Recovery Point: 2009/1/9 15:46:54:281:773:8
```

Figure 333

### 12.3.2.13 Listcommonpoint

This command is used to identify a common recovery point across a bunch of replication pairs. The syntax is

```
cdpcli.exe --listcommonpoint
```

```
C:\Program Files\InMage Systems>cdpcli --listcommonpoint
Common Recovery Point: 2008/12/23 5:42:20:693:0:0
```

Figure 334:

## lopattern

This option provides statistics of the I/O distribution of the target storage to check whether the target volume is bottleneck. This option reads the retention database and print out the I/O profile. The syntax is

```
Cdpcli.exe --iopattern --vol=<path name>
```

```
F:\InMage Systems>cdpcli.exe --iopattern --vol=H:

Io Profile:
size    %Access %Read %Random Delay Burst Alignment Reply
512B    1      100    100    0      1      1      sector none
512B    1      0      0      0      1      1      sector none
512B    1      100    0      0      1      1      sector none
512B    2      0      100    0      1      1      sector none
1KB     0      100    100    0      1      1      sector none
1KB     0      0      0      0      1      1      sector none
1KB     0      100    0      0      1      1      sector none
1KB     3      0      100    0      1      1      sector none
2KB     2      100    100    0      1      1      sector none
2KB     2      0      0      0      1      1      sector none
2KB     2      100    0      0      1      1      sector none
```

Figure 335

## 12.3.2.14 Displaystatistics

This command is used to display target volume, differentials pending on CX, differentials pending on target, current RPO, apply rate and apply time of the differentials to the target volume.

```
cdpcli.exe --displaystatistics --vol=<volumename>
```

```
C:\Program Files\InMage Systems>cdpcli.exe --displaystatistics --vol=F:

\\.\F: is a symbolic link to F:

-----
##### REPLICATION STATISTICS #####
-----
Target Volume Name:      F:
Diffs pending in CX:    0
Diffs pending in Target: 198946196
Current RPO (secs):      0
Apply rate (Bytes/sec):  63
Apply time (secs):       3157876

C:\Program Files\InMage Systems>_
```

Figure 336:

## 12.3.3 Cdpcli on Linux/ Solaris

### 12.3.3.1 Validate

This command is used to check the health of the retention logs. Navigate to the VX agent installation folder and issue the following command

Syntax

```
./cdpcli --validate --vol=<target_volume>
```

Or

```
./cdpcli --validate --db=<path of the retention db for this replication pair>
```

```
[root@Lin-DR bin]# pwd
/usr/local/Xiotech/Vx/bin
[root@Lin-DR bin]# ./cdpcli --validate --vol=/dev/mapper/volume-v1

Note:
  Validation of the database requires exclusive access to the database.
  No other application/service will be able to access the database
  while validation is in progress

                                cdpcli --validate
-----
verifying data file :/mnt/retention/0b8c64df4b/8fb864a8-165b-4178-a4d0-2791a0a1d
f6e/dev/mapper/volume-v1/cdpv1__diffsync__000443a94f7d700a.dat
```

Figure 337:

### 12.3.3.2 Showsummary

The showsummary switch will display

- Total number of consistency tags and their type
- Time range of retention logs
- Space occupied on disk
- Location of the retention logs
- total number of retention logs

The syntax is

```
Cdpcli --showsummary --vol=<target volume>
```

Or

```
Cdpcli --showsummary --db=< path of the retention db for this replication pair>
```

```
[root@Lin-DR bin]# ./cdpcli --showsummary --vol=/dev/mapper/volume-v1

Database:                /mnt/retention/0b8c64df4b/8fb864a8-165b-4178-a4d0-2791
a0a1df6e/dev/mapper/volume-v1/cdpv1.db
Version:                  1
Revision:                 1
Log Type:                 Roll-Backward
Disk Space (app):         18230 bytes
Total Data Files:         3
Recovery Time Range(GMT): 2008/1/14 6:24:7:746:446:3 to
                          2008/1/14 6:30:6:244:946:2
Consistency Event Summary:
-----
Application              Num. Events
-----
FS                        2
USERDEFINED               2
```

Figure 338:

### 12.3.3.3 Listevents

The listevents switch will display a complete list of vacp consistency tags issued from the source. This command displays

- Timestamp of the consistency tag
- Accuracy of the tag
- Application name
- Event name

The syntax is

```
./cdpcli --listevents --vol=<target_volume>
```

Or

```
./cdpcli --listevents --db=< path of the retention db for this replication pair>
```

```
[root@Lin-DR bin]# ./cdpcli --listevents --vol=/dev/mapper/volume-v1
```

No.	TimeStamp(GMT)	Accuracy	Application	Event
1	2008/1/14 6:30:6:244:946:2	Approximate	FS	FileSystem478b016e
2	2008/1/14 6:30:6:244:946:2	Approximate	USERDEFINED	Tag_Two
3	2008/1/14 6:29:55:659:555:5	Approximate	FS	FileSystem478b0163
4	2008/1/14 6:29:55:659:555:5	Approximate	USERDEFINED	Tag_One

Total Events:4

```
[root@Lin-DR bin]# █ cdpcli --listevents
```

Figure 339:

### App

Using the “--app” switch will display the number of consistency tags for a specified application. The syntax is

```
./cdpcli --listevents --vol=<target_volume> --app=FS
```

```
[root@target bin]# ./cdpcli --listevents --vol=/dev/mapper/s-vol3 --app=FS
```

No.	TimeStamp (GMT)	Accuracy	Application	Event
1	2008/12/23 0:0:28:522:718:3	Exact	FS	FileSystem49502a1c
2	2008/12/23 0:0:24:938:263:2	Exact	FS	FileSystem49502a18
3	2008/12/22 23:52:55:390:604:8	Exact	FS	FileSystem49502857

Total Events:3

Figure 340:

You may use the switch “--eventnum” to perform a recovery to the specific event number. This switch is to be used in combination with the “--app” switch. When the “--app” switch is not specified by default the application is set to “USERDEFINED”

## --Event

This command is used to display the information about a consistency tag.

```
./cdpcli --listevents --vol=<target_volume> --event=<tag name>
```

```
[root@target bin]# ./cdpcli --listevents --vol=/dev/mapper/s-vol3 --event=tag3
-----
                                cdpcli --event
-----
No.   TimeStamp (GMT)              Accuracy      Application Event
-----
1     2008/12/23 0:0:24:938:263:2  Exact        USERDEFINED tag3
Total Events:1
```

Figure 341:

## --at

This command displays the consistency tags applicable for a specified time. The syntax is

```
./cdpcli --listevents --vol=<target_volume> --at=<time to be displayed>
```

```
[root@target bin]# ./cdpcli --listevents --vol=/dev/mapper/s-vol3 --at=
2008/12/23 0:0:28:522:718:3
-----
                                cdpcli --at
-----
No.   TimeStamp (GMT)              Accuracy      Application Event
-----
1     2008/12/23 0:0:28:522:718:3  Exact        USERDEFINED tag4
2     2008/12/23 0:0:28:522:718:3  Exact        FS           FileSystem49502a1c
Total Events:2
[root@target bin]# ./cdpcli --listevents --vol=/dev/mapper/s-vol3 --beforetime
=2008/12/23 0:0:28:522:718:3
```

Figure 342:

## --Beforetime

This command is used to display the list of events that occur before a specific time. The syntax is

```
./cdpcli --listevents --vol=<target_volume> --beforetime =<time>
```

```
[root@target bin]# ./cdpcli --listevents --vol=/dev/mapper/s-vol3 --beforetime
=2008/12/23 0:0:28:522:718:3
-----
                                cdpcli --beforetime
-----
No.   TimeStamp (GMT)              Accuracy      Application Event
-----
1     2008/12/23 0:0:24:938:263:2  Exact        FS           FileSystem49502a18
2     2008/12/23 0:0:24:938:263:2  Exact        USERDEFINED tag3
3     2008/12/22 23:52:55:390:604:8 Exact        FS           FileSystem49502857
4     2008/12/22 23:52:55:390:604:8 Exact        USERDEFINED tag2
Total Events:4
```

Figure 343:

## --Aftertime

This command is used to display the list of events that occur after a given time. The syntax is

```
./cdpcli --listevents --vol=<target_volume> --aftertime =<time>
```

```
[root@target bin]# ./cdpcli --listevents --vol=/dev/mapper/s-vol3 --aftertime
=2008/12/22 23:52:55:390:604:8
```

cdpcli --aftertime				
No.	TimeStamp (GMT)	Accuracy	Application Event	
1	2008/12/23 0:0:24:938:263:2	Exact	USERDEFINED	tag3
2	2008/12/23 0:0:24:938:263:2	Exact	FS	FileSystem49502a18
3	2008/12/23 0:0:28:522:718:3	Exact	USERDEFINED	tag4
4	2008/12/23 0:0:28:522:718:3	Exact	FS	FileSystem49502a1c
Total Events:4				

Figure 344:

### 12.3.3.4 Snapshot

To perform a snapshot, access the target host's console and navigate to the VX installation folder to issue the following command.

syntax

```
./cdpcli --snapshot --dest=<where it has to be copied> --source <target volume in the
replication pair> --mountpoint =<where it has to be mounted>
```

```
[root@Lin-DR bin]# ./cdpcli --snapshot --dest=/dev/mapper/volume-v2 --source=/de
v/mapper/volume-v1 --mountpoint=/v6
```

cdpcli --snapshot	
Note:	
1. Specifying a volume which is part of any replication	
as either source or destination volume is not allowed	
unless the Replication agent service is stopped.	

Figure 345:



## Snapshot Pairs

You may perform a one snapshot per replication pair through the “**snapshotpairs**” switch. For e.g. two snapshots may be performed for two replication pairs through this command.

Syntax:

```
./cdpcli --snapshot --snapshotpairs=<target volume1, snapshot volume1,  
mountpoint1; target volume2, snapshot volume2, mountpoint2>
```

```
[root@target bin]# ./cdpcli --snapshot --snapshotpairs=/dev/mapper/s-vol  
/dev/mapper/s-vol132, /vol132;
```

**Snapshot Pairs**  
Note:  

1. Specifying a volume which is part of any replication  
as either source or destination volume is not allowed  
unless the Replication agent service is stopped.
2. All data on destination volume would be overwritten by  
this operation.
3. Aborting this process during the operation may leave the  
source/destination volume in dismounted or locked state.  
Please check the troubleshooting reference to restore the  
volume state.

Figure 346:

### 12.3.3.5 Recover

Time based recovery can be performed by using the switch recover, this can take a snapshot back in time on to a specified disk mentioned

The syntax is

```
./Cdpcli --recover --source <source volume for snapshot> --dest= <Snapshot volume> --  
{event=<Event> or at=<Time>} --mountpoint=<Where to mount>
```

```
[root@Lin-DR bin]# ./cdpcli --recover --source=/dev/mapper/volume-v1 --dest=/dev/  
/mapper/volume-v2 --event="Tag_One" --mountpoint=/R4  
  
Selected Recovery Point: 2008/1/14 6:29:55:659:555:5 Event based  
Note:  
1. Specifying a volume which is part of any replication
```

Figure 347:

```
[root@Lin-DR bin]# ./cdpcli --recover --source=/dev/mapper/volume-v1 --dest=/dev/  
/mapper/volume-v2 --at=2008/1/14 6:25:00:00 --mountpoint=/R3  
  
Note: The specified time 2008/1/14 6:25:00:00 is not a crash consistent point fo  
r all the volumes.  
Selected Recovery Point: 2008/1/14 6:25:0:0:0:0 cdpcli --recover  
Note:
```

Figure 348:

```
./Cdpcli --recover --recoverypairs <target volume1, snapshot volume1, mountpoint1;  
target volume2,snapshot volume2, mountpoint2> --{event=<Event> or --at=<Time> or --  
recentcrashconsistentpoint or --lastcrashconsistentpoint or --  
lastfsconsistentpoint or --lastappconsistentpoint}
```

```
[root@Lin-DR bin]# ./cdpcli --recover --recoverypairs="/dev/mapper/volume-v2,/de  
v/mapper/volume-v1, /home/rec1;/dev/mapper/volume-v3,/dev/mapper/volume-v4,/home  
/rec2" --recentcrashconsistentpoint recentcrashconsistentpoint  
  
Selected Recovery Point: 2008/1/15 0:16:59:512:450:4  
Note:  
1. Specifying a volume which is part of any replication  
as either source or destination volume is not allowed
```

Figure 349: recovery pairs



#### Notes:

Snapshot option will give a physical snapshot of the protected volume at the time when the command is issued.

Recover option will give a physical snapshot of the protected volume at some point back in time. This can be based time or an event.

Specifying the mountpoint option is mandatory to view the snapshot

### 12.3.3.6 Rollback

Using the switch rollback, will perform a target volume rollback. This can be done based on time and an event. Performing a rollback will break the replication pair.

The syntax for event based target volume rollback is

```
./cdplci --rollback --dest=<target volume> --event=<name of the event> --  
mountpoint=<where the rolled back volume shall be mounted>
```

```
[root@Lin-DR bin]# ./cdplci --rollback --dest=/dev/mapper/volume-v1 --event="Tag  
_One" --mountpoint=/R5  
  
Selected Recovery Point: 2008/1/14 6:29:55:659:555:5  
Note:  
1. Specifying a volume which is part of any replication  
as either source or destination volume is not allowed  
unless the Replication agent service is stopped
```

**Rollback**

Figure 350: Event based rollback

The syntax for time based target volume rollback is

```
./cdplci --rollback --dest=<target volume> --time=<time to which the target volume is to  
be rolledback > --mountpoint=<where the rolled back volume shall be mounted>
```

```
[root@Lin-DR bin]# ./cdplci --rollback --dest=/dev/mapper/volume-v1 --at=2008/1/  
14 7:43:45:00:00 --mountpoint=/r1  
  
Note: The specified time 2008/1/14 7:43:45:00:00 is not a crash consistent point  
for all the volumes.  
Selected Recovery Point: 2008/1/14 7:43:45:0:0:0
```

Figure 351: Time based rollback

```
./cdplci --rollback --rollbackpairs=<target volume 1, mount point 1 ; target volume 2,  
mount point 2> {--at=<time where the rollback has to be attained> or --event=<event to be  
rolled back to> or --recentcrashconsistentpoint or --lastcrashconsistentpoint or  
--lastfsconsistentpoint or --lastappconsistentpoint}
```

```
[root@Lin-DR bin]# ./cdplci --rollback --rollbackpairs="/dev/mapper/volume-v2, /  
home/rec1;/dev/mapper/volume-v3,/home/rec2" --recentcrashconsistentpoint  
  
Selected Recovery Point: 2008/1/15 0:33:45:396:532:8  
Note:  
1. Specifying a volume which is part of any replication  
as either source or destination volume is not allowed  
unless the Replication agent service is stopped.
```

**recentcrashconsistentpoint**

Figure 352: Rollbackpairs to recentcrashconsistentpoint

## deleteretentionlog

This command deletes the retention logs after the volume is rolled back. The syntax is

```
./cdpcli --rollback --dest =<Target volume> --mountpoint= <mount point> --event=
<tag name> --deleteretentionlog=<yes or no>
```

```
[root@LIN-DR bin]# ./cdpcli --rollback --dest=/dev/mapper/volume-v1 --mountpoint
=/v1 --event="Tag_1" --deleteretentionlog=yes
/dev/mapper/volume-v1:Rollback Started.
/dev/mapper/volume-v1:Rollback InProgress.

UnMounting /home/vsnap1 ...
Trying to find processes accessing /home/vsnap1 ...
Performing unmount operation on /home/vsnap1 ...

Trying to find processes accessing /dev/vs/cli1 ...
Removing entries for the device /dev/vs/cli1 from /etc/fstab ...
Removal of Mountpoints corresponding to /dev/vs/cli1 succeeded
Device /dev/vs/cli1 removal succeeded.
Cx is not updated with the deletion of the vsnap /dev/vs/cli1
Deleting VSNAP /dev/vs/cli1 of Target Volume /dev/mapper/volume-v1 ,since it is
in RW Mode now...
Replication agent service (svagents) is running
```

## deleteretentionlog

Figure 353

### 12.3.3.7 Hide

To hide a volume use the switch “hide” followed by the volume letter

The syntax is

```
./cdpcli --hide <volume to be hidden>
```

```
[root@Lin-DR bin]# ./cdpcli --hide /dev/mapper/volume-v1

Note:
While Replication status is in resync, hide or unhide operat
target volume should not be performed.
```

## cdpcli --hide

Figure 354:

### 12.3.3.8 Unhide\_ro

To expose a protected volume (target volume) in read only mode use the switch “**unhide\_ro**”

The syntax is

```
./cdpcli --unhide_ro <volume to be exposed> --mountpoint=<where it needs to be mounted>
```

```
[root@Lin-DR bin]# ./cdpcli --unhide_ro /dev/mapper/volume-v1 --mountpoint=/v3 -  
-filesystem=ext3
```

**cdpcli --unhide\_ro**

Note:  
While Replication status is in resync, hide or unhide operations on  
target volume should not be performed.

Checking for Replication agent service (svagent) status ...  
Replication agent service (svagents) is running

Figure 355:

### 12.3.3.9 Unhide\_rw

To expose the protected volume (target volume) in read write mode use the switch “**unhide\_rw**”

The syntax is

```
./cdpcli --unhide_rw < volume to be exposed> --mountpoint=<where it needs to be  
mounted>
```

```
[root@Lin-DR bin]# ./cdpcli --unhide_rw /dev/mapper/volume-v1 --mountpoint=/v4 -  
-filesystem=ext3
```

**cdpcli --unhide\_rw**

Note:  
While Replication status is in resync, hide or unhide operations on

Figure 356:

### 12.3.3.10 Virtual volume

This creates a file on the target system which can be mounted and used as a local disk (can also serve as a target for a replication pair). The command and its syntax are very similar to its windows counterpart. Using a virtual volume will require you to create a sparse file first and then creating a device.

#### Creating a sparse file

Syntax

```
./cdpcli --virtualvolume --op=createsparsefile --filepath=<full path where the sparse file should be created> --size=<Size in MB>
```

```
[root@Lin-DR bin]# ./cdpcli --virtualvolume --op=createsparsefile --filepath=/home/vvol --size=50
```

**Create sparse file**

Sparse file /home/vvol of size 50 MB created succesfully

Figure 357:

#### Creating a device

Syntax

```
./cdpcli --virtualvolume --op=createvolume --filepath=<full path of sparsefile>
```

```
[root@Lin-DR bin]# ./cdpcli --virtualvolume --op=createvolume --filepath=/home/vvol
```

**Mount Volume**

The devicefile /dev/volpack0 has been created succesfully for /home/vvol

Figure 358:

#### Listing virtual volumes

Getting a list of virtual volumes

```
"./cdpcli --virtualvolume --op=list"
```

```
[root@Lin-DR bin]# ./cdpcli --virtualvolume --op=list
```

Following is the list of virtual volumes mounted in the system

```
1) /dev/volpack0
```

Figure 359: list of virtual volumes on Linux

#### Removing a virtual volume

Syntax

```
./cdpcli --virtualvolume --op=unmount --drivename =<name of the drive>
```

```
[root@Lin-DR bin]# ./cdpcli --virtualvolume --op=unmount --drivename=/dev/volpack0
```

**Remove virtual volume**

The volpack device /dev/volpack0 has been unmounted succesfully

Figure 360: To remove only one volume

Using "--op=unmountall" will remove all virtual volumes

**Table 30**

Flags	Description
--op=<createsparsefile>	Creates a sparse file
--op=<mount >	Mounts the sparse file
--op=<unmount>	To unmount a virtual volume
--op=<unmountall>	To unmount all virtual volumes
--op=<list>	To display list of all virtual volumes
--filepath=<path of the file>	Path of the file to create or mount a sparse file
--drivename=<available drive name>	Drive name for mount or unmount a virtual volume

### 12.3.3.11 Vsnap

Virtual snapshots are of three kinds

- Read only virtual snapshot
- Read write virtual snapshot
- Read write tracking virtual snapshot

The read only virtual snapshots do not support any writes and all writes are discarded

The read write virtual snapshot retains writes onto in the form of log files stored under another volume. Once unloaded the writes too are discarded

The read write tracking virtual snapshot are much advanced than the other two since the writes made to the disk can be retained and applied to another virtual snapshot.

Each of these snapshots can be based on time or an event. So the combinations increase.

#### Point in time virtual snapshot

This is a plain snapshot that does not require CDP retention option to be enabled on the replication pair  
The syntax is

```
./cdpcli --vsnap --target=<target volume> --virtual=<where it has to be mounted as virtual> --op=mount
```

```
[root@Lin-DR bin]# ./cdpcli --vsnap --target=/dev/mapper/volume-v1 --virtual=/home/ro --op=mount
Creating Point-in-time Virtual Snapshot...
Mounting the virtual volume /home/ro
Virtual volume /home/ro mounted successfully, VsnapId: 1200301902
This operation took 0 milliseconds
```

Figure 361:



## Event based Read only virtual snapshot.

For event based virtual snapshots the syntax is

```
./cdpcli --vsnap --target=<target volume> --virtual=<virtual volume where it has to be mounted> --event=<event_name> --op=mount
```

When you intend to use the file system or application tags, you will need to specify the “--app=<USERDEFINED/ FILESYSTEM>” switch

Optionally db path, pre and post scripts can also be used here.

```
[root@Lin-DR bin]# ./cdpcli --vsnap --target=/dev/mapper/volume-v1 --virtual=/home/ro1 --op=mount

Creating Point-in-time Virtual Snapshot...
Mounting the virtual volume /home/ro1
Virtual volume /home/ro1 mounted successfully, VsnapId: 1200302695

This operation took 0 milliseconds
[root@Lin-DR bin]# █
```

Figure 362:

## Event based read write virtual snapshot

For event based read/write virtual snapshots the syntax is

```
./cdpcli --vsnap --target=<target volume> --virtual= <virtual volume> --flags=rw --event=<event_name> --op=mount --datalogpath=<datalogpath>
```

```
[root@Lin-DR bin]# ./cdpcli --vsnap --target=/dev/mapper/volume-v1 --virtual=/home/ro2 --flags=rw --event="Tag_One" --datalogpath=/home/datarw --op=mount

Creating Recovery Virtual Snapshot...
Recovery Time: 2008/1/14 8:22:7:396:175:9 Event Name: Tag_One
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume /home/ro2
Virtual volume /home/ro2 mounted successfully, VsnapId: 1200302855

This operation took 0 milliseconds
█
```

Figure 363:



### Notes:

For time based virtual snapshots --time option is used.  
For time based physical snapshots --at option is used

## Time based read only virtual snapshot

For time based virtual snapshot use the syntax

***./cdpcli --vsnap --target=<target volume> --virtual=<mountpoint where vsnap is to be mounted> --flags=<ro for read only(default), rw for read write> --op=<mount> --time=<time falling between the retention time range>***

```
[root@Lin-DR bin]# ./cdpcli --vsnap --target=/dev/mapper/volume-v1 --virtual=/vsnapro --time=2008/1/14 8:00:00:00 --op=mount
```

Creating Recovery Virtual Snapshot... **Time based vsnap**  
Recovery Time: 2008/1/14 8:0:0:0:0:0  
Parsing Retention Logs and Generating the map. Please wait.  
Mounting the virtual volume /vsnapro  
Virtual volume /vsnapro mounted successfully, VsnapId: 1200303486  
This operation took 0 milliseconds

Figure 364:

## Time based Read Write Virtual Snapshot

This is similar to a time based read only virtual snapshot except for the “Flags” for read only its “ro” and for read write its “rw”

```
[root@Lin-DR bin]# ./cdpcli --vsnap --target=/dev/mapper/volume-v1 --virtual=/vsnaprw --time=2008/1/14 8:00:00:00 --op=mount --flags=rw --datalogpath=/home/rwdata
```

Creating Recovery Virtual Snapshot... **Time based RW vsnap**  
Recovery Time: 2008/1/14 8:0:0:0:0:0  
Parsing Retention Logs and Generating the map. Please wait.  
Mounting the virtual volume /vsnaprw  
Virtual volume /vsnaprw mounted successfully, VsnapId: 1200304119  
This operation took 0 milliseconds

Figure 365:

## Working with Read Write Tracking Virtual Snapshots

This section includes creating read write tracking virtual snapshot, unloading a virtual snapshot, unloading a virtual snapshot but keeping the track logs intact and applying track logs to a new virtual snapshot.

For read write tracking virtual snapshot the syntax is

```
cdpcli --vsnap --op=mount--target=<target volume letter> --virtual=<virtual volume drive letter or mount point> --event=<event name> --datalogpath=<folder where vsnap writes are stored> --flags=rwt
```

```
[root@Lin-DR bin]# ./cdpcli --vsnap --target=/dev/mapper/volume-v1 --virtual=/vsnaprw --time=2008/1/14 8:00:00:00 --op=mount --flags=rwt --datalogpath=/home/rwd/ata
```

Creating Recovery Virtual Snapshot... **Creating RWT Vsnap**  
Recovery Time: 2008/1/14 8:0:0:0:0:0  
Parsing Retention Logs and Generating the map. Please wait.  
Mounting the virtual volume /vsnaprw  
Virtual volume /vsnaprw mounted successfully, VsnapId: 1200304234

This operation took 0 milliseconds  
[root@Lin-DR bin]# █

Figure 366: Created a vsnap (read write tacking) on Linux

Open up the virtual volume and create some files on it.

Unmount the virtual volume and keep the track logs intact by using the command

```
./cdpcli --vsnap --op=unmount --virtual=<virtual volume> --flags=nodelete
```

```
[root@Lin-DR bin]# ./cdpcli --vsnap --op=unmount --virtual=/vsnaprw --flags=nodelete
```

UnMounting /vsnaprw ... **Unload RWT vsnap**  
shutting down all processes accessing /vsnaprw ...  
Performing unmount operation on /vsnaprw ...

shutting down all processes accessing /dev/vsnap0 ...  
Removing entries for the device /dev/vsnap0 from /etc/fstab ...  
UnMount /dev/vsnap0 succeeded.  
/vsnaprw Unmounted Successfully  
[root@Lin-DR bin]# █

Figure 367: unmount vsnap on Linux with logs intact

## Create another read write Virtual Snapshot

```
[root@Lin-DR bin]# ./cdpcli --vsnap --target=/dev/mapper/volume-v1 --op=mount --
flags=rw --event="Tag_One" --virtual=/home/rwl --datalogpath=/logs
Creating Recovery Virtual Snapshot... Event based RW vsnap
Recovery Time: 2008/1/14 8:22:7:396:175:9 Event Name: Tag_One
Parsing Retention Logs and Generating the map. Please wait.
Mounting the virtual volume /home/rwl
Virtual volume /home/rwl mounted successfully, VsnapId: 1200304551
This operation took 0 milliseconds
```

Figure 368:

Applying old writes of X: to new virtual volume, the syntax is

**`./cdpcli --vsnap --op=<applytracklog> --vsnapid=<vsnap id of the old volume on which files were created> --target=<new virtual volume> --datalogpath=<datalog path of old virtual volume>`**

```
[root@Lin-DR bin]# ./cdpcli --vsnap --op=applytracklog --vsnapid=1200304234 --ta
rget=/home/rwl --datalogpath=/home/rwdata
UnMounting /home/rwl ...
shutting down all processes accessing /home/rwl ... Apply tracklogs
Performing unmount operation on /home/rwl ...

shutting down all processes accessing /dev/vsnap0 ...
Removing entries for the device /dev/vsnap0 from /etc/fstab ...
UnMount /dev/vsnap0 succeeded.
Started applying track logs. This operation may take some time to complete. Abor
ting this operation in the middle may leave the target volume in invalid state.

Applying Logs, Percentage Completed: 100%
Performing post apply track logs operations. Please wait....
Applying Track Logs Completed Successfully.
```

Figure 369: Applying Old Logs to New Vsnap on Linux

To delete track logs the syntax is

**`./cdpcli --vsnap --op=deletelogs --vsnapid=<vsnapid to which the tracklogs belong> --
datalogpath=<location of datalogs>`**

```
[root@Lin-DR bin]# ./cdpcli --vsnap --op=deletelogs --vsnapid=1200304234 --datal
ogpath=/home/rwdata
Completed Successfully Deleting vsnap logs
[root@Lin-DR bin]# █
```

Figure 370:

## To Apply Tracks of an Existing Vsnap to another Vsnap

The syntax is

```
./cdpcli --vsnap --op=<applytracklog> --virtual=<original vsnap with writes> --target=<new vsnap target to which the writes will be applied> --datalogpath=<location of original track logs>
```

```
[root@Lin-DR bin]# ./cdpcli --vsnap --op=applytracklog --virtual=/home/rwl --target=/r23 --datalogpath=/logs
UnMounting /r23 ...
shutting down all processes accessing /r23 ...
Performing unmount operation on /r23 ...

shutting down all processes accessing /dev/vsnap1 ...
Removing entries for the device /dev/vsnap1 from /etc/fstab ...
UnMount /dev/vsnap1 succeeded.
Started applying track logs. This operation may take some time to complete. Aborting this operation in the middle may leave the target volume in invalid state.

Applying Logs, Percentage Completed: 100%
Performing post apply track logs operations. Please wait....
Applying Track Logs Completed Successfully.
```

Figure 371:

## List of virtual snapshots

The syntax for listing all virtual snapshot is

```
./cdpcli --vsnap --op=list
```

```
[root@Lin-DR bin]# ./cdpcli --vsnap --op=list

Following is the list of virtual volumes mounted in the system

      1) /dev/vsnap0           /home/rwl
      2) /dev/vsnap1           /r23
[root@Lin-DR bin]#
```

Figure 372:

## Unmount

```
./cdpcli --vsnap --op=unmount --virtual=<Virtual snapshot volume>
```

```
[root@Lin-DR bin]# ./cdpcli --vsnap --op=unmount --virtual=/home/rwl
UnMounting /home/rwl ...
shutting down all processes accessing /home/rwl ...
Performing unmount operation on /home/rwl ...

shutting down all processes accessing /dev/vsnap0 ...
Removing entries for the device /dev/vsnap0 from /etc/fstab ...
UnMount /dev/vsnap0 succeeded.
/home/rwl Unmounted Successfully
```

Figure 373:

You may use the “**--op=unmountall**” to unmount all the virtual snapshots on the target host

### 12.3.3.12 Recentcrashconsistentpoint

This switch is used to perform recovery operations such as rollback, vsnap, snapshot etc to a common consistency point for all existing target volumes within the target host. For e.g. when the target host contains three target volumes, using this switch will perform a recovery to a recent common consistency point among the three target volumes. When a “**--timerange**” switch is used in combination with the “**recentcrashconsistentpoint**” then the search for the common point will be restricted to the time range

The syntax for “**recentcrashconsistentpoint**” is

```
./cdpcli --vsnap --vsnappairs="<target volume 1,vsnap1;target volume 2,vsnap2>" --  
recentcrashconsistentpoint --op=mount
```

```
[root@imits182 bin]# ./cdpcli --vsnap --vsnappairs="/dev/sda9,/home/doc3;/dev/sda10,/home/  
doc4" --recentcrashconsistentpoint --op=mount  
  
Retention DB Path /home/12/89logs/2b867a355b/ed86285b-72a5-4193-a91c-ed7d4155b2e1/dev/sda9  
being used for storing vsnap metadata for the pair /dev/sda9 -> /home/doc3  
Retention DB Path /home/13/810logs/364b5afb49/ed86285b-72a5-4193-a91c-ed7d4155b2e1/dev/sda  
10 being used for storing vsnap metadata for the pair /dev/sda10 -> /home/doc4  
Selected Recovery Point: 2009/1/9 13:47:25:674:870:1  
Corresponding I/O Sequence Point: 54693  
  
Recovery Time: 2009/1/9 13:47:25:674:870:1  
  
/home/doc3      100%  
  
Vsnap device /dev/vs/cli165 created successfully, VsnapId: 167
```

Figure 374

### Recentfsconsistentpoint

Take Recovery Vsnap on the specified volumes to the last common file system consistent point in the specified time range. If the time range is not specified, the most recent common file system consistent point is chosen.

The syntax is

```
./cdpcli --vsnap --vsnappairs="<target volume 1,vsnap1;target volume2,vsnap2>" --  
recentfsconsistentpoint --op=mount
```

```
[root@imits182 bin]# ./cdpcli --vsnap --vsnappairs="/dev/sda9,/home/doc3;/dev/sda10,/home/  
doc4" --recentfsconsistentpoint --op=mount  
  
Retention DB Path /home/12/89logs/2b867a355b/ed86285b-72a5-4193-a91c-ed7d4155b2e1/dev/sda9  
being used for storing vsnap metadata for the pair /dev/sda9 -> /home/doc3  
Retention DB Path /home/13/810logs/364b5afb49/ed86285b-72a5-4193-a91c-ed7d4155b2e1/dev/sda  
10 being used for storing vsnap metadata for the pair /dev/sda10 -> /home/doc4  
Selected Recovery Point: 2009/1/9 13:34:13:735:263:3
```

Figure 375

## Recentappconsistentpoint

Take Recovery Vsnap on the specified volumes to the last common application consistent point in the specified time range. If the time range is not specified, the most recent common application consistent point is chosen.

The syntax is

```
./cdpcli --vsnap --vsnappairs="target volume 1,vsnap1;target volume2,vsnap2" --  
recentappconsistentpoint --op=mount
```

```
[root@imits182 bin]# ./cdpcli --vsnap --vsnappairs="/dev/sda9,/home/doc3;/dev/sda10,/home/  
doc4" --recentappconsistentpoint --op=mount  
  
Retention DB Path /home/12/89logs/2b867a355b/ed86285b-72a5-4193-a91c-ed7d4155b2e1/dev/sda9  
being used for storing vsnap metadata for the pair /dev/sda9 -> /home/doc3  
Retention DB Path /home/13/810logs/364b5afb49/ed86285b-72a5-4193-a91c-ed7d4155b2e1/dev/sda  
10 being used for storing vsnap metadata for the pair /dev/sda10 -> /home/doc4  
Selected Recovery Point: 2009/1/9 13:34:13:735:263:3  
Acquiring read lock on /dev/sda9. attempt:1
```

Figure 376

### 12.3.3.13 Listcommonpoint

This command is used to identify a common recovery point across a group of replication pairs within the same target host. The syntax is

```
./cdpcli --listcommonpoint
```

```
[root@target bin]# ./cdpcli --listcommonpoint  
  
Common Recovery Point: 2008/12/22 23:37:26:385:835:0  
Corresponding I/O Sequence Point: 1259597  
[root@target bin]#
```

Figure 377

#### 12.3.3.14 lpattern

This option provides statistics of the I/O distribution of the target storage to check whether the target volume is bottleneck. This option reads the retention database and print out the I/O profile. The syntax is

```
./cdpcli --iopattern -vol=<path name>
```

```
[root@target bin]# ./cdpcli --iopattern --vol=/dev/mapper/s-vol3

Io Profile:
size      %Access %Read %Random Delay Burst Alignment Reply
1KB       25    100    100    0    1    sector none
1KB       25     0     0    0    1    sector none
1KB       25   100     0    0    1    sector none
1KB       25     0   100    0    1    sector none
[root@target bin]#
```

Figure 378

#### 12.3.3.15 Display Statistics

This command displays the statistics of a given target volume. The option displays the following

- Target volume name
- Differentials pending in CX
- Differentials pending in target
- Current RPO
- Apply rate
- Apply time

```
./cdpcli --displaystatistics --vol=<path name>
```

```
[root@target bin]# ./cdpcli --displaystatistics --vol=/dev/mapper/s-vol3

-----
##### REPLICATION STATISTICS #####
-----
Target Volume Name:      /dev/mapper/s-vol3
Diffs pending in CX:     47805390
Diffs pending in Target: 58314797
Current RPO (secs):      0
Apply rate (Bytes/sec):  140
Apply time (secs):       416534

Display Statistics
```

Figure 379



## 12.4 Generating reports through bwreport.pl

[Bandwidth Report](#)s as explained on page 183 may also be generated through the CLI tool on both windows and Linux CX servers. To generate bandwidth reports from CLI, access the respective CX server's command prompt and navigate to the CX installation path and issue the following commands

Table 31

Switch	Description
-Host	hostname to generate the report for
[-daily]	daily report
[-week]	weekly report
[-monthly]	report for the current month
[-yearly]	report for the current year
[-from=dd/mm/yyyy]	specify a start date range in dd/mm/yyyy
[-to=dd/mm/yyyy]	specify a end date range in dd/mm/yyyy
[-mode=daily]	specify the report mode (hourly, daily, monthly)
[-file=filename]	specify a filename to send the output to

### 12.4.1 Daily report

Given below is an example for a daily report

```
./bwreport.pl -host=<hostname> -daily
```

```
[root@localhost bin]# ./bwreport.pl -host=IMIT14.INMAGE.IN -daily
-----
BANDWIDTH REPORT [ 2009-01-17 to 2009-01-17 ]
-----
  DATE      HOUR      IN          OUT          MAX          SUM
-----
  01/17/2009  00      14.80 KB     23.41 KB     23.41 KB     38.21 KB
  01/17/2009  01      13.56 KB     21.15 KB     21.15 KB     34.71 KB
  01/17/2009  02      10.10 KB     16.42 KB     16.42 KB     26.52 KB
  01/17/2009  03          0.00 B          0.00 B          0.00 B          0.00 B
  01/17/2009  04          0.00 B          0.00 B          0.00 B          0.00 B
  01/17/2009  05          0.00 B          0.00 B          0.00 B          0.00 B
  01/17/2009  06       1.18 KB       1.76 KB       1.76 KB       2.93 KB
-----
  TOTAL              39.63 KB     62.72 KB     62.72 KB    102.35 KB
-----
[root@localhost bin]#
```

## 12.4.2 Weekly report

To generate a weekly report use the following command

```
./bwreport.pl -host=<hostname> -weekly
```

```
[root@localhost bin]# ./bwreport.pl -host=IMIT14.INMAGE.IN -weekly
```

BANDWIDTH REPORT [ 2009-01-11 to 2009-01-17 ]				
DATE	IN	OUT	MAX	SUM
01/11/2009	0.00 B	0.00 B	0.00 B	0.00 B
01/12/2009	0.00 B	0.00 B	0.00 B	0.00 B
01/13/2009	0.00 B	0.00 B	0.00 B	0.00 B
01/14/2009	0.00 B	0.00 B	0.00 B	0.00 B
01/15/2009	0.00 B	0.00 B	0.00 B	0.00 B
01/16/2009	16.77 KB	27.72 KB	27.72 KB	44.49 KB
01/17/2009	43.02 KB	67.98 KB	67.98 KB	110.99 KB
TOTAL	59.79 KB	95.70 KB	95.70 KB	155.48 KB

```
[root@localhost bin]#
```

## 12.4.3 Monthly and yearly report

Similarly to generate a monthly or yearly report use the `-monthly` or `-yearly` switch respectively e.g. commands

```
./bwreport.pl - host=<hostname> -monthly
./bwreport.pl -host=<hostname> -yearly
```

## 12.4.4 Custom report

To generate a custom report use the following command

```
./bwreport.pl -host=<hostname> -from=<define the starting date for the report> -to=<define the end date for the report>
```

## 12.5 ICAT

### 12.5.1 Introduction to ICAT

Icat is a command line tool used to archive content to multiple archival repositories at the same time. This utility is used in combination with an event based scheduled snapshot to periodically update content archival repositories with latest data.

This solution is deployed in three phases

**Configure:**

Ensure that the FX and VX agents on the production and DR servers are up and running.  
Configure the “ICAT.conf” file as described in the section Configure the icat.conf configuration file on page 6 file keeping in mind the path which the ICAT functionality reads from.

**Protect**

Set the VX replication pair from the production server to the DR server. Please refer to the Hitachi Dynamic Replicator Administration Guide for further details

**ICAT Functionality**

Schedule an event based snapshot, preferably a virtual snapshot for faster execution. Specify the ICAT utility in the post script of the schedule snapshot.  
Finally set the FX job for ICAT consistency, you should find two templates for ICAT consistency, one for windows platform and the other for Linux platforms.

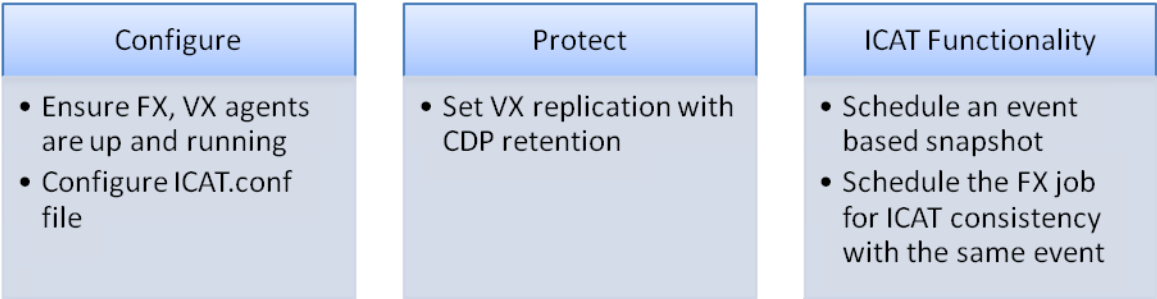


Figure 380

## 12.5.2 How ICAT works

When the FX job for ICAT consistency starts it issues a consistency tag on the production server. Once this consistency tag reaches the DR server, an event based snapshot is fired (ensure that you specify the same tag name in the FX consistency job and the scheduled snapshot). When the snapshot is complete, the ICAT command is executed.

The ICAT functionality reads from the snapshot and updates the archival repositories.

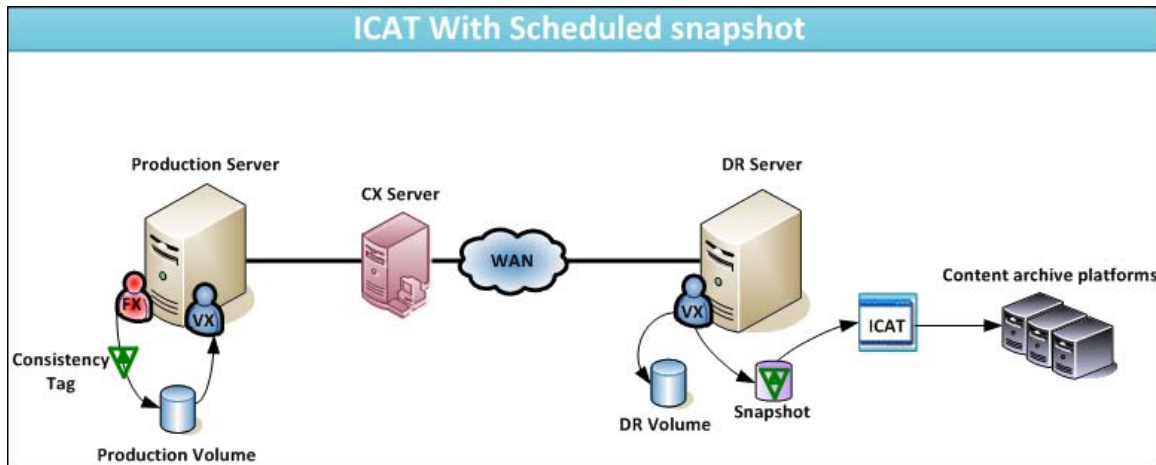


Figure 381

## 12.5.3 Before you begin

- Scout is deployed and actively protecting at least one volume on the production server
- Consistency tags are being issued on the production volumes at regular intervals through FX job.
- Ensure that no firewall is blocking any of the Scout components
- Determine on the volume on which a snapshot will be performed
- Determine on the name of the consistency tag

## 12.5.4Configure

Often icat may require several inputs as part of the command. Executing it through CLI at times may be time consuming and error prone. We recommend that you use the “icat.conf” configuration file as an input to the icat.exe. This will simplify the command to a large extent, an example is below

```
<FX installation path>icat.exe --config=icat.conf
```

### Configure the icat.conf configuration file

This configuration file is found under the FX agent installation folder and bundled with both the VX and FX agents. There are six major sections in this configuration file.

#### Remote Office:

“**Servername**” = Ensure that you enter the source server for the VX replication pair.

“**Branchname**” = Enter the name of the location where the source server is placed. A separate folder is created on the archival repository under which all the data will be copied to. For e.g.  
“/branchname/servername/volname/<folder with date and time>/<all data>”

“**Sourcevolume**” = Enter the volume that is being replicated

```
[remoteoffice]
servername = Production_server
branchname = Location_Alfa
sourcevolume = Volume
```

#### Archiverepository

“**Transport**” = By default http is used, you may also use cifs or nfs. Ensure that you enter only one of them.

“**Nodes**” = Enter the IP address of the archival repository. You may enter multiple node’s IP address separated by a comma. When the HTTP port of the node is different, ensure that you mention the port number as well. For e.g. nodes=10.0.164.21:564.

“**Dnsname**” = When the node IP is unknown, you may choose to enter the DNS name. You may separate multiple DNS names by a comma. By giving the domain name here, icat will replicate to all the archival repositories within the domain.

“**Rootdir**” = the archival repository contains two folders by default the data will be replicated to the fcfs\_data folder

```
[archiverepository]
Transport = http
Nodes = 10.0.164.21
Dnsname =
Rootdir = fcfs_data
```

## Tunables

“Retryonfailure” = the number of times icat will try when failed.

“Retrylimit” = the number of times icat will try before reporting a failure

“Retryinterval” = the number of seconds icat will wait before trying again.

“Exitonretryexpiry” = By default it is 0 thus skips the file and continues with the next one

“Maxconnects” = Number of threads to be spawned

“Tcprecvbuffer” = TCP receive window size for optimized bandwidth utilization

“Tcpsendbuffer” = TCP send window size for optimized bandwidth utilization

“Lowspeedlimit” = Number of changes expected per second

“Lowspeedtime” = Number of seconds waited before lowspeedlimitlimit is reached.

“Connectiontimeout” = Number of seconds of no data transfer before exiting.

```
[tunables]
retryonfailure=1
retrylimit=5
retryinterval=0
exitonretryexpiry=0
maxconnects=10
tcprecvbuffer = 36782
tcpsendbuffer = 78340
lowspeedlimit = 7489
lowspeedtime = 10
connectiontimeout = 180
```

## Content.source

There may be many content sources, you will need to copy the five lines as shown below, however ensure that you use different numbers for e.g. content.source2

- **“Directoryname”** = Folder on the DR server (target host of the VX replication) which will be replicated to the archival repository. Specify the directory path on the snapshot volume as the path here.
- **“Excludelist”** = the list of folders that are ignored by icat while replicating from the directoryname. You may separate multiple directories by a comma here.
- **“Filefilter”** = pattern=<file patterns> and/or date=<file last modified date> and/or size=<size of the file> for multiple patterns values should be separated by comma delimiter and, it allows = and != operators.  
date format should be YYYY-MM-DD and it allows =, >, <, >=, <= operators  
size should be given in terms bytes and it allows =, >, <, >=, <= operators

**“Include”** = Enter **“true”** to replicate all the files matching the filefilter set. You may enter **“false”** to ignore the filefilter.

```
[content.source1]
directoryname= H:\vsnap
excludelist=
filefilter=pattern=
include= true
```

## Config

- **“Fromlastrun”** = Enter the value as 0 when you need to perform incremental backup from the last time ICAT has run. Enter 1 to replicate all files.
- **“Forcerun”** = When the value is set to 1 ICAT resumes archival process even though resume information is not found.
- **“Overwrite”** = To overwrite existing files enter the value as 0 and 1 to ignore existing files.
- **“Autogendestdirectory”** = Enter the value as 0 to generate the folder name as specified under the [remote office]. Enter 1 to use an alternative folder name in the place of “branchname/servername/volname”
- **“Targetdirectory”** = This is used when **“Autogendestdirectory”** is set to 1. Enter the name of the folder to be used in the place of “branchname/servername/volname”.
- **“Logfilepath”** = All log files are created under this folder on the host where icat runs.
- **“Loglevel”** = 0 to 7

```
[config]
fromlastrun = 0
forcerun = 1
overwrite = 0
autogendestdirectory = 1
targetdirectory = mystestfil21
logfilepath = C:\testicat\
loglevel = 7
```

Table 32: Types of Loglevel

Debug Log Level	Severity of Messages logged into the log file
0	Logging is disabled
1	Only FATAL messages are logged
2	Only FATAL and severe messages are logged
3	Only FATAL, severe and error messages are logged
4	Only FATAL, severe, error and warning messages are logged
5	Only FATAL, severe, error, warning & info messages are logged
6	Only FATAL, severe, error, warning, info & debug messages are logged
7	All messages will be logged



### Caution:

Refrain from setting the loglevel as 7 as this may consume system resources thus degrading the machine/services performance.



## Delete

- **“Maxlifetime”** = The list of files icat attempted to copy irrespective of the result (success or failure) are maintained in its database for the specified number of days. Once the number of days is crossed all entries of the files are deleted from the database.
- **“Maxcopies”** = A file cannot occur for more than the specified number of times in the icat database

```
[delete]
maxlifetime = 30
maxcopies = 10
```

## Using through command line

You may also run the ICAT utility through command line. Access the DR server's command prompt and navigate to the FX agent installation folder to find the icat.exe. The syntax as show below

```
icat.exe {--nodes [ip:port] + | --dnsname }[--transport] [--uid] [--gid] {--
autogendestdirectory=1 --branchname --servername --sourcevolume | --
autogendestdirectory=0 --targetdirectory | --targetdirectory | --branchname -
--servername --sourcevolume} {--directory [--excludelist] [--include] --
namepattern --comp --op=and/or --date --comp --op=and/or --size}+ [--
overwrite=1][--retryonfailure][--maxretries] [--retryinterval] [--
exitonretryexpiry=1] [--logpath][--fromlastrun][--forcerun] [--sendtcpbuffer]
[--recvtcpbuffer][--connectiontimeout][--lowspeedlimit][--lowspeedtime][--
connects][--lifetime][--copies][--rootdir][--loglevel]
```

Table 33

Mandatory switches	Dependent switches	Optional switches
Nodes or dns name		Transport, UID, GID
--autogendestdirectory=1,	--branchname --servername --sourcevolume	
--autogendestdirectory=0	--targetdirectory	[--branchname --servername --sourcevolume]
--directory		--excludelist, --include, --namepattern, --comp, --op=and/or --date --comp --op=and/or --size
--resume (used only to resume a failed ICAT)		
--config (used only for passing config file)		
--rootdir		

## 12.5.5 Resume mode

A resume mode is used when the normal ICAT operation fails or to perform incremental backup to the same folder on the archival repository. Adding the “**--resume**” switch to the ICAT command will:-

- Continue the operation from the point where the normal mode failed.
- Perform incremental backup when the normal mode was successful

For e.g. *<FX agent install path>***icat.exe --resume --config=** *<full path of the Icat.conf file>*



### Notes:

Additionally FX template “ICAT consistency for Windows” and “ICAT consistency for Linux” may be used to issue consistency tags on the production volumes which in turn fires a snapshot and ultimately the ICAT utility.

You will need to use the ICAT.exe as part of the snapshot post script to achieve certain level of automation.

Please refer to the ICAT usage document for detailed steps

## 12.6 Application Support

### 12.6.1 Application.exe

“**Application.exe**” is used for discovery, verifying consistency, failover and failback of applications, while “**dns.exe**” is used for DNS server failover and failback.

**Table 34: application.exe options**

Flags	Description
-discover	This switch uses -app <application name> and the result will display the list of volumes and other application specific information.
-verifyconsistency	This switch will also use -s <source name> -t <target name> -tag <tag name>.
-failover	This switch will use -planned/unplanned -s <source name> -t <target name> -tag. For a failback the same command is used with an addition of -ip switch -ip <actual source ip address>.
-apptag	Using this switch along with failover command will issue a consistency tag before performing a failover.
-retainDuplicateSPN	If this option is specified then after failover/failback using application.exe/exfailover.exe duplicate SPN entries may be observed in Active Directory which may affect Kerberos Password Authentication between Outlook & Exchange servers after failover/failback
-app	Name of the application.
-planned	To perform a planned failover
-unplanned	To perform an unplanned failover
-virtualserver	Used to perform a failover of source SQL/Exchange virtual server to a standalone server.
-mta	Exchange virtual server's name hosting exchange MTA resource. This has to be specified when the source exchange cluster has multiple exchange virtual servers.
-IP	This switch is used while performing a failback to restore the original IP address of the source server.
-tag	This switch is used on the source server to generate a tag. The same switch is used on the target host to recover to the specified tag.
-s	Name of the source host.
-t	Name of the target host

### 12.6.1.1 SQL failover example

The following picture shows how SQL server is fail over is performed. Using a switch of `-apptag` will issue another tag on the source volume, but the failover will be performed to the `-tag` switch.

Syntax

**Application** `-failover -planned -app sql2005 -s <source name> -t <target name> -tag none`

```
C:\Program Files\InMage Systems>Application.exe -failover -planned -app sql2005
-s cluster1 -t imits070 -builtin -tag NONE
cluster1 IP address = 192.168.71.30
imits070 IP address = 10.0.1.70
```

Figure 382:

The output of the command should be executed on the target host to complete the failover.

```
*****IMPORTANT INFORMATION*****
Please run the following command on the Target Host to complete failover
Application.exe -failover -planned -app sql2005 -s cluster1 -t imits070 -builtin
-tag FileSystem46d7b0ba
*****
C:\Program Files\InMage Systems>
```

Figure 383:

For failback an extra switch of `-ip <actual source IP address>` is added

### 12.6.1.2 Discover syntax

Syntax:

**Application.exe** `-discover -app <application name>`

### 12.6.1.3 Verifyconsistency syntax

Syntax:

**Application.exe** `-verifyconsistency -app <application name> -s<source name> -t <target name>`

## 12.6.2DNS.exe

Table 35: dns.exe options

Flags	Description
-failover	This will performs a failover, other keys used within this are -s, -host,-ip, -dnsserverip -dnsdomain -user -password. While some are optional others are required, this will depend on the setup in place.
-failback	Performs a failback and requires the same switches as that of failover
-s	Source host name
-t	Target host name
-ip	New ip address that is to be assigned
-host	Source host name
-dnsserverip	Dns server IP address (used if there are more than one DNS servers)
-dnsdomain	Domain of the DNS server
-user	User name
-password	password

Dns.exe performs DNS failover. The syntax is given below

**Dns -failback -host <hostname> -ip <new ip address> -<DNS server IP>**

```
C:\Program Files\InMage Systems>dns -failback -host bes1 -ip 10.0.100.100
Connecting to CX to determine if host : bes1 is part of a cluster...
Using DNS name : bes1 For DNS failback
bes1 FQDN = bes1.careermakers.in
***** DNS record for bes1 now modified to point to 10.0.100.100 successfully
DNS cache flushed successfully.
```

Figure 384:

## 12.6.3 Winop.exe

Winop.exe is a support tool used by advanced users or support team. The following table shows major switches and their associated functions.

**Table 36**

Switch	Functionality
AD	To perform AD replication
SPN	Host SPN, exchange specific SPN entries can be added/removed
NETBIOS	To add/delete/change the netbios name
Security	Used as part of the File server solution when the FX agent is run with local user privileges. By using this switch, you may store domain credentials in the registry to use it at a later time.
Cluster	This is used on clustered environment to find out the active node.
Mapdriveletter	This is used for the ESX solution to get a list of the volumes and their respective GUIDs.

### AD: To perform AD replication

**Table 37**

Switch name	Importance	Description
Replicate	Mandatory	This switch is used in combination with the optional switches to replicate AD.
UpdateAllDnsServers	Mandatory	This switch is used to update all the DNS servers within the domain
UpdateDnsServer	Mandatory	This switch is used to update a specific DNS server
Dc	Optional	The "DC" switch is used to specify the attribute of the AD
ZoneUpdateFromDs	Optional	This switch determines the DNS server with the updated information after the application failover
Domain	Optional	You may execute the winop.exe command through a different user account by specifying that user's credentials through the domain, user and password switches.  Use the -domain switch to pass the domain name. Ensure that you use the user and password switch as well.
User	Optional	This switch is used to pass the user name
password	Optional	This switch is used to pass the password.

## Examples for AD only update

To replicate updated AD information to all the ADs within the domain

### Example

```
WinOp.exe AD -replicate DEFAULT
```

To replicate a defined attribute or a specific domain controller to all ADs in the domain use the following command

### Example

```
WinOp.exe AD -replicate "DC=Schema,DC=Configuration,DC=mydomain,DC=COM"
```

To replicate AD by specifying the source and target DCs

### Example

```
WinOp.exe AD -replicate "DC=Schema,DC=Configuration,DC=mydomain,DC=COM" -dc SALES_DC
```

When you do not have administrator privileges you may pass the domain administrator credentials through the command line

### Example

```
WinOp.exe AD -replicate "DC=mydomain,DC=COM" -domain mydomain.com -user administrator -password mycred
```

To replicate specific attribute to all ADs

```
WinOp.exe AD -replicate DEFAULT; "DC=Schema, DC=Configuration, DC=MyDomain, DC=COM"
```

To replicate specific attribute to specific domain controller

```
WinOp.exe AD -replicate "DC=mydomain, DC=COM"; "DC=Schema, DC=Configuration, DC=mydomain, DC=COM" -dc SALES_DC
```

## Examples for DNS only update

DNS servers are updated after the application failover. The closest DNS server is first updated and then the rest of the DNS servers are updated using the following command

```
WinOp.exe AD -UpdateAllDnsServers
```

When you know the DNS server that is updated after the application failover, use the following command to update the rest of the DNS servers.

```
WinOp.exe AD -UpdateAllDnsServers -ZoneUpdateFromDs Inmange.net;DevZone
```

To update a given set of DNS servers, use the following command

```
WinOp.exe AD -UpdateDnsServer DC1,DC2,DC3
```

To update a set of DNS servers by a specific DNS server use the following command

```
WinOp.exe AD -UpdateDnsServer DC1,DC2,DC3 -ZoneUpdatefromDs Inmange.net;DevZone
```

## SPN: To add/ remove Exchange SPN entries

Table 38: winop.exe SPN flags

Flags	Description
Add	This switch is used to add Exchange specific EPN entries to the desired host
Delete	This switch is used to delete Exchange specific SPN entries from the desired host
Addhost	Used to add a host to the active directory
Deletehost	Used to delete the host from the active directory

The –add switch is used to add exchange SPN entries to the specified host. Given below is the syntax

**Winop.exe –add <host name>**

```
C:\Program Files\InMage Systems>WinOp.exe spn -add imits035
The search result size for filter: (&(objectClass=Computer)(CN=imits035)) is 1
*****Processing Computer DN CN=IMITS035,CN=Computers,DC=bit32,DC=org
exchangeMDB/imits035
exchangeMDB/imits035.bit32.org
exchangeRFR/imits035
exchangeRFR/imits035.bit32.org
SMTPSVC/imits035
SMTPSVC/imits035.bit32.org
HTTP/imits035.bit32.org
Successfully set servicePrincipalName attribute in computer account: imits035
C:\Program Files\InMage Systems>
```

Figure 385

The –delete switch is used to delete exchange specific SPN entries from the desired host. Given below is the syntax

**Winop.exe –spn –delete <host name>**

```
C:\Program Files\InMage Systems>WinOp.exe spn -delete imits035
The search result size for filter: (&(objectClass=Computer)(CN=imits035)) is 1
*****Processing Computer DN CN=IMITS035,CN=Computers,DC=bit32,DC=org
exchangeMDB/imits035
exchangeMDB/imits035.bit32.org
exchangeRFR/imits035
exchangeRFR/imits035.bit32.org
SMTPSVC/imits035
SMTPSVC/imits035.bit32.org
HTTP/imits035.bit32.org
Successfully set servicePrincipalName attribute in computer account: imits035
```

Figure 386



The `-addhost` is used to add a host entry to the AD. Given below is the syntax

Syntax

`Winop spn -addhost <name of the host>`

```
C:\Program Files\InMage Systems>winop spn -addhost imits035
The search result size for filter: (&(objectClass=Computer)(CN=imits035)) is 1
****Processing Computer DN CN=IMITS035,CN=Computers,DC=bit32,DC=org
HOST/imits035
HOST/imits035.bit32.org
Successfully set servicePrincipalName attribute in computer account: imits035
```

Figure 387

The `-deletehost` is used to delete a host entry from the AD. Given below is the syntax

Syntax

`Winop spn -deletehost <name of the host>`

```
C:\Program Files\InMage Systems>winop spn -deletehost imits035
The search result size for filter: (&(objectClass=Computer)(CN=imits035)) is 1
****Processing Computer DN CN=IMITS035,CN=Computers,DC=bit32,DC=org
HOST/imits035
HOST/imits035.bit32.org
Successfully set servicePrincipalName attribute in computer account: imits035
```

Figure 388

## Netbios: To add/remove/change netbios name of the host

Winop netbios flags

Table 39: winop.exe netbios flags

Flags	Description
<b>-remote</b>	Is used to specify a remote machine (default being local machine)
<b>-add</b>	Used to add netbios name to the netbios name table
<b>-delete</b>	Used to remove netbios name from the netbios name table
<b>-changeto</b>	Used to change the name of the computer (requires a reboot)
<b>-domain</b>	Optional. Will require <b>-user</b> and <b>-password</b> along with it. Used to explicitly specify the name of the domain or server containing the user account
<b>-user</b>	Optional. Requires <b>-domain</b> and <b>-password</b> along with it. Used to mention the name of the user
<b>-password</b>	Optional. Requires <b>-domain</b> and <b>-password</b>

Syntax:

**Winop.exe NETBIOS -remote <Computer Name> -add <Name> -delete <Name> -changeto <Name> -domain <DomainName> -user <user name> -password <password>**

You may use the “**-add**” and “**-delete**” switches are used to add and delete the netbios name of the host respectively.

```
C:\Program Files\InMage Systems>WinOp.exe netbios -delete test2
Successfully deleted test2 from NetBios name table
C:\Program Files\InMage Systems>WinOp.exe netbios -add test2
Successfully added test2 to NetBios name table.
NOTE: Add operation succeeds if the server name specified
is added to at least one transport.
```

Figure 389

You may also change the netbios name using the “**-changeto**” switch as shown in the picture below.

```
C:\Program Files\InMage Systems>WinOp.exe netbios -changeto test2
Successfully changed computer name from  to test2
C:\Program Files\InMage Systems>hostname
IMITS035
C:\Program Files\InMage Systems>_
```

Figure 390

## Security: To store domain administrator credentials

The FX agent needs to be running with domain administrator privileges to perform a failover or failback. When the FX agent does not have domain administrator privileges you will need to store a domain administrator's credentials using the following command

**Winop security -encrypt**

```
C:\Program Files\InMage Systems>winop security -encrypt
Enter Domain name:inmage
Enter Domain User:test
Enter Password:
Reenter Password:
C:\Program Files\InMage Systems>
```

Figure 391

You will need to enter the domain name, domain administrator name and password. This will encrypt and store the credentials in the windows registry.

## Cluster: To perform identify the active node

The cluster switch is used to identify the active node on a cluster environment. The example is shown below:

**Winop -cluster checkactivenode <virtual server name>**

## Mapdriveletter: Used for ESX solution

The “mapdriveletter” is used to collect and store drive information for restoring it on the target ESX server

```
C:\Program Files\InMage Systems>winop mapdriveletter -s
Persisting Drive Letter Information to the ClusUtil.log File during -s operation:
-----
Volume Guid:  \\?\Volume{d5f71e74-404e-11de-9796-005056c00008}\      Volume Mount Point:  E:\
Volume Guid:  \\?\Volume{46c69e95-150c-11de-b63a-005056c00008}\      Volume Mount Point:  J:\
Volume Guid:  \\?\Volume{e655667b-04ab-11de-9c41-005056c00008}\      Volume Mount Point:  M:\
Volume Guid:  \\?\Volume{f12b432f-04c5-11de-9cfc-005056c00008}\      Volume Mount Point:  N:\
Successfully persisted all the Drive Letters information.
C:\Program Files\InMage Systems>winop mapdriveletter -t
Mapping Correct Drive Letters Information to the ClusUtil.log File during -t operation:
-----
WARNING: If any new volumes are added and if they clash with the
original cofiguration of drive letters,then those volumes will be Un-Mounted!

No Drive Letters are changed from the previous system's Drive Letter Configuration.
```


Figure 392

Using the “-t” switch will restore drive mapping on the target guest (on the ESX server)

## 12.6.4 Tweaking FailoverServices.conf

This file is located under the “**consistency**” folder (under VX installation path). It contains three sections Exchange, SQL2000 and SQL2005. Use this file to list all services that are dependent on exchange, SQL 2000 and SQL 2005. Doing this will enable the Scout to stop and start the dependent services before doing the same for the application service. If dependent services (if they exist) are not listed then the dependent services will not let the application service stop.

If there are any services dependent on exchange then mention them under the “[**exchange**]” section, under start and stop keys. This ensures that these services are stopped and started while performing failover.



**Figure 393:**

FX job can be used to maintain the same file on the target to preserve the list of dependencies post failover.

# Part 8: Performance Tuning

Chapter 6, 7, 8, and 9 are covered in this section.

## **Chapter 6**

Tuning the FX agent for faster execution

## **Chapter 7**

This chapter describes various options that can be fine tuned on the source and target side VX agents for better performance.

## **Chapter 8**

Tuning the CX server for better performance

## **Chapter 9**

This chapter describes creating bandwidth policies specific to the CX server.

## 13 Tuning FX

### 13.1 Tuning the FX for Faster Execution

FX replication process depends on various factors such as number of files, file sizes, available bandwidth etc. One could use different FX options to speed FX replication process, however the trade-offs need to be well understood before making these changes.

FX replication performance can take a hit when there are large numbers of files [typically around > 100 thousand]. During the FX replication process, the agent needs to determine the changed file list [w.r.t. target], this operation can tend to be costly when the number of files increases. The motivation of this section is to explain various options and trade-offs to further speed up the replication

#### 13.1.1 Options for Faster Execution

##### 113. Tuning File-listing options.

This approach focus on controlling the way file-list is generated. There are two different mechanism used to determine changed files.

- **“Checksum approach”**: checksum of each file on source and comparing them with corresponding target checksum
- **“File Timestamp and Size”**: Timestamp and File size (faster than checksum)



##### Notes:

Since checksum is a CPU consuming task, other options such as “file size” and “timestamps” can be used to further boost up the FX execution

##### 114. Additional replication options

- Reducing activity logging operations

##### 115. Setting TCP send and receive buffer size

You may choose to tweak two dword type registry keys “**TcpSendBufSize**” and “**TcpRecvBufSize**” under “**HKLM/SOFTWARE/SV Systems/FileReplicationAgent**” to set the send and receive buffer values of TCP for the FX agent. For Linux agents you may edit the “**config.ini**” file under the FX installation folder to find the keys as “**#TcpSendBufSize = 0**” and “**#TcpRecvBufSize = 0**”. Remove the “**#**”, set the desired values and save the file. You will need to restart the FX agent service after this.

### 13.1.1.1 Performance Tuning

The following are the different options that can be looked at which can have an impact on the performance, while configuring a job.

- **“Checksum”**: If this option is selected, FX performs checksum of all the files using “128-bit MD4 checksum” algorithm before transferring. It also performs the checksum on the target to determine whether the files are in sync or not. Calculating checksum is CPU intensive and requires I/O to be performed. For faster FX execution, disable the option **“Always perform checksum”**

File/Directory Options	
<input checked="" type="radio"/>	Copy the source directory to a subdirectory of the target directory
<input checked="" type="radio"/>	Copy the contents of the source directory directly into the target directory
<input type="checkbox"/>	Always perform checksum
	Checksum block size: 8192
<input type="checkbox"/>	Whole files (no incremental checks)

Figure 394:

- **“size-only”**: Choosing the option **“Ignore files with same size”** will force FX to ignore same sized files across source and target folders. The other option **“Ignore files with same size and timestamp at destination”** will ignore files with the same size and timestamp across source and target directories.

Inclusion/Exclusion Options	
<input checked="" type="checkbox"/>	Update only (Do not overwrite newer files)
<input type="checkbox"/>	Only update files that already exist at the destination
<input type="checkbox"/>	Ignore files that already exist at the destination
<input checked="" type="checkbox"/>	Ignore files with same size and timestamp at destination
<input type="checkbox"/>	Ignore files with same size
<input type="checkbox"/>	Exclude files matching pattern: Separated by ;
<input type="checkbox"/>	Include subset of exclude list matching pattern: Separated by ;

Figure 395:

- **“Verbose”**: This option will determine the level of FX logging> Selecting **“Very low verbosity”** will only log critical activity leading to lesser IO>

Feedback Options	
<input type="radio"/>	Very high verbosity
<input type="radio"/>	High verbosity
<input type="radio"/>	Medium verbosity
<input type="radio"/>	Low verbosity
<input checked="" type="radio"/>	Very low verbosity
<input type="radio"/>	Quiet
<input checked="" type="checkbox"/>	Show file transfer statistics (Necessary for trending information)

Figure 396:

## 13.2 Auto Delete FX Logs on Source / target

Every time an FX job executes it is considered as a separate instance and a log file is created under the FX agent installation folder. There is a mechanism to delete these log files after a period of seven days. To enable automatic FX log file deletion:

On windows: Open the registry editor then navigate to "**HKEY\_LOCAL\_MACHINE\Software\SV Systems\FileReplicationAgent**" then set the registry value of "**EnableFrLogFileDeletion**" to 3255 (decimal) on both the source and target FX hosts

On Linux platforms: Open the "**config.ini**" file under the FX installation folder (default path being "**/usr/local/Inmage/Fx**") and set the value as 3255 on both the source and target FX hosts.



## 14 Tuning VX

### 14.1 Source side Driver Tweaks

#### Driver:

Driver is a low level component that acts as a bridge between the hardware and the Agent service.

#### File writer:

A logical unit within the driver that reads from the memory, sorts data and directs to threads based on pre-defined logic to write the data to disk.

File writers are specific to a particular disk write and additionally there can also be common file writers

#### Threads:

Threads are sub components of file writers; they acquire data from file writer and write it to disk as chunks of predefined size (4mb by default)

#### 14.1.1 Working Process

When the initial copy is complete and the replication pair enters into “**differential sync**” mode, the driver takes over by intercepting writes to the disk (Source side).

This is a three step process

- First the driver transfers the writes to the memory then
- The writes are sent to File Writers, each file writer has two threads if it is a uni processor system and four threads if it is a dual processor system.
- Threads create files of 4mb each, which are later picked up by the agent service.

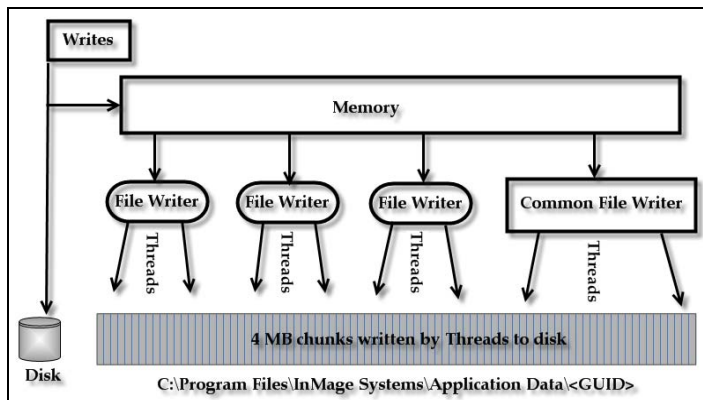


Figure 397

The file writer accumulates and assigns a bunch of writes to its threads. This procedure can be tweaked to enhance performance

## 14.1.2 Registry Values

The registry values are stored under

“HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet001\Services\involut\Parameters”

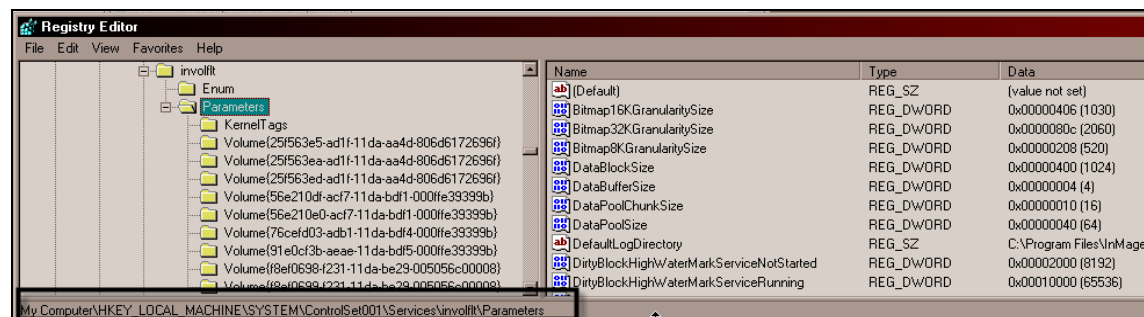


Figure 398

Volume filtering is performed in two modes

116. **“Data mode”**: All the disk writes are tracked into the main memory in the form of metadata (Offset and Length of data) along with data itself.

117. **“Meta data mode”**: This returns the location and length of data written on the disk (metadata) to agent service. In case of low physical memory, the driver enters into bitmap mode, these saved changes are read back when there is free memory available. (Frequent disk access leads to slower performance). Certain driver thresholds determine the mode in which the driver operates

**VolumeDataFiltering**: This enables disk writes to be tracked in Data Mode and once the memory usage threshold is reached, based on whether file based tracking is enabled it will write data to Disk files. If file based tracking is not enabled then it is automatically shifted to Meta data mode.

**VolumeDataFilteringForNewVolumes**: The driver will also monitor newly created volumes. By default this is enabled (value 1). By default this monitors in data mode and once the threshold is reached it switches to Meta data mode. To disable it change the value to 0.

**NumberOfThreadsPerFileWriter**: The number of threads can be increased, two for uni processor and four for dual processor. Default value being 0.

**NumberOfCommonFileWriters**: This is a user created special file writer for one or more specific replication pairs, Default value being 1

**VolumeFileWriterId**: User specified ID for a file writer. This is used if a dedicated writer with the specified id is to be associated with a specific volume or group of volumes.

**FreeThresholdForFileWrite**: File writer would start queuing data to disk when the amount of total space left is less than or equal to this value

**VolumeThresholdForFileWrite:** File writer would start queuing data to disk when the amount of space used per volume is more than this value

**VolumeDataLogDirectory:** This value is created under another key (guid). Specifies where the files have to be written. By default the files are written under <installation path>/Application data/<guid>. As the guid changes for each replication so does the physical location.

**VolumeOutOfSyncTimestampInGMT:** Specifies the volume out of sync time stamp in seconds (GMT)

**VolumeFileWriterThreadPriority:** This has to be configured per replication pair. This value is common for all writers and specifies the thread priority for file writer threads. Individual writers for specific volume can be set to have different priority. Value can range from 0 to 19 (0 being the lowest and 19 being the highest). More than one replication pair can have the same thread priority

**VolumeDataToDiskLimitInMB:** By default threads create 4 MB files (each containing writes). Collective space occupied by all these files can be configured through this. Default value is 512MB. Although the agent service will continuously transfer them out of the source a limit is set so that free space is ensured on the source. This is configured per replication pair.

**VolumeDataFiles:** Disabling this (changing to 0) will not create data files (4mb files) on the disk, rather they are created in the memory and once the limit is reached it is switched into meta data mode.

## 14.2 For production environments with aggressive data writes

You may choose to tweak the driver's global registry settings under  
"HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\involflt\Parameters"

**DataPoolSize:** Memory allocated from paged pool that is used to store the change data coming to the driver

Default value: 64 MB for 32 bit and 256 MB for 64 bit platform

**MaxNonPagedPoolInMB:** This value is a portion of memory allocated from non paged pool.

Default value: 12 MB for 32 bit and 32 MB for 64 bit platform

**VolumeDataSizeLimit:** This is the maximum limit that is used to determine the data pool size which may be used by any particular volume

Default value: 32 MB for both 32 and 64 bit platforms

## 14.3 Tunable for Solaris

“inm\_dmit” tool is used to modify and retrieve common and volume specific attribute values. Inm\_dmit file is located in the VX installation path -- i.e., `/InMage/Vx/bin` (This path can be different for different users depending on VX installation path)

To get a help on how to use inm\_dmit tool, navigate to the VX installation path and execute the command “inm\_dmit”.

### 14.3.1 Common Tunable

To retrieve common attributes’ value, execute the following command in the vx installation path:

```
inm_dmit --get_attr <attribute name>
```

Example:

```
inm_dmit --get_attr DataPoolSize
```

To modify common attributes’ value, execute the following command in the vx installation path:

```
inm_dmit --set_attr <attribute name> <attribute value>
```

Example:

```
inm_dmit --set_attr DataPoolSize 128
```

The list of the common attributes which can be modified with inm\_tool is as follows:-

- DataPoolSize
- DefaultLogDirectory
- FreeThresholdForFileWrite
- VolumeThresholdForFileWrite
- DirtyBlockHighWaterMarkServiceNotStarte
- DirtyBlockLowWaterMarkServiceRunning
- DirtyBlockHighWaterMarkServiceRunning
- DirtyBlockHighWaterMarkServiceShutdown
- DirtyBlocksToPurgeWhenHighWaterMarkIsReached
- MaximumBitmapBufferMemory
- Bitmap512KGranularitySize
- VolumeDataFiltering
- VolumeDataFilteringForNewVolumes
- VolumeDataFiles
- VolumeDataFilesForNewVolumes
- VolumeDataToDiskLimitInMB
- VolumeDataNotifyLimit
- SequenceNumber
- MaxDataSizeForDataModeDirtyBlock
- VolumeResDataPoolSize

- MaxDataPoolSize



#### Notes:

Don't specify MB while setting DataPoolSize, that means just pass 128 not 128MB.

### 14.3.2 Volume Specific Tunable

To retrieve volume specific attributes' value, execute the following command in the vx installation path:

```
inm_dmit --get_attr volume <volume name> <attribute name>
```

Example:

```
inm_dmit --get_attr volume /dev/dsk/clt9d0s2 BitmapOnProtectedVolume
```

To modify volume specific attributes' value, execute the following command in the vx installation path:

```
inm_dmit --set_attr volume <volume name> <attribute name> <attribute value>
```

Example:

```
inm_dmit --set_attr volume /dev/dsk/clt9d0s2 BitmapOnProtectedVolume 1
```

The list of the volume specific attributes which can be modified with inm\_tool is as follows:-

- VolumeFilteringDisabled
- VolumeBitmapReadDisabled
- VolumeBitmapWriteDisabled
- VolumeDataFiltering
- VolumeDataFiles
- VolumeDataToDiskLimitInMB
- VolumeDataNotifyLimitInKB
- VolumeDataLogDirectory
- VolumeBitmapGranularity
- VolumeResyncRequired
- VolumeOutOfSyncErrorCode
- VolumeOutOfSyncErrorStatus
- VolumeOutOfSyncCount
- VolumeOutOfSyncTimestamp
- VolumeOutOfSyncErrorDescription
- BitmapOnProtectedVolume
- VolumeFilterDevType
- VolumeNblks
- VolumeBsize
- VolumeResDataPoolSize

## 14.4 Tuning target side agent

The following are the registry values that can be tweaked on the target side during differential sync to enhance performance. All the values are located under “**HKEY\_LOCAL\_MACHINE\SOFTWARE\SV Systems\VxAgent**”. These values can be edited manually through windows registry editor (regedit.exe)

### **MaxFastSyncApplyThreads**

Functionality:

This parameter is set on target to increase the number of threads used while performing initial and Resync on the target. More threads result in faster resync process. However increase this value only if the host has more than 4 CPUs. General guide line is to make the value equal to the number of CPUs. In case of single CPU machines the default is 4.

**Default value:** 4

### **MaxRunsPerInvocation**

**Functionality:** Target side VX Agent makes few runs before it exits. In each run it

- Acquires the file list from the CX server
- Applies them to the target
- Execute any ready snapshot requests

The value here determines how many runs will be performed before the exit

**Type:** Dword

**Default value:** 10

**Other values:** 1 to any value (0 being infinite). If you set this value to zero, target site VX agent would run forever until there is a replication configuration change or svagent service shutdown.

Larger value lets the agent complete its task but it runs for a longer time.

### **MaxMemoryUsagePerReplication**

**Functionality:** This is used to set a limit on memory that is used by a replication pair for storing differential files until they get applied to the target volume.

**Type:** Dword

**Default value:** 32 MB

**Min value:** 2 \* maximum of (MaxInMemoryCompressedFileSize and MaxInMemoryUnCompressedFileSize)

### **MaxInMemoryCompressedFileSize:**

**Functionality:** This determines the file size limit of a compressed file that should be downloaded to memory. If a compressed differential file size on CX is larger than this value, it will be downloaded to the application cache directory

**Type:** Dword

**Default value:** 4MB

**Other values:** If 0 (is also the min value) is set then all the compressed files from CX would be downloaded to application cache directory

## **MaxInMemoryUnCompressedFileSize**

**Functionality:** This determines the file size limit of an uncompressed file that should be downloaded to memory. If an uncompressed differential file size on CX is larger than this value, it will be downloaded to application cache directory.

**Functionality**

**Type:** Dword

**Default value:** 8 MB

**Other values:** If 0 (is also the min value) is set then all the uncompressed files from CX will be downloaded to application cache directory

## **CompressionChunkSize**

**Functionality:** All the differential files are uncompressed in chunks on the target. This registry value determines the uncompressed chunk size.

**Type:** Dword

**Default value:** 1 MB

## **CompressionBufSize**

**Functionality:** This registry value determines the size of buffer allocations for storing uncompressed output. We allocate memory in multiples of this registry value. In general this should be set to a value such that we do not have to do realloc in most cases.

**Type:** Dword

**Default value:** 4 MB

## **SequenceCount and SequenceCountInMsecs**

**Functionality:**

These two registry values are used in case of volume groups. It will prevent any one volume in the volume group from getting too far ahead as well as prevent any one volume from monopolizing the cache.

“**SequenceCount**” determines the no. of files that can be downloaded per volume without waiting for other volumes in the volume group. On reaching the “**SequenceCount**”, file download will be stopped and resumed only when download file count on all the volumes have reached the “**SequenceCount**”.

“**SequenceCountInMsecs**” determines the amount of data in milliseconds granularity that can be downloaded per volume without waiting for others in the volume group.

**Type:** Dword

**Default value for SequenceCount:** 8

Default value for **SequenceCountInMsecs**: **SequenceCountInMsecs** = 60000 i.e. 60 seconds

The **SequenceCount** determines the delay between replication pairs in a volume group; larger number indicates greater distance between them.

## **FastSyncHashCompareDataSize**

**Functionality:**

This value determines the amount of communication data sent from the source to the CX server, the value is set to 1024 KB. The same should be on the target side, else the resync will pause.

**Type:** Dword

**Default value:** 512 kb for older agents, 1024 for versions 3.5.2 and above

The following are the configuration values that may be tweaked on the target side during differential sync to enhance performance.

You will need to restart the agent service restart after any changes

All the values are located in

Linux: <InMage Vx Install Path>/etc/drscout.conf

Windows: <InMage Vx Install Path>/Application Data/etc/drscout.conf

#### **CacheDirectory:**

##### **Functionality:**

This is the intermediate directory for downloading differentials from cx to target before applying it to the target volume

**Type:** string

##### **Default value:**

Linux: /home/svsystems/var

Windows: <InMage Vx InstallPath>/Application Data

**Recommended Value:** Set this on a separate disk/partition other than system drive. This directory also needs to be excluded from virus scanners to prevent replication from slowing down. Use junctions/soft links to allow using separate cache directory per pair.

#### **MaxDiskUsagePerReplication:**

**Functionality:** This determines the max disk space that can be consumed per replication pair by the cache directory. The cumulative disk space allocated for all the replication pairs should be less than the available space.

**Type:** integer

**Units:** bytes

**Default value:** 512MB

**Recommended value:** Set this equal to (or greater than) peak data change so that all changes can be downloaded to the target server and minimize data loss in case of a disaster.

#### **MaxMemoryUsagePerReplication**

**Functionality:** This is used to set maximum memory that can be used by a replication pair for storing differential files until they are written to the cache directory. Once the memory utilization crosses this threshold, differentials will be downloaded directly to application cache directory till the memory usage for the pair again drops back.

**Type:** integer

**Units:** bytes

**Default value:** 64 MB

**Recommended value:** Increase this value to get better wan bandwidth utilization. Note, the combined memory usage for all the replication should be below the available physical memory leaving at least 1 gb memory for rest of processes. In case of windows, the combined memory usage for all the replication pairs should be below 3 Gb.

Min value: 2 \* maximum of (MaxInMemoryCompressedFileSize and MaxInMemoryUnCompressedFileSize)



**NWThreadsPerReplication:**

**Functionality:** This determines the no. of parallel downloads from CX server to the target per replication pair.

**Type:** integer

**Units:** none

**Default value:** 2

**Recommended value:** Increase this if the network bandwidth utilization is low. See IOThreadsPerReplication before changing this parameter.

**IOThreadsPerReplication:**

**Functionality:** This determines the no. of parallel writes from memory to the application cache per replication pair.

**Type:** integer

**Units:** none

**Default value:** 2

**Recommended value:** Increase this if the changes are getting accumulated into memory and writes are becoming bottleneck.

**MaxRunsPerInvocation:****Functionality:**

Target side VX Agent makes few runs before it exits. In each run it

- Gets file list from the cache directory
- Applies them to the target
- Execute any ready snapshot requests

The value here determines how many runs will be performed per invocation.

**Type:** integer

**Units:** bytes

**Default value:** 1

**Range:** 1 to any value (0 being infinite). If you set this value to zero, target site VX agent would run continuously until there is replication configuration change or svagent service is shutdown. Increasing this value reduces the process creation overheads and also allows for applying the changes from cache directory to target volume even if CX server is unavailable.

**TcpSendWindowSize and TcpRecvWindowSize:****Functionality:**

The TCP Window is a buffer that determines how much data can be transferred before the server (cx) waits for acknowledgement.

**Type:** integer

**Units:** bytes

**Default value:** 1048576

**Recommended value:** needs to be large enough to fit the maximum\_available\_bandwidth x maximum\_anticipated\_delay. Use <http://www.speedguide.net/bdp.php> to calculate the product and use the rounded value.

**CMRetryDelayInSeconds:****Functionality:**

This determines the delay value between attempts by cache manager on encountering an error.

**Type:** integer

**Units:** seconds

**Default value:** 30

**Recommended value:** 30 seconds

**CMMaxRetries:**

**Functionality:** This determines max. no of retries by cache manager on encountering an error while downloading a file before it exits.

**Type:** integer

**Units:** number

**Default value:** 10

**Recommended value:** 10

**CacheMgrExitTime:**

**Functionality:** This determines max. Amount of time for which cache manager continuously run before exiting for next invocation.

**Type:** integer

**Units:** seconds

**Default value:** 16200

**Recommended value:** none.

**MaxInMemoryUnCompressedFileSize****Functionality:**

This determines the max size of an uncompressed file that should be downloaded to memory. If an uncompressed differential file size on CX is larger than this value, it would be downloaded directly to application cache directory.

**Type:** integer

**Units:** bytes

**Default value:** Linux: 16MB, Windows: 8MB

**Recommended value:** Set this equal to (or greater than) average differential file size.

**Disable value:** If 0 (is also the min value) is set then all the uncompressed files from CX would be downloaded to application cache directory

**MaxInMemoryCompressedFileSize:**

**Functionality:** This determines the max size of a compressed file that should be downloaded to memory. If a compressed differential file size on CX is larger than this value, it would be downloaded directly to application cache directory.

**Type:** integer

**Units:** bytes

**Default value:** Linux: 16MB, Windows: 4MB

**Recommended value:** Set this equal to (or greater than) average differential file size.

**Disable value:** If 0 is set then all the compressed files from CX would be downloaded to application cache directory

**CompressionChunkSize**

**Functionality:** All the differential files are uncompressed in chunks on the target. This value determines the uncompress chunk size.

**Type:** integer

**Units:** bytes

**Default value:** 1 MB

**Recommended value:** 1 MB

**CompressionBufSize**

**Functionality:** This determines the size of buffer allocations for storing uncompressed output. We allocate memory in multiples of this setting. In general this should be set to a value such that we do not have to do reallocations in most cases.

**Type:** integer

**Units:** bytes

**Default value:** 4 MB

**Recommended value:** Set this equal to (or greater than) average differential file size.

**MinCacheFreeDiskSpace and MinCacheFreeDiskSpacePercent:**

**Functionality:** These values control the amount of unused free disk space on the application cache volume for other applications sharing the volume. Minimum of either is used.

**Type:** integer

**Units:** bytes

**Default value:**

**MinCacheFreeDiskSpace** = 1 GB

**MinCacheFreeDiskSpacePercent** = 25

**Recommended value:** if the cache volume is used only as intermediate storage for differentials, set this to a low value. (5% of cache volume size)

## 14.5 Tuning TCP Window Sizes

TCP Windows Tuning on Source/Target Servers, the defaults would suffice for servers with in a local LAN [Gig Network]

**Parameters:** `TcpSendWindowSize` and `TcpRecvWindowSize`

**Location:** `Drscout.conf`

Linux: `/home/svsystems/var`

Windows: `<InMage Vx InstallPath>/Application Data/etc`

**Functionality:** The TCP Window is a buffer that determines how much data can be transferred before the server (cx) waits for acknowledgement.

**Type:** integer

**Units:** bytes

**Default value:** 1048576

It is observed that with 1 Gbps network link and 0.1ms latency 1048576 bytes results in better data transfer rates.

**Recommended value:** needs to be large enough to fit the `maximum_available_bandwidth x maximum_anticipaded_delay`.

**Use:** <http://www.speedguide.net/bdp.php> to calculate the product and use the rounded value.

## 15 Tuning CX

### 15.1 Using a Lower Encryption

For a higher encryption between CX and the target access console of the CX server and navigate to the folder “/etc”.

By default a 256 bit encryption is used, to use a lower encryption you can edit the file “proftpd.conf” under “/etc”. To use a 128 bit encryption remove the # symbol before “#TLSCipherSuite MEDIUM” then save the file and restart the service.

```
# Are clients required to use FTP over TLS
TLSRequired off
Remove the # symbol to uncomment
#cipher selection for secure mode transfers
# Different intensity cipher can be selected
#please uncomment below line if you need
#TLSCipherSuite MEDIUM
#please uncomment below line if you need
#TLSCipherSuite HIGH
# Server's certificate
TLRSACertificateFile /home/svsystems/etc/
TLRSACertificateKeyFile /home/svsystems/
# CA the server trusts
TLSCACertificateFile /home/svsystems/etc/
# Authenticate clients that want to use FTP
TLSVerifyClient off
</IfModule>
```

Figure 399



#### Notes:

Do not uncomment both the lines

### 15.2 TCP Window tuning for Linux based CX

Add the following lines to “/etc/sysctl.conf” and then run “/sbin/sysctl -p”

```
# increase TCP max buffer size settable using setsockopt()
net.core.rmem_max = 16777216
net.core.wmem_max = 16777216
# increase Linux autotuning TCP buffer limits
# min, default, and max number of bytes to use
# set max to at least 4MB, or higher if you use very high BDP paths
net.ipv4.tcp_rmem = 4096 87380 16777216
net.ipv4.tcp_wmem = 4096 65536 16777216
```

## 16 Bandwidth Shaping

You can define the bandwidth shaping policies to regulate outgoing data traffic rates to the target. This helps in controlling the bandwidth usage by Scout. Bandwidth Shaping is applicable to Volume Replication only.

You can specify the following information for bandwidth policy:

- Policy name, description, bandwidth shaping information, cumulative bandwidth, and bandwidth shaping schedule.
- Percentage bandwidth allocation between CX and Target machines called Bandwidth Shaping pairs.
- Sharing of unused bandwidth between two bandwidth shaping pairs.



Notes:

You can configure the bandwidth policies only after setting up the VX replication process.

### 16.1 Preparing for bandwidth shaping

Access the CX configuration server's command prompt to verify if TCP segmentation offload is enabled by running the "**ethtool -k eth0**" command. A sample command is shown as below

```
[root@imits198 tmp]# ethtool -k eth0
Offload parameters for eth0:
Cannot get device udp large send offload settings: Operation not supported
rx-checksumming: on
tx-checksumming: on
scatter-gather: on
tcp segmentation offload: on
udp fragmentation offload: off
generic segmentation offload: off
[root@imits198 tmp]#
```

Observe that the TCP segmentation offload is set to on. Proceed to disable it through the command "**ethtool -K eth0 tso off**"

You may verify it again through the "**ethtool -k eth0**" command. Always ensure that the "**TCP segmentation offload**" is set to off for the bandwidth shaping to function

## 16.2 Creating bandwidth shaping pairs

The fine-grained cumulative bandwidth shaping (provisioning) is based on Bandwidth Shaping pairs. Bandwidth shaping pair is a unique combination of CX and participating Target host in the volume replication.

118. On the “System” menu click on the “Bandwidth Shaping” link. The “Configure Bandwidth Utilization” window appears. The list of process servers, NICs on the process server and target VX agents are displayed.



### Notes:

The Traffic Control displays the CX-PS server and Target server’s IP Address pairs.  
Even windows process server’s NIC is shown as eth0  
Windows CX-PS server with multiple NICs do not support bandwidth shaping.

System: Bandwidth Shaping

v5.50.1.DIT.1578.1 (DAILY\_5-50-1\_DIT\_1578\_Apr\_28\_2009\_InMage).

Logged in as 'admin@10.0.7.21' - Logout

Server Time: May-8-2009 09:21:04

Monitoring

Bandwidth Shaping

CX Settings

Agent Settings

License Management

Agent Heartbeat

Process Server Traffic Load Balancing

Process Server Failover

Remote

CX

Versions and Patches

User Documents

Logs

Agent Installers

RX Settings

Configure Bandwidth Utilization

	Process Server Name	Process Server NIC	Bandwidth	Link Speed	Traffic Control	Policy Status
					Process Server --> VX Agent	
	vm-ps_7_23.localdomain	eth0 [10.0.7.31]	WAN	100Mbps	10.0.7.31 --> 10.0.7.11	NOT ENABLED
	vm-ps_7_23.localdomain	bond0 [10.0.7.34]	WAN	100Mbps	10.0.7.34 --> 10.0.7.12	NOT ENABLED
	cs_n_ps	eth0 [10.0.7.21]	WAN	100Mbps	10.0.7.21 --> 10.0.7.12 10.0.7.21 --> 10.0.7.31	ENABLED

Submit

Reset

Enable Policy

Figure 400: Enabling the global bandwidth shaping option

119. Select the desired process server’s NIC –target VX agent entry and click on “Submit” to create a bandwidth policy for that PS-VX pair.

120. Click on “Create Policy” on the new screen

Bandwidth Configuration - WAN

v5.50.1.DIT.1578.1 (DAILY\_5-50-1\_DIT\_1578\_Apr\_28\_2009\_InMage).

Logged in as 'admin@10.0.7.21' - Logout

Server Time: May-8-2009 09:27:56

Monitoring

Bandwidth Shaping

CX Settings

Agent Settings

License Management

Agent Heartbeat

Process Server Traffic Load Balancing

Process Server Failover

Remote

CX

Versions and Patches

User Documents

Logs

Agent Installers

RX Settings

Create Policy

Existing Policies for cs\_n\_ps [10.0.7.21]

Policies		Schedule			Action		
ID	Policy Name	Time	Criteria	Priority	View	Edit	Delete

Action

Back

Figure 401

121. Enter the name and description of the bandwidth policy.

122. Enter the “**Cumulative Bandwidth**” in kbps. This bandwidth will be shared between all the replication pairs handled by the NIC on the process server. Recall that you have selected the process server’s NIC-target VX pair in the previous steps.

**Create Bandwidth Policy WAN**  
v5.50.1.DIT.1578.1 (DAILY\_5-50-1\_DIT\_1578\_Apr\_28\_2009\_InMage).

Monitoring | **Bandwidth Shaping** | CX Settings | Agent Settings | License Management | Agent Heartbeat | Proc  
CX | Versions and Patches | User Documents | Logs | Agent Installers | RX Settings

**Bandwidth Policy for cs\_n\_ps [10.0.7.21]**

Policy Name:  ☐ Default Policy

Description:

Cumulative Bandwidth (Kbps):

Target	Shape Bandwidth	Allocation(%)	Bandwidth (Kbps)
10.0.7.12	<input checked="" type="checkbox"/> Include	<input type="text" value="55"/>	<input type="text" value="550"/>
10.0.7.31	<input checked="" type="checkbox"/> Include	<input type="text" value="45"/>	<input type="text" value="450"/>

☒ Share Unused Bandwidth      Bandwidth Allocated:  %

**Set Schedule**

From:  :   To:  :

☒ Daily  
☐ Weekly  
☐ Monthly  
☐ Yearly

☐ Every  day(s)  
☒ Every Day

**Figure 402: Creating new bandwidth policy**

123. Select the desired VX agents that need to be part of the bandwidth policy then allocate the % of the cumulative bandwidth to each of the VX agents. Optionally you may choose to enable the “**Share Unused Bandwidth**” option.

124. The “**Set Schedule**” is not shown for the “**Default Policy**”, Select the desired schedule and click on “**Save**”.



**Notes:**

Select the **Share Unused Bandwidth** checkbox if you want to share the unused bandwidth between the two bandwidth shaping pairs.

When a month does not have a given day then there will be no affect even if it is used i.e. If you select the 30<sup>th</sup>, there will be no affect for the month of February.



125. A message is displayed on the CX UI indicating that the policy is created. Click on “Next” to continue.

**Create Bandwidth Policy Confirmation WAN**  
v5.50.1.DIT.1578.1 (DAILY\_5-50-1\_DIT\_1578\_Apr\_28\_2009\_InMage).

Monitoring | Bandwidth Shaping | CX Settings | Agent Settings | License Management | Agent Heartbeat | Process Server Traffic Load Balancing | Process Server Failover | Remote

CX | Versions and Patches | User Documents | Logs | Agent Installers | RX Settings

**Policy Confirmation**

Message	cs_n_ps_To_Tgt [7] was Created
Action	Success, Click on forward for list of Bandwidth Policies
Potential Policy Conflicts	None

Next

Figure 403:

126. You may view, edit, or even delete the policy by selecting the appropriate option and clicking on the “Action” button.

**Bandwidth Configuration - WAN**  
v5.50.1.DIT.1578.1 (DAILY\_5-50-1\_DIT\_1578\_Apr\_28\_2009\_InMage).  
Logged in as 'admin@10.0.7.21' - Logout  
Server Time: May-8-2009 09:32:04

Monitoring | Bandwidth Shaping | CX Settings | Agent Settings | License Management | Agent Heartbeat | Process Server Traffic Load Balancing | Process Server Failover | Remote

CX | Versions and Patches | User Documents | Logs | Agent Installers | RX Settings

Create Policy

**Existing Policies for cs\_n\_ps [10.0.7.21]**

ID	Policy Name	Time	Criteria	Priority	View	Edit	Delete
7	cs_n_ps_To_Tgt [default]	12:00 A.M.to 11:59 P.M.	Applicable when no policies are defined or no other policy can be enforceable	0			

Action Back

Figure 404

127. The policy may be enabled or disabled through the screen shown below.

**System: Bandwidth Shaping**  
v5.50.1.DIT.1578.1 (DAILY\_5-50-1\_DIT\_1578\_Apr\_28\_2009\_InMage).  
Logged in as 'admin@10.0.7.21' - Logout  
Server Time: May-8-2009 09:32:36

Monitoring | Bandwidth Shaping | CX Settings | Agent Settings | License Management | Agent Heartbeat | Process Server Traffic Load Balancing | Process Server Failover | Remote

CX | Versions and Patches | User Documents | Logs | Agent Installers | RX Settings

**Configure Bandwidth Utilization**

	Process Server Name	Process Server NIC	Bandwidth	Link Speed	Traffic Control	Policy Status
					Process Server --> YX Agent	
	vm-ps_7_23.localdomain	eth0 [10.0.7.31]	WAN	100Mbps	10.0.7.31 --> 10.0.7.11	NOT ENABLED
	vm-ps_7_23.localdomain	bond0 [10.0.7.34]	WAN	100Mbps	10.0.7.34 --> 10.0.7.12	NOT ENABLED
	cs_n_ps	eth0 [10.0.7.21]	WAN	100Mbps	10.0.7.21 --> 10.0.7.12 10.0.7.21 --> 10.0.7.31	ENABLED

Submit Reset Enable Policy

Figure 405

## 16.2.1 Policy Conflicts

Policy conflicts occur, when two or more policies are eligible for policy enforcement at the same time. To resolve these conflicts, Scout assigns priorities by default, which is based on the associated scheduling option. You can also prioritize the policies manually, based on your requirements. For better illustrations, consider the following policies as examples

- Policy 1: *Daily Policy* - **Every 1 day**  
From: 06:00 AM to 11:50 P.M.
- Policy 2: *Weekly Policy* – **Every 1 Week on Every Monday, Tuesday**  
From: 09:00 AM to 6:50 P.M.
- Policy 3: *Monthly Policy* – **Every 1 Month on October 24th**  
From: 10:00 AM to 3:00 P.M.

Consider current data & time to be: October 24<sup>th</sup>, Tuesday, 1:00 P.M. – All the policies will be applicable for enforcement as they meet the policy scheduling criteria.

Bandwidth shaping modules detect these conflicts when policies are created or edited, in case of such conflicts, the policy with the highest priority is enforced, and in the above case it would policy-3. For more details on priorities refer to the subsequent section. The order of priority enforced by Scout is:

**Table 40: Bandwidth policy and priority**

Policy Type	Priority
Default Policy	0
Daily-Periodic Ex. Every [1] day/s	1
Daily-Specific Ex. Every Day	2
Weekly	3
Monthly-Periodic Ex. Day 1 of every 2 month(s)	4
Monthly-Specific Ex. The first day of every 2 month(s)	5
Yearly-Periodic Ex. Every January 21	6
Yearly-Specific Ex. The last Wednesday of September	7

The conflicts in the policy are detected when they are created or edited. In case of such conflicts, the policy with the highest priority is enforced.



### Notes:

The default policy is applied only if the bandwidth shaping option is enabled and no other bandwidth policies can be enforced.

You can define only one default policy

The default policy does not have any schedule attached to it.

By default, all outgoing traffic originating from the ports 20, 21, 873, 989, 990 are throttled

Do not try to change other BPM related parameters, this may disrupt the bandwidth shaping feature.

## 16.3 How Policies are enforced

- The basic rule is the policy with highest assigned priority will be selected.
- If there are multiple policies with the same priority, then the policy creation time will be used to resolve the policy.
- In case, there are no-policies enforceable and default policy is configured, then default policy will be enforced automatically

## 16.4 Effect of Stopping VX Replication Pairs on the Bandwidth Policies

Whenever the VX Replication is stopped, the bandwidth policies get readjusted. Consider the following illustrations

**Condition:** If all the VX replication pairs are stopped

**Effect:** all bandwidth policies get deleted.

**Condition:** If there are multiple VX replication pairs (with same target host), one replication pair is stopped

**Effect:** no change in the bandwidth policies.

**Condition:** If there are multiple VX replication pairs (each with a different target host), one of the replication pair is stopped

**Effect:** the bandwidth allocated to the stopped pair is allocated proportionately to the remaining pairs.

## 16.5 Advanced Configuration

Bandwidth shaping feature, by default should require additional configuration apart from creating bandwidth policies. Advanced hardware related configurations can be done using “**amethyst.conf**” file located in “**/home/svsystems/etc**”.

Some of the possible hardware related configurations are

- Configuring to change or use a particular network device on CX Server, by default the device “eth0” is specified.

```
BPM_NTW_DEVICE = "eth0"
BPM_LNK_SPEED  = "100Mbit"
```

- Shaping all outgoing traffic between CX and Outpost.

```
BPM_SRC_PORTS  = ""
BPM_DEST_PORTS = ""
```

## **Hitachi Data Systems**

### **Corporate Headquarters**

750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
Phone: 1 408 970 1000  
[www.hds.com](http://www.hds.com)  
[info@hds.com](mailto:info@hds.com)

### **Asia Pacific and Americas**

750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
Phone: 1 408 970 1000  
[info@hds.com](mailto:info@hds.com)

### **Europe Headquarters**

Sefton Park  
Stoke Poges  
Buckinghamshire SL2 4HD  
United Kingdom  
Phone: + 44 (0)1753 618000  
[info.eu@hds.com](mailto:info.eu@hds.com)



**MK-98DF8212-00**