



# **Disaster Recovery as a Service, Release 1.0 Design and Implementation Guide**

September 11, 2013



CCDE, CCENT, CCSI, Cisco Eos, Cisco Explorer, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco TrustSec, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco:Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLYNX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1002R)

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

*Disaster Recovery as a Service, Release 1.0 Design and Implementation Guide*  
© 2013 Cisco Systems, Inc. All rights reserved.



# CONTENTS

## Preface v

Documentation, Support, and Security v

---

## CHAPTER 1

### Overview 1-1

DRaaS: Business Drivers 1-3

DRaaS: Technical Challenges 1-3

DRaaS: Host-Based Replication As Preferred Choice 1-5

Value of Cisco DRaaS Architecture for Service Providers 1-8

Value of Cisco DRaaS for Enterprises 1-10

---

## CHAPTER 2

### System Architecture 2-1

DRaaS 1.0 System Architecture 2-1

System High Level Architecture 2-1

End-to-End Architecture 2-4

DRaaS Operational Workflows 2-5

VMDC Cloud Infrastructure 2-8

VMDC 2.3 Architecture 2-9

VMDC 2.3 Network Containers 2-13

Modifications in VMDC Network Containers for DRaaS 2-17

VMDC Orchestration using BMC CLM 2-20

CLM 3.1 SP1 Architecture 2-20

Container Pre-Provisioning for DR Services 2-24

Available Replication Types 2-26

Hypervisor-based vs. Guest OS vs. Storage 2-27

InMage Software Architecture 2-28

InMage Overview 2-28

Components 2-31

How InMage ScoutCloud Works 2-37

Component Flow 2-38

Deployment Considerations 2-39

Journal Sizing 2-39

Storage 2-41

Compression 2-43

External Cisco Products 2-43

- DR Vendor 2-44
- Encryption 2-44
- Compute 2-45
  - Oversubscription 2-45
- Key Findings 2-45
  - Concurrency 2-45
  - Limitations 2-45

**CHAPTER 3**

**Implementation and Configuration 3-1**

- Master Target—Enterprise and Service Provider 3-1
  - Master Target OS-Specific Information 3-2
  - Master Target Deployment Limit 3-2
  - Volume Requirements 3-3
    - Retention Volume Sizing 3-3
    - Cache Volume Sizing 3-5
- InMage Scout Server Details 3-6
  - Scout Server Storage and Compute Implementation 3-6
  - Scout Server Network Implementation 3-7
  - Scout Server Network Bandwidth 3-8
  - Scout Server Replication Options 3-10
- InMage vContinuum 3-12
- InMage Agent Configuration 3-13
  - CX UI for Linux Bulk Install and Upgrade 3-14
  - vContinuum for Windows Bulk Install and Upgrade 3-16
- Multi-Tenant Portal—RX Server 3-17
  - Multi-Tenant Portal Version Compatibility 3-17
  - Multi-Tenant Portal User Accounts 3-18
  - CX Server Registration with Multi-Tenant Portal 3-20
  - Multi-Tenant Portal Rebranding 3-21
- Summary Tables of Components for All Tenants 3-21
  - InMage Components for Service Provider 3-22
  - InMage Component for Enterprise Tenants 3-23
- VMDC 2.3 3-24
  - VMDC 2.3 Integrated Compute and Storage Stack 3-25
    - UCS Implementation 3-26
    - ESXi Implementation 3-28
    - Nexus 1000V 3-30
    - VSG Implementation 3-32
  - Mapping DR Components to VMDC 2.3 Containers 3-32

Tenant Configuration	3-33
IPsec	3-33
Out of Band Management Portal Access	3-34
VMDC Container Modifications	3-35
Connectivity across the WAN	3-39
Storage Configuration	3-42
SAN Implementation Overview	3-42
VNX5500 Configuration Overview	3-45
BMC Cloud Lifecycle Management	3-49

**CHAPTER 4**

<b>Disaster Recovery Workflow</b>	<b>4-1</b>
Protection Workflows	4-2
Setting up Virtual-to-Virtual (V2V) Protection Plan	4-2
Setting up Physical-to-Virtual (P2V) Protection Plan	4-18
Offline Sync	4-26
Offline Sync Export	4-27
Offline Sync Import	4-33
Recovery Workflows	4-37
Failback Protection Workflows	4-46
Virtual Failback Protection	4-46
Virtual-to-Physical (V2P) Failback Protection	4-59
Prepare the Physical Server	4-60
Prepare the USB Flash Drive	4-63
Create the V2P Failback Plan	4-68
Recover the Physical Server	4-73
Resume Protection Workflows	4-78
DR Drill Workflows	4-84

**CHAPTER 5**

<b>Monitoring, Best Practices, Caveats, and Troubleshooting</b>	<b>5-1</b>
InMage Resource Monitoring	5-1
Bandwidth Monitoring	5-2
Scout Server Health Monitoring	5-3
RPO and Health Monitoring	5-5
I/O Monitoring	5-7
Implementation Best Practices	5-8
Caveats	5-12
Troubleshooting	5-15
VMware Tools	5-19

vContinuum Logging 5-20  
Configuration and Processing Server Logging 5-21  
InMage Tools 5-22

---

**APPENDIX A**

**Characterization of Replication Process A-1**  
Replication with Compression Disabled A-3  
Replication with Compression Enabled A-4  
Comparison of Compression Enabled and Disabled A-5

---

**APPENDIX B**

**Extending a Linux Volume B-1**

---

**APPENDIX C**

**Acknowledgements C-1**

---

**APPENDIX D**

**Glossary D-1**



# Preface

Chapter content in this document is described in [Table 1](#).

**Table 1** *Document Organization*

Chapter	Discussion
<a href="#">Preface</a>	Document organization
<a href="#">Chapter 1, “Overview ”</a>	Introduction to the DRaaS 1.0 System.
<a href="#">Chapter 2, “System Architecture”</a>	Describes high level architecture, end-to-end architecture, and DRaaS workflows in the DRaaS 1.0 System.
<a href="#">Chapter 3, “Implementation and Configuration”</a>	Details the master target, InMage Scout Server details, InMage Agent configuration, InMage vContinuum, Multi-tenant Portal - RX Server, summary tables of components for all tenants, VMDC, and BMC Cloud Lifecycle Management within the DRaaS 1.0 System.
<a href="#">Chapter 4, “Disaster Recovery Workflow”</a>	Includes the protection, recovery, failback, resume, and DR drill workflows in the DRaaS 1.0 System.
<a href="#">Chapter 5, “Monitoring, Best Practices, Caveats, and Troubleshooting”</a>	Includes procedures for operational maintenance, best practices, caveats, and troubleshooting techniques in the DRaaS 1.0 System.
<a href="#">Chapter A, “Characterization of Replication Process”</a>	Describes the characterization of a replication process.
<a href="#">Chapter B, “Extending a Linux Volume”</a>	Describes how to extend a Linux volume.
<a href="#">Chapter C, “Acknowledgements”</a>	List of authors of this document.
<a href="#">Chapter D, “Glossary”</a>	List of terms and acronyms.

## Documentation, Support, and Security

Specific information about DRaaS-related software can be obtained at the following locations:

**InMage (contact InMage for access)**

- [InMage documentation](#)

## **VMDC 2.2**

- [Cisco VMDC 2.2 Design Guide](#)
- [Cisco VMDC 2.2 Implementation Guide](#)
- [Cisco VMDC Documentation on Cisco.com Design Zone](#)
- [Cloud Ready Infrastructure Smart Solutions Kits Accelerate Design and Deployment of Unified DC](#)

## **VMDC 2.3**

- [VMDC 2.3 Design Guide](#)
- [VMDC 2.3 Implementation Guide](#)
- [VMDC 2.3 Test Results Report](#)
- [SP Cloud Smart Solutions with VMDC](#)
- [Cloud Service Assurance for VMDC Design and Implementation Guide](#)
- [Cloud Orchestration for VMDC with BMC Cloud Lifecycle Management 3.1 SP1 Design and Implementation Guide](#)

For information on obtaining documentation, submitting a service request, and gathering additional information, refer to the monthly [What's New in Cisco Product Documentation](#), which also lists all new and revised Cisco technical documentation, at:

Subscribe to the What's New in Cisco Product Documentation as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.





# CHAPTER 1

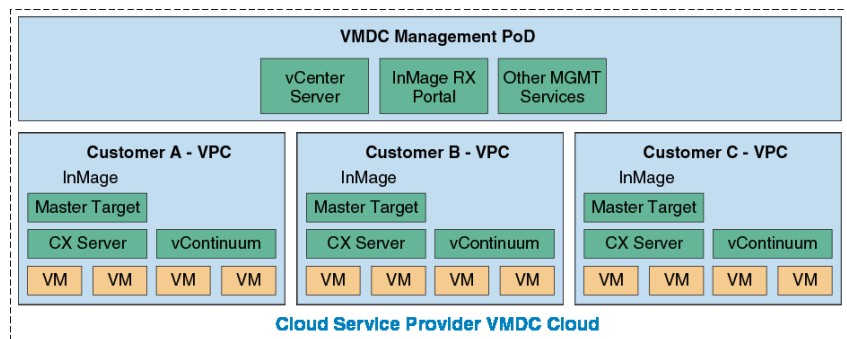
## Overview

Cisco Disaster Recovery as a Service Solution (DRaaS) architecture described in this document is designed to provide a new set of related capabilities allowing Virtualized Multi-Tenant Data Center (VMDC)-based service providers (SP) to enhance their addressable market, financial performance, and differentiation vs. commodity cloud solutions. Many of Cisco VMDC-based SPs seek better monetization of their existing VMDC investments through layered services that are synergistic with the advanced networking capabilities delivered by VMDC. These SPs demand new, easily deployable services both to keep pace with the innovation of commodity/public cloud providers such as Amazon Web Services (AWS) and to address portions of the market that are not well served by commodity cloud solutions.

The key end user consumable services being enabled by this system architecture is to enable a SP to offer disaster recovery for both physical and virtual servers from a customer data center to a SP virtual private cloud (VPC). The DRaaS System primarily targets SMBs and enterprises. The global DRaaS and cloud-based business continuity is expected to grow from \$640.84 million in 2013 to \$5.77 billion by 2018, at a CAGR of 55.20%.


The traditional disaster recovery (DR) system constitutes a substantial portion of expenses annually. With the "pay as you go" model of the cloud-based DR system, the impact of downtime can be minimized through replication. DR can start up applications once the disaster is identified. In addition to recovery, cloud-based DR incorporates business continuity. Implementation of DRaaS with a virtualized cloud platform can be automated easily and is less expensive, since DR cost varies before and after a disaster occurs. The key requirements for DRaaS are Recovery Point Objective (RPO), Recovery Time Objective (RTO), performance, consistency, and geographic separation (Figure 1-1).

**Figure 1-1** What is Disaster Recovery as a Service?




The market presents a strong opportunity for the SPs to take advantage of the demand for DRaaS services as illustrated by Figure 1-2.

**Figure 1-2 Strong Market Demand for DRaaS**



Further investigation of the global demand patterns for DRaaS indicates that the market opportunity and interest is equally spread across the enterprise, mid-market, and SMB segments as summarized in [Figure 1-3](#).

**Figure 1-3 Global DRaaS Demand by Segment**




This chapter includes the following major topics:

- [DRaaS: Business Drivers, page 1-3](#)
- [DRaaS: Technical Challenges, page 1-3](#)
- [DRaaS: Host-Based Replication As Preferred Choice, page 1-5](#)
- [Value of Cisco DRaaS Architecture for Service Providers, page 1-8](#)
- [Value of Cisco DRaaS for Enterprises, page 1-10](#)

## DRaaS: Business Drivers

Increased regulatory pressure drives the need for disaster recovery (DR) and business continuity plans and presents a hierarchy of requirements for the implementation of these solutions (geographic restrictions, regulatory compliance, etc.). Enterprises are constantly faced with budget constraints that prevent infrastructure duplication. Building DR infrastructure is a contextual business activity that requires a degree of specialization with IT skillsets or resources that are significantly harder to build without sufficient scale. Under these circumstances a growing desire exists to consume DR as a service, allowing incremental deployment and growth as budget becomes available.

**Figure 1-4 Cisco's DRaaS Blueprint Solution**




## DRaaS: Technical Challenges

The selection of a specific technology and implementation for the DRaaS is a highly complex decision with technology challenges that need to be adequately explored and analyzed prior to choosing an appropriate technology. The following questions arise in the choice of the DRaaS implementation:

- How do we replicate data, databases, and virtual machines?
- What technology of replication do we use?
- What are our RTO/RPO requirements for the various applications requiring Disaster Recovery?
- How should we monitor what is being done during the testing and recovery events?
- How should we perform failover when needed either by a test or a disaster event?
- How should we virtual machines and databases be rebuilt?
- How can we ensure the consistency of databases and applications?
- How can we redirect traffic, reconfigure the Domain Name Services, etc.?
- How should we perform failback after a recovery event?
- How should our organization staff for Disaster Recovery and testing?

- How can our organization afford Disaster Recovery (which a cost and not a revenue generating activity)?

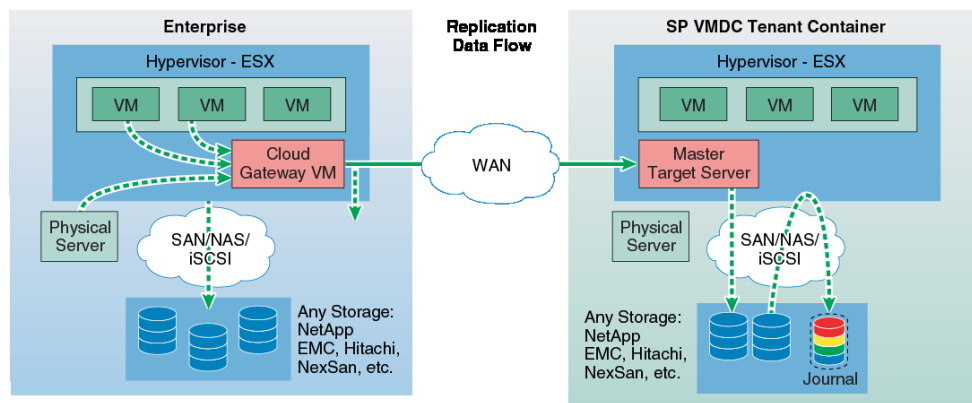
**Figure 1-5 DRaaS Technical Challenges**



### Challenges with Traditional Storage-based Replication

The use of traditional storage-based replication requires an identical storage unit on the DR site from the same vendor. The storage array-based replication software is not application aware and needs additional intervention at the host level to achieve application consistency. Multiple points of management are required while performing DR and this introduces complexity in protecting and recovering workloads. The traditional storage-based replication approaches lack granularity and can replicate all virtual machines (VM) or none that are residing on a logical unit number (LUN). Replication of data happens between LUN pairs that need to be identical and this restricts the ability to failover a single VM residing on the LUN.

**Figure 1-6 Any-to-Any Replication**



Traditional storage replication approaches need additional functionality to take snapshots or clones of the target LUN to perform disaster recovery drills without interrupting data replication. Otherwise, replication has to be stopped for DR drills. Storage array-based replication does not support continuous data protection natively and data cannot be protected from logical failures.

VMware Site Recovery Manager (SRM) is an orchestration and runbook automation tool that streamlines the workflows for failovers and recovery of workloads. SRM leverages storage-based replication or vSphere replication to provide DR. [Table 1-1](#) shows a comparison of the VMware approach to DR with the Cisco DRaaS approach.

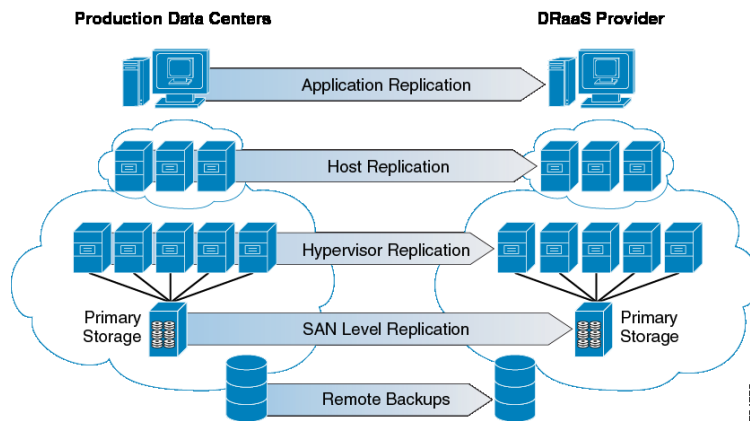
**Table 1-1** VMware Disaster Recovery Solution Comparison

SRM with Storage Replication	SRM with vSphere Replication	Cisco Solution
<ul style="list-style-type: none"> <li>• Supports only vSphere-to-vSphere replication.</li> <li>• Needs to have similar storage arrays on customer and DR sites.</li> <li>• Involves a complex configuration to provide point-in-time copies for recovery.</li> <li>• Issues with incompatibility, as SRM coordinates with multiple components (Storage Array software, SRAs, Multipath software, vSphere versions).</li> <li>• Needs storage configuration or reorganization before SRM is ready to use.</li> <li>• Limitation with N:1 replication and expensive to set up.</li> <li>• No multi-tenant portal</li> </ul>	<ul style="list-style-type: none"> <li>• Supports only vSphere-to-vSphere replication.</li> <li>• Does not provide point-in-time copies for recovery.</li> <li>• Limited ESXi version support (only supports vCenter 5.1 and above and ESXi 5.x and above).</li> <li>• RPO cannot be less than 15 minutes.</li> <li>• Limitations with N:1 replication and scalability: <ul style="list-style-type: none"> <li>– Simultaneous VM failover - between 10 - 80.</li> <li>– Site Pairing - 10 Sites only per vCenter/ SRM pair.</li> <li>– Limited to 500 VMs.</li> </ul> </li> <li>• Lack of cloning capability at DR site for performing DR drills.</li> <li>• No multi-tenant portal.</li> </ul>	<ul style="list-style-type: none"> <li>• Supports Any-to-vSphere replication.</li> <li>• Provides continuous data replication with multiple point in time copies for recovery.</li> <li>• Supports N:1 replication with any number of source sites.</li> <li>• Provides multi-tenant portal for customers.</li> <li>• Supports any-to-any replication with any storage type and vendor.</li> <li>• Supports near zero RPO and RTO.</li> </ul>

## DRaaS: Host-Based Replication As Preferred Choice


Several options exist in the choice of technology for the implementation of DRaaS, which is associated with varying levels of cost, complexity, and operational models. A summary of technology options for the implementation is presented in [Figure 1-7](#).

**Figure 1-7 Many Approaches to DRaaS**




The host-based replication technology is the recommended implementation for Cisco's DRaaS System architecture. It is delivered in partnership with InMage ScoutCloud product offering because of the value and the differentiation it provides delivering DR services for physical-to-virtual (P2V) and virtual-to-virtual (V2V) workloads.

**Figure 1-8 Host-Based Replication/Recovery Process**



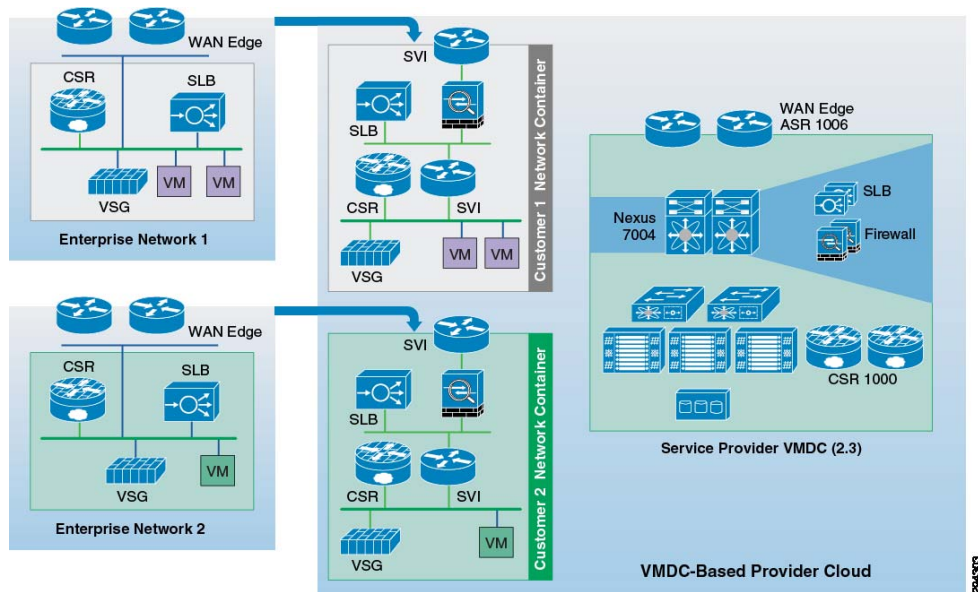
The Cisco DRaaS System, which offers architecture based on the VMDC 2.3 infrastructure architecture, provides P2V and V2V DR and business continuity capabilities. Cisco VMDC-based cloud provides secure multi-tenancy, services orchestration, high availability, and modularity.

Figure 1-9 DRaaS Offering



Layer 2 Extensions and IP mobility using Overlay Transport Virtualization (OTV) and Lisp to support partial failovers and active-active scenarios are targeted to be addressed as part of future capabilities of VMDC architecture. The solution presents heterogeneous, storage, and infrastructure-agnostic data replication capabilities for the creation and offer of DR solution offerings. The system offers continuous data protection (CDP)-based recovery with the ability to roll back to any point in time. The system provides guaranteed application consistency for most of the widely-used applications.

Figure 1-10 Host-based DRaaS on VMDC Architecture



# Value of Cisco DRaaS Architecture for Service Providers

DRaaS offers the following value to SPs:

- **Increased Customer Relevance:** Not all of the customers requiring DR services want Infrastructure as a Service Offering (IaaS). Offering DRaaS provides better alignment with a typical IT buyer's focus. Leverage of DRaaS offerings by SPs provide them an opportunity to differentiate from commodity and over-the-top IaaS providers.
- **Bigger, More Profitable Deals:** DR instances command a premium and provide improved margins due to lack of commoditization. DR deals are typically larger compared to IaaS deals for SPs and generate higher margins. DRaaS offerings create reduced capital expenditures on compute resources and lower operating expenses on licensing due to oversubscription opportunities.
- **Strong Services Growth:** DRaaS offerings present a strong ability to attach additional services with the offerings and creates a pipeline of revenue from new and existing customers through new and improved monetization via services growth. Additional monetization opportunities present themselves through possibilities for hybrid services.

## Cisco DRaaS Approach vs. Backup-based Disaster Recovery

One commonly encountered question is how do the backup-based disaster recovery approaches compared to Cisco's recommendation for DRaaS architecture for SPs. [Table 1-2](#) shows the key considerations and a comparison of the approaches.

**Table 1-2 Comparison of Cisco DRaaS vs. Backup-based DR**

	Managed backup using Cloud Storage	Backup-based Cloud Recovery using Snapshots	Cisco Approach
<b>Use Case</b>	Backup to cloud: Cloud storage for backups	Disaster recovery: SP-managed disaster recovery	Disaster recovery: SP or customer self-managed disaster recovery
<b>Pros</b>	Customers have ability to store data offsite without shipping tapes or having a secondary site to host data	Makes use of existing backup and virtualization tools for recovery	SP managed or enterprise self managed Single solution for protecting both physical and virtual environments Automated recovery
<b>Cons</b>	<ul style="list-style-type: none"> <li>• Does not ensure continuity of operations. Provides data availability only.</li> <li>• Impacts performance of application during backup window.</li> <li>• No automated recovery</li> </ul>	<ul style="list-style-type: none"> <li>• No P2V capability, protection for only virtual environments</li> <li>• Performance impact on production applications during snapshot creating</li> <li>• No automated recovery</li> </ul>	
<b>RPO/RTO</b>	Very high	High	Near Zero
<b>Continuous Data Protection (CDP)</b>	N/A; works based on traditional backups	Near CDP, cannot achieve real CDP. Depends on the frequency of snapshots.	Real CDP, provides multiple point in time copies for an extended period of time.



**Service Provider Tenant Operating Models**

Cisco DRaaS presents a model that clearly delineates the responsibilities of the SPs providing the DRaaS services and the end customer guidance on the ownership and expectations in the system offering.


**Table 1-3 Well-Defined Tenant/SP Operational Responsibilities Model**

Responsibility	Service Provide	Tenant
Provide standby recovery environment (compute, network)	X	
Configure standby recovery environment with replication/ recovery plans for protected servers and network elements	X	
Recover/ boot protected servers to recovery environment with pre-defined VLAN/ IP address mapping and network topology	X	
Provide recovery to a specific point in time using CDP technology to create a bootable VMDK; boot associated VMs	X	
Ensure reachability of running VMs over pre-defined recovery network	X	
Validate application configuration and functionality		X
Provide notification of changes requiring recovery plan updates - VLANs, IPs, added/ removed volumes, new servers		X
Participate in annual recovery tests/ drills (no production impact)	X	X
Declare disaster		X

**SP Monetization of Cisco DRaaS**

Figure 1-11 is a financial model that presents the monetization opportunity for SPs associated with the deployment of the Cisco DRaaS System architecture.

**Figure 1-11 Monetization Opportunity for SPs**




# Value of Cisco DRaaS for Enterprises

DRaaS provides the following value for Enterprises:

- **Recovery Time Is Key:** Enterprises frequently lack the knowledge to select and deploy the optimal DR tools for their needs. Current enterprise tools for low RPO/RTO tend to be cost prohibitive for widespread deployment.
- **Reduced Cost and Impact of Disaster Recovery Testing:** DR exercises present a significantly high cost and are a "distraction factor" to the normal business operation. The use of DRaaS allows enterprises to focus on application validation without being distracted by rack, stack, and recover activities with their infrastructure and IT services. It also presents a potential opportunity to better leverage the DR environment.
- **Accelerated Implementation:** The use of DRaaS presents an easier framework for implementation of business continuity plans and test execution and provides end customers with the ability to grow over time from a limited scope. An equivalent DRaaS solution to replace one that is provided and managed through a SP's robust offerings would be extremely time consuming to build for enterprises on their own as they include self-service, monitoring, and service assurance capabilities as a holistic offer from SPs.
- **Better Odds of Success:** The use of specialized SP offerings eliminate the need for a strong DR competency and addressed the difficulty associated with hiring and retaining talent for DR. The DRaaS is a niche technology that requires a significantly large scale to gain the required specialized experience. Globalization means many organizations cannot use traditional primary and secondary model of dedicated infrastructures for DR and business continuity operations.

**Figure 1-12** Why Enterprises Choose DRaaS





## CHAPTER 2

# System Architecture

---

This chapter includes the following major topics:

- [DRaaS 1.0 System Architecture, page 2-1](#)
- [VMDC Cloud Infrastructure, page 2-8](#)
- [VMDC Orchestration using BMC CLM, page 2-20](#)
- [Available Replication Types, page 2-26](#)
- [Deployment Considerations, page 2-39](#)
- [Key Findings, page 2-45](#)

## DRaaS 1.0 System Architecture

This section includes the following topics:


- [System High Level Architecture, page 2-1](#)
- [End-to-End Architecture, page 2-4](#)
- [DRaaS Operational Workflows, page 2-5](#)

## System High Level Architecture

This section describes the high level architecture of the DRaaS System. The system provides disaster recovery for customer physical/virtual servers by deploying recovery VMs in the VMDC 2.3-based container on the provider side.

[Figure 2-1](#) illustrates the high level architecture of DRaaS System.

Figure 2-1 DRaaS High Level Architecture



The physical system architecture consists of the following building blocks:

### Provider Cloud

The provider cloud within the DRaaS System will be based on VMDC 2.3. The VMDC 2.3 design is based on the earlier VMDC 2.2 design, with changes to optimize the design for lower cost, fewer layers, and increased tenancy scale. The Cisco VMDC System provides vPC-based L3 hierarchical virtual routing and forwarding (VRF)-Lite DC design, multi-tenancy, secure separation, differentiated service tiers, and high availability in a data center environment. It also provides secure separation between replicated workloads and provides shared network services for customers in DRaaS.

The VMDC 2.3 architecture works with Vblock, FlexPod, or any other integration stack. Integrated stacks can be added as required to scale the SP cloud environment.

Based on the customer's production environment and needs, a specific tenancy model can be selected to provide similar services in the cloud-matching production environment. VMDC architecture and deployment models will be covered in detail in this chapter.

### Enterprise Data Center

The DR solutions should address enterprise customer requirements for various vertical industries and geographies. The enterprise data center design is therefore expected to vary from customer to customer. The intent of the DRaaS System is to keep the enterprise DC architecture generic so as to provide the greatest coverage. While the DC architecture is almost irrelevant and the solution supports heterogeneous replication across any-to-any infrastructure, a typical three tier (core/aggregation and access) DC architecture is suggested in the system.

### WAN Connectivity

The WAN connectivity design principles provided by VMDC are maintained and supported without requiring any additional components and technologies. The replicated data between the enterprise and SP data center can be encrypted with the help of Cisco technologies like IPsec VPN based on Cisco ASA firewalls. Optionally, for low cost implementation to support a small number of servers, inflight replicated data encryption can be provided by InMage partner software.

To support partial failover of customer's environment, technologies like Overlay Transport Virtualization (OTV) can be used for L2 extension between the customer's data center and the cloud. L2 connectivity allows customers to use the same IP from enterprise network in the cloud without the need to change for accessing workloads in the cloud after recovery.

### Partner Solution for Providing Disaster Recovery

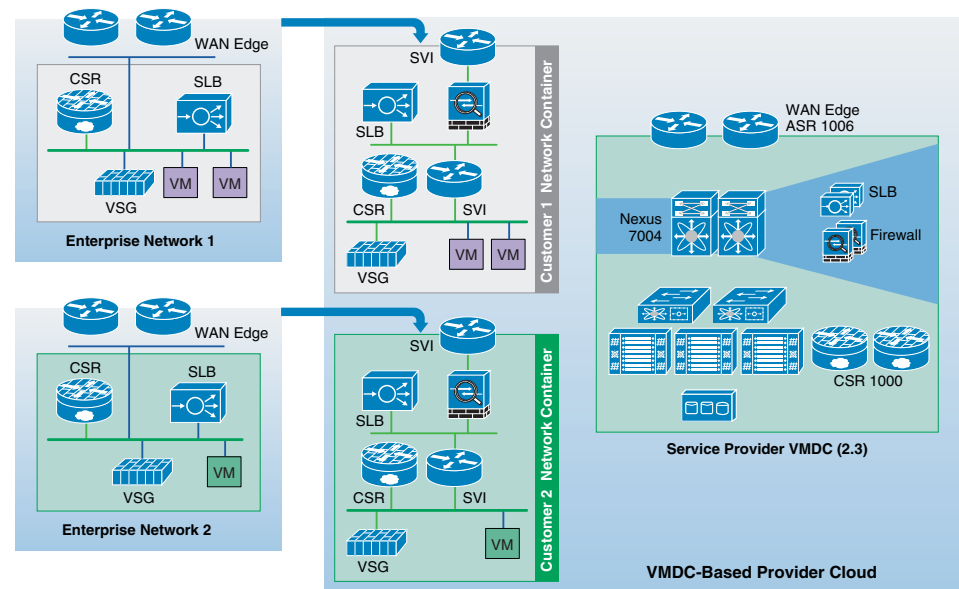
Data replication and recovery of the production servers will be provided by InMage ScoutCloud technology. InMage ScoutCloud is a software-based replication and recovery technology, which can protect both physical and virtual servers into the cloud. InMage ScoutCloud is being integrated into the DRaaS System, providing the following functionality:

- Heterogeneous data replication
- Continuous data protection
- Application consistency
- Recovery automation
- DR Drill

### System Logical Topology

Figure 2-2 covers the logical topology of the DRaaS System.

**Figure 2-2** DRaaS Logical Topology



As shown in Figure 2-2, each customer will have a dedicated network container created on the SP VMDC cloud. The network containers will be created based on the necessary security and network services required by the enterprise customers. Any network topology on the customer's data center can be matched on the VMDC cloud using network containers. Predefined containers provide examples for different types of deployments. Automated provisioning and management logic for each customer type is pre-defined in the management and orchestration software. Customers can choose from existing models or define their own customized models. The production workloads from each enterprise data center will be replicated to the corresponding network container on the VMDC cloud and will be available for recovery purposes.


## End-to-End Architecture

The DRaaS System addresses the following design principles and architectural goals:

- Secure Multi-Tenancy
- Secure, modular, and highly available cloud
- Continuous Data Protection (CDP)
- Physical-to-Virtual (P2V) and Virtual-to-Virtual (V2V) Disaster Recovery
- Near zero RPO and RTO-capable DRaaS
- Automated run book automation
- Self-Service Multi-Tenant Portal

By utilizing the architecture above, DRaaS in a multi-tenant environment can be supported as shown in [Figure 2-3](#).

**Figure 2-3** End-to-End Architecture



In a multi-tenant environment, each customer is mapped as a separate VMDC tenant where the necessary network security is provided and traffic segregation is maintained. Figure 2-3 depicts the end-to-end architecture of the DRaaS System based on VMDC.

With the deployment of lightweight components as shown in Figure 2-3 and utilizing the network security provided by VMDC architecture, customers can replicate their data into a secure cloud environment for recovery.

Data changes are collected from the production servers as they occur, directly in memory before they are written to disk, and sent to a software appliance within an enterprise data center. Because of this approach, absolutely no additional I/O load is induced on production servers due to replication. The appliance is responsible for further offloading compute-intensive tasks from production systems, such as compression, encryption, WAN acceleration, and consolidated bandwidth management.

The system provides CDP for the customer's production servers. The customers will be able to recover their environments to any point in time before the disaster occurred. The servers are not only protected from the physical disasters, but also from logical disasters due to CDP.

Application consistency is enforced at regular intervals through VSS integration on Windows and native application-specific mechanisms on Linux and Solaris systems. Application consistency is also enforced at the guest level in virtual environments such as VMware ESX, Xen Server, and Hyper-V. These application-consistent points are tagged by a bookmark and archived as part of the CDP data. They can be leveraged to perform application consistent recoveries within stringent recovery time objectives.

The following use cases are covered as part of the DRaaS System and will be discussed in more detail in the following sections:

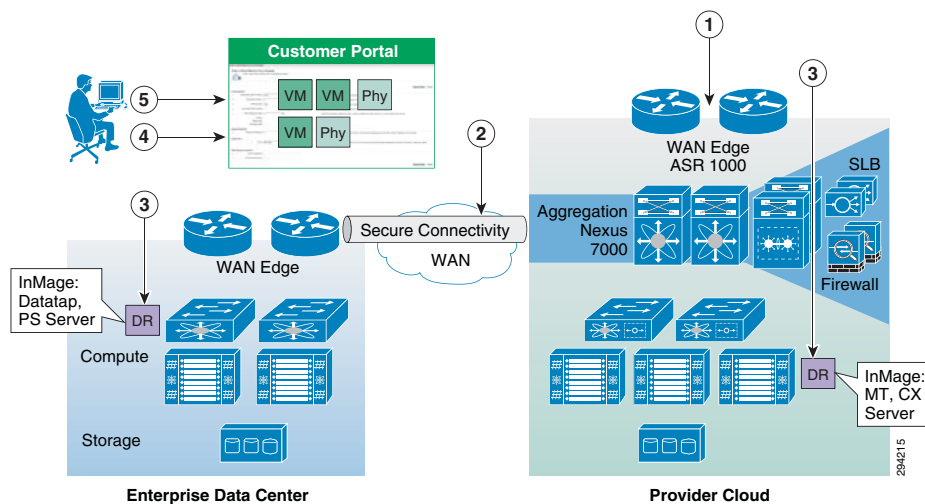
- [Protection Workflows, page 4-2](#)
- [Recovery Workflows, page 4-37](#)
- [Failback Protection Workflows, page 4-46](#)
- [Resume Protection Workflows, page 4-78](#)
- [DR Drill Workflows, page 4-84](#)

## DRaaS Operational Workflows

Following are the workflows for protecting and recovering the customer's production workloads into the cloud. The workflows describe the process of creating the network containers for customers within the SP cloud, replication of workloads into the network containers, and recovery of workloads in the event of a disaster.

The workflow in [Figure 2-4](#) is used for protection and failover scenarios.

**Figure 2-4** New Customer Protection Workflow




- 
- Step 1** Based on the customer requirements, deploy a VMDC Network Container using BMC.
- Step 2** Secure IPsec connectivity is manually set up between the Enterprise and the VMDC-based cloud provider setup.
- Step 3** At both enterprise and SP data centers, deploy and configure the necessary DR components.

- Step 4** Use the InMage management wizard to select the machines to be protected and set up the recovery plans.
- Step 5** Allow customers to monitor the status of DR and RPO/RTO utilizing the Partner Product portals.

The workflow in case of a failure scenario is shown in [Figure 2-5](#).

**Figure 2-5 Failure Scenario**



- Step 1** When the customer DC goes down, customer declares a disaster and communicates to SP what VMs to restore and what checkpoints to use. SP can use the recovery plan (which could be preconfigured), which details the list of protected VMs, the startup order, and any custom steps.
- Step 2** SP logs into the DR product portal and brings up the required VMs in its environment. Customers with self-service capabilities will be able to recover VMs in the cloud themselves using the self-service portal.
- Step 3** Customer works with its DNS provider to direct the client traffic to the SP DC. If the customer is utilizing a Global Site Selector (GSS)-based DNS solution or has a L2 extension, this step will be automatic or not required.
- Step 4** When the Enterprise DC is back up, customer works with the SP during a maintenance window to bring up the VMs in customer DC, failback the VMs from SP to enterprise, and update the DNS so that the client traffic is re-routed to the customer DC.

### Network Deployment Considerations to Support Recovery Environment

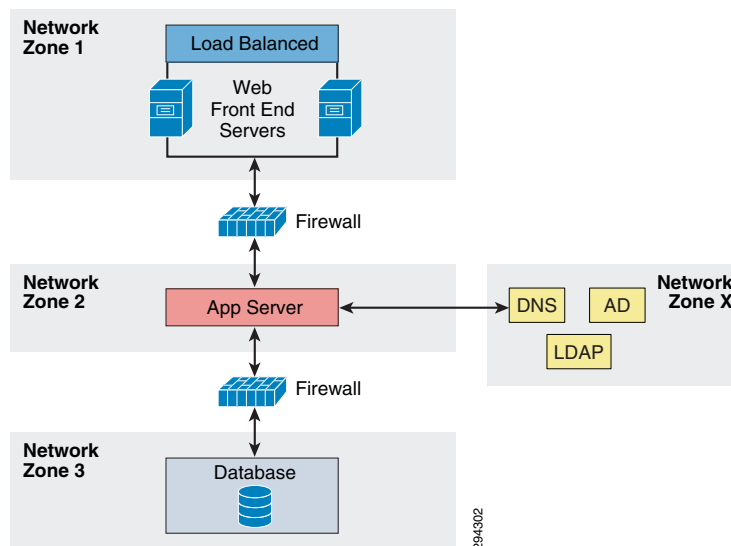
[Table 2-1](#) shows the considerations in matching the networks between the enterprise's and SP's VPC. Logically, the enterprise network will consist of VLANs and network services, including firewall rules and load balancing. Based on the requirements of enterprise, which depend on the type of applications that are protected, network containers can be created on the VMDC to meet those requirements.



**Table 2-1 Network Containers Available on VMDC**

Container	VLANs	Network Services
Gold	3	Tenant firewall, intra-tenant firewall, and load balancer
Silver	3	Load balancer
Bronze	1	Intra-tenant firewall, load balancer
Copper	1	Intra-tenant firewall

The typical deployment of a multi-tiered application running in the enterprise is shown in [Figure 2-6](#).

**Figure 2-6 Application Deployment Example**

The following is the onboarding procedure of a customer running the application shown above:

- The enterprise IT admin needs to coordinate with the SP to have the network container created on the VMDC, based on the requirements and dependencies of the application being protected. The options of creating the network container and maintaining consistency on the SP side are as follows:
  - The container is pre-configured by SP with the necessary VLANs and network services. The firewall rules and the load balancing rules pre-configured based on the pre-determined IPs of recovery servers on VMDC.
  - The container is preconfigured and the firewall and load balancing rules are configured dynamically by the SP using BMC orchestration or manually through CLI during the failover process of the servers. Any changes done with the network services after replication has been set up on the enterprise data center have to be communicated to the SP. This ensures network consistency during recovery of the servers. Optionally, the SP can enable the Enterprise customer to manage the firewall and load balancing services on the VMDC cloud. This can be done by providing access to the BMC orchestrator or to the specific devices directly for modifications.
- For the application to run properly on the VMDC cloud after recovery, all the components of the application from different tiers needs to communicate with each other. All the servers needs to be brought up in an order based on the application dependencies.

- The other important dependency is the IP addressing. Two types of configurations are done on the servers within an application for intra-component communication:
  - Configuration based on IP address
  - Configuration based on DNS names

Legacy application servers configured based on IP address can run seamlessly as long as they have the same IPs on the VMDC cloud. This may or may not be the case for all the customers. Customers who have different network subnets available on the VMDC need to reconfigure the servers to have the new IPs part of the application configuration. The task mentioned about can be performed by a SP administrator in a managed recovery use case or by the customer after the servers are available on the VMDC cloud.

The re-IPing of the servers can be eliminated if the servers with in the applications are using DNS names for communicating, in which case the DNS entries can be updated to reflect the new IPs. An SP can also perform the modification of DNS entries if the customer is using SP DNS servers. Optionally, DNS entries can be modified automatically using scripts.

In cases of the customer protecting the enterprise DNS servers, the DNS servers can be brought online during the recovery of application and based on the new IP addresses the configuration can be updated by the customer manually or can be automated.

#### Use of Cisco GSS

To accelerate the disaster recovery service and the dynamic distribution of the workload between the primary and secondary data centers, Cisco provides different network services to optimize the access and the distribution of the user traffic to the remote sites using a Global Site Load Balancing (GSLB) solution. This global GSLB solution for traditional L3 interconnection between sites relies on three major technologies:

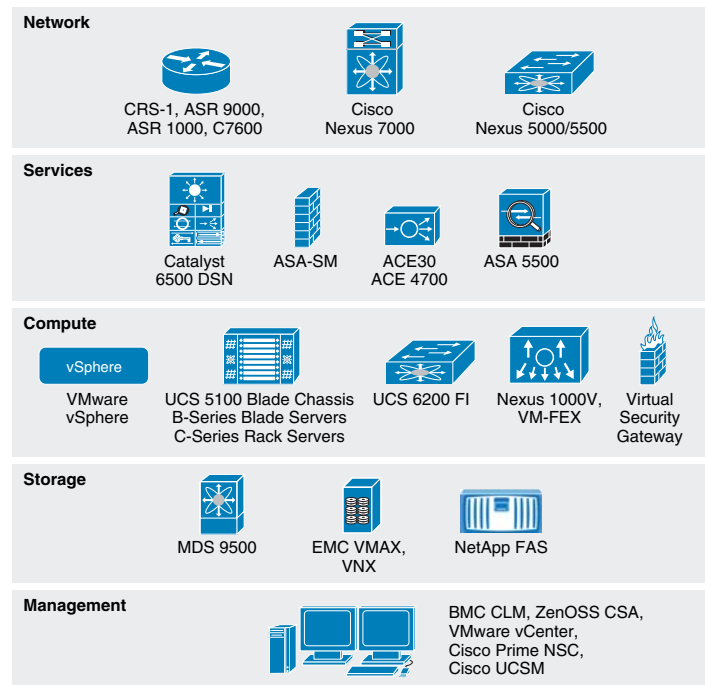
- Intelligent Domain Name System (DNS): A DNS known as the Global Site Selector (GSS) redirects the requests from end-users to the physical location where the application is active.
- HTTP Traffic Redirection between Sites: In case of resource unavailability, the local Server Load Balancing (SLB) device will return an HTTP redirection message type (HTTP status code 3xx) to the end-user so that the web browser of the client can be automatically and transparently redirected to the elected backup data center where resources and information are available.
- Route Health Injection (RHI): RHI provides a real-time, very granular distribution of user traffic across multiple sites based on application availability. This method is initiated by an SLB device that will inform the upward router about the presence or absence of selected applications based on extremely accurate information. This information is usually related to the status of the services that it supports. Therefore, the redirection of the user request to a remote site occurs in real time.

## VMDC Cloud Infrastructure

The VMDC System is the Cisco reference architecture for IaaS cloud deployments. This Cisco cloud architecture is designed around a set of modular DC components consisting of building blocks of resources called PoDs, or Points of Delivery. These PoDs comprise the Cisco UCS, SAN and NAS storage arrays, access (switching) layers, and aggregation (switching and routing) layers connecting into the DSN-based services layer or connecting directly to service appliances; and multiple 10 GE fabric using highly scalable Cisco network switches and routers. The VMDC system is built around the UCS, Nexus 1000V, Nexus 5000 and Nexus 7000 switches, Multilayer Director Switch (MDS), ASR 1000, ASR 9000, ASA 5585-X or Adaptive Security Appliance Services Module (ASASM), Catalyst 6500 DSN, ACE, Nexus 1000V VSG, VMware vSphere, EMC VMAX, VNX and NetApp FAS storage arrays.

Cloud service orchestration is provided by the BMC Cloud Lifecycle Management (CLM) suite and cloud service assurance is provided by the ZenOSS Cloud Service Assurance (CSA) suite. [Figure 2-7](#) provides a synopsis of the functional infrastructure components comprising the VMDC system.

**Figure 2-7 VMDC Infrastructure Components**



This section includes the following topics:

- [VMDC 2.3 Architecture, page 2-9](#)
- [VMDC 2.3 Network Containers, page 2-13](#)
- [Modifications in VMDC Network Containers for DRaaS, page 2-17](#)

## VMDC 2.3 Architecture

The VMDC System utilizes a hierarchical network design for high availability and scalability. The hierarchical or layered DC design uses redundant switches at each layer of the network topology for device-level failover that creates highly available transport between end nodes using the network. DC networks often require additional services beyond basic packet forwarding, such as SLB, firewall, and intrusion prevention. These services might be introduced as modules populating a slot of one of the switching nodes in the network or as stand-alone appliance devices. Each service approach also supports the deployment of redundant hardware to preserve the HA standards set by the network topology. This layered approach is the basic foundation of the VMDC design to provide scalability, performance, flexibility, resiliency, and service assurance. VLANs and VRF instances are used to provide tenant isolation within the DC architecture, and routing protocols within the VRF instances are utilized to interconnect the different networking and service devices. This multilayered VMDC architecture is comprised of core, aggregation, services, and access layers. This architecture allows for DC modules to be added as demand and load increases. It also provides the flexibility to create different logical topologies utilizing device virtualization, the insertion of service devices, and traditional L3 and L2 network configurations.

The VMDC 2.3 System is the latest released version of the VMDC architecture, with VMDC 2.2 being the previous release. Architecturally, VMDC 2.3 is based on VMDC 2.2 (and 2.0), but with several optimizations to reduce cost and footprint and increase tenancy scale. The key differences between VMDC 2.3 and 2.2 include:

- VMDC 2.3 includes an ASR 1000 as the DC Edge (PE) router, while VMDC 2.2 uses the ASR 9000.
- VMDC 2.3 includes a collapsed core/aggregation layer, while VMDC 2.2 includes a separate Nexus 7000 core layer and Nexus 7000 aggregation layers.
- VMDC 2.3 includes an ASA 5585-X for the perimeter firewall, while VMDC 2.2 uses the ASA5585-X or ASASM module on Catalyst 6500 DSN.
- VMDC 2.3 includes an ACE 4710 for Server Load Balancing, while VMDC 2.2 uses the ACE-30 module on the Catalyst 6500 DSN.
- VMDC 2.2 optimizes the Enhanced Gold, Silver, and Bronze network containers to consume fewer resources on the platforms, compared to VMDC 2.3.
- VMDC 2.3 utilizes the ACE 4710 in One-Arm mode, while VMDC 2.2 uses the ACE30 in Two-Arm mode.

**Note**

---

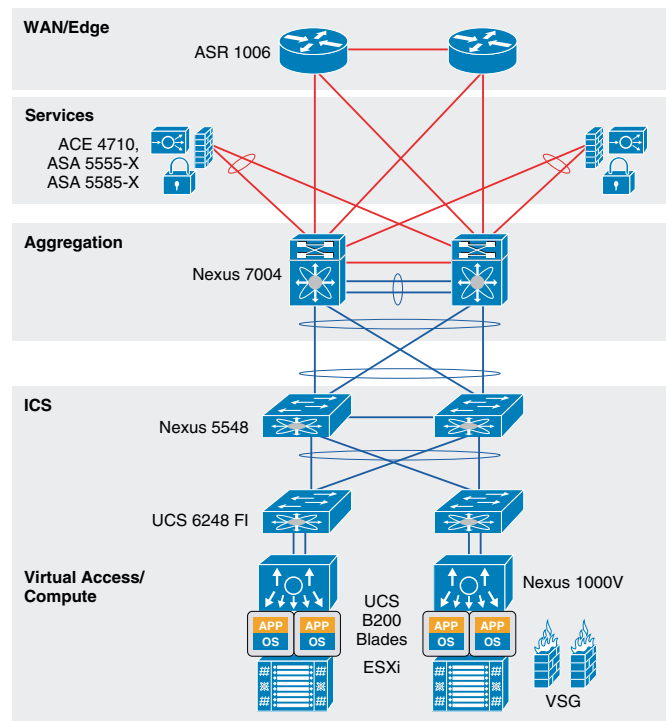
For detailed information on VMDC 2.3 System architecture, refer to the following documents:

- [VMDC 2.3 Design Guide](#)
- [VMDC 2.3 Implementation Guide](#)

For information on the previous VMDC 2.2 System architecture, refer to the following documents:

- [VMDC 2.2 Design Guide](#)
  - [VMDC 2.2 Implementation Guide](#)
- 

[Figure 2-8](#) provides a representation of the VMDC 2.3 physical architecture.

**Figure 2-8 VMDC 2.3 System Architecture****VMDC 2.3 Modular Components**


The VMDC System architecture provides a scalable solution that can address the needs of Enterprise and SP cloud data centers. This architecture enables customers to select the design that best suits their immediate needs while providing a solution that can scale to meet future needs without retooling or redesigning the DC. This scalability is achieved using a hierarchical design with two different modular building blocks, Point of Delivery (PoD), and ICS.

**Point of Delivery (PoD)**

The modular DC design starts with a basic infrastructure module called a PoD. A PoD is a repeatable, physical construct with predictable infrastructure characteristics and deterministic functions. A PoD identifies a modular unit of DC components and enables customers to add network, compute, and storage resources incrementally. This modular architecture provides a predictable set of resource characteristics (network, compute, and storage resource pools, power and space consumption) per unit that are added repeatedly as needed.

In this design, the aggregation layer switch pair, services layer nodes, and one or more Integrated Compute and Storage (ICSs) are contained within a PoD. The PoD connects to the WAN/PE layer device in the DC, in the VMDC 2.3 architecture, and connects to the core layer in previous VMDC 2.2 and 2.0 architectures. To scale a PoD, providers can add additional ICSs and can continue to scale in this manner until the PoD resources are exceeded. To scale the DC, additional PoDs can be deployed and connected to the core layer devices. [Figure 2-9](#) illustrates how PoDs can be used to scale compute, network, and storage in predictable increments within the DC.


Figure 2-9 VMDC 2.3 PoDs for Scaling the Data Center



ICS

The second modular building block utilized is a generic ICS based on existing models, such as the VCE Vblock or Cisco/NetApp FlexPod infrastructure packages. The VMDC architecture is not limited to a specific ICS definition, but can be extended to include other compute and storage stacks. An ICS can include network, compute, and storage resources in a repeatable unit. In this guide, the access layer switch pair, storage, and compute resources are contained within an ICS. To scale a PoD, customers can add additional integrated compute stacks and can continue to scale in this manner until the PoD resources are exceeded. Figure 2-10 illustrates how integrated compute stacks can be used to scale the PoD.

Figure 2-10 VMDC 2.3 ICS for Scaling the Data Center




## VMDC 2.3 Network Containers

The VMDC 2.3 solution defines a reference three-tier Infrastructure as a Service (IaaS) model of Gold, Silver, and Bronze tiers. These service tiers define resource and service levels for compute, storage, and network performance. This is not meant to be a strict definition of resource allocation, but to demonstrate how differentiated service tiers could be built. These are differentiated based on the following features:

- **Network resources.** Differentiation based on network resources and features:
  - **Application tiers.** Service tiers can provide differentiated support for application hosting. In some instances, applications may require several application tiers of VMs (web, application, database). VMDC 2.3 Gold and Silver services are defined with three application tiers on three separate VLANs to host web, application, and database services on different VMs. The Bronze service is defined with one VLAN only so if there are multi-tiered applications, they must reside on the same VLAN or potentially on the same VM (Linux, Apache, MySQL, PHP, Perl, or Python (LAMP)/Windows Apache, MySQL, PHP, Perl or Python (WAMP) stack). All three services, Gold, Silver, and Bronze, are defined with separate VRF instances to provide security and isolation.
  - **Stateful services.** Tenant workloads can also be differentiated by the services applied to each tier. The Gold service is defined with an ASA 5585-X virtual firewall context, ACE 4710 Virtual Server Load Balancer (vSLB) context, and secure remote access (IPSec VPN and SSL-VPN) on the ASA 5555-X. The Silver tier is defined with an ACE vSLB. The Bronze tier is defined with no services on ASA or ACE. All three services include the Nexus 1000V Virtual Security Gateway (VSG) for compute firewall services.
  - **Quality of Service (QoS).** Bandwidth control during periods of network congestion can be a key differentiator. QoS policies can provide different traffic classes to different tenant types and prioritize bandwidth by service tier. The Gold tier supports VoIP/real-time traffic, call signalling and data class, while the Silver, Bronze, and Copper tiers have only data class. Additionally, Gold and Silver tenants are given bandwidth guarantee with Gold getting more bandwidth (2x) than Silver.
- **VM resources.** Service tiers can vary based on the size of specific VM attributes, such as CPU, memory, and storage capacity. The Gold service tier is defined with VM characteristics of four vCPUs and 16 GB memory. The Silver tier is defined with VMs of two vCPUs and 8 GB, while the Bronze tier VMs have one vCPU and 4 GB.
- **Storage resources.** To meet data store protection, RPOs, or RTOs, service tiers can vary based on provided storage features, such as redundant array of independent disks (RAID) levels, disk types and speeds, and backup and snapshot capabilities. The Gold service is defined with 15k FC disks, Silver tier on 10k FC disks, and Bronze tier on SATA disks.

Figure 2-11 shows a representation of a VMDC 2.3 Gold service tier network container.

**Figure 2-11 VMDC 2.3 Expanded Gold Network Container**



The network container is a logical (virtual) segment of the shared (common) physical network resources (end-to-end through the DC) that represents the DC network domain carrying tenant traffic. The physical infrastructure is common to all tenants, but each network device (routers, switches, firewalls, and so forth) is virtualized such that each tenant's virtual network container is overlaid on the common physical network.

The Gold tenant gets two network (and compute/storage) zones to place workloads into. Each zone has its own set of VLANs, VRF instances, and firewall/load balancer contexts. Figure 2-11 shows a logical representation of a two-zone VMDC 2.3 Expanded Gold network container.

This Gold service tier provides the highest level of sophistication by including secure remote access, firewall, and load balancing to the service. The vFW (on the ASA 5585-X60) provides perimeter security services, protecting tenant VMs. The vSLB (ACE 4710 appliance) provides load balancing across VMs in each tier of the tenant. The ASA 5555-X provides virtualized secure remote access (IPsec-VPN and SSL-VPN) to tenant VMs from the Internet. The ACE and ASA service module/appliance are utilized in routed (L3) virtual mode in the VMDC 2.3 design. The Gold service tier also includes the Nexus 1000V VSG for providing virtual security services to the VMs. The Gold service provides higher QoS SLA and three traffic classes - real-time (VoIP), call signaling, and premium data.

The two zones can be used to host different types of applications, to be accessed through different network paths. The two zones are discussed below.

- **PVT Zone:** The Private Zone (PVT) and its VMs can be used for cloud services to be accessed through the customer MPLS-VPN network.
  - The customer sites connect to the provider MPLS-core and the customer has their own MPLS-VPN (Cust-VRF).
  - The VMDC DC ASR 1000 PE connects to the customer sites through the MPLS-VPN (Cust-VRF in Figure 2-11).
  - This Cust-VRF is extended through the VMDC network to the Nexus 7004 aggregation switch.
  - On the agg/access Nexus 7004, the Cust-VRF connects to the ASA Cust-vFW, and then is connected back into a Cust-PVT-VRF on the Nexus 7004 agg/access device (VRF sandwich to insert service nodes), and then to the compute layer on the UCS.



- For the VMDC 2.3 Gold tenant, the PVT zone is defined with three server VLANs.
- In addition, each tenant is assigned a separate Nexus 1000V VSG instance. The tenant is defined as an ORG in the VSG (PNSC), with the three VLANs placed into separate VSG sub-zones.
- The VSG is used to provide security policies to monitor and protect traffic between the VLANs (sub-zones).
- **DMZ Zone:** The VMDC 2.3 Gold container supports a DMZ Zone for tenants to place VMs into a DMZ area, for isolating and securing the DMZ workloads from the PVT workloads, and also to enable users on the Internet to access the DMZ-based cloud services.
  - The ASR 1000 PE WAN router is also connected to the Internet and a shared (common) VRF (usually global routing table) exists for all Gold tenants to connect to (either encrypted or unencrypted).
  - Encrypted (SSL or IPsec Remote Access VPN) traffic is sent to an ASA 5555-X, and based on the VPN policy, is mapped to a particular tenant and the corresponding tenant VPN VLAN.
  - The tenant VPN VLAN then connects to the tenant DMZ-vFW (different vFW context on the ASA 5585-X than the tenant PVT-vFW), then to the tenant DMZ-VRF (different VRF on the Nexus 7004 agg/access than the tenant PVT-VRF), and then to the Compute layer for the DMZ Zone.
  - Similarly, unencrypted traffic from the Internet, based on the destination VM/VIP address, is sent to the tenant DMZ-vFW, then to the DMZ-vSLB, DMZ-VRF, and the DMZ Compute Zone.
  - The DMZ Zone can be used to host applications like proxy servers, Internet-facing web servers, email servers, etc. The DMZ Zone consists of one server VLAN in this implementation.

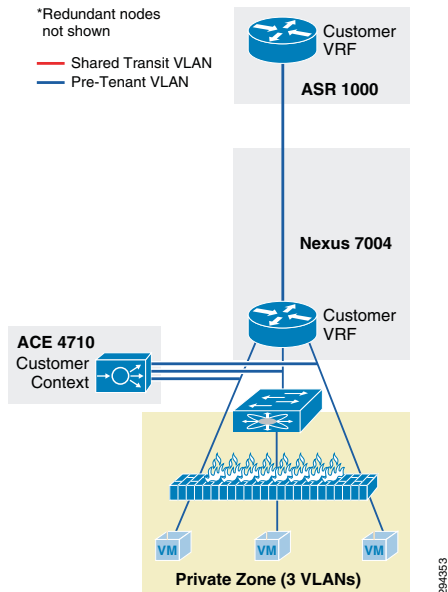
In VMDC 2.3, a Gold tenant can choose to have only the PVT Zone, or both the PVT and DMZ Zones. If the tenant has both PVT and DMZ Zones, then the Gold tenant will consume three VRF instances (Cust, Cust-PVT, and Cust-DMZ) on the Nexus 7004 Agg, two VFW instances, two vSLB instances, two VSGs, and four server VLANs. To facilitate traffic flows between the DMZ and PVT Zones (for example, proxy or web servers in the DMZ Zone, application and database servers in the PVT Zone), the DMZ-vFW and PVT-vFW are interconnected. Configuring appropriate security policies (routing, NAT, firewall rule, ACLs) on the DMZ-vFW and PVT-vFW can allow or disallow communication between the two zones.

Load-balanced traffic for all tiers of Gold tenants is implemented using the ACE 4710, which has one interface in each of the tiers.

- MPLS-VPN to PVT Zone
- Unsecured (clear) Internet to DMZ Zone
- Secure (Remote Access SSL/IPsec VPN) Internet to DMZ Zone
- DMZ to PVT Zone
- MPLS-VPN to DMZ Zone
- PVT to Internet Zone is via an HTTP proxy hosted in the DMZ Zone

Figure 2-12 is a representation of a VMDC 2.3 Silver network container.

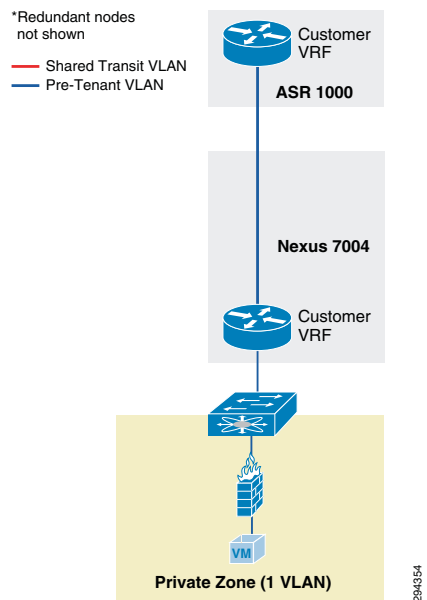
**Figure 2-12 VMDC 2.3 Silver Network Container**



The Silver service tier includes one VRF instance per Silver tenant and three server VLANs (three-tiered applications) for each tenant. The Silver service includes a load-balancing service for more sophistication over the Bronze tier. The vLB (ACE 4710 appliance) provides load balancing across VMs in each tier of the tenant. The ACE service load balancer is utilized in one arm, routed (L3), virtual mode in the VMDC 2.3 design, and one context is used per Silver tenant. The context has links on each of the server VLANs and works in one-arm mode. The Silver service tier also includes the Nexus 1000V VSG to provide virtual security services to the VMs. The Silver service provides medium QoS SLA and one traffic class, premium data.

Figure 2-13 is a representation of a VMDC 2.3 Bronze network container.

**Figure 2-13 VMDC 2.3 Bronze Network Container**



The Bronze service tier includes one VRF instance and one server VLAN for each tenant. The Bronze service is the least sophisticated tier and does not include any perimeter security services. The Bronze service tier does include the Nexus 1000V VSG for providing virtual security services to the VMs. The Bronze service provides lower QoS SLA and one traffic class, standard data.

**Note**

Additionally, VMDC 2.3 also defines a Copper network container, which has the similar characteristics as Bronze, but has only Internet-based access and no L3VPN-based access. The Copper container also uses a shared perimeter firewall (ASA vFW context) for all tenants. However, the VMDC 2.3 Copper network container has not been validated with the DRaaS System.

## Modifications in VMDC Network Containers for DRaaS

The VMDC 2.3-based infrastructure and Gold, Silver, or Bronze network containers (specific container used by a tenant based on FW, SLB services needed) can be used for DR services, but the following modifications need to be made:

- Utilize a new ASA context per tenant for IPsec-VPN services to encrypt the communication between InMage control servers in the DR site and the Enterprise site. This ASA context needs to be added whether the tenant is using Gold, Silver or Bronze container on the DR site. This context will logically reside close to the server VLANs.

**Note**

In the case of Silver or Bronze VMDC containers, no existing ASA context is being used in the network container for firewall or VPN services. Therefore inserting this ASA context for InMage VPN purposes will be a new addition to the network container. In the case of the VMDC Gold container, an ASA context (on the multi-context ASA5585X) is utilized for perimeter firewall services, and a shared ASA (single-context ASA5555) is utilized for remote access VPN purposes. However, these existing ASA contexts in the VMDC Gold container cannot be used for the InMage VPN purposes since they logically sit in a different part of the network container. This new ASA context for the tenant can be created on the existing ASA5585-FW device (if enough capacity for contexts and throughput exists) or a new ASA device can be utilized. It is recommended to use a new physical ASA device (ASA5555 or ASA5585 based on VPN throughput needed) for the InMage VPN purposes. Thus, the VMDC 2.3 infrastructure for DRaaS would have three separate physical ASA devices: one each for FW, RA-VPN, and one for Inmage Site-Site VPN. The VMDC 2.3 Gold container for DRaaS would have three logical ASA devices: one per-tenant context for FW, one shared/global ASA for RA-VPN, and one per-tenant context for InMage Site-Site VPN

- In the case of Gold containers, the tenant ASA context performing perimeter firewall services needs to have a security policy (ACL) configured to permit the IPsec-VPN traffic from the ENT site to the DR site. This ACL should be specific to allow IPsec traffic only between the IPsec tunnel endpoints (local ASA Site-Site VPN endpoint in the DR site, and remote VPN endpoint in the ENT site) used to encrypt the InMage traffic.
- Create a new VLAN for Bronze container to host the InMage control servers. To insert an ASA context to encrypt the InMage traffic, we need to create an outside and inside interface on the ASA context. Since the VMDC 2.3 Bronze container is defined with only one VLAN, we need to define a new VLAN to host the InMage servers. The first VLAN (where recovered VMs will be placed) will serve as the outside interface for the ASA VPN context and the second (new) VLAN (where the InMage servers are placed) will serve as the inside interface for the ASA VPN context.

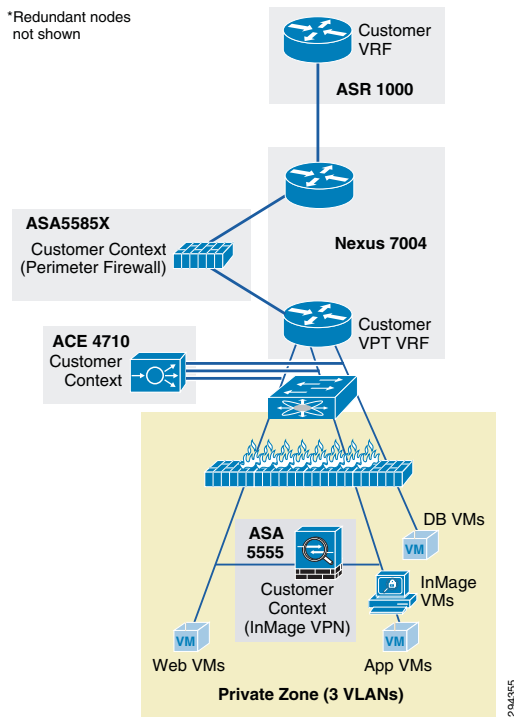


**Note**

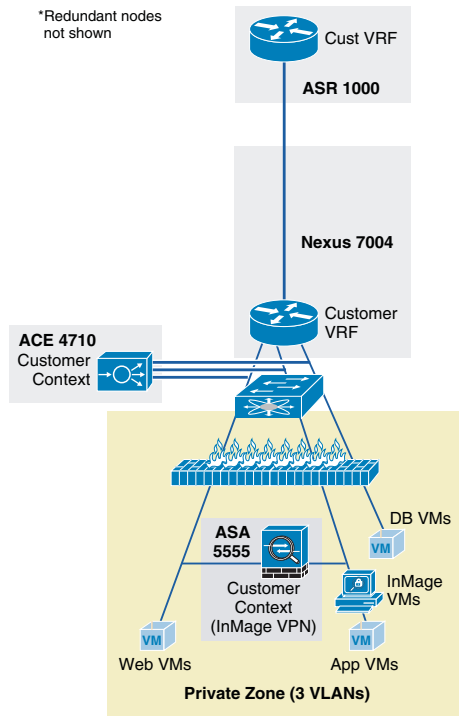
For the VMDC 2.3 Gold and Silver containers, three server VLANs already support three-tier applications so there is no need to create a new VLAN to host the InMage servers or for InMage VPN purposes. Instead, the InMage servers can be hosted in the second VLAN (App tier). The first VLAN (Web tier, where recovered Web-tier VMs will be placed) will serve as the outside interface for the ASA VPN context, and the second VLAN (App tier, where the recovered App-tier VMs and also the InMage servers will be placed) will serve as the inside interface for the ASA VPN context.

Figure 2-14, Figure 2-15, and Figure 2-16 show logical representations of the modified VMDC 2.3 Gold, Silver, and Bronze network containers for DRaaS.


**Figure 2-14 Modified VMDC 2.3 Gold Container for DRaaS**



**Figure 2-15 Modified VMDC 2.3 Silver Container for DRaaS**



**Figure 2-16 Modified VMDC 2.3 Bronze Container for DRaaS**



# VMDC Orchestration using BMC CLM

The Cisco-BMC cloud management architecture for VMDC is designed to meet the growing needs of today's data center and cloud deployments. BMC Cloud Lifecycle Management (CLM) provides an end-to-end automated lifecycle management solution for cloud-based IT hosting environments.

The architecture focuses on the planning, governance, provisioning, operation, administration, and maintenance of cloud services, the runtime environments and infrastructure resources needed to sustain them, and the management services that comprise CLM.

The VMDC 2.3 architecture and network containers have been validated to be orchestrated by CLM 3.1 Service Pack 1 (SP1). CLM 3.1 SP1 includes all of the elements that are essential to enabling a VMDC 2.3-based cloud environment:

- **Self-service Portal and Service Catalog.** Provides the ability to order and track deployed services.
- **Service delivery automation.** Automates provisioning of services. CLM can also provide usage metering of services, by using additional BMC components.
- **Resource management.** Provisions and manages resources as per-service needs. This includes network, compute, and storage resources.
- **Operational process automation.** Automates operational processes such as user management, service desk integration, and alerting. Capacity management and service level management can also be provided by additional BMC components like BMC Capacity Optimization (BCO) and BMC ProactiveNet Performance Management (BPPM).

CLM 3.1 SP1 enables onboarding and pooling of resources for compute, storage, and networking, and creation of policies to manage those pools. It provides functionality to provision network containers, physical servers, and virtual server instances. It also provides the ability for end users, through a portal, to place service requests to create and manage server instances. CLM 3.1 SP1 is fully multi-tenant/multi-service aware. It can support simultaneous use of the cloud environment by multiple tenants that can request, deploy, and operate services independently.



## Note

For detailed information on using BMC CLM 3.1 SP1 for orchestrating VMDC 2.3 architecture, refer to the following document:

- [Orchestrating VMDC 2.3 with BMC CLM 3.1 SP1 Design & Implementation Guide](#)

This section includes the following topics:

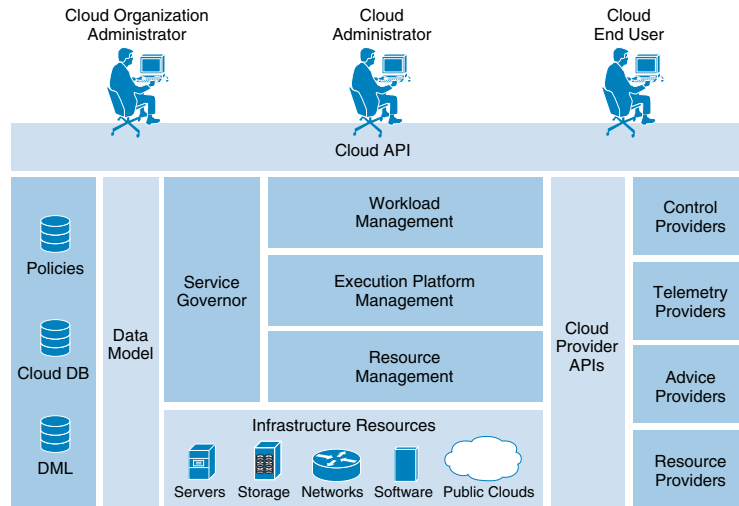
- [CLM 3.1 SP1 Architecture, page 2-20](#)
- [Container Pre-Provisioning for DR Services, page 2-24](#)

## CLM 3.1 SP1 Architecture

CLM 3.1 SP1 is a general-purpose, one-size-fits-all management solution for cloud hosting environments. CLM 3.1 SP1 can manage environments that reside entirely on-premise or off-premise and hybrid environments that are hosted partially on-premise and off premise. CLM 3.1 SP1 can manage hosting environments that use physical or virtual compute, storage, and network resources. CLM 3.1 SP1 can now manage multi-hypervisor environments that include Microsoft Hyper-V and VMware vSphere. It can also manage environments that use cloud resources, including resources and services offered by other Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) clouds.

CLM 3.1 SP1 is fully multi-tenant aware. It can support simultaneous use of the cloud by multiple tenants that can request, deploy, and operate multiple services independently. CLM 3.1 SP1 has an architecture that provides the foundation for scaling the cloud and for configuring multiple data centers.

**Figure 2-17 CLM 3.1 SP1 Architecture**



### User Roles

CLM 3.1 SP1 supports three different classes of users:

- **Cloud Administrator.** A Cloud Administrator is an IT professional responsible for the full lifecycle of the cloud environment, including initial planning, deployment and configuration, and continued administration, operation, and maintenance. The Cloud Administrator uses the Administration console for most tasks.
- **Cloud Organization (Tenant) Administrator.** A Cloud Organization (Tenant) Administrator is responsible for managing a subset of the cloud that is provisioned for a particular Tenant or Service.
- **Cloud End User.** Cloud End Users request services made available to them by the Cloud Administrator through the BMC My Services console. Cloud End Users can request virtual as well as physical resources, view and manage their commissioned resources, monitor the health of their commissioned resources, and decommission resources.

### Consoles

CLM 3.1 SP1 has three consoles that cloud users can use to manage the VMDC cloud infrastructure:

- **BMC CLM My Cloud Services console.** This enables users and administrators to request, deploy, and operate Service Offerings from the Service Catalog.
- **BMC CLM Administration console.** This enables Cloud Administrators to manage the cloud and the services that it hosts.
- **BMC Tenant Administration console.** This is an enhanced My Cloud Services console. This console offers additional administration capabilities to Tenant Admins, such as, capabilities around network management.

### Service Catalog

The Service Catalog contains the Service Offerings that are available for consumption by cloud users. Cloud Administrators maintain the Service Catalog by creating, modifying, and deleting Service Offerings. They can also control which offerings in the Service Catalog are available to each Tenant.

### Cloud Database

The Cloud DB contains operational state and configuration information about the objects managed by the cloud. These managed objects include the Service Offering Instance (SOI), virtual cloud resources, and physical and virtual infrastructure resources.

### Product Catalog and Definitive Media Library

The Product Catalog and Definitive Media Library (DML) list all software that can be provisioned in the cloud. The Product Catalog does not store the software itself. Instead, it contains a unique reference to each piece of software, while the software itself remains in native repositories such as the BMC Server Automation (BSA) software repository or the Amazon AWS Amazon Machine Images (AMI) repository. The Product Catalog also contains software metadata, such as software and hardware requirements pertaining to software provisioning, as well as other data used during software configuration. Cloud Administrators create and maintain entries in the Product Catalog by using interfaces provided by the Product Catalog.

### Cloud Blueprints

Cloud blueprints define cloud services and resources that can be provisioned in the VMDC infrastructure. CLM 3.1 SP1 uses the following cloud blueprints:

- **Service blueprints** describe the functional structure of a given Service Offering, including its functional components and communication paths. They also define how a Service Offering is to be deployed under different circumstances. Each Service Offering in the Service Catalog has a Service Blueprint that is used for its instantiation. When creating a Service Blueprint, the user can define the service and how it is deployed:
  - *Service definitions* of applications or server instances specify the topology (number of tiers), configuration, operating systems, and software packages that need to be provisioned to "stand up" an application or server.
  - *Service deployment definitions* for each Service Blueprint specify a set of one or more ways in which the blueprint could be instantiated when it is provisioned.

For example, in a blueprint for an application, one service is related to three deployments, Small, Medium, and Large, that are mapped to a Service Offering in the Service Catalog. The Small deployment definition for the application might use a single Resource Set that consists of one VM to support all three tiers, web, business logic, and database. In contrast, the Large deployment definition might distribute the application component to three different Resource Sets, each corresponding to a different application tier.

- **PoD blueprints** define the physical topology of the VMDC infrastructure.
- **Network container blueprints** define the logical segmentation of the VMDC cloud infrastructure.

### Infrastructure Resources

Cloud infrastructure resources represent physical or virtual DC resources that host Service Offerings in the cloud. All compute, storage, network, and software resources that are part of the VMDC infrastructure are considered to be infrastructure resources. In the VMDC 2.3 system, the following components comprise the infrastructure resources:

- ASR 1006 (WAN Router)



- Nexus 7004 (DC Aggregation layer)
- Nexus 5548 (DC Access layer)
- ACE 4710 (Server Load Balancer appliance)
- ASA 5585-X (Firewall appliance)
- UCS B-series blade and C-series rack servers
- UCS 6248 (Fabric Interconnect)
- Nexus 1000V (Distributed Virtual Switch)
- Virtual Security Gateway (Compute Firewall)
- NetApp FAS and EMC VMAX, VNX (NAS and SAN devices)
- VMware Virtual Center
- VMware ESXi clusters and hosts
- Microsoft Hyper-V 2012
- Microsoft System Center Virtual Machine Manager 2012

### Cloud Providers

Cloud providers are software programs that act as element managers for different types of resources, platforms, and services consumed by the cloud. At installation, CLM 3.1 SP1 includes the following providers:

- **BMC Server Automation (BSA)**. BSA is a control and resource provider for various types of infrastructure compute resources, such as physical servers, VMs, VM clusters, and virtual cluster resource pools.
- **BMC Network Automation (BNA)**. BNA is a control and resource provider for network resource providers, such as IP addresses, routers, firewalls, load balancers, and VLANs
- **Amazon EC2**. This provider facilitates integration with the Amazon EC2 cloud.
- **vCloud Director**. This provider facilitates integration with the VMware vCloud Director-based cloud environment.

### Cloud Workload Manager

The cloud Workload Manager (WLM) instantiates Service Offerings selected from the Service Catalog. It also administers and maintains those services based on policies defined in the cloud policy DB.

### Cloud Platform Manager

The cloud Platform Manager provisions, operates, administers, and maintains runtime platform instances.

### Cloud Resource Manager

The cloud Resource Manager (RM) manages cloud infrastructure resources, including onboarding, Organization, assignment, and allocation.

### Cloud Service Governor

The cloud Service Governor orchestrates workload management, platform management, and resource management operations based on policies and rules defined by Cloud Administrators to provision hosted Service Offerings. The Service Governor makes placement decisions on where resources need to be

placed, based on resource management/tracking and service policies/tagging. The cloud Service Governor also distributes and subsequently enforces operational policies across CLM 3.1 SP1 components.

**CLM 3.1 SP1 Northbound API**


The solution architecture also defines a set of APIs that can be used by third-party application developers to enhance the CLM 3.1 SP1 implementation. The API can be used to do the following tasks:

- Integrate with CLM 3.1 SP1
- Automate repetitive processes
- Create a customized portal in CLM 3.1 SP1

The API is a model-based, object-oriented RESTful web service that features an easy-to-use interface facilitated by standard HTTP requests and response messages that carry JSON-format payloads.

Figure 2-18 illustrates the BMC CLM 3.1 components and interactions with infrastructure components.

**Figure 2-18 CLM 3.1 SP1 Components and Interactions**



## Container Pre-Provisioning for DR Services

To enable the DRaaS, the VMDC-based network containers need to be pre-provisioned in the DR site. Depending on the services (firewalls, load balancers) needed, the DR customer can choose to use the VMDC 2.3 Gold, Silver, or Bronze Network containers. BMC CLM can be used to automatically provision the VMDC Gold, Silver, or Bronze network containers on a VMDC 2.3 infrastructure. This capability for orchestrating VMDC 2.3 Gold, Silver, or Bronze network containers is available out-of-the-box in BMC CLM 3.1 SP1.



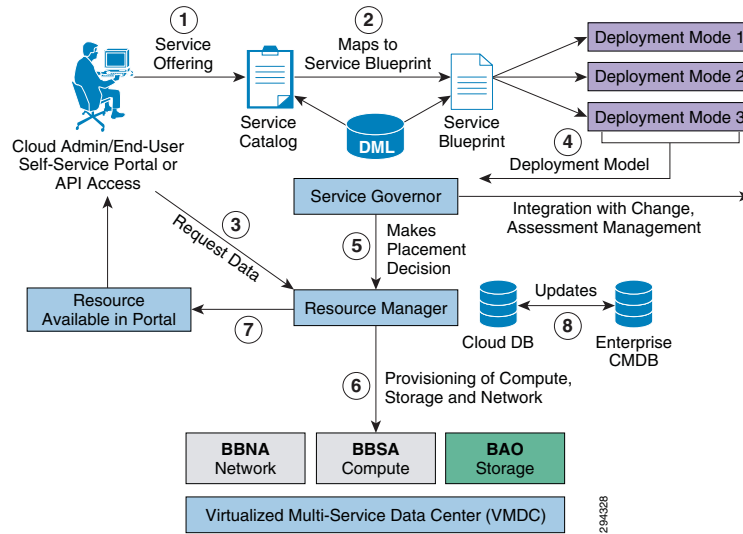
**Note**

For more information on using BMC CLM 3.1 SP1 to orchestrate VMDC 2.3 network containers, refer to the following document:

[Cloud Orchestration for VMDC with BMC Cloud Lifecycle Management 3.1 SP1 Design and Implementation Guide](#)

Figure 2-19 illustrates the high-level CLM orchestration workflow.

Figure 2-19 CLM 3.1 SP1 End-to-End New Request Flow



The workflow steps needed for container pre-provisioning and preparing the DR service are listed:

- Step 1** VMDC 2.3-based infrastructure, topology, and base infrastructure configurations are created in the DR site.
- Step 2** BMC CLM 3.1 SP1 are installed in the DR site. The VMDC 2.3 PoD and Network Container Blueprints are on-boarded into CLM. The CLM portal is set up and running.
- Step 3** Tenant (ENT customer requiring DR services) uses the BMC CLM portal to request the Network Container from the Service Catalog. The tenant can choose Gold, Silver, or Bronze network container.
- Step 4** CLM Cloud Admin (DR Provider) approves the change request, and the requested Network Container for DR services is created on the VMDC infrastructure for the tenant. Appropriate Compute and Storage resources are set aside for the Tenant.
- Step 5** Necessary modifications are made in the VMDC network container to facilitate the DR service are made. This includes setting up any additional VLANs, creating the ASA context for IPsec VPN tunnels (for secure InMage communication from customer site to DR site, etc., as documented in VMDC container. Adding VLANs can be done by the Cloud Admin through the CLM portal. Creating the ASA context for VPN, and connecting to tenant's VMDC network container and VLANs has to be done manually by the DR site network administrator.
  - a. Install the InMage control plane servers into the tenant's VMDC network container in the appropriate VLANs. This task can be done through CLM by the CLM Cloud Admin, if the InMage applications are created as Service Offerings in the CLM Service Catalog. Otherwise, these InMage applications have to be manually installed by the DR site server administrator.
  - b. Set up the InMage DR components, storage and other necessary steps as documented in Chapter 3 Implementation and Configuration. These InMage DR components have to be set up at both the Enterprise and Provider DR sites.
  - c. Verify IPsec connectivity between the DR components in the Enterprise and SP DR sites.
  - d. Verify DR functionality and create protection and recovery plans through InMage.

As part of the VMDC network container creation, BMC CLM also creates the firewall and load balancer services for the relevant containers. For example, in the VMDC 2.3 Gold container, perimeter FW services are provided by an ASA context, load balancing services are provided by an ACE context, and back-end or compute FW services are optionally provided by a VSG. When BMC CLM is used to create a VMDC 2.3 Gold network container for DR services, the ASA context, ACE context, and VSG are created for the tenant container. CLM also provisions some base security rules on the ASA and VSG. In addition, the CLM portal can be used to provision more specific FW security rules through network paths. These can be done by the Cloud Admin or by the Tenant Admin, and the security policies can be done at a VLAN (IP subnet and protocol/port) level or at a VM (specific IP address and protocol.port) level. In addition, CLM portal can be used (by Cloud Admin or Tenant Admin) to create load balancing virtual servers (VIPs) for specific protocols on the tenant ACE context, and associate servers or server pools (VMs) to the VIPs.

In the DRaaS scenario, after a failure, the recovered VMs will be brought online in the VMDC network container in the provider DR site. As such, these VMs need to have associated security policies on the ASA (and VSG if tenant is using compute FW), and associated load balancing policies on the ACE context. The following steps need to be followed to accomplish this once the workflow steps described above for creating the DR service have been completed.

---

**Step 1** Network Container Pre-provisioning

- a. Follow the workflow steps described above to pre-provision the tenant network container through CLM in the DR site, install the DR components on the ENT and DR sites, and identify VMs to be protected/recovered.

**Step 2** vFW Policies

- a. Use the CLM portal to create Network Paths to allow/deny specific traffic flows to/from the VMs that will be recovered on the DR site. These Network Paths will be translated into appropriate ASA and VSG security policies and provisioned on the tenant ASA and VSG by CLM.
- b. The Network Paths and security policies can be based on VM subnets or specific VM IPs. It is recommended to use security policies based on VM subnets, so as to minimize the changes necessary when additional VMs get added into the protection plans (within the same VLAN or subnet).
- c. It is recommended to configure these Network Paths and underlying security policies before DR declaration, so that all services are in place when VMs get recovered to the DR site. Post-recovery, the CLM portal can again be used to tweak the Network Paths and security policies as needed.

**Step 3** vSLB Policies

- a. Use the CLM portal to create Virtual Servers (VIP) and associate to real servers (VMs), and define the protocols to be load balanced and the mechanisms for probing the real servers. CLM will provision these policies on the tenant ACE context. The real servers have to reflect the specific IP addresses to which the VMs to be load balanced.
  - b. It is recommended to configure these SLB policies before DR declaration, so that all services are in place when VMs get recovered to the DR site. Post-recovery, the CLM portal can again be used to tweak the load balancing policies as needed.
- 

## Available Replication Types


This section includes the following topics:

- [Hypervisor-based vs. Guest OS vs. Storage](#)
- [InMage Software Architecture, page 2-28](#)

## Hypervisor-based vs. Guest OS vs. Storage

Figure 2-20 shows the different types of replication technologies that can be used for disaster recovery purposes.

**Figure 2-20 Disaster Recovery Replication Technologies**



### Storage Array Level Replication

The most popular replication method used by most of the organizations today is storage array-level replication. Array-based replication is expensive and lacks granularity. You need to purchase from a single storage vendor the exact type, brand, and model number of a storage array on both the source and target side of your DR solution. You need to budget for exactly the same storage class and tier. One of those storage arrays will stay dormant until a recovery situation requires it to be active. An array-based solution typically replicates an entire volume even if there is only one VM in the volume that needs to be replicated. It does not provide the flexibility of replicating a single VM. It also requires multiple points of management while performing disaster recovery tasks and needs a separate run book management tool along with the storage array management console.

### Hypervisor-Based Replication

Hypervisor-based replication is a good option for organizations who has all of their environment virtualized. The agent that captures the changes on the production servers sits at the hypervisor layer. Since hypervisor-based replication is "VM-aware," it is possible to select the VMs that need to be replicated, while saving storage space at the secondary site by avoiding replicating the ones that don't. Hypervisor-based replication allows you to be much more granular in what you protect, and it also allows you to group VMs by defining protection groups. And it can be managed from virtualization management suites like VMware's vCenter or Microsoft's System Center. The main limitation of hypervisor-based replication is that it's specific to a hypervisor and using the same solution physical environments cannot be protected.

### Guest OS/Host Based Replication

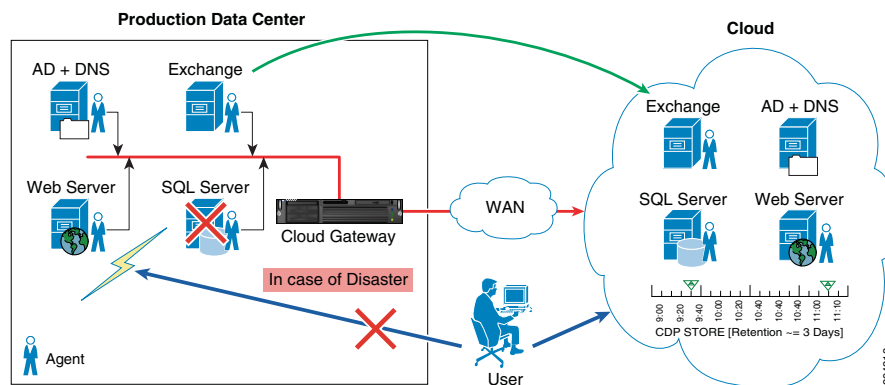
Many enterprises use host-based replication because it is relatively inexpensive. The process involves installing a replication agent onto the operating systems of the servers to be replicated. This agent processes and replicates I/O traffic on any storage systems (NAS, DAS, SAN, etc.) to a secondary

replication target system, which use storage of any type, from any vendor. It saves money compared to array-based replication because licensing host-based replication software is much less expensive than for most array-based replication systems. Also, there's no need to go to the expense of purchasing a second storage array that's identical to the primary one. SPs can deploy any storage type in their cloud while offering DRaaS. This allows them to offer DRaaS to customers using any storage and infrastructure.

Though many organizations are embracing virtualization, most organizations are still not 100% virtualized and still have critical and legacy applications running on physical environments. Using host-based replication, both physical and virtual environments can be protected and the solution is agnostic to the server, operating system, and storage. The DRaaS system uses host-based replication for its simplicity and for providing greater coverage of protecting physical and virtual environments.

Figure 2-21 shows OS/Host-based replication in the DRaaS System.

**Figure 2-21 OS/Host-Based Replication**



## InMage Software Architecture

This section includes the following topics:

- [InMage Overview](#), page 2-28
- [Components](#), page 2-31
- [How InMage ScoutCloud Works](#), page 2-37
- [Component Flow](#), page 2-38

## InMage Overview

### InMage ScoutCloud Enables Recovery as a Service

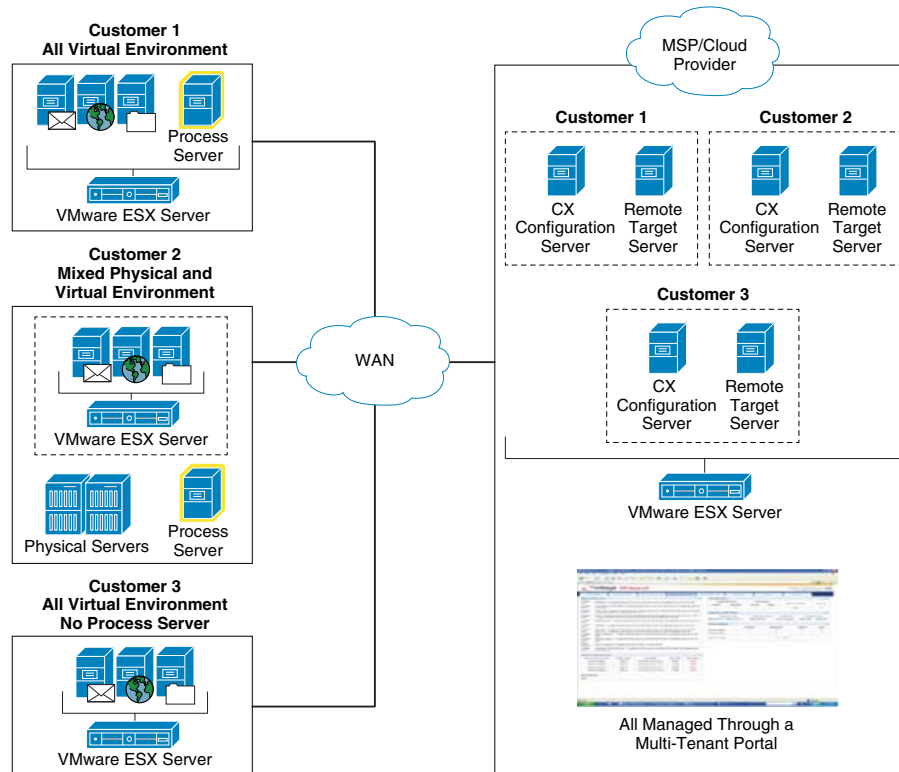
The InMage ScoutCloud platform addresses the growing market for cloud-based disaster recovery products, also referred to as the Recovery as a Service (RaaS) market. InMage ScoutCloud leverages next generation recovery technologies including disk-based recovery, CDP, application snapshot API integration, asynchronous replication, application awareness, and WAN optimization. These next generation recovery technologies are wrapped up in a single product offering, enabling MSPs and cloud providers to have the fastest time-to-market when offering customers a near zero RPO and RTOcapable RaaS with:

- Best-in-class data protection.

- A comprehensive P2V and V2V recovery engine that supports all applications.
- A provisioning manager that automates provisioning of recovery for VMs and associated storage combined with a full-fledged multi-tenant portal.

Figure 2-22 shows the InMage ScoutCloud architecture in a DRaaS environment.

**Figure 2-22 InMage ScoutCloud Architecture**



### InMage ScoutCloud Concepts

**Continuous Data Protection (CDP):** CDP refers to a technology that continuously captures or tracks data modifications by saving a copy of every change made to your data, essentially capturing every version of the data that you save. It allows you to restore data to any point in time. It captures the changes to data and sends them to a separate location. CDP-based solutions can provide fine granularities of restorable objects ranging from crash-consistent images to logical objects such as files, mail boxes, messages, and database files and logs.

Traditional backups require a schedule and restore data to the point at which it was backed up. CDP does not need a schedule because all the data changes on the primary server are tracked and sent to a secondary server asynchronously.

Most CDP solutions save byte or block-level differences rather than file-level differences. This means that if you change one byte of a 100 GB file, only the changed byte or block is saved. CDP technology has the following fundamental attributes:

- Data changes of primary server are continuously captured or tracked.
- All data changes are stored in a separately located secondary server.
- It enables data recovery in much lesser time as compared to tape backup or archives.

**Disaster Recovery (DR):** DR is the process of preparing for recovery or continuation of technology infrastructure critical to an organization after a natural or human-induced disaster. DR solution using CDP technology replicates your data to a separately located secondary server. In case of disaster, you can get immediate access to a primary server's data, which is up-to-the minute of disaster.

**Application Protection Plan:** An efficient Application Protection Plan can protect customer's critical applications from natural as well as human-interfered disaster. Every individual application of an organization should have a unique protection plan where the application can have single or multiple protections; i.e., the application can be protected locally for backup purpose or it can be protected to remote locations for DR purposes.

**Replication Stages:** InMage ScoutCloud replicates drive level data in three stages:

- **Resyncing (Step I):** In this step, data at the primary server is replicated to the secondary server. This is done only once for each drives that you want to replicate to a secondary server drive.
- **Resyncing (Step II):** All data changes during Resyncing (Step I) are replicated to the secondary server in this step.
- **Differential Sync:** Differential Sync is a continuous process where any change in the primary server volume is copied to the Secondary server volume simultaneously.

**Consistent Data:** In case of DR, the restored data should be consistent with the original data. To ensure the consistency of backup data, the consistent tags/bookmarks are issued at the primary server at periodic intervals of time or on demand.

**Journal/Retention or CDP Logs:** The retention or CDP logs store information about data changes on primary server within a specified time period on a separately located secondary server. This timeframe is referred to as the retention window. Consistent points are stored as bookmarks/tags in retention window. An application can be rolled back to one of the bookmarks/tags in this retention window. Alternately, an application can be rolled back to any point in time of this retention window. Applications that are rolled back to any of the bookmarks/tags in this retention window will only be consistent. Three types of retention policy are associated with this retention window:

- **Time-based:** The data in the retention window will be overwritten after the specified time period.
- **Space-based:** The data in the retention window will be overwritten once the size is exhausted.
- **Time and space-based:** The data in the retention window will be overwritten once the time specified or space specified qualifies first.

**Sparse Retention:** For long term data retention purposes, the sparse policy is used, which helps to save disk space on retention volumes and makes it possible to afford a wider retention window. Depending on the type of policy enforced, the retention window is maintained by discarding older data changes within the retention log files to make rooms for new data changes.

**Failover:** This is the process of switching production server to secondary server. The failover process can be a planned or an un-planned operation. The planned failover is used for periodic maintenance or software upgrades of primary servers wherein the data writes to primary server are stopped. An unplanned failover happens in case of actual failure of the primary server.

**Failback:** This is the process of restoring the primary server from the secondary server after a planned or un-planned failover. A failover operation is usually followed by a failback operation. In this failback process, the data writes on the secondary server are also restored to the primary server. Scout also supports fast failback where the data changes of the secondary server are not applied to the primary server while restoring.

**Snapshot:** A snapshot is an exact replica of a primary server's data as it existed at a single point in time in retention window. The two types of snapshot are Physical Snapshot and Virtual Snapshot:

- For Physical Snapshot, you can take a snapshot on a physical volume. It requires the intended snapshot volume to be equal or larger than the Secondary server volume (in the replication pair).



- For Virtual Snapshot, you can take a snapshot on a virtual volume. It is also known as "vsnap," which requires minimal system resources and are faster in loading or unloading. These snapshots can be accessed in one of following modes:
  - Read-Only: As the name indicates, read only snapshots are for informative purposes and are not capable of retaining writes on to them.
  - Read-Write: Read/write virtual snapshots retains writes on to them; this is done by maintaining an archive log on some part of the local disk as specified.
  - Read-Write Tracking: Read/write tracking 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.

**Application Consistency:** Application Consistency ensures the usability of the application when DR copies of the application's primary server data are used in place of the original data. An application can be rolled back to any bookmark/tag in the retention window. Consistency bookmarks are of the following three types:

**Application bookmarks:** This bookmark ensures consistency at the application level. This is issued after flushing the application buffers to the disk.

**File System bookmarks:** This bookmark ensures consistency of the data at the file system level. This is issued after flushing the file system cache to the disk.

**User-defined bookmarks:** This is a user-defined name for a bookmark which is associated with application bookmark or a file system bookmark or both. These are human readable bookmarks unlike the application or file system bookmarks, which are used by the DR administrators to recover the data.

## Components

This section includes the following topics:

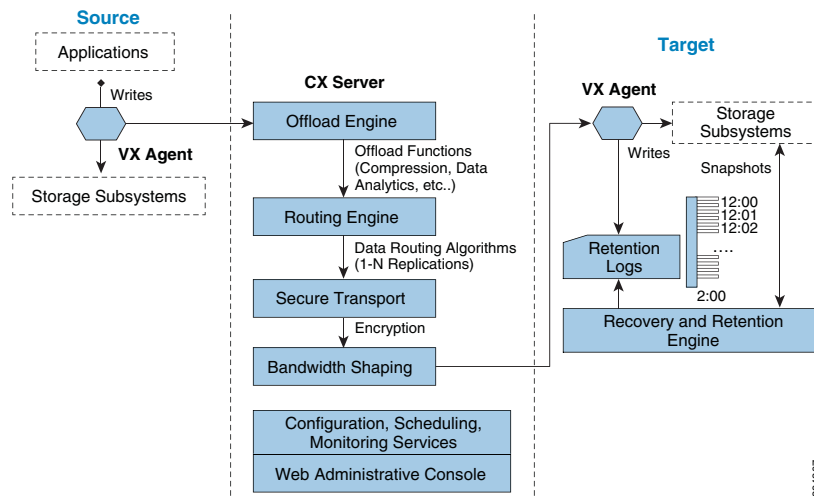
- [Unified Agent, page 2-31](#)
- [Master Target, page 2-33](#)
- [CX Server, page 2-33](#)
- [RX Server, page 2-35](#)
- [Management Console, page 2-36](#)
- [Self Service, page 2-37](#)

### Unified Agent

Unified Agent (aka VX Agent) is a lightweight agent that is installed on to each VM or physical server protected. It offloads the data changes to the CX appliance. Unified Agent is installed automatically by the vContinuum wizard.

[Figure 2-23](#) shows the VX Agent Theory of Operation.

Figure 2-23 VX Agent Theory of Operation




### VX Agent - Responsible for Volume Level Replication and CDP

- Sentinel: The Sentinel software is installed on the protected hosts. It is responsible for keeping track of the data changes that occur. It is also necessary for the Initial Sync and Resync processes.
- Outpost Agent: The Outpost Agent software is installed on the CDP/DR hosts. It is responsible for keeping the CDP/DR host volumes in sync with the replicated host volumes. It is necessary for the Initial Sync and Resync processes. It includes the snapshot functionality.

Figure 2-24 shows the VX Agent Architecture.

Figure 2-24 VX Agent Architecture



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.

### VX Components

Once the VX Agent is installed, a service named "svagents" is created. Svagents (Svagents.exe) is the Windows service that is responsible for launching and managing all other user space components of the VX Agent. This service runs two threads "dataprotection.exe" and "s2.exe":

- 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.

## Master Target

A dedicated VM created on secondary vSphere server to act as a target for replication is called master target (MT). It is used as a target for replicating disks from primary VMs and MT contains the retention (CDP) data. Retention data is the log of prior changes using which you can recover a VM to prior point in time or to a prior application consistent point.

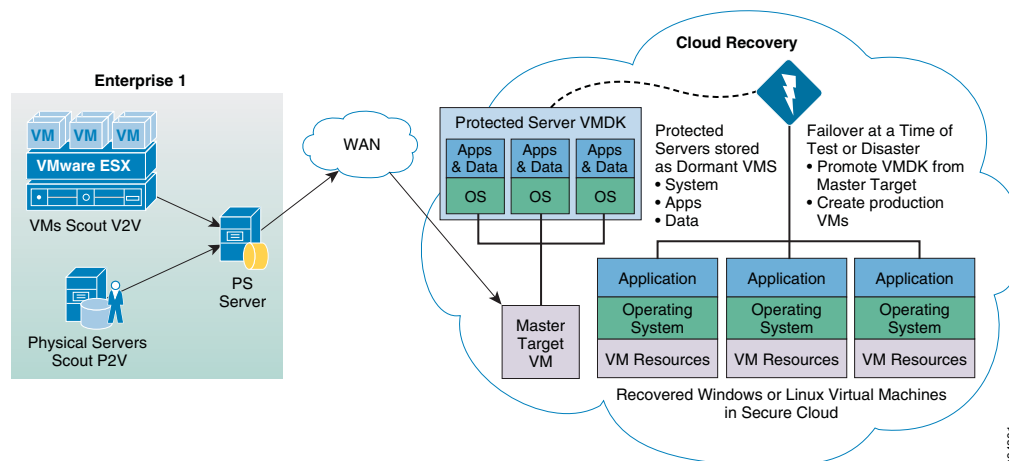
When the initial replication plan is set up for a server or group of servers. The data volumes on the production servers are created as VMDKs on the target site and get mounted to the MT server for writing data. In the event of a disaster, the VMDKs gets released by the MT server and will get mounted to the actual recovery servers.

The MT should be of the same OS family as that of primary servers. If primary VMs are Windows, MT has to be Windows. For Linux primary VMs, MT must be a Linux VM.

Win2k8R2 is recommended to protect Windows VMs. You can have more than one master target on secondary vSphere servers. To perform failback protection and failback recovery, a MT VM is required on the primary vSphere server. In case of failback, replication is set in reverse direction from recovered secondary VMs back to MT on the primary vSphere server.

Figure 2-25 shows the MT functionality.

**Figure 2-25 Master Target Functionality**



## CX Server

CX Server is the combination of Configuration Server (CX-CS) and Process Server (CX-PS).

The Scout PS/CS server is an appliance that performs the data movement between primary and secondary servers. It offloads various CPU intensive tasks from the primary server, such as bandwidth management, caching, and compression. It is also used to monitor protection plans by the vContinuum wizard.

CX Server is:

- Responsible for data offload and WAN optimization functions such as:
  - Compression
  - Securing the data over WAN
  - Data Routing
  - Bandwidth Optimization

- Provides centralized UI for configuration and monitoring.
- Provides centralized error reporting using logs, SNMP, and email alerts.
- A mandatory component for all environments.

CX-PS server is deployed on the Enterprise data server which receives data from the production servers and sends to the target site.

CX-CS server is deployed on the cloud SP's datacenter dedicated per customer to manage the replication and recovery.

Figure 2-26 shows the CS/PS server theory of operation.

Figure 2-26 CS/PS Server Theory of Operation


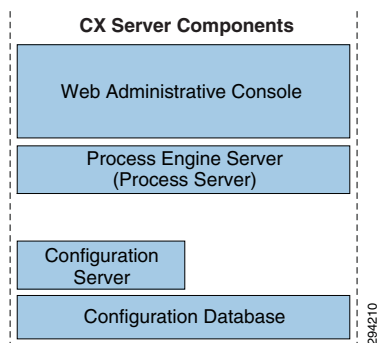


Figure 2-27 shows the CS/PS server architecture.

Figure 2-27 CS/PS Server Architecture



The CS/PS server has the following components:

- MySQL
- Common tables
- Hosts
- logicalVolumes
- srcLogicalDestinationLogicalVolume

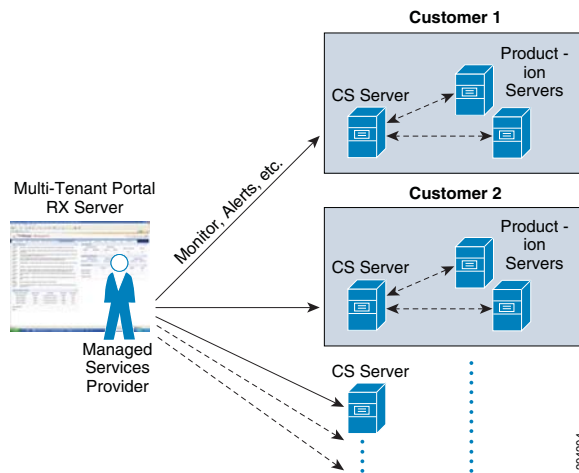
- Scripting
- Apache
- Configuration
- Transport Server
- HTTP/FTP
- Tmanagerd
- Executive functions
- Monitor.pl
- Keeps the GUI up to date
- Gentrends.pl
- Graphing
- Amethyst.conf
- Bootstrap configuration

## RX Server

RX is the multi-tenant portal that enables the management of all customer services through a single portal and provides:

- Centralized monitoring across all customers.
- Fully re-brandable ready to use customer-facing dashboard.
- A full-fledged API stack for deeper integration into partner portals.
- Replication health and statistics for each CS server.
- License statistics for each CS server.
- Alerts and notifications.
- Provision to create logical groups for multiple CS servers to enforce policies on multiple CS servers in one shot.
- Custom reports on bandwidth usage for each CS server.

Figure 2-28 RX: Multi-Tenant Portal




## Management Console

The management console/GUI Wizard is a Windows 32 bit-based GUI wizard that progresses through the protection and recovery steps:

- In the case of Windows CX, it is installed along with the CX server.
- In the case of Linux CX, the Wizard has to be installed on Windows 2008 R2, Windows7, XP or Vista desktop. The vContinuum wizard can be installed on the MT in the case of Windows.

vContinuum is stateless and does not have the current information regarding the replication status; it talks to the CX-CS server to get this information.

Figure 2-29 vContinuum



The vContinuum wizard helps the cloud provider to perform the following tasks:

- Push agents to source production servers
- Create protection plans to protect servers
- Modify existing protection plan.
- Perform DR drill.
- Resume protection
- Failback

- Offline Sync

## Self Service

Self Service can be enabled for the customers in the following ways:

- **RX Portal:** RX multi-tenant portal also allows the customers to perform recovery of their environments. This can be controlled by enabling the recovery option for the customer user account within the RX.
- **vContinuum:** vContinuum is a dedicated component deployed one per customer on the SP cloud. The SP can provide access to the vContinuum GUI to their customers who wants to have total control of the disaster recovery process. vContinuum allows customers to perform all the operations required for protecting and recovering the workloads into cloud.


## How InMage ScoutCloud Works

InMage ScoutCloud is based on CDP technology that gives it granular DR capabilities to meet most stringent DR requirements. InMage ScoutCloud can be configured to support long distance DR requirements as well as operational recovery requirements and supports heterogeneous servers running on Windows, Linux, or UNIX. ScoutCloud supports a web browser-based management UI that allows all management operations for both application and data recovery across different production servers and applications to be tracked and managed using a common management paradigm. A CLI to do the same is available as well. Management capabilities are protected through the use of a multi-level security model.

InMage ScoutCloud replicates a production server's data to one or more secondary servers that can be either local or remote, which can also be virtual or physical systems. ScoutCloud can be deployed into existing environments without disrupting your business continuity.

To understand how ScoutCloud works, let's look at a basic configuration with a single primary server and multiple secondary servers communicating to CX-CS server through a single CX-PS. The CX-CS is deployed in the primary server LAN network component whose failure and/or replacement do not impact production server's operation. The VX and FX component of ScoutCloud are deployed on your primary server, which utilizes negligible resource on your primary server. They asynchronously send writes as they occur on primary server to CX-CS. The VX and FX component of Scout are deployed on your secondary server, as well, to communicate continuously with CX-CS.

**Figure 2-30 Scout Environment**



ScoutCloud protects data by setting replication between primary server drive/file and secondary server drive/file. The replication process at drive level happens through stages. At the beginning of the replication process, a baseline copy of primary server's drive that you want to protect is created at the

secondary server. This step is known as Resyncing (Step I). Data changes during this step are sent to the secondary server Resyncing (Step II). Thereafter, Scout captures and sends only the changes in primary server drive. This stage is known as Differential Sync. This differential sync is a continuous process which is archived through VX agents. Scout supports fast resync, where the replication process starts directly from differential sync instead of replication stages. Unlike drive level replication, file/ folder level replication between primary and secondary server are one time activity, which is archived through FX agents.

For maintenance activities on the primary server or actual failure of the primary server, Scout switches the primary server to secondary server through failover. A failover operation is always followed by a failback operation; i.e., restoring the primary server from the secondary server. Scout uses CDP technology to replicate data, so that it can restore data to any point in time. To ensure the consistency of primary server drive data, the consistent tags/bookmarks are issued at the primary server at periodic intervals of time or on demand. The secondary server can be rolled back to any of the consistency bookmarks to ensure consistency of backup/DR data.

InMage ScoutCloud also supports the storing of Snapshots (exact replica of primary server's drive data as existed in single point in time) on physical or virtual volumes. These snapshots are stored as per consistency bookmarks applied on the secondary server.

## Component Flow


This section includes the following topics:

- [Process Server to Master Target, page 2-38](#)
- [Process Server to Master Target \(Reverse Protection\), page 2-38](#)

### Process Server to Master Target

Figure 2-31 shows the flow of data from protected servers to the PS server at the Enterprise to the MT at the SP.

**Figure 2-31 Data Flow: Protected Servers to the PS Server**




The Agent running on the source-protected servers collects data from the servers as it is created and sends to the local PS server, which would then send the data to the MT server residing at the SP premise where the data will be stored.

The PS server is also responsible for compressing and encrypting the data before sending over to the MT server. It's also capable of caching the data for extended periods of time in any WAN failure scenarios.

### Process Server to Master Target (Reverse Protection)

Figure 2-32 shows the data flow from the recovered servers after a failover into the Enterprise data center.



**Figure 2-32 Data Flow: Reverse Protection**

The data flow is similar to the scenario of protecting a server. The changed data from the recovered server is collected by the Datatap agent on the server and is sent to the CX/PS Scout Server on the SP side. The server compresses and encrypts data and in turn sends the data to the MT server on the Enterprise side. A MT server is required on the Enterprise side for the failback scenario.

## Deployment Considerations

This section includes the following topics:

- [Journal Sizing, page 2-39](#)
- [Storage, page 2-41](#)
- [Compression, page 2-43](#)
- [Encryption, page 2-44](#)
- [Compute, page 2-45](#)

## Journal Sizing

Journal volume holds the data changes happening on the production servers. All the changes during a retention period will be stored in the journal volume for recovery purposes. This gives the ability to recover the production servers to any point in time within the retention period.

InMage uses the copy-on-write method to capture changes in to the journal. After the initial copy of data to the target site, InMage tracks the changing blocks on the recovery volumes. The original data that is being written to is copied into the journal for point in time recovery.

Before a write is allowed to a block, copy-on-write moves the original data block to the journal space and updates the original block.

Capacity and performance sizing for journal is a very important consideration. The journal should have correct performance characteristics to handle the total write performance of the data being protected. The journal will have the same amount of writes that will occur at the source protection site. All the changed blocks have to be written into the journal.

### Journal Volume Capacity Considerations

The two important considerations for capacity planning of journal are:

- The amount of change rate of the source servers.
- The retention window required.

Retention window is the period during which the changes will be stored. When setting up a CDP configuration, the administrator will create a baseline copy of the data to be protected (effectively the state of the data at initial synchronization), take the average change rate of the data per day into account, determine how many hours, days, or weeks of writes they want to retain, and size the retention log accordingly. The log operates in a circular manner in the sense that it will keep all writes within the

defined time period but discard all writes older than that. For example, if the administrator defines a retention window of 3 days, then the retention log will retain all writes for the first 3 days of operation, and then begin to discard writes that are older than 3 days as newer writes are logged.

For recovery purposes, administrators generally want the most recent data. In cases where data corruption was the problem, or where root cause analysis will be performed, one or more older recovery points may be desired, but generally even these points are no older than 24 hours. Many customers establish a retention policy of one to four weeks, balancing the availability of granular recovery points against the size of the retention log. With a retention window of only a couple of days, you may not have access to certain older points, such as a quarterly close, that may be of interest because it may already have aged out of the retention window. In the case of the quarterly close, you probably aren't interested in any of the points around it, just that point which marked the state of the database of record when the quarterly close was completed and before any new transactions from the next quarter came in. It's probably not worth it in terms of storage capacity to extend the size of the retention log out so that older points like this could be retained, but it would be nice if certain older points could be saved without saving all the rest of the data. This is where the concept of sparse retention policies comes into play. The default retention log policy is "retain all writes for x time", but multiple policies with different retention periods could be defined.

For example, four policies might be defined:

- For the most recent 24 hours, keep all writes and bookmarks.
- For data that is between 24 and 48 hours old, keep a recovery point every 4 hours plus all bookmarks.
- For data that is older than 48 hours, keep only bookmarks of a certain type.
- Keep no data longer than 4 weeks.

A policy like that initially keeps all data, but then begins pruning unneeded metadata (for potential recovery points that are being discarded) to reclaim retention log space as data ages. Once the metadata for a particular point is discarded, that point can no longer be retroactively created.

By being able to define different retention period policies within the same CDP timeline, you get the recovery granularity needed for recent data while still being able to save certain key older points for easy access without having to use too much storage.

#### Journal Volume Performance Considerations

For every write on the Source Production Volume, three I/Os exist on the target side:

- Write to the Replica volume
- Read from the Replica volume
- Write to the Journal volume

The pattern of the IO, therefore, is 1\*Read and 2\*Write on the target side, which is split between Journal and Replica Volumes. The exact breakdown of IO type is:

- Journal Volume = 1 sequential write
- Replica Volume = 1 random read and 1 random write (Production Volume IO Pattern)

The Journal size is therefore dependent on the retention period that a customer wants to recover their workloads from and should be able to support the write IOPS from the production site of the customer.

## Storage

Storage is the main component in the DRaaS System. Proper storage sizing and deployment is very critical for delivering optimized service to customers. The following storage efficiency feature is recommended at the SP recovery site:

- **Thin Provisioning:** Thin provisioning is a good method for optimizing utilization of available storage. It relies on on-demand allocation of blocks of data versus the traditional method of allocating all the blocks up front. This method eliminates all the unused space, which helps avoid poor utilization rates. The best practice is to enable thin provisioning at the storage level or at the hypervisor level to avoid management challenges. In the DRaaS System, as InMage is capable of creating VMs using thin provisioning in the cloud, it is recommended to implement it on the hypervisor layer.

The following storage efficiency features are specific to EMC VNX when using vBlock as the ICS:

- **FAST Cache:** FAST Cache technology is an extension of your DRAM cache where it allocates certain flash drives to serve as FAST Cache. The benefit is that hotter data from applications running inside the VM will be copied to FAST Cache. Hence, these applications will see improved response time and throughput since the I/O is now serviced from flash drives. In DRaaS environments, FAST Cache will be useful during concurrent customer site failovers and during the on-boarding of new customers. In general, FAST Cache should be used in cases where storage performance needs to improve immediately for I/O that is burst-prone in nature.
- **FAST VP:** Data has a lifecycle. As data progresses through its lifecycle, it experiences varying levels of activity. When data is created, it is typically heavily used. As it ages, it is accessed less often. This is often referred to as being temporal in nature. FAST VP is a simple and elegant solution for dynamically matching storage requirements with changes in the frequency of data access. FAST VP segregates disk drives into the following three tiers: **Extreme Performance Tier** — Flash drives; **Performance Tier** — Serial Attached SCSI (SAS) drives for VNX; and **Capacity Tier** — Near-Line SAS (NL-SAS) drives for VNX platforms.
  - You can use FAST VP to aggressively reduce TCO and/or to increase performance. A target workload that requires a large number of Performance Tier drives can be serviced with a mix of tiers and a much lower drive count. In some cases, an almost two-thirds reduction in drive count is achieved. In other cases, performance throughput can double by adding less than 10 percent of a pool's total capacity in flash drives.
  - FAST VP and FAST Cache can be used together to improve storage system performance. Customers with a limited number of flash drives can create FAST Cache and storage pools consisting of performance and capacity drives. For performance, FAST Cache will provide immediate benefits for any burst-prone data, while FAST VP will move warmer data to performance drives and colder data to capacity drives.
  - FAST Cache is storage system aware where storage system resources are not wasted by unnecessarily copying data to FAST Cache if it is already on flash drives. If FAST VP moves a slice of data to the extreme performance tier, FAST Cache will not promote that slice into FAST Cache - even if the FAST Cache criteria is met for promotion.
  - When initially deploying flash drives in a storage system, use them for FAST Cache. FAST Cache will track I/Os smaller than 128 KB and requires multiple cache hits to 64 KB chunks. This will initiate promotions from performance or capacity drives to Flash Cache and as a result, I/O profiles that do not meet this criteria are better served by flash drives in a pool or RAID group.

The following storage efficiency features are specific to NetApp when using FlexPod as an integrated stack within VMDC:

- **Flash Cache:** Flash Cache speeds access to data through real-time intelligent caching of recently read user data and NetApp metadata. It is effective for random read-intensive workloads, including databases, e-mail, and file services. The combination of intelligent caching and NetApp data storage efficiency technologies enables the virtual storage tier, which promotes hot data to performance media in real time without moving the data, allowing you to scale performance and capacity while achieving the highest level of storage efficiency in the industry.
- **Flash Pool:** Flash Pool is a technology that allows flash technology in the form of solid-state disks (SSDs) and traditional hard disk drives (HDDs) to be combined to form a single Data onTap aggregate. When SSD and HDD technologies are combined in a Data onTap aggregate, the NetApp storage system takes advantage of the latency and throughput benefits of SSD while maintaining the mass storage capacity of HDD.
  - A Flash Pool is built from a Data onTap aggregate in a two-step process. Essentially, it is the addition of SSDs into an aggregate to provide a high-bandwidth, low-latency location that is capable of caching random reads and random overwrites. \*\* The feature does not require a license and works with any NetApp SSDs and one type of HDD per Flash Pool. That is, SSD and SAS performance drives can be combined to make a Flash Pool or SSD and SATA capacity drives can be combined to make a Flash Pool. You cannot combine SSD, SAS, and SATA into a single Flash Pool.
  - As a key component of the NetApp Virtual Storage Tier, Flash Pool offers a real-time, highly efficient implementation of automated storage tiering. Fine-grain promotion of hot data elements, combined with data deduplication and thin cloning, enables optimal performance and optimal use of Flash technology.
- **Deduplication:** NetApp deduplication is an integral part of the NetApp Data onTap operating environment and the WAFL file system, which manages all data on NetApp storage systems. Deduplication works "behind the scenes," regardless of what applications you run or how you access data, and its overhead is low.
  - NetApp deduplication is a key component of NetApp's storage efficiency technologies, which enable users to store the maximum amount of data for the lowest possible cost.
  - NetApp deduplication is a process that can be triggered when a threshold is reached, scheduled to run when it is most convenient, or run as part of an application. It will remove duplicate blocks in a volume or LUN.
- **Steady State Replication:**
  - FAST VP from EMC
  - Flash Pool from NetApp.
  - During the steady state replication, the target storage will have the information about the I/ O characteristics and about the data blocks.
- **Summary:**
  - Flash Cache and FAST Cache are useful in dealing with unpredicted I/O needs that can be observed during the recovery of multiple customer environments during a disaster.
  - Flash Pool and FAST VP are useful efficiency features which helps the SP to use storage space efficiently during steady state replication scenario. Warmer data gets moved to the faster drives and cold data gets moved to the capacity disks automatically.
  - Deduplication and thin provisioning reduces the total storage foot print required to support customer workloads.

## Compression

The efficient way to replicate data from one site to another is to compress the data before it is sent over to the WAN Network. This helps in reducing the WAN bandwidth required for data replication. This can be accomplished by using a dedicated external device or can be done by using the components of InMage.

The advantages of going with both these approaches include:

- Use of external device provides better handling of data compression and management as it will be used only for this functionality. This offloads the load on process server from InMage which does the compression.
- Compression consumes a lot of CPU resources and will effect the ability of the process server in performing other tasks.
- Troubleshooting for events will become easier.

For more information on the utilization of resources when compression is enabled, Refer to [Appendix A, “Characterization of Replication Process”](#).

The InMage process server can perform compression of data. This is a good option for customers who do not want to have a dedicated device for this functionality and this would be an ideal choice for customers who have fewer servers being protected.

This section includes the following topics:

- [External Cisco Products, page 2-43](#)
- [DR Vendor, page 2-44](#)

## External Cisco Products

Network links and WAN circuits can have high latency and/or packet loss as well as limited capacity. WAN optimization devices can be used to maximize the amount of replicated data that can be transmitted over a link.

A WAN Optimization Controller (WOC) is an appliance that can be placed in-line or out-of-path to reduce and optimize the data that is to be transmitted over the WAN. These devices are designed to help mitigate the effects of packet loss, network congestion, and latency while reducing the overall amount of data to be transmitted over the network. In general, the technologies utilized in accomplishing this are TCP acceleration, data deduplication, and compression. WAN and data optimization can occur at varying layers of the OSI stack, whether they be at the network and transport layer, the session, presentation, and application layers, or just to the data (payload) itself.

Cisco wide area application services (WAAS) devices can be used for data optimization. The WAAS system consists of a set of devices called wide area application engines (WAE) that work together to optimize TCP traffic over your network. Cisco WAAS uses a variety of transport flow optimization (TFO) features to optimize TCP traffic intercepted by the WAAS devices. TFO protects communicating devices from negative WAN conditions, such as bandwidth constraints, packet loss, congestion, and retransmission. TFO includes optimization features such as compression, windows scaling, Selective ACK, increased buffering, BIC TCP, and TCP Initial Window Size Maximization.

Cisco WAAS uses Data Redundancy Elimination (DRE) and LZ compression technologies to help reduce the size of data transmitted over the WAN. These compression technologies reduce the size of transmitted data by removing redundant information before sending the shortened data stream over the WAN. By reducing the amount of transferred data, WAAS compression can reduce network utilization and application response times.

When a WAE uses compression to optimize TCP traffic, it replaces repeated data in the stream with a much shorter reference and then sends the shortened data stream out across the WAN. The receiving WAE uses its local redundancy library to reconstruct the data stream before passing it along to the destination. The WAAS compression scheme is based on a shared cache architecture where each WAE involved in compression and decompression shares the same redundancy library. When the cache that stores the redundancy library on a WAE becomes full, WAAS uses a FIFO algorithm (first in, first out) to discard old data and make room for new.

## DR Vendor

InMage supports compression of data before it sends the data over to the WAN. This helps in reducing the need for high WAN bandwidth to carry data traffic. WAN optimization features like data compression will help customers to achieve low RPOs.

InMage collects data changes from production servers in real time, and places them in memory before they are written to disk. They changes are sent to a software appliance called the process server and then transferred to the secondary site. The server offloads compute intensive tasks from the production systems, such as compression, encryption, WAN acceleration, and consolidated bandwidth management.

The process server does the compression of data and sends the data over to the MT server residing on the SP's cloud.

Customers who do not want to bear additional cost for the WAN optimization devices can leverage InMage for achieving low RPOs.

## Encryption

Encryption of data-in-transit and data-at-rest is the best method to enforce the security and privacy of data, regardless of where it resides. Data-in-transit encryption is necessary to keep the data secure while in transit. The network connection between sites must be secure and the data must be protected. The use of IPsec or SSL to encrypt WAN connections ensures that no visibility occurs at the packet level if any of the datagrams are intercepted in transit.

Encryption of data-in-transit between the sites can be accomplished in two ways:

- **InMage technology** is capable of encrypting data in flight, the CX-PS server on the enterprise encrypts the data before it sends it over to MT over the WAN. Enabling the encryption will secure the data transmission between CX-PS and a secondary server. Since the encryption is performed on the CX-PS server, any performance impact will be limited to the CX-PS server.
- The other option is to use **Cisco ASA** for encrypting data between the Enterprise and the SP's data centers. The Cisco ASA 55xx Series is a purpose-built platform that combines superior security and VPN services for enterprise applications. The Cisco ASA 55xx Series enables customization for specific deployment environments and options, with special product editions for secure remote access (SSL/IPsec VPN).

The Cisco ASA 55xx Series SSL/IPsec VPN Edition uses network-aware IPsec site-to-site VPN capabilities. This allows customers to securely extend their networks across low-cost Internet connections to the service provider cloud.

Encryption of data-at-rest can add further security to the storage environment on the cloud SP's data center. Any external key manager can be used in conjunction with SAN fabrics and storage arrays to secure data-at-rest.

## Compute

This section includes the following topic:

- [Oversubscription, page 2-45](#)

## Oversubscription

DRaaS utilizes shared resources on the recovery site. Since resources at failover site sit idle most of the time, DR enables high over-subscription ratios, making it ideal for cloud environments.

The SP can have fewer compute resources compared to the customer's production environments. The compute within the SP cloud is based on Cisco UCS servers, which can be rapidly deployed with the help of the service profiles to meet any unexpected or rare scenario where all the customers fail over to the cloud. In this scenario, new UCS servers can be deployed and added to the existing compute clusters for additional compute resource needs.

Every server that is provisioned in the Cisco UCS is specified by a service profile, which is a software definition of a server and its LAN and SAN network connectivity. In other words, a service profile defines a single server and its storage and networking characteristics. Service profiles are stored in the Cisco UCS 6xxx Series Fabric Interconnects. When a service profile is deployed to a server, UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the service profile. This automation of device configuration reduces the number of manual steps required to configure servers, network interface cards (NICs), host bus adapters (HBAs), and LAN and SAN switches.

## Key Findings

This section includes the following topics:

- [Concurrency, page 2-45](#)
- [Limitations, page 2-45](#)

## Concurrency

A maximum of 40 VMDKs/RDMs per MT based on best practice. Refer to “[Implementation Best Practices](#)” section on page 5-8 for details.

Maximum supported change rate per day per Scout Process Server is 1TB. Deploy additional process servers to support additional change rate. Refer to [Table 3-1 \(Scout Server Storage and Compute Implementation, page 3-6\)](#) for details.

InMage relies on vCenter limits for throttling and concurrent operations. Refer to <http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf> for additional details.

## Limitations

The following are current limitations for the DRaaS 1.0 System:

- For recover later plans, outside of the Multi-Tenant Portal, no mechanism exists to view VMs included in a recovery plan. The assumption is the recovery plan should map exactly to a protection plan. If the protection plan changes, a new corresponding recovery plan should be created.
- Storage vMotion of MT is not supported. Compute vMotion is supported.
- MT OS type needs to match the OS of the protected servers. Also a single protection plan can only span across a single MT. It is not possible to migrate a protection plan from one MT to another.
- Offline sync requires a deployment of MT in the primary enterprise data center. MT has to be a VM that gets shipped to a SP's secondary data center.
- Smallest configuration retention window from the vContinuum server is 1 day.
- The RX portal can have up to a 10 minute lag in display. Resource consumption for storage and compute is not available from the RX portal.
- It is not possible to set VM replication priority within a protection group. It is possible to limit available bandwidth at a MT level.
- Infrastructure masking for tenants resources relies on VMware vCenter permissions.





# CHAPTER 3

## Implementation and Configuration

---

Typical InMage deployments require control servers to be deployed in both the Enterprise and SP data centers. Proper sizing of each component, based on change rate, is required to achieve desired recovery point objective (RPO), recovery time objective (RTO), and retention window. This chapter will provide sizing details along with deployment recommendations for each component. The following major topics will be discussed:

- [Master Target—Enterprise and Service Provider, page 3-1](#)
- [InMage Scout Server Details, page 3-6](#)
- [InMage vContinuum, page 3-12](#)
- [InMage Agent Configuration, page 3-13](#)
- [Multi-Tenant Portal—RX Server, page 3-17](#)
- [Summary Tables of Components for All Tenants, page 3-21](#)
- [VMDC 2.3, page 3-24](#)
- [BMC Cloud Lifecycle Management, page 3-49](#)

The above contents are not organized based on deployment order, but are rather general guidelines for each component as a standalone entity. Actual deployment order may vary depending on the method of onboarding the tenant and the SP method of operation.

## Master Target—Enterprise and Service Provider

The InMage disaster recovery solution creates replication pairs between a primary VM and a dedicated VM on a secondary ESX server called "master target" or "MT." The MT may act as the protection target for multiple physical or virtual servers. Each machine (physical or virtual) can only be mapped to a single MT and a single protection plan. A single protection plan can also only span across a single MT. Depending on the number of servers under protection, each tenant will have at least one MT. Based on the type of use case, a tenant may require MTs in both the enterprise and SP:

- Enterprise to SP Recovery Protection: At least one instance of MT required in the SP.
- SP to Enterprise Failback Protection: At least one instance of MT is required in the Enterprise.

MTs can be deployed on Linux or Windows virtual servers, but must be of the same OS family as primary servers being protected. A number of factors, which are described in the following topics, determine deployment and sizing of the MT:

- [Master Target OS-Specific Information, page 3-2](#)
- [Master Target Deployment Limit, page 3-2](#)

- [Volume Requirements, page 3-3](#)

## Master Target OS-Specific Information

The OS that matches that of the primary server needs to be deployed on the MT.

- If Primary VMs are running Windows, the MT needs to be Windows based. Windows 2008 R2 is recommended for Windows-based MTs.
- Otherwise, if the primary server is running Linux, the Linux MT needs to be deployed. The CentOS-based MT is recommended.

Although detailed OS commands vary greatly when setting up a Windows MT from Linux, most of the high level planning steps are very similar. InMage recommends both the Linux and Windows MT to be configured with three volumes:

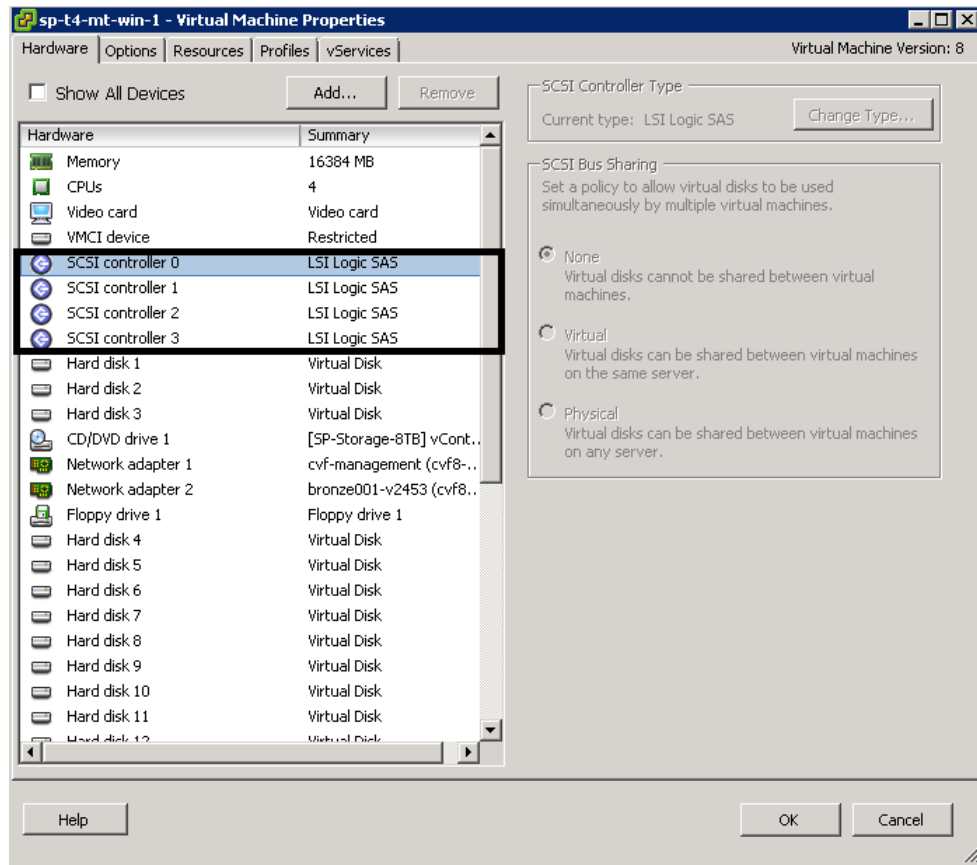
- OS boot volume
- Cache
- Retention

Both Cache and Retention are required for journal. Space required for each volume is based on change rate. Refer to [“Volume Requirements” section on page 3-3](#) for additional details.

## Master Target Deployment Limit

As discussed previously, the vContinuum will create a shadow copy of each disk under protection and dual mount the disk locally on the shadow VM and the MT. Each time a block changes at the source machine, the MT is the target of the changed block. As such, to prevent consistency problems on virtual disks, a lock is placed against all protected volume/VMDKs. Since the MT is a VM running in the vSphere environment, vCenter limits apply to each MT. vCenter limits each host to 4 SCSI controllers; each SCSI controller is capable of supporting 15 VMDKs. Refer to [Figure 3-1](#).

Figure 3-1 Master Target Deployment Limit



Due to the VMware SCSI controller limits, a single MT can support a maximum of 60 VMDKs. Additionally, InMage requires three volumes on the MT for OS, Cache, and Retention. Assuming DR Drill functionality is not required, a MT can support a maximum of 57 VMDKS and therefore protect 57 primary VMs. If DR Drill functionality is required, InMage recommends to not exceed 40 VMDKS per MT. If the number of primary VMs that need protection exceeds 57 (or 40 with DR Drill), then additional MTs must be deployed. Refer to “[Implementation Best Practices](#)” section on page 5-8 for details.

## Volume Requirements

To properly size a deployment, it is important to ensure the InMage MT has sufficient disk space required to support the desired change rate, journal history and intermediate caching. These are described in the following topics:

- [Retention Volume Sizing, page 3-3](#)
- [Cache Volume Sizing, page 3-5](#)

## Retention Volume Sizing

Retention policies effect the sizing of retention volume. The three types of retention policies are:

- **Time based:** The data in the retention window will be overwritten after the specified time period.

- **Space based:** The data in the retention window will be overwritten once the size is exhausted.
- **Time and space based:** The data in the retention window will be overwritten once the time specified or space specified qualifies first.

To ensure the SP meets and maintains SLA, time-based retention is often implemented. In a time-based deployment, two factors determine the size of retention volume:

- Daily Data Change Rate
- Retention Window

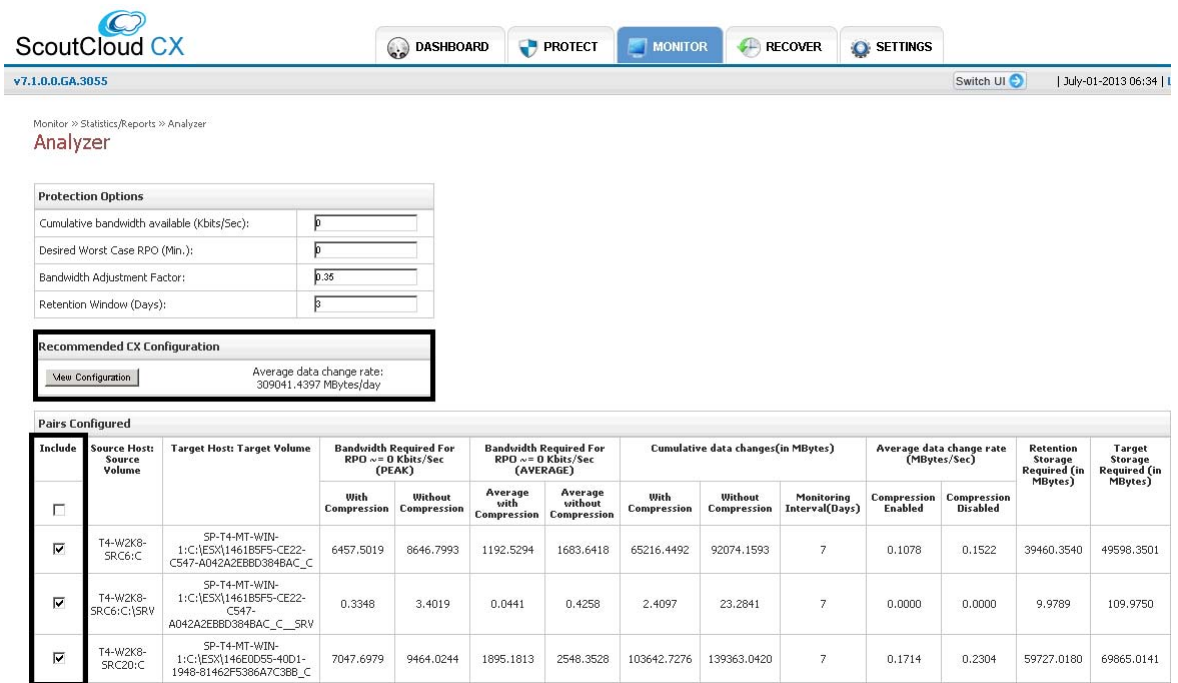
Retention window is the time period information about data changes on the primary server is stored. An application can be rolled back to any point in time or application-consistent bookmarks/tags in this window. If the data change rate is expected to be constant over a 24-hour window, the following simple relationship can be used to derive the size of retention drive:

**Retention Drive Size = (Daily Change Rate) \* (Number of hours of restore point or retention window)**

As an example, if an enterprise expects 750GB of data change in an 24-hour window and a retention window/journal history of six hours, the expected retention size will be (750G / 24 hours) \* (6 hours) = 187G (200G rounded up).

If the data change rate follows a traditional bell curve, where lots of change occurs during normal operation and little change occurs after hours, when sizing the retention drive based on time and space, it is important to ensure the retention drive is based on the peak data change window. In a production environment, the daily record of change rate profile is stored in the CX server. This information can be used to fine tune the size of the retention drive as needed. Figure 3-2 is a sample screen capture of the change rate profile screen:

Figure 3-2 Change Rate



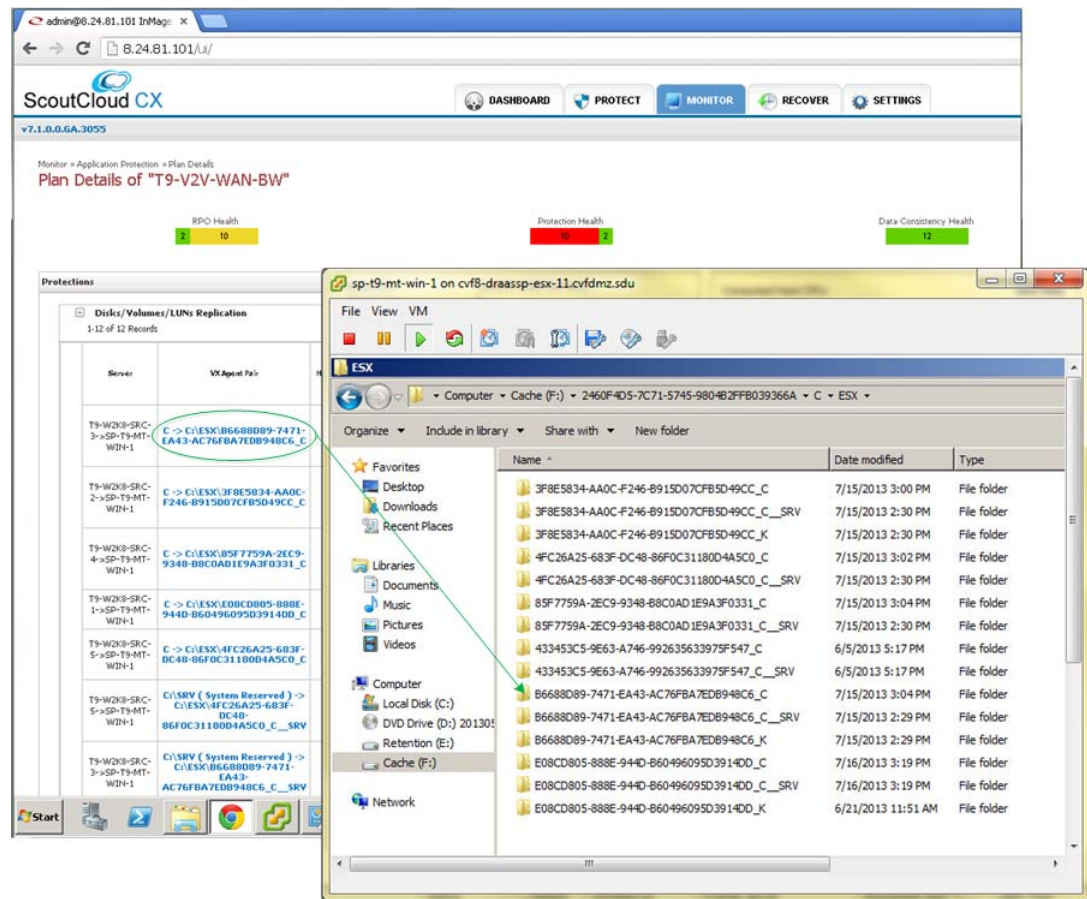
The highlighted box indicates the average change rate per day across all volumes under protection. Average change rate can be removed or added from the data change profile with the enable/disable check box next to "Individual volume."

## Cache Volume Sizing

The cache volume is used by the MT to store new uncompressed incoming data from the primary site prior to be processed and written to journal. For each source disk/volume under protection, a separate and dedicated folder is created on the cache volume on the MT.

Figure 3-3 shows the Cloud CX UI and the vCenter console for the MT named SP-T9-MTWIN-1. In the CX UI, several servers and VX Agent Pairs are shown. The VX Agent Pair associates a local drive on the primary server with a folder on the MT. The naming convention for the folders uses a host ID for the primary server with the volume name appended. Looking at the first row, changes for the C: drive on the T9-W2K8-SRC-3 primary server are cached on the F:\2460F4D5-7C71-5745-9804B2FFB039366A\C\ESX\B6688D89-7471-EA43-AC76FBA7EDB948C6\_C folder on the SP-T9-MT-WIN-1 MT.

Figure 3-3 Cache Volume Sizing



As a design guideline, 500MB of disk space per Virtual Machine Disk (VMDK) under protection should be reserved on the cache volume. The total size of the cache volume is a function of total number of volumes under protection:

$$\text{Size of Cache Volume} = (\text{Total number of volumes under protection}) * (\text{500MB per volume})$$

We expect the cache volume utilization to stay relatively low when no resource bottlenecks exist.

# InMage Scout Server Details

InMage Scout server is a multi-purpose server that, depending on deployment, can be referred to as one of the following:

1. **Process Server (PS):** Caches copies of block changes on primary server(s) residing in enterprise data center and sends them to master targets residing in the SP secondary site. Optionally, compression and encryption of block changes can be performed prior to being transmitted over the WAN link.
2. **Central Configuration Server (CX-CS):** Allows an operator to perform tasks through a web-based UI, such as fault monitoring (RPO violation, agent heartbeat), configuration, accounting, and performance monitoring. This server also generates reports, trend graphs, e-mail, and SNMP trap alerts.
3. **Dual role (PS + CX):** Similar to the CX-CS server, but adds the PS functionality in the server provider server to enable failback protection for the secondary servers back to the primary servers.

This section includes the following topics:

- [Scout Server Storage and Compute Implementation, page 3-6](#)
- [Scout Server Network Implementation, page 3-7](#)
- [Scout Server Network Bandwidth, page 3-8](#)
- [Scout Server Replication Options, page 3-10](#)

## Scout Server Storage and Compute Implementation

To understand Scout Server sizing, it is important to understand the multi-stage processing required to set up a protection plan:

- **Resyncing (Step I):** Baseline copy of primary server's drive under protection is created at the secondary server.
- **Resyncing (Step II):** Data changes during Step I are sent to the secondary server.
- **Differential Sync:** Once Step II completes, only block changes are sent to the secondary server.

To minimize WAN overhead, Scout also supports fast resync, where the replication process starts directly from differential sync instead of replication stages. Fast resync does a block-by-block comparison between storage volume at the primary and secondary sites. Data is sent over the WAN only if the two blocks are different. The primary case for fast resync is if the source went down abruptly causing some changes in memory to be lost.

As shown in [Figure 3-4](#), for each volume protected, disk space is reserved on the PS cache drive for the syncing stages using the following fields:

- **Resync File Threshold (MB):** For the resync steps, which defaults to 2G.
- **Differential Files Threshold (MB):** For differential sync, based on the expected or observed change rate, which defaults to 8G.

**Figure 3-4** Pair Settings

Pair Settings											
Visible	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 type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Fast	30	2	2048	8192	CX-PS Based	

However, the 8G threshold can also be changed if the change rates for a single volume are higher. A single volume from a DB server, for example, could generate a large number of changes. It is a good practice to monitor cache utilization proactively and adjust thresholds based on application change rate. The main intent behind increasing these thresholds is to cache data on the PS instead of primary servers in the event there is a WAN bottleneck. If no bottlenecks exist, the default settings should work for most application deployment scenarios and provide enough storage to cache changes for a 6-8 hour WAN outage.

Total disk requirements on the processing server, based on default settings, will be the number of volumes \* (8G + 2G), which should not exceed the size of the cache partition. [Table 3-1](#) is the general CPU/memory/disk sizing recommendation by InMage based on change rate:

**Table 3-1** CPU/Memory/Disk Sizing Recommendation

Data Change Rate	CPU	Memory	Boot Volume Disk Type and Memory	Cache Disk
< 300GB	1 Quad Core	8 GB	RAID 1+0 15K disk, 40GB	RAID 1+0 10K/15K disk, 400GB
< 700GB	2 Quad Core	16 GB	RAID 1+0 15K disk, 40GB	RAID 1+0 10K/15K disk, 790GB
< 1 TB	2 Quad Core	32 GB	RAID 1+0 15K disk, 40GB	RAID 1+0 10K/15K disk, 790GB

The above recommendation assumes data encryption and compression as well as potential WAN outage of up to six hours. The CPU requirement could be reduced if data encryption and compression are offloaded to an external appliance. Consult with InMage for alternative deployment considerations.

## Scout Server Network Implementation

“[Component Flow](#)” section on page 2-38 documented data flows between the enterprise and SP among various InMage components that enable CDP and automatic recovery upon disaster declaration.

The primary and secondary server's VX agents communicate configuration and status information to the CX server over the HTTP protocol using standard port 80. The stateful firewall contains predefined rules for both HTTP and FTP/FTPS; i.e., no need to open a range of ports specifically exists. Also, the secondary server VX agents can open connections to port 873 (RSYNC) of the CX server when using resync.

The primary and secondary server agents use FTP/FTPS as the data transfer protocol to send and receive data from the CX server. The FTP protocol comes in two flavors, Active and Passive. InMage defaults to Passive FTP. Active FTP uses fixed ports 20 and 21 on the server side (CX server). In Passive FTP, clients initiate all connections; thus, no client-side firewall settings need to be configured. The clients, however, do need to be able to access port 21 and all ports greater than 1024 on the server side (CX server). It is possible to limit the range of ports to be opened up; this setting is controlled by `/etc/proftpd.conf`. To achieve this, simply create an entry for PassivePorts directive in your `proftpd.conf`:

```
PassivePorts 60000 65535 # allowed ports
```

The FX agents communicate configuration and status information to the CX server over the HTTP protocol. The data transfer protocol by default is a single socket connection to port 874 of the primary or secondary server.

In summary, InMage recommends that network traffic on the following ports not be blocked by hardware- or software-based firewall(s):

**Table 3-2 Network Ports**

Component	Traffic Type (Port)
Source Host	<ul style="list-style-type: none"> <li>• HTTP (80)</li> <li>• HTTPS (443)</li> <li>• Unencrypted data transfer (9080)</li> <li>• Encrypted data transfer (9443)</li> </ul>
CX Configuration Server	<ul style="list-style-type: none"> <li>• SMTP (25)</li> <li>• HTTP (80)</li> <li>• HTTPS (443)</li> <li>• MySQL (3306)</li> <li>• Unencrypted data transfer (9080)</li> <li>• Encrypted data transfer (9443)</li> </ul>
Target Host	<ul style="list-style-type: none"> <li>• HTTP (80)</li> <li>• HTTPS (443)</li> <li>• VX replication data traffic (873)</li> <li>• FX replication (874)</li> <li>• Unencrypted data transfer (9080)</li> <li>• Encrypted data transfer (9443)</li> </ul>
Optional	<ul style="list-style-type: none"> <li>• SNMP (162)</li> </ul>
Process Server	<ul style="list-style-type: none"> <li>• MySQL (3306)</li> <li>• Unencrypted data transfer (9080)</li> <li>• Encrypted data transfer (9443)</li> </ul>

## Scout Server Network Bandwidth

InMage Bandwidth Manager controls and shares bandwidth from the processing server to the master targets. Available WAN bandwidth, along with data change rate, directly impacts target RPO. To determine WAN bandwidth required during peak and off peak hours, InMage offers the ability to profile a single or multiple disk volumes before or after enabling block level replication. To enable profiling before setting up block level replication, data changes can be sent to a local CX-CS server or "InMageProfiler" for analysis of the data rather than sending data changes across the WAN. The objective is to simplify the task of initial sizing of the WAN for disaster recovery. Profiling as part of a regular block level data replication is intended to be regular ongoing capacity management, performed by the cloud admin, to adjust available WAN bandwidth as application usage pattern changes.



To gain valuable insight into data change rate and compressibility on the primary server(s) before setting up block level replication, first install the scout agent on the primary server. Configure a CXCS and CX-PS, aka "InMageProfiler," in the same LAN as the primary server with the following procedure:

- Step 1** From CX UI, navigate to **Protect > Protection Plans > Create Protection Plan > Setup Profiling**.
- Step 2** Provide the protection plan name and choose the type of profiling as **Individual Volume Profiling**.
- Step 3** Select a desired volume and select next to profile. All the hosts are listed under Host Drives. Expand the hosts to select the desired volume and click **Next**.
- Step 4** In the second step, select the target as InMageProfiler. The Replication Options are optional and the CDP retention option will not be available. Click **Next**.

The **Monitor > Volume Protection** page can be used for monitoring:

**Figure 3-5 Monitor/Volume Protection Page**

Monitor >> Volumes >> Reports  
Replication Reports

Statistics Reports Settings

Pair Details							
Primary Server	Primary Volume	Remote Server	Target Volume	Process Server	Replication Pool	Fast Resync Unmatched %	Agent Log
T4-W2K8-SRC2	C	InMageProfiler	P	t4-ps-1.t4.sdu [ 6.126.101.6 ]	24	N/A	N/A

Health Report [ Jul 01, 2013 - Jul 05, 2013 ]											
T4-W2K8-SRC2											
T4-W2K8-SRC2 (C) - PROTECTED <span style="float: right;">Change Rate   RPO   Retention   Health</span>											
Date	Data changes (in MBytes)		Retention Window (Days)		RPO		No. of hours RPO not met	Data Flow Controlled (Hours)	Retention log reset?	Available Consistency Points	Protection Coverage
	With Compression	Without Compression	Policy	Available	Threshold	Max					
Jul 04, 2013	3051.25	4511.41	0	0	30 min	7.23 days	0.03	0	NO	0	100%
Jul 05, 2013	8094.66	11524.09	0	0	30 min	0.26 min	0	0	NO	0	100%
<b>Total:</b>	11145.91	16035.5	N/A	N/A	N/A	N/A	0.03	0	N/A	N/A	100%

Data change profiling as part of a regular block level data replication is enabled by default. Navigate to **Monitor > Analyze Profile Results**.

The network planner can input the following four metrics into the protection options window:

- Proposed available bandwidth
- Desired RPO (minutes)
- Bandwidth adjustment factor
- Retention window (days)

InMage is based on historical data change metrics profile if proposed RPO is achievable during peak and average usage pattern. Based on profile result, the network planner can further adjust the WAN bandwidth or reduce the desired RPO expectations.

**Figure 3-6 Analyze Profiling Results**

Monitor » Statistics/Reports » Analyzer

**Analyzer**

Protection Options	
Cumulative bandwidth available (Kbits/Sec):	<input type="text" value="1000"/>
Desired Worst Case RPO (Min.):	<input type="text" value="1"/>
Bandwidth Adjustment Factor:	<input type="text" value="0.35"/>
Retention Window (Days):	<input type="text" value="3"/>

Recommended CX Configuration	
<input type="button" value="View Configuration"/>	Average data change rate: 497026.2771 MBytes/day

Result			
Data Change	Compression	Bandwidth Required (Kbits/sec)	Rpo Achieved
Peak	Yes	1443.6268	No
	No	1979.2029	No
Average	Yes	746.9172	Yes
	No	1060.3227	No

By default, all volumes under protections are included. Simply unselect a volume(s) to exclude it from the final data change profiling / analysis.

**Figure 3-7 Selecting Volumes for Analysis**

Recommended CX Configuration	
<input type="button" value="View Configuration"/>	Average data change rate: 533310.4646 MBytes/day


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)
			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"/>	T4-W2K8-SRC6:C	SP-T4-MT-WIN-1:C:\ESX\1461B5F5-CE22-C547-A042A2E8BD384BAC_C	483116.8089	646908.6852	215926.9877	308570.9205	157836.4816	225556.5593	7	0.2610	0.3729	96667.0968	106805.0929
<input checked="" type="checkbox"/>	T4-W2K8-SRC6:C:\SRV	SP-T4-MT-WIN-1:C:\ESX\1461B5F5-CE22-C547-A042A2E8BD384BAC_C__SRV	25.0516	254.5150	6.9271	66.1408	5.0635	48.3471	7	0.0000	0.0001	20.7202	120.7163

Unselecting the volume does not operationally impact the volume, remove the volume from data protection, or prevent data migration, regular backup, etc. It simply excludes the volume from data profiling.

## Scout Server Replication Options

This section outlines the configurable replication options available on the Scout Server. The settings described below are configured on a per-volume basis. [Figure 3-8](#) is a screen capture taken from a protection deployment.

Figure 3-8 Volume Replication Configuration Options

v7.1.0.0.GA.3055 Switch UI 

Monitor » Volumes » Settings  
**Replication Settings**

Statistics Reports **Settings**

Pair Details							
Primary Server	Primary Volume	Remote Server	Target Volume	Process Server	Replication Pool	Fast Resync Unmatched %	Agent Log
T2-LX-SRC-3	/dev/sda	SP-T2-MT-LX-1	/dev/mapper/36000c290836457ccda08f44077f05bcf	t2-ps-1 [ 6.126.99.201 ]	6	N/A	N/A

Pair Settings										
Visible	Resync	Profiling Mode	Secure CX-PS to Destination	Secure Source to CX-PS	Resync Mode	RPD Threshold	Replication Pool (1-24)	Resync Files Threshold (MB)	Differential Files Threshold (MB)	Compression Enable
<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Fast	0	5	2048	192	CX-PS Based

Restart Resync Accept Changes Reset

Select a different Process Server

SP-T2-PS-1 (8.24.71.101) t2-ps-1 (6.126.99.201)	Number of Pair Configured 1
--	--------------------------------

Accept Changes

Automatic Resync

Start between hours  :  and  :  after waiting  minutes

Accept Changes

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	/mnt/retention/Retention_Logs	80	256.00	Purge older logs

Edit Disable Retention

1. Secure Transport: Refer to “IPsec” section on page 3-33.
2. Batch Resync: Refer to “Scout Server Storage and Compute Implementation” section on page 3-6 regarding Resync. This setting controls number of replication errors that can be resynced simultaneously in a protection plan. Typically, this setting should be set based on WAN bandwidth and storage. If WAN throughput is low, it is recommend to limit the number of concurrent resyncs. Default setting is 3 when setting up a protection plan from vContinuum.
3. Automatic Resync: Used when a replication pair is required to address data inconsistencies automatically. A resync is required if any inconsistency exists between the primary and secondary. For example, the primary server reboots causing data in memory to be lost. Resync is required to ensure data consistency. When the Automatic Resync Option is enabled and data inconsistency occurs, the replication pair waits for a certain period of time (by default, 30 minutes) before performing a forced resync within the Start Between Hours time frame. This ensures data consistency and minimizes manual intervention.
4. Use Compression. The three options are:
  - No Compression: Use this option to replicate data without compression.
  - CX-PS based: Use this option to enable data compression on the CX-PS.
  - Source based: Use this option to compress data on the primary server before sending it to the CX-PS.

5. Resync and Differential Thresholds: Refer to [Scout Server Storage and Compute Implementation, page 3-6](#).
6. RPO Threshold: Worst case RPO before an email alerts are sent and RPO alarms are raised.
7. Bandwidth restriction per MT.

## InMage vContinuum

InMage vContinuum is used to set up application discovery, protection, and failback for VMware vSphere VM as well for physical servers. Together with Scout processing (PS), Scout CX, and MT, vContinuum protects VMs as well as physical servers by replicating them to a secondary ESX/ESXi server and recovering them on the secondary VMware vSphere environment when needed. When integrated with application VSS writers, application backup and recovery can be application consistent instead of crash consistent. Refer to InMage's compatibility matrix for information on supported operating system versions.

[Table 3-3](#) displays the platforms that are supported by vContinuum.

**Table 3-3** *Continuum-Supported Platforms*

Platform	Version Number
vSphere	ESX 3.5, 3.5 U2, 4.0, 4.1 ESXi 3.5, 3.5 U2, 4.0, 4.1, 5.0, 5.0 U1, 5.1
Guest OS	Windows 2003, 2008, 2008 R2, 2012 SLES 9.x, 10.x, 11.x CentOS 4.x, 5.x, 6.x RHEL 4.x, 5.x, 6.x
vCenter	vCenter 4.0, 4.1, 5.0, 5.0 U1, 5.1



### Note

- To provide failover from enterprise to SP, secondary vSphere (SP) version should be either the same or higher than the source (enterprise) vSphere server. To perform a failback from SP to Enterprise, enterprise vSphere version should be either the same or higher than the SP vSphere. vSphere server may need to be upgraded if failback is required.
- For new installations, InMage recommends:
  - Secondary ESXi Platform: ESXi 5.1
  - MT Platform for Windows: Windows 2008 R2 Enterprise Edition
  - MT Platform for Linux: CentOS 6.2 64-bit
  - CX Scout OS: CentOS 6.2 64-bit
  - vContinuum: Windows 2008 R2

Deploying the vContinuum wizard can occur in two major ways:

- In the case of the Linux-only tenant, the wizard has to be installed on a dedicated VM running Windows 2008 R2, Windows7, or XP desktop.
- In the case of Windows, the vContinuum wizard can be installed on the MT or on a dedicated VM running Windows 2008 R2, Windows7, or XP.

Running the vContinuum wizard on top of the MT server reduces the number of touch points and deployment footprint. This is the deployment model implemented in the Cisco Cloud Validation Facility.

## InMage Agent Configuration

InMage agents comes in various flavors depending on the guest OS. The Unified Windows Agent covers most releases for Windows 2003, 2008, and 2012, as well as some Windows XP, Vista, 7, and 8 releases. The specific Windows edition support is highlighted in [Figure 3-9](#).

**Figure 3-9 InMage Windows Edition Support**

Windows Guest Operating Systems								
GUEST OS			EDITION					
OS Version	Bit	Release	Web	Standard	Enterprise	Data Center	Professional	
Windows 2003	32 bit	Base*	✓	✓	✓			
		SP1*	✓	✓	✓			
		SP2	✓	✓	✓			
		R2 SP1		✓	✓			
		R2 SP2		✓	✓			
	64 bit	Base*			✓	✓		
		SP1*			✓	✓		
		SP2			✓	✓		
		R2 SP1			✓	✓		
		R2 SP2			✓	✓		
Windows 2008	32 bit	SP1	✓	✓	✓	✓		
		SP2	✓	✓	✓	✓		
	64 bit	SP1	✓	✓	✓	✓		
		SP2	✓	✓	✓	✓		
		R2	✓	✓	✓	✓		
Windows 2012#	64 bit	Base		✓		✓		
Windows XP	64	SP2					✓	
Windows Vista	32	Base			✓			
	64	Base			✓			
Windows 7	32	Base					✓	
	64	Base					✓	
Windows 8	64	Base					✓	



### Note

- If Windows 2003 (Base, SP1) source machines have applications that require application quiesce. It is then strongly suggested to upgrade to Windows 2003 SP2 to overcome the VSS-related errors.
- Storage Space is not supported. Protect Windows 2012 VM with ReFS Filesystem, requires matching MT.



### Note

UEFI with Dynamic disk will not work.

- Windows XP 32 bit VM on ESX cannot be protected in V2V workflow, but it can be protected using P2V workflow.

- Windows XP 32 /64 recovered VM will not have network drivers installed automatically; the user needs to install manually.

Unified Agent support for Linux is specific to the distribution and release. RHEL 5/6, CentOS 5/6, and SUSE 10/11 are supported. RHEL 5 U3 / CentOS 5 U3 and older versions require network changes to be manually configured after recovery to secondary data center. For Linux physical-to-virtual protection, GRUB bootloader must be used on the source server. Refer to [InMage\\_vContinuum\\_Compatibility\\_Matrix](#) for complete details.

On the Windows platform, the VSS framework is used by vContinuum and Scout to enforce application consistency. The InMage agent interfaces with the VSS writer to quiesce the application and take application consistent snapshots. Refer to [InMage\\_vContinuum\\_Compatibility\\_Matrix](#) for a complete list of certified applications across various operation systems.

There are two ways to install InMage agents onto the primary servers:

1. **Manual Installation:** Network operator can pre-load the Scout agent on each server:
  - Interactive install
  - Command line install (Silent mode)
  - File-based install (Linux)
2. **Bulk Push:** Network operator can push the Scout agent to each server using a UI:
  - CX UI (Linux)
  - vContinuum (Windows)

The first option works well in a smaller deployment in which there are few primary servers.

InMage\_Scout\_Standard\_User\_Guide has excellent step-by-step installation procedures, which will not be repeated in this document. We do want to point out that for Linux agent install, by default, the installer checks for a free space of 2 GB on /root. This can be suppressed by adding the extra switch -k (space in KBs) or -m (space in MBs) or -g (space in GBs).

Bulk push works better for a medium or large deployment as it automates a great number of repetitive clicks and configuration settings that a cloud operator is required to enter. The choice of UI depends on the primary server OS platform. For Linux bulk install, the CX UI is used and for Windows bulk install the vContinuum is used.

This section includes the following topics:

- [CX UI for Linux Bulk Install and Upgrade, page 3-14](#)
- [vContinuum for Windows Bulk Install and Upgrade, page 3-16](#)

## CX UI for Linux Bulk Install and Upgrade

Depending on distribution of Linux, a distribution-specific agent installer needs to first be loaded into the Software Repository. To upload the desired installers, navigate through **Settings > Agent Installers > Software Repository**. Click **Browse** to upload software.

### *Figure 3-10 Home Installer 1*

Once installers are loaded, navigate through **Settings > Agent Installers > Manage Agent Installation/Upgrade**. Click **Install Agent** to provide install requirements. This is a four-step process:

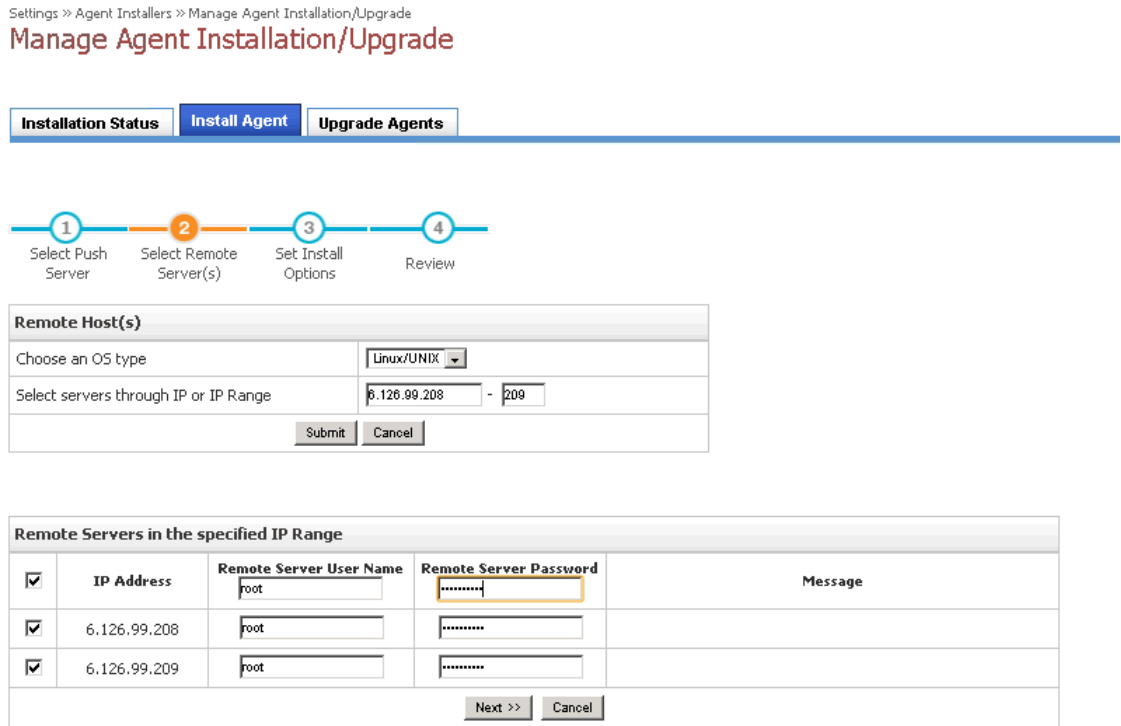
- Step 1** Select **Push Server**. Select the CX server as the push server.
- Step 2** Select **Remote Server**. Enter the IP address range of the primary server(s). Click **Submit** to enter the primary server username/password. Enter the username/password at the top row if the same username/password will be used for all servers; otherwise, enter the username/password next to individual servers.

**Figure 3-11 Home Installer 2**



- Step 3** Set **Install Options**. This is where installation directory, agent type, and CX IP/port is entered. Installation directory defaults to /usr/local/InMage/, Unified Agent (both file and volume agent), and source Scout Agent.

**Figure 3-12 Home Installer 3**



- Step 4** Click **Run** after you have reviewed information for accuracy.

**Note**

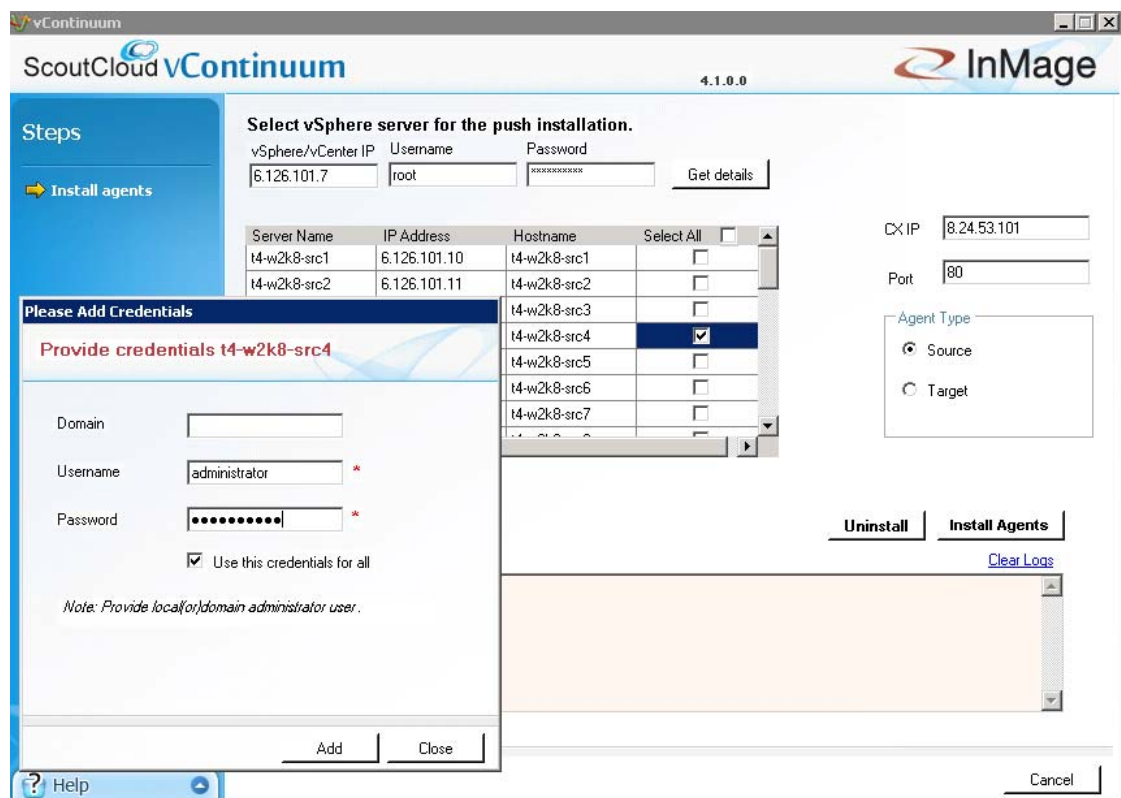
For Linux install, ensure SSH access to the primary server has been enabled.

## vContinuum for Windows Bulk Install and Upgrade

Agent update/install for Windows primary servers is supported from the tenant vContinuum server. vContinuum and CX server pre-bundles compatible windows source agent unlike Linux bulk install. A separate upload into software repository is not required. To install agents, access the vContinuum UI from the vContinuum server and select **Push Agent**. This is a four-step process:

- Step 1** Provide vCenter information.
- Step 2** Select server(s) to install the agent. Currently, the UI does not indicate agent status. It is not possible to determine which server has agent already been installed.
- Step 3** Provide credentials for the server. Default setting is to apply the same credentials to all servers.

**Figure 3-13 vContinuum: Provide Credentials**



- Step 4** vContinuum validates if servers can be reached with provided credentials. Click **Install Agent** to begin the agent install.

A number of outstanding enhancements exist to simplify the user interface. Refer to “[Monitoring, Best Practices, Caveats, and Troubleshooting](#)” section on page 5-1 for details.

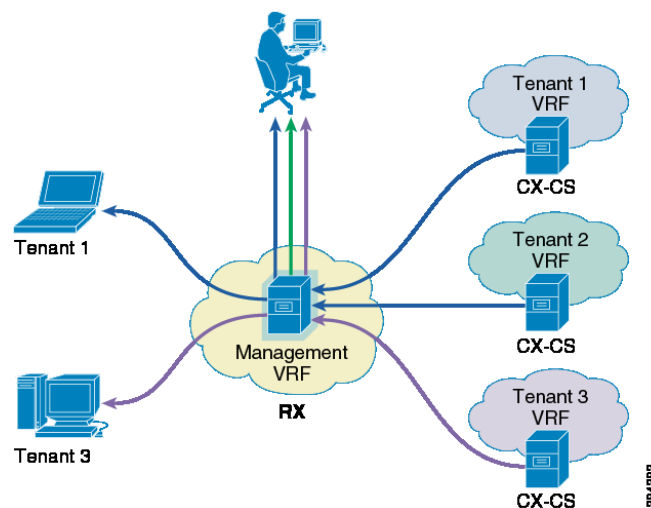


## Multi-Tenant Portal—RX Server

The RX server is a multi-tenant portal for end customers to monitor the protected servers. It also enables the SP to consolidate administration to a single interface rather than a dedicated CX UI for each customer/CX server, greatly simplifying monitoring. The RX server enables:

- Centralized monitoring for all CS servers and source servers for all tenants.
- Central user account management.
- Consolidated reports of protection health of servers, servers details, usage statistics, RPO, and alerts.
- Publish news and announcements by the SP to users.

**Figure 3-14 Multi-Tenant Portal**



As displayed in [Figure 3-14](#), the RX server is an aggregation point for statistics coming from tenant-specific CX-CS servers. Based on role-based access control (RBAC), a tenant user can be limited to monitor statistics from a single CX-CS server while a Managed Services Provider (MSP) user can be assigned to monitor multiple deployments of CX-CS from a single interface.

This section includes the following topics:

- [Multi-Tenant Portal Version Compatibility](#), page 3-17
- [Multi-Tenant Portal User Accounts](#), page 3-18
- [CX Server Registration with Multi-Tenant Portal](#), page 3-20
- [Multi-Tenant Portal Rebranding](#), page 3-21

## Multi-Tenant Portal Version Compatibility

The 7.1 RX server can be either installed on a 64 bit ScoutOS version 5 and Update5, or a ScoutOS version 6 and update 2. When integrating RX with the CX-CS, RX version is always backwards compatible with the CX-CS version. As a rule of thumb, RX should be always higher or equal to the CX-CS version. [Table 3-4](#) is a compatibility table with ScoutOS 6.2 CX.

**Table 3-4** Compatibility Table

	ScoutOS 5.5 RX	ScoutOS 6.2 RX
ScoutOS 5.5 CX	Yes	Yes
ScoutOS 6.2 CX	No	Yes

7.1 RX is compatible with 7.1 GA CX without additional updates/hotfixes. Compatibility with an earlier version of CX release may require additional updates/hotfixes. Refer to InMage Scout Cloud RX release notes for details.

Although the number of users accessing the multi-tenant portal and CX pairing will vary, a minimum of 150G of disk, 1 vCPU, and 4G of memory should be allocated to the VM hosting the RX server.

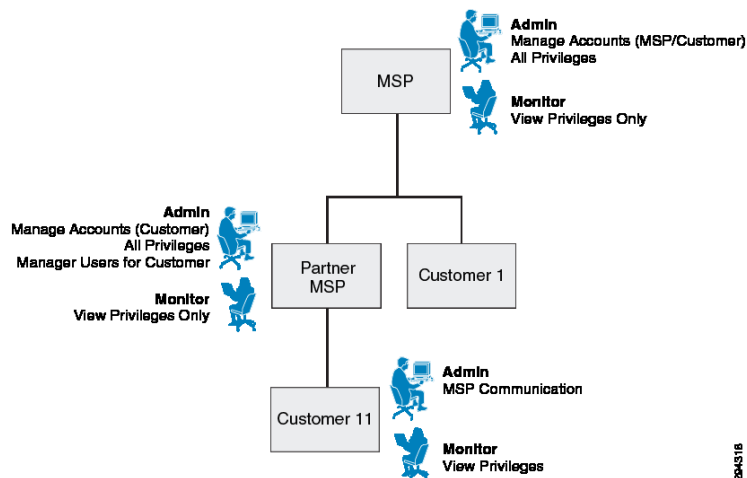
## Multi-Tenant Portal User Accounts

Figure 3-15 lists out various types of user accounts that can be created in the multi-tenant portal. There are two major types of accounts: MSP and Customer:

The root MSP account is created by default with the RX installation; it has full access to all features and data for all customers. Additional sub-MSP accounts can be created in a managed services/white labeling deployment. Account data view would be limited to their own customers.

Customer accounts can be created when new tenants are added.

Both root and sub-MSP accounts can create additional customers and associate the corresponding CXCS server for multi-tenant portal access.

**Figure 3-15** Multi-Tenancy Overview

Users can be set up as monitor user (read only) or administrator (read/write). The ability to perform self-serviced recovery from the RX portal can be allowed per tenant based on the recovery setting. To enable a user for self-serviced recovery, first navigate to Accounts > Customer Accounts > Edit customer. Simply select the check box next to Allow to Perform Recovery Options, as shown in [Figure 3-16](#).

**Figure 3-16 Self Service Recovery Setting**

Accounts » Customer Accounts » Edit Customer

### Edit Customer Account

**Company Details**

Tenant 9	Reference ID: 9 Contract ID: 9 Contract Start Date: 05/22/2013 Contract Expiry Date: 05/01/2014
----------	--

**User Details**

Full Name	User Name	User Type	Creation Date	Account Status	Action
Tenant 9	tenant9	Administrator	2013-05-22 15:20:17	Active	✎

**Assigned CS Servers**

IP Address	Host Name
8.24.81.101	sp-t9-ps-1

**Support Contact**

**Recovery Settings**

Allow to Perform Recovery Operations    (
  Allow Hardware and Network Configuration )

Support: tenant9@cs.sdu

**Recovery Settings**

Allowed to Perform Recovery Operations with Advanced Options (Hardware and Network Configuration)

The CVD validation focused on SP-managed deployment with direct tenants. [Figure 3-17](#) is a summary view.

Figure 3-17 Summary Tenant View

InMage Systems							
Company	Reference ID	Contract ID	Contract Start Date	Contract Expiry Date	Recovery Allowed?	Account Status	View Dashboard
Tenant 1	1	1	05/22/2013	05/01/2014	✓	✓	
Tenant 10	10	10	05/22/2013	05/01/2014	✓	✓	
Tenant 11	11	11	05/22/2013	05/01/2014	✓	✓	
Tenant 12	12	12	05/22/2013	05/01/2014	✓	✓	
Tenant 2	2	2	05/22/2013	05/01/2014	✓	✓	
Tenant 3	3	3	05/22/2013	05/01/2014	✓	✓	
Tenant 4	4	4	05/22/2013	05/01/2014	✓	✓	
Tenant 5	5	5	05/22/2013	05/01/2014	✓	✓	
Tenant 6	6	6	05/22/2013	05/01/2014	✓	✓	
Tenant 7	7	7	05/22/2013	05/01/2014	✓	✓	
Tenant 8	8	8	05/22/2013	05/01/2014	✓	✓	
Tenant 9	9	9	05/22/2013	05/01/2014	✓	✓	

As displayed in Figure 3-17, in an SP-managed scenario, all tenants are mapped under the default/root MSP account, InMage Systems. CX-CS servers are mapped to each company, and multiple users can exist per company. When a root MSP user logs in, a summary list of tenants along with tenant status, dashboard, and self-recovery status is displayed.

## CX Server Registration with Multi-Tenant Portal

You can point the configuration server to Scout Cloud RX in the following two ways:

- From the RX UI, also known as Pull Method, where Cloud RX pulls the required data from all the registered CX-CS servers OR
- From the CX-CS UI, also known as Pull/Push Method. This is when CX-CS is behind firewall and RX cannot initiate connection to the CX-CS, CX-CS can push data to RX using the Push method.

To register CX-CS with Scout Cloud RX through RX UI:

- 
- Step 1** Navigate to Settings > CS Server.
  - Step 2** Enter the CX IP address or IP Range along with the HTTP port and click Discover.
  - Step 3** CX Servers in the specified IP range are displayed along with the details such as Server IP, Status, and Host Name.
  - Step 4** Enter user name, password, and alias name for CX. This ensures secured CS registration.
  - Step 5** Select the customer account. Select the CS Server and click Register with RX.
- 



### Note

Discover / CX-CS registration is available to MSP admin users only.

To register CX with Scout Cloud RX through CX-CS UI:

- Step 1** Navigate to Settings > RX Server.
- Step 2** Enter the RX IP address and CX alias name. Synchronize data mode defaults to Allow CX to push data to RX.
- Step 3** CX server will be mapped to the "Unassigned CS Servers" pool as shown in [Figure 3-18](#).

**Figure 3-18** Newly Added CX-CS Server In Unassigned Pool

Registered CS Servers										
UnAssigned CS Servers										
	CS Server IP	CS Server Name	CS Alias Name	CS Version	Registration Date	Last Synchronization Time	Synchronization Interval (minutes)	Data Synchronization Method	Communication Type	Action
<input type="checkbox"/>	8.24.85.101	sp-t10-ps-1	sp-t10-ps-1		2013-07-06 15:32:24		5	PUSH	HTTP	

- Step 4** Select the checkbox and click Action to assign the CX-CS server to Customer. Once assigned, the CS server is moved from the unassigned pool to a specific customer as shown in [Figure 3-19](#).

**Figure 3-19** Newly Added CX-CS Server Assigned to Customer

Tenant 10										
	CS Server IP	CS Server Name	CS Alias Name	CS Version	Registration Date	Last Synchronization Time	Synchronization Interval (minutes)	Data Synchronization Method	Communication Type	Action
<input type="checkbox"/>	8.24.85.101	sp-t10-ps-1	sp-t10-ps-1	7.1.0.0	2013-07-06 15:32:24		5	PUSH	HTTP	



**Note**

If a CX-CS server is configured with dual NICs, one is to communicate with remote Scout agents and the other is to communicate with RX. Use the push method instead of pull when registering with the RX server. This is a known limitation and will be addressed in future releases of InMage software.

## Multi-Tenant Portal Rebranding

InMage's Scout Product provides the ability to support user interface customizations as required by its partners. This rebranding allows changing the user interface to carry the partner's company and product name. Also, it allows customizing color and graphics for uniformity across partner products.

Refer to InMage\_Scout\_Cloud\_RX\_Branding document for additional details.

## Summary Tables of Components for All Tenants

This section includes the following topics:

- [InMage Components for Service Provider, page 3-22](#)
- [InMage Component for Enterprise Tenants, page 3-23](#)

## InMage Components for Service Provider

Based on sizing discussion in earlier sections, twelve tenants were configured based on various change rates. [Table 3-5](#) summarizes disk / CPU and memory configuration required for InMage control servers.

**Table 3-5** *Disk/CPU/Memory Configuration (Service Provider)*

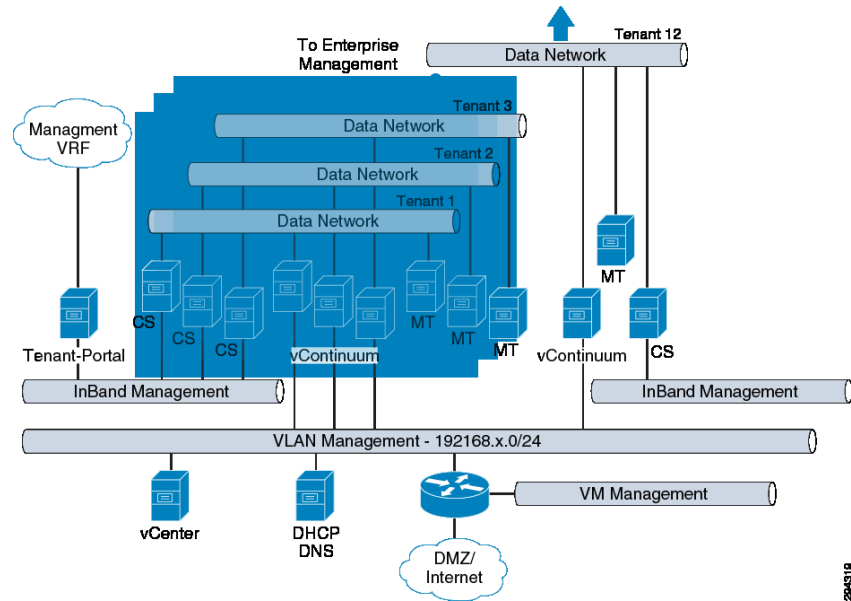
Tenant	CX-PS (GB)	vCPU	RAM (GB)	MT Retention (GB)	MT Cache	vCPU	RAM (GB)
1	800	8	32	2X 300	2X40	2X4	2X16
2	250	4	8	300	40	4	16
3	800	8	32	300	40	4	16
4	800	8	16	300	40	4	16
5	800	8	16	300	40	4	16
6	800	8	16	300	40	4	16
7	800	8	16	300	40	4	16
8	800	8	16	300	40	4	16
9	250	4	8	50	40	2	8
10	250	4	8	50	40	2	8
11	250	4	8	50	40	2	8
12	250	4	8	50	40	2	8

- Tenants 1 and 3 are sized based on 1TB change rate per day.
- Tenants 4, 5, 6, 7 and 8 are sized based on 750GB change rate per day.
- Tenants 2, 9, 10, 11, and 12 are sized based on 350G change rate per day.

MT retention is based on time (9 hours) and space capacity. With the exception of Tenant 2, all tenants are Windows-based tenants running Windows 2008 R2. vContinuum is installed directly on the MT. Tenant 2 is a Linux-only tenant, running CentOS 6.3; as such, a dedicated vContinuum server running windows 2008R2 with 1 vCPU, 3G memory, and 50G disk was utilized. With the exception of Tenants 9 - 12, all MTs are deployed with 4 vCPU, 16G memory, and 300G retention drive.

Deployment topology and logical connectivity between InMage control servers are captured in [Figure 3-20](#).

Figure 3-20 Deployment Topology and Logical Connectivity between InMage Control Servers



From the SP VPC, communication with enterprise management is via inband data NIC part of VMDC server VLAN. A common out-of-band (OOB) management VLAN was extended from the Cloud Validation Facility management POD to the VMDC 2.3 infrastructure. OOB management VLAN is used to communicate with shared resources residing in the management PoD such as vCenter, Active Directory, vCenter Database, and DHCP server.

## InMage Component for Enterprise Tenants

Twelve Enterprise tenants were configured to mirror the SP setup. Table 3-6 summarizes Disk / CPU and memory configuration required for enterprise InMage component. Retention policy, vCPU, and RAM requirements mirror that of the SP site based on 1TB, 750GB, and 350GB change rate/day. MT is deployed in the enterprise for failback protection. Enterprise master target is expected to be idle and not needed until failback after a disaster recovery. It can be pre-provisioned or installed during time of failback. Each tenant will have a dedicated vCenter server managing ESXi resources. Tenant 1 and Tenant 2 have both physical and virtual servers. Both P2V and V2V protection are needed. To generate the desired data change rate, each enterprise is assigned a set of machine machines (Table 3-6).

Table 3-6 Disk/CPU/Memory Configuration (Enterprise)

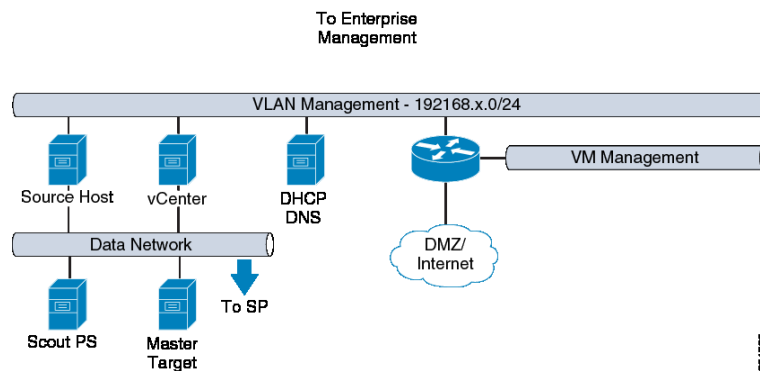
Tenant	PS (GB)	vCPU	RAM (GB)	MT Retention (GB)	MT Cache (GB)	vCPU	RAM (GB)	vCenter disk (GB)	vCPU	RAM (GB)	Total VM
1	800	8	32	2 X 300	2 X 40	2X 4	2 X 16	100	2	4	60
2	250	4	8	300	40	4	16	100	2	4	10
3	800	8	32	300	40	4	16	100	2	4	50
4	800	8	16	300	40	4	16	100	2	4	30
5	800	8	16	300	40	4	16	100	2	4	30
6	800	8	16	300	40	4	16	100	2	4	30
7	800	8	16	300	40	4	16	100	2	4	30

Table 3-6 Disk/CPU/Memory Configuration (Enterprise) (continued)

Tenant	PS (GB)	vCPU	RAM (GB)	MT Retention (GB)	MT Cache (GB)	vCPU	RAM (GB)	vCenter disk (GB)	vCPU	RAM (GB)	Total VM
8	800	8	16	300	40	4	16	100	2	4	30
9	250	4	8	50	40	2	8	100	2	4	7
10	250	4	8	50	40	2	8	100	2	4	7
11	250	4	8	50	40	2	8	100	2	4	8
12	250	4	8	50	40	2	8	100	2	4	8

Deployment topology and logical connectivity between InMage Enterprise control servers are captured in Figure 3-21.

Figure 3-21 Deployment Topology and Logical Connectivity between InMage Enterprise Control Servers



From the Enterprise VPC, communication with the SP management is via the inband data NIC part of VMDC server VLANs. A common OOB management VLAN was extended from the Cloud Validation Facility Management PoD to the VMDC 2.2-based enterprise VPC. OOB management VLAN is used to communicate with shared resources residing in the Management PoD such as vCenter, Active Directory, vCenter Database, and DHCP server.

## VMDC 2.3

This section includes the following topics:

- [VMDC 2.3 Integrated Compute and Storage Stack, page 3-25k](#)
- [Mapping DR Components to VMDC 2.3 Containers, page 3-32](#)
- [Tenant Configuration, page 3-33](#)
- [Connectivity across the WAN, page 3-39](#)
- [Storage Configuration, page 3-42](#)



## VMDC 2.3 Integrated Compute and Storage Stack

The VMDC 2.3 release uses modular blocks for compute and storage, generically referred to as Integrated Compute and Storage (ICS) stacks. Several stacks can be attached to a PoD, providing compute and storage scale. With VMDC 2.3, three ICS stacks can be connected to the Nexus 7004, with 4x 10G links per aggregation switch, for a total of 80 Gbps to the ICS switch layer. Refer to the Cisco VMDC 2.3 Design Guide at <http://www.inwats.cisco.com/publications/viewdoc.php?docid=6637> for more discussion of the scaling factors.

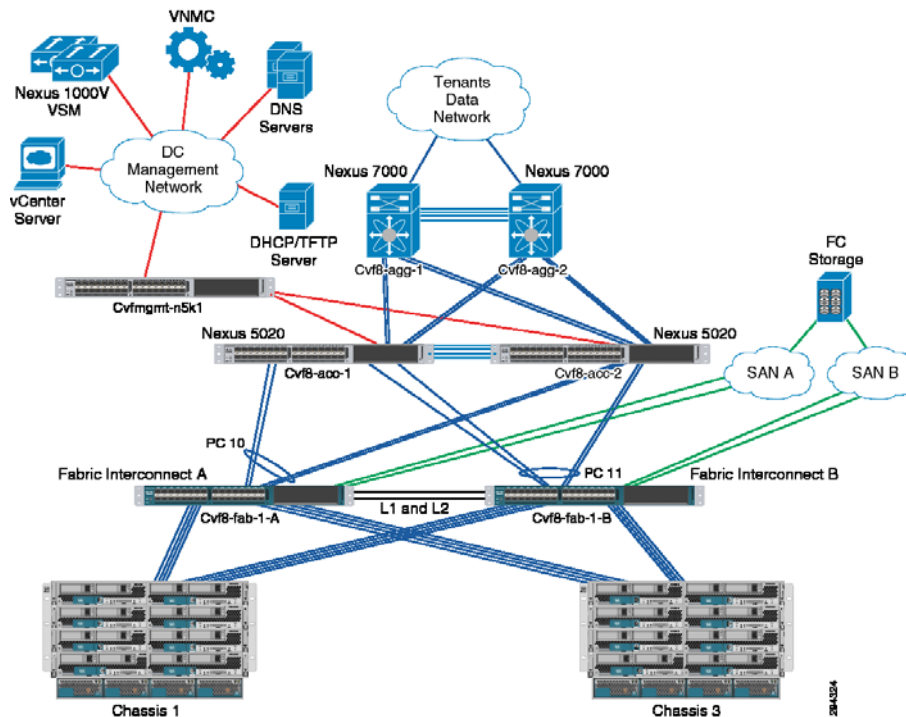
In our implementation, a smaller footprint ICS was built as listed in [Table 3-7](#)

**Table 3-7 ICS Footprint**

Tenant Type	Number of Tenants	Number of VLANs per Tenant	Number of VMs
Gold	4	3	30
Silver	2	3	60
Bronze	6	1	160
Total	12	24	250

The ICS design uses the VNX 5500 as the SAN storage. The details of the ICS buildout are covered in [Figure 3-22](#).

**Figure 3-22 ICS Buildout**



This section includes the following topics:

- [UCS Implementation, page 3-26](#)

- [ESXi Implementation, page 3-28](#)
- [Nexus 1000V, page 3-30](#)
- [VSG Implementation, page 3-32](#)

## UCS Implementation

[Table 3-8](#) shows UCS components for implementation.

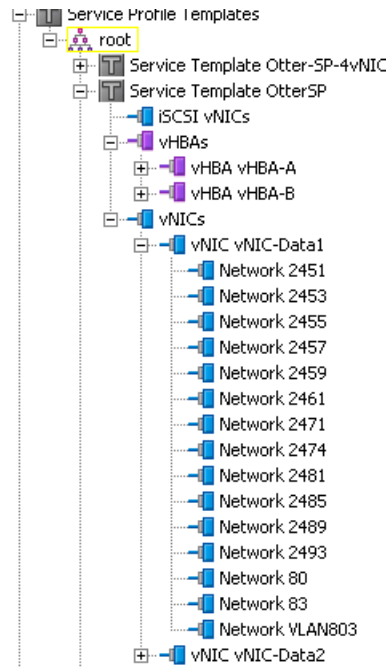
**Table 3-8 UCS Implementation**

Component	Product Name	Quantity
Fabric Interconnect (FI)	Cisco UCS 6120	2
Chassis	Cisco UCS 5108	2
I/O Module	Cisco UCS 2104XP	4
Blade Server	Cisco UCS B200 M3 (2 x 8 cores CPU, 196GB Memory)	3
Blade Server	Cisco UCS B200 M2 (2 x 6 cores CPU, 96GB Memory)	9
Adapter	Cisco UCS VIC 1240	3
Adapter	Cisco UCS M81KR	9

Two UCS 5108 chassis are connected to a pair of UCS 6120 FIs. Each chassis has four server links to each FI. The UCS FIs are configured in End Host (EH) mode into a cluster to provide active/standby management plane redundancy for the UCSM, active/active for data forwarding. The uplinks on the FIs are bundled into port-channels to the upstream Nexus 5000 switch. Both management and data traffic are carried in the same port-channel. Nexus 5000 switches with Fibre Channel (FC) links for access to SAN storage. Each UCS blade is configured with two vHBAs for access to SAN storage via SAN-A and SAN-B for storage multipathing.

The UCSM service-profile in [Figure 3-23](#) is used on the UCSM.

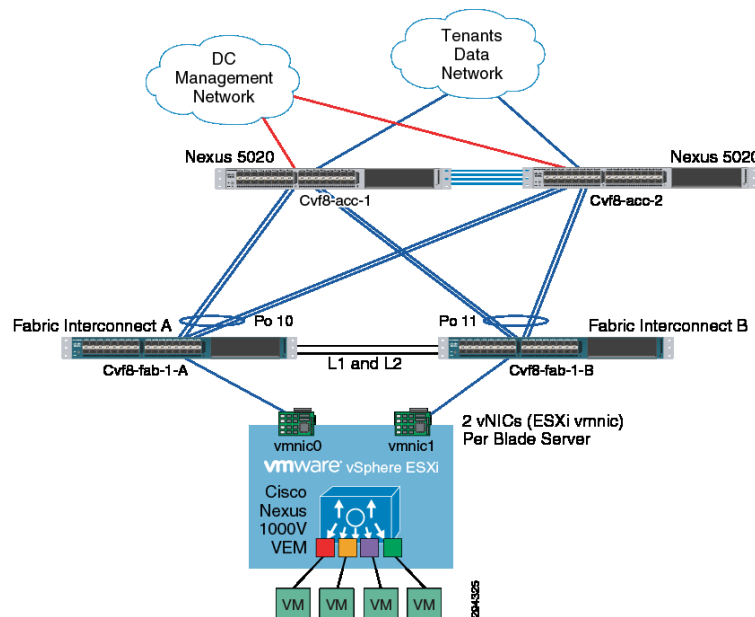
Figure 3-23 UCSM Service-Profile



Each service-profile is associated with a respective server pool and blades from both chassis are made available to the server pool.

The UCS Fabric Interconnects are connected to a pair of Nexus 5000 for redundancy. Both FIs are configured with the same set of VLANs. All VLANs are trunked to all available uplinks and the same uplink port channel is used to carry both management and tenant data traffic.

Figure 3-24 UCS Implementation



## ESXi Implementation

A VMware vSphere Cluster is a grouping of servers with similar characteristics and can be used as one unified compute resource. VMware vSphere Cluster is the foundation used to achieve a pooling of resources, HA, and Distributed Resource Scheduling.

Refer to the following documents for more information on VMware HA and VMware Distributed Resource Scheduler:

- HA Deepdive at <http://www.yellow-bricks.com/vmware-high-availability-deepdiv/>
- Distributed Resource Scheduler Deepdive at <http://www.yellow-bricks.com/drs-deepdive/>

The following set of clusters were created based on the recommendations provided in VMDC 2.3 Design Guide at [http://www.cisco.com/en/US/docs/solutions/Enterprise/Data\\_Center/VMDC/2.3/design\\_guide/VMDC\\_2.3\\_DG\\_1.html#wp1361952](http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.3/design_guide/VMDC_2.3_DG_1.html#wp1361952).

**Table 3-9 Host per Cluster**

Cluster Type	Number of Hosts	Memory	CPU
Bronze	5	900GB	207 GHz
Silver	3	500 GB	129 GHz
Gold	4	700 GB	165 HZ

It is recommended to size the number of hosts per cluster based on capacity and HA requirements of each individual implementation. A minimum of three hosts is recommended to provide non-disruptive host maintenance without loss of VMware HA protection. Cluster size varies from 3 ESXi hosts to 5 ESXi depending on workload type.

The following set of VMware HA parameters is necessary to define capacity requirements and isolation response:

- Admission Control
- Number of Host Failures Permitted
- Restart Priority
- Isolation Response
- Host Monitoring Status

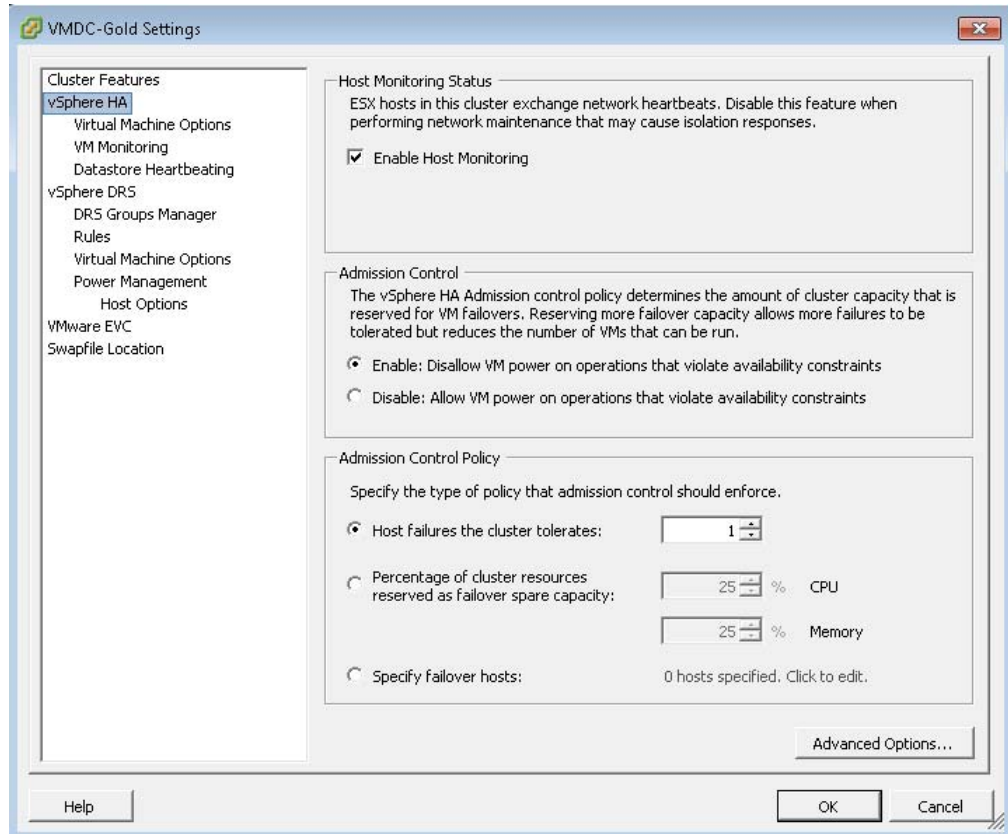
Based on the recommendations provided in VMDC , all three clusters are assigned the same HA parameters. [Table 3-10](#) provides the sample settings for the Gold Cluster.

**Table 3-10 Sample Settings for the Gold Cluster**

Category	Setting
Admission Control	Enabled: Do not power on VMs that violate availability constraints
Number of Host Failures Permitted	1 host failures cluster tolerates
Restart Priority	Medium
Isolation Response	Leave VM powered on
Host Monitoring Status	Enabled

[Figure 3-25](#) shows the Gold Cluster parameters.

Figure 3-25 Gold Cluster HA -1



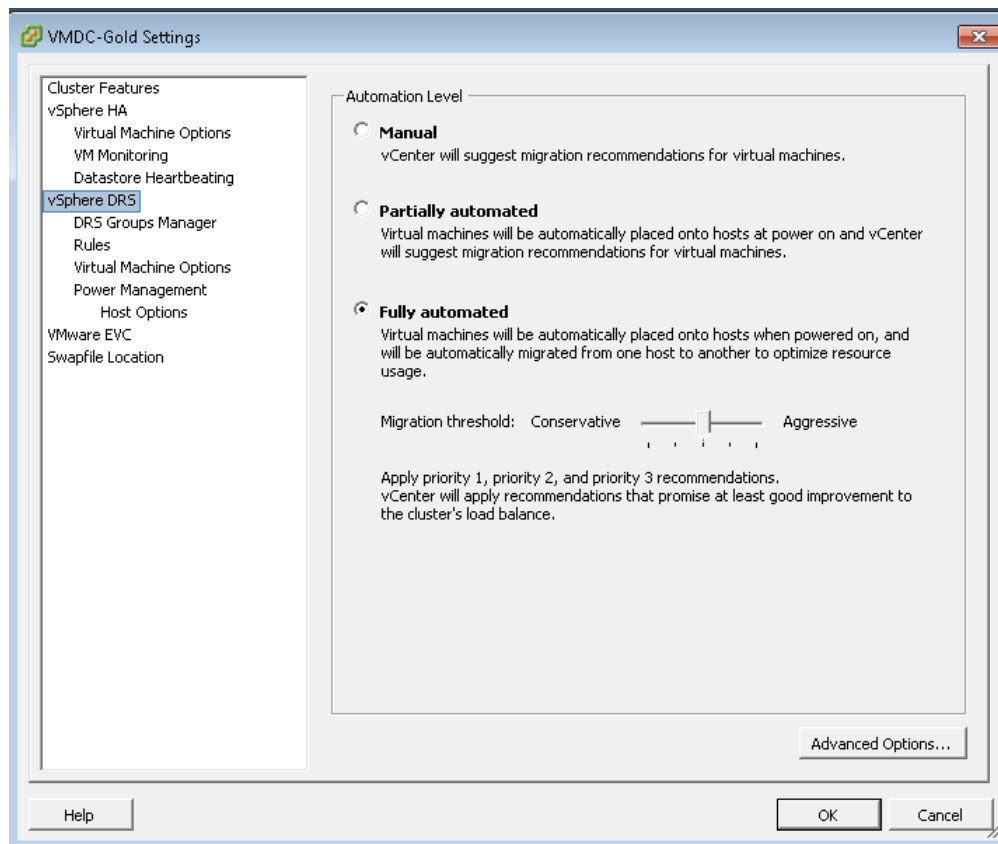
The VMware Distributed Resource Scheduler functions by monitoring the VM (CPU and memory) loads in a virtual computer cluster and, if necessary, moves the VMs from one physical ESX server to another in an attempt to load balance the workload. Distributed Resource Scheduler works in one of three modes: fully automatic, partially automatic, or manual. Based on the recommendations provided in the VMDC 2.3 Design Guide at [http://www.cisco.com/en/US/docs/solutions/Enterprise/Data\\_Center/VMDC/2.3/design\\_guide/VMDC\\_2.3\\_DG\\_1.html#wp1361952](http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.3/design_guide/VMDC_2.3_DG_1.html#wp1361952), all three clusters are assigned the same Distributed Resource Scheduler parameters.

Table 3-11 Distributed Resource Scheduler Parameters

Distributed Resource Scheduler Mode	Fully Automated
Migration Threshold	3 stars

Figure 3-26 shows the Gold Distributed Resource Scheduler setting.

Figure 3-26 Gold Cluster HA -1



vSphere supports cluster sizes of up to 32 servers when HA and/or DRS features are utilized. In general practice, however, the larger the scale of the compute environment and the higher the virtualization (VM, network interface, and port) requirement, the more advisable it is to use smaller cluster sizes to optimize performance and virtual interface port scale. Therefore, in large VMDC deployments, cluster sizes are limited to eight servers; in smaller deployments, cluster sizes of 16 or 32 can be utilized. Gold, Silver, and Bronze compute profiles are created to represent Large, Medium, and Small workload types. Gold has one vCPU/core and 16 GB RAM, Silver has 0.5 vCPU/core and 8 GB RAM, and Bronze has 0.25 vCPU/core and 4 GB of RAM.

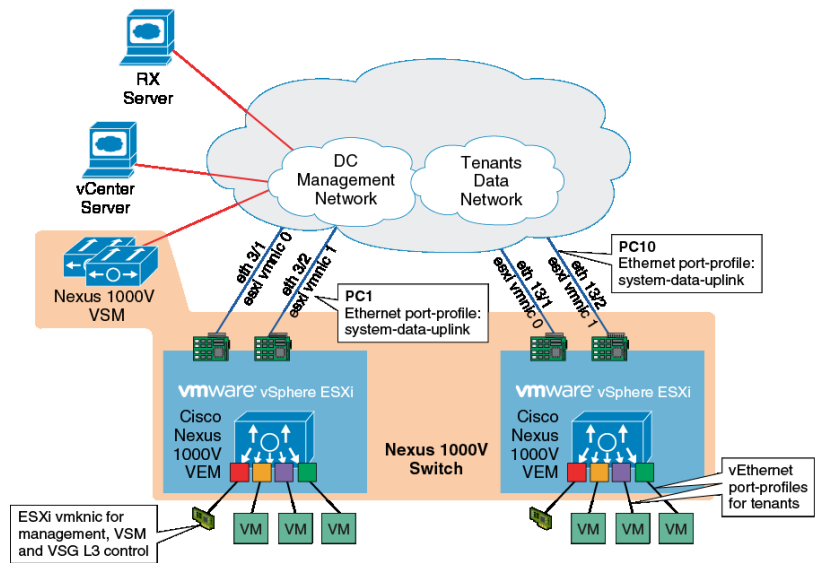
While the VMDC 2.3 architecture works with Vblocks and FlexPods, the system has been validated with VNX 5500.

## Nexus 1000V

The Nexus 1000V Series Switches provide a comprehensive and extensible architectural platform for VM and cloud networking. In our implementation, all networking needs of the VMs are provided by Nexus 1000V Series Switches and it is implemented identically to the VMDC 2.3 specification.

Without reduplicating VMDC 2.3 documentation, [Figure 3-27](#) is a summary description of the implementation.

Figure 3-27 Summary Description of Nexus 1000v Implementation



The Nexus 1000V VSM is configured in L3 SVS mode. In L3 SVS mode, VSM encapsulates the control and packet frames into User Datagram Protocol (UDP) packets. The VSM uses its mgmt0 interface to communicate with the VEMs. The VEMs are located in a different IP subnet from the VSM mgmt0 interface. On each VEM, the vmk0 vmkernel interface is used to communicate with the VSM. The following configuration shows the VSM svcs-domain configuration:

```
svs-domain
domain id 1
control vlan 1
packet vlan 1
svs mode L3 interface mgmt0
```

The UCS is configured with EHV mode and has the upstream L2 switches performing the split between the management and customer production data domains. Each ESXi/VEM host is configured with two NICs (also referred to as the ESXi VM Network Interface Card (VMNIC) or UCS vNIC), carrying both management network and tenants' data network (for UCS Fabric A - fabric B redundancy). On the Nexus 1000V, the following configuration shows the Ethernet port-profile configuration:

```
port-profile type ethernet system-data-uplink
vmware port-group
switchport trunk allowed vlan 80,83,2451,2453,2455,2457,2459,2461,2471
switchport trunk allowed vlan add 2474,2481,2485,2489,2493
switchport mode trunk
switchport trunk native vlan 83
channel-group auto mode on mac-pinning
no shutdown
system vlan 83
max-ports 32
state enabled
```

When the ESXi host is added to the Nexus 1000V DVS, the vmnic0 and vmnic1 interfaces are attached to the system-data-uplink Ethernet uplink port profile. In this implementation, the vmknic ESXi kernel interfaces (vmk0) are also managed by the Nexus 1000V. The following shows the configuration used for the ESXi management.

```
port-profile type vethernet esxi-mgmt-vmknic
capability l3control
vmware port-group
switchport mode access
pinning id 0
```

```

switchport access vlan 83
no shutdown
system vlan 83
max-ports 32
state enabled

```

Refer to VMDC 2.3 (Nexus 1000V Series Switches at [http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Data\\_Center/VMDC/2.3/implementation\\_guide/VMDC2.3\\_IG2.html#wp2272450](http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Data_Center/VMDC/2.3/implementation_guide/VMDC2.3_IG2.html#wp2272450)) for additional details.

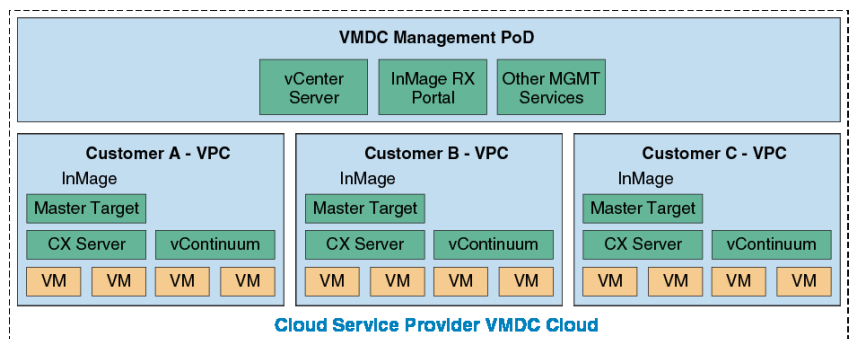
## VSG Implementation

For information about VSG implementation, please see: [http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Data\\_Center/VMDC/2.3/implementation\\_guide/VMDC2.3\\_IG5.html#wp2277285](http://www.cisco.com/en/US/partner/docs/solutions/Enterprise/Data_Center/VMDC/2.3/implementation_guide/VMDC2.3_IG5.html#wp2277285)

## Mapping DR Components to VMDC 2.3 Containers

As previously discussed in “InMage Components for Service Provider” section on page 3-22, a common management PoD is deployed to host shared services (example: vCenter, DHCP, and DNS). The multi-tenant portal (RX) is part of the management PoD. The remaining InMage specific components are co-located with the production ICS cluster corresponding to service tier (Gold / Silver / Bronze). Refer to [Figure 3-28](#).

**Figure 3-28** VMDC Management PoD



InMage RX server will be deployed with dual NICs:

- Management NIC: Used to communicate with the Configuration Server (CX).
- OOB Management NIC: Server VLAN part of a dedicated VMDC 2.3 container used for OOB tenant access. Refer to “[Out of Band Management Portal Access](#)” section on page 3-34.

InMage CX-CS servers will be deployed with dual NICs:

- Management NIC: Used to communicate with the Multi-Tenant Portal (RX).
- Data NIC: VMDC 2.3 Server VLAN: Used for Scout Agent, MT, and PS server registration and configuration updates.

Master Target: Two deployment scenarios are based on OS:

- InMage Windows MT is deployed with dual NICS because vContinuum is colocated with the MT:
  - Management NIC; used by vContinuum to communicate with shared vCenter in the SP Cloud.
  - Data NIC; used by vContinuum to communicate with tenant vCenter in the enterprise private cloud., and to receive data changes from the enterprise private cloud.



- InMage Linux MT will be deployed with a single NIC: – Data NIC:
  - Used to receive data change from Enterprise Private Cloud - Communication with vCenter is from vContinuum.

Virtual Machines - VM settings are configured during recovery, refer to [Chapter 4, “Recovery Workflows,”](#) for details. Based on the VMDC specification, each VM will have two NICs:

- Management NIC—Protected by the VSG and accessible by a cloud orchestration system, such as BMC or CIAC.
- Data NIC—VMDC 2.3 Server VLAN.

## Tenant Configuration

This section includes the following topics:

- [IPsec, page 3-33](#)
- [Out of Band Management Portal Access, page 3-34](#)
- [VMDC Container Modifications, page 3-35](#)

## IPsec

When sending data across the WAN from a primary site to a secondary site, securing the data transmission is critical. Data may be left vulnerable if encryption between the CX-PS server to a MT in the secondary site is not enabled. Two encryption points are possible in the architecture:


- Primary Server: Encrypt data from the source primary server to local InMage CX-PS.
  - This option requires the InMage agent (DataTap) on the primary server under protection to encrypt the data before sending it to the processing server. Scout PS server does not support protection of a device when it is encrypted outside of InMage. This approach will consume CPU/memory and potentially other resources on the production server. InMage recommends to not use this option unless deployment topology prevents a processing server from being deployed.
- Processing Server: Encrypt data from InMage CX-PS to MT before transmitting the packet out of the enterprise WAN.
  - This option is preferred as it assumes encryption is performed by an external device, not the production server. A number of ways are possible for accomplishing this. InMage supports data encryption within the processing server or external encryption appliance such as the Cisco ASA. Either option will secure the data transmission; however, it is important to remember if the processing server is encrypting traffic across the WAN. Performance penalties can occur on the processing server when compared to unencrypted transmissions.

To simplify capacity management and operational support, external ASA appliances were used in our implementation. ASA pairs deployed in multi-context mode are used to secure the communication between the Enterprise LAN and SP LAN across a L3 MPLS VPN. The benefits of ASA are:

- Multi-context support: Tenant separation and delegation of role / responsibility.
- Capacity monitoring and management: Single pair of ASA VSXs deploying additional vCPU/ MEM to each processing server to support encryption.
- Operational Monitoring: Single tunnel per customer vs. encryption setting per disk/volume under protection

All traffic needing encryption between the enterprise and service are redirected to the IPsec tunnel based on static routing. ASA appliance will encrypt the entire frame and re-encapsulate it with an additional IPsec header:

**Figure 3-29** IPsec Header




Refer to “VMDC Container Modifications” section on page 3-35 Updates for details on routing changes.

## Out of Band Management Portal Access

The topology in Figure 3-30 is used to validate the functions and features of the InMage Multi-tenant Portal.

**Figure 3-30** Out of Band Management Portal Access Topology



All DRaaS users share a VMDC 2.3 Gold container for management access. A Gold container is used as it offers the guaranteed resources as well as flexibility in network container resources such as firewall and server load balancing servers.

The Enterprise route targets are imported on the Cisco ASR 1000 for MPLS connectivity across the WAN.

For example:

```
vrf definition mgmt-gold-pub rd 800:800
route-target export 800:800 route-target import 800:800 route-target import 2451:2451
route-target import 3003:3003 route-target import 2471:2471 route-target import 3008:3008
route-target import 2474:2474 route-target import 3015:3015 route-target import 2453:2453
route-target import 3019:3019 route-target import 2455:2455 route-target import 3020:3020
route-target import 2457:2457 route-target import 3023:3023 route-target import 2459:2459
route-target import 3027:3027 route-target import 2461:2461 route-target import 3029:3029
route-target import 3482:3482 route-target import 3010:3010 route-target import 3040:3040
route-target import 3042:3042 route-target import 3034:3034
```

The ASR 1000 peers downstream with the Cisco Nexus 7000 aggregation to learn specific prefixes required for management access.

Traffic flows from the Enterprise to the Cisco ASR 1000 via the MPLS core, and the Cisco ASR 1000 follows the path of the VMDC 2.3 Gold container until it arrives on the Cisco Nexus 7000 aggregation public VRF. From the Cisco Nexus 7000 aggregation, traffic is sent across the firewall to the load balancer for server load balancing (if needed).

## VMDC Container Modifications

This section includes the following topics:

- [VMDC Gold, page 3-35](#)
- [VMDC Silver, page 3-36r](#)
- [VMDC Bronze, page 3-38](#)

### VMDC Gold

Changes to the VMDC Gold container were implicit to the overall architecture. No modifications or changes were made to the baseline VMDC Gold container; server VLANs were neither introduced nor removed. Non-InMage/DRaaS-related traffic streams follow identical network paths as documented in the VMDC 2.3 Design and Implementation guide (<http://www.in-wats.cisco.com/publications/viewdoc.php?docid=6637>). Traffic sourced from the VPC will first be routed to the private VRF default gateway based on default routing. Traffic will be forwarded from the private VRF to the public VRF across the vFW. Once traffic arrives in the public VRF, it will be L3 routed to the ASR1K (PE) towards the L3 VPN.

Traffic to and from the InMage server from the primary enterprise site needs to be forwarded to the IPsec tunnel for encryption or decryption.

- This was accomplished by setting up a static site-to-site tunnel between the SP and enterprise, placing the IPsec-inside interface on Gold server VLAN 1 and the outside interface on Gold server VLAN 2.
- InMage traffic from the enterprise LAN will be sent to the ASA-outside interface residing on server VLAN 2.
- Once received traffic is decrypted by the ASA, it will be forwarded to InMage servers residing on server VLAN 1 via the IPsec-inside interface.
- InMage traffic from the SP to Enterprise is accomplished by adding static routes on the CX and MT server pointing to the inside interface of the ASA.
- Once received traffic is encrypted by the ASA, it will be forwarded out of the ASA-outside interface residing on server VLAN 2 towards the HSRP address of the private VRF interface.
- Once encrypted traffic is received in server VLAN 2, normal VMDC traffic flow occurs.

See [Figure 3-31](#) for details.

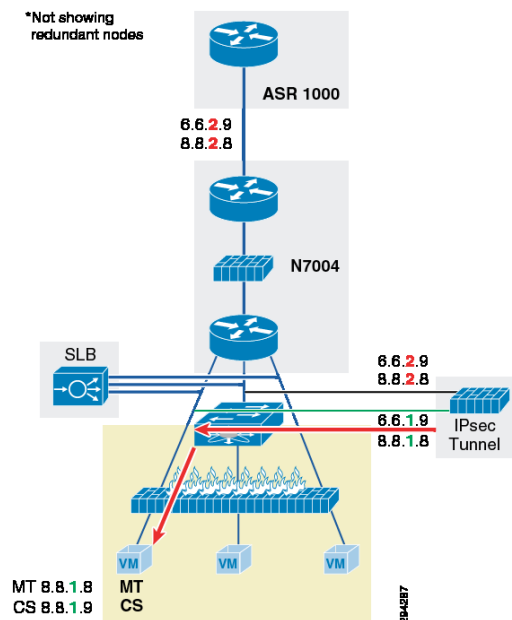
**Figure 3-31 V2V Traffic Flow (Gold Container)**

Figure 3-31 is an example of V2V Traffic flow between the Enterprise InMage PS server and the SP MT:

- From 6.6.1.9 (Enterprise InMage processing server)
- Destination 8.8.1.8 (SP InMage MT)
- Tunnel Source—Enterprise ASA IPsec-outside interface: 6.6.2.9
- Tunnel Destination—SP ASA IPsec-outside interface: 8.8.2.8

#### Enterprise

1. PS receives changes from data tap on primary server.
2. PS sends traffic (6.6.1.9, 8.8.1.8) to IPsec router.
3. IPsec router encrypts the traffic and sends re-encapsulated packet with (6.6.2.9, 8.8.2.9) to SP.

#### Service Provider

1. IPsec-encrypted traffic with (6.6.2.9, 8.8.2.9) is received by ASR 1000.
2. ASR 1000 forwards (6.6.2.9, 8.8.2.9) to the aggregation Nexus 7000 public VRF.
3. (6.6.2.9, 8.8.2.9) is routed from Nexus 7000 public VRF to Nexus 7000 private VRF via the vFW.
4. (6.6.2.9, 8.8.2.9) is send to ASA IPsec-outside interface for decryption.
5. ASA decrypts the packet, strips header (6.6.2.9, 8.8.2.9) and sends re-encapsulated packet with (6.6.1.9, 8.8.1.8) to InMage MT.

#### VMDC Silver

No modifications or changes were made to the baseline VMDC Silver container. As with the Gold container, server VLANs were neither introduced nor removed. Non-InMage/DRaaS-related traffic streams follow identical network paths as documented in the VMDC 2.3 Design and Implementation Guide (<http://www.in-wats.cisco.com/publications/viewdoc.php?docid=6637>). Traffic sourced from the VPC will first be routed to public VRF default gateway. Once traffic arrives in the public VRF, it will be L3 routed to the ASR 1000 (PE) towards the L3 VPN.

- Traffic to and from InMage server from the primary enterprise site needs to be forwarded to the IPsec tunnel for encryption or decryption.
- This was accomplished by setting up a static site-to-site tunnel between the SP and enterprise, placing the IPsec-inside interface on the Silver server VLAN 1 and the outside interface on the Silver server VLAN 2.
- InMage traffic from the enterprise LAN will be sent to the ASA-outside interface residing on server VLAN 2.
- Once received traffic is decrypted by the ASA, it will be forwarded to InMage servers residing on server VLAN 1 via the IPsec-inside interface.
- InMage traffic from SP to Enterprise is accomplished by adding static routes on the CX and MT server pointing to the inside interface of the ASA.
- Once received traffic is encrypted by the ASA, it will be forwarded out of the ASA-outside interface residing on server VLAN 2 towards the HSRP address of the private VRF interface.
- Once encrypted traffic is received in server VLAN 2, normal VMDC traffic flow occurs. Refer to [Figure 3-32](#) for details.

**Figure 3-32 V2V Traffic Flow (Silver Container)**




Figure 3-32 is an example of V2V traffic flow between the Enterprise InMage PS server and the SP MT:

- From 6.6.1.9 (Enterprise InMage processing server).
- Destination 8.8.1.8 (SP InMage MT).
- Tunnel Source—Enterprise ASA IPsec-outside interface: 6.6.2.9.
- Tunnel Destination—SPO ASA IPsec-outside interface: 8.8.2.9.

### Enterprise

1. PS receives changes from data tap on the primary server.
2. PS sends traffic (6.6.1.9, 8.8.1.8) to IPsec router.
3. IPsec router encrypts the traffic and sends re-encapsulated packet with (6.6.2.9, 8.8.2.9) to SP.

**Service Provider**

1. IPsec-encrypted traffic with (6.6.2.9, 8.8.2.9) is received by ASR 1000.
2. ASR 1000 forwards (6.6.2.9, 8.8.2.9) to the aggregation Nexus 7000 customer VRF.
3. (6.6.2.9, 8.8.2.9) is sent to ASA IPsec-outside interface for decryption.
4. ASA decrypts the packet, strips header (6.6.2.9, 8.8.2.9) and sends re-encapsulated packet with (6.6.1.9, 8.8.1.8) to InMage MT.

**VMDC Bronze**

Unlike Gold and Silver, we had to make some modifications to the baseline VMDC Bronze container to support encryption with IPsec. This is because unlike Gold and Silver container, Bronze container has only a single server VLAN. It wasn't possible to set up a site-to-site tunnel with only a single server VLAN; an additional IPsec-outside interface is required. This interface could connect directly to the ASR or to the aggregation 7004. Connecting the IPsec interface to the ASR introduced a number of fundamental design changes; it was much simpler to introduce a dedicated SVI interface on the 7004 and extend it to the ASA. The benefits of this approach are the following:

- VMDC Alignment: Minimal changes to VMDC baseline container models.
- Operation Consistency: IPsec configuration is nearly identical among all VMDC container types.
- Single Point of Resource Management: Only VLAN resources on the N7k are required. Does not introduce any changes to ASR or ASA.

Non-InMage/DRaaS-related traffic streams follow identical network paths as documented in the VMDC 2.3 design and implementation guide (<http://www.in-wats.cisco.com/publications/viewdoc.php?docid=6637>). Traffic sourced from the VPC will first be routed to customer VRF default gateway. Once traffic arrives in the customer VRF, it will be L3 routed to the ASR 1000 (PE) toward the L3 VPN.

Traffic to and from InMage server from the primary enterprise site needs to be forwarded to the IPsec tunnel for encryption or decryption.

- This was accomplished by setting up a static site-to-site tunnel between the SP and enterprise, placing the IPsec-inside interface on the Bronze server VLAN 1 and the outside interface on a newly created SVI interface between the ASA and aggregation Nexus 7004.
- InMage traffic from the enterprise LAN will be sent to an ASA-outside interface.
- Once received traffic is decrypted by the ASA, it will be forwarded to InMage servers residing on the server VLAN 1 via the IPsec-inside interface.
- InMage traffic from SP to enterprise is accomplished by adding static routes on the CX and MT server pointing to the inside interface of the ASA.
- Once received traffic is encrypted by the ASA, it will be forwarded out of the ASA-outside interface. Encrypted traffic follows normal VMDC traffic path once received by Nexus 7004.

See [Figure 3-33](#) for details.

Figure 3-33 V2V Traffic Flow (Bronze Container)




Figure 3-33 is an example of V2V traffic flow between the Enterprise InMage PS server and the SP MT:

- From 6.6.1.9 (Enterprise InMage processing server).
- Destination 8.8.1.8 (Service provider InMage MT).
- Tunnel Source—Enterprise ASA IPsec-outside interface: 6.6.2.9.
- Tunnel Destination—SP ASA IPsec-outside interface: 8.8.2.8.

#### Enterprise

1. PS receives changes from data tap on the primary server.
2. PS sends traffic (6.6.1.9, 8.8.1.8) to IPsec router.
3. IPsec router encrypts the traffic and sends re-encapsulated packet with (6.6.2.9, 8.8.2.9) to SP.

#### Service Provider

1. IPsec encrypted traffic with (6.6.2.9, 8.8.2.9) is received by ASR 1000.
2. ASR1K forwards (6.6.2.9, 8.8.2.9) to the aggregation Nexus 7000 customer VRF.
3. (6.6.2.9, 8.8.2.9) is sent to ASA IPsec-outside interface for decryption.
4. ASA decrypts the packet, strips header (6.6.2.9, 8.8.2.9) and sends re-encapsulated packet with (6.6.1.9, 8.8.1.8) to InMage MT.

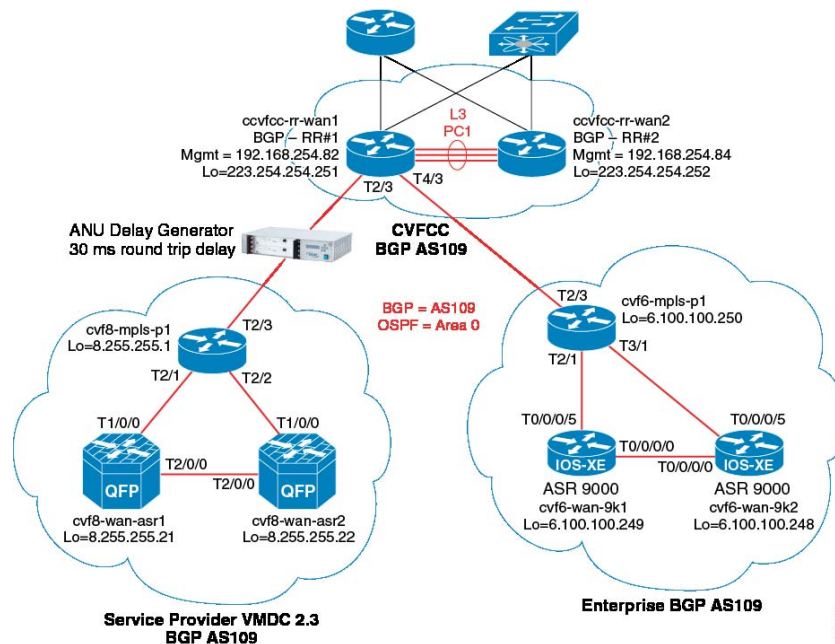
## Connectivity across the WAN

The enterprise is connected to the SP through a MPLS-VPN (L3VPN providing connectivity through an MPLS core) network for data plane connectivity. This VPN network provides connectivity in the data path between the private and public clouds for users and applications in the Enterprise to access applications in the public cloud. The same VPN network is utilized to provide in-band management connectivity between the Enterprise and the SP. Thus, the control plane (management) connectivity

between the InMage processing server on the Enterprise and the CX and MT on the SP side is carried in the same path (in-band) as the data plane connectivity. This model allows the Enterprise and SP management applications (InMage) to have either public or private IP addressing.

Figure 3-34 shows the WAN topology deployed; resources below the PE router are not drawn. Refer to earlier sections for details. Based on VMDC recommendation, a pair of ASR 1006s was deployed as PE routers in the VMDC 2.3 topology, PE routers connected to P routers running MPLS using 10 GE. At the Enterprise site, VMDC 2.2 container was utilized to simulate enterprise tenants. A pair of ASR 9000s was deployed as the PE router.

Figure 3-34 Connectivity across the WAN



MPLS VPN route is a combination of route distinguisher (RD) and actual prefix. RD is a unique identifier used to distinguish the same prefix from different customer. We define it at PE router for particular VRF. Prefix combined with RD and actual IPv4 prefix is called vpnv4 prefix and is carried by MP-BGP. BGP-extended community support is required to carry vpnv4 prefixes and labels. To carry vpnv4 prefixes, we configured iBGP in the core network. It either requires full mesh connections between routers or route reflectors (RR) deployment. We went with the second option using RR. The following is the configuration on the RR:

```
router bgp 109
bgp router-id 223.254.254.251
bgp log-neighbor-changes
neighbor RRCC peer-group
neighbor RRCC remote-as 109
neighbor RRCC description ibgp-to-RR
neighbor RR peer-group
neighbor RR remote-as 109
neighbor RR description CVFCC-WAN-6k1-to-CVFCC-WAN-6k2
neighbor RR update-source Loopback0
neighbor 6.100.100.248 remote-as 109
neighbor 6.100.100.248 peer-group RRCC
neighbor 6.100.100.249 remote-as 109
neighbor 6.100.100.249 peer-group RRCC
neighbor 6.100.100.250 remote-as 109
neighbor 6.100.100.250 peer-group RRCC
```



```

neighbor 8.255.255.1 remote-as 109
neighbor 8.255.255.1 peer-group RRCC
neighbor 8.255.255.21 remote-as 109
neighbor 8.255.255.21 peer-group RRCC
neighbor 8.255.255.22 remote-as 109
neighbor 8.255.255.22 peer-group RRCC
neighbor 223.254.254.252 remote-as 109
neighbor 223.254.254.252 peer-group RR
!
address-family ipv4
neighbor RRCC send-community
neighbor RRCC route-reflector-client
neighbor RR send-community both
neighbor 6.100.100.248 activate
neighbor 6.100.100.249 activate
neighbor 6.100.100.250 activate
neighbor 8.255.255.1 activate
neighbor 8.255.255.21 activate
neighbor 8.255.255.22 activate
neighbor 223.254.254.252 activate no auto-summary
no synchronization
network 223.254.0.0 mask 255.255.0.0 exit-address-family
!
address-family vpnv4
neighbor RRCC send-community both
neighbor RRCC route-reflector-client
neighbor RRCC next-hop-self
neighbor RR send-community both
neighbor 6.100.100.248 activate
neighbor 6.100.100.249 activate
neighbor 6.100.100.250 activate
neighbor 8.255.255.1 activate
neighbor 8.255.255.21 activate
neighbor 8.255.255.22 activate
neighbor 223.254.254.252 activate exit-address-family

```

PE routers are configured to import prefix from remote PE as well as export local prefix. The following is an example of a Gold tenant:

```

vrf definition tenant11-gold-pub rd 3486:3486
route-target export 3486:3486
route-target import 3486:3486
route-target import 3040:3040
!
address-family ipv4
exit-address-family
!

```

A WAN can be distributed over a large geographical region. The data packets need more time to travel through the network. WAN latency, in general, is orders of magnitudes higher than latency in local area networks. TCP throughput and application performance are directly impacted by latency. Many applications are slow or do not work at all for users far away from the data centers. In our implementation, an Anue delay generator was inserted between the P router and the BGP RR to inject round trip delay of 30ms between the primary Enterprise data center and the SP secondary site.

[Figure 3-35](#) is a screen capture of Anue configuration.

Figure 3-35 Anue Configuration

#	Name	Enabled	Delay	Policing	Shaping	Drop	Modify	Corrupt	CRC Corrupt	Reorder	Duplicate	Jitter	Other	Bandwidth (Mbps)	
														Tx	Rx
0	Default	✓	15.000ms	—	—	—	—	—	—	—	—	—	✓	462.256 79170 Pkts	461.183 78992 Pkts
1	<a href="#">Profile #1</a>														

The following formula can be used to calculate effective TCP throughput based on round trip delay:

$$\text{Throughput} = \text{Window Size} / \text{RTT}$$

All InMAGE components are running with default window size based on host OS.

Using the standard 64KB TCP window size of a Windows machine:

$$65536 * 8 \text{ bits} / 0.030 \text{ seconds} = 17.4 \text{ Mbps maximum possible throughput}$$

Linux window size is based on:

```
net.core.rmem_max = 131071
net.core.rmem_default = 124928
```

$$131071 * 8 \text{ bits} / 0.030 \text{ seconds} = 34.9 \text{ Mbps maximum possible throughput}$$

Based on the above, the TCP stream between the Enterprise process server and SP MT is limited to 17.4 Mbps.

## Storage Configuration

In the current implementation, the VNX 5500 is used to provide the storage needs of the solution. The VNX 5500 is based on a unified storage architecture and provides Storage Area Network (SAN) and Network-Attached Storage (NAS) capabilities on a single platform. In this solution, only SAN is utilized. The Nexus 5000 is the FC switch that connects server blades and storage to provide SAN capability.


This section presents the following topics:

- [SAN Implementation Overview, page 3-42](#)
- [VNX5500 Configuration Overview, page 3-45](#)

## SAN Implementation Overview

This section explains the Fibre Channel over Ethernet (FCoE) connection from servers to the FI and Fibre Channel (FC) connectivity to carry SAN traffic from the FI to the MDS (storage switch) to VNX 5500 Filers. [Figure 3-36](#) shows an overview of the SAN infrastructure.

**Figure 3-36 Storage Infrastructure Overview**



Features of FC configuration in the data center follow:

- Each blade server has two vHBAs that provide server to storage SAN connectivity. This is to provide server level host bus adapter (HBA) fabric redundancy.
- Storage traffic from server blades to FIs is FCoE. Each virtual SAN (VSAN) is mapped to a unique VLAN that carries storage traffic from server to FI.
- Each FI is mapped to one VSAN. In this case, FI-A (CVF8-FAB-1-A) carries all VSAN88 traffic and FI-B (CVF8-FAB-1-B) carries all VSAN89 traffic.
- FCoE, by default, maps FC traffic to a no-packet drop class using the system QoS policy. This assures that during congestion storage traffic will not be dropped.

Figure 3-37 shows the list of VSANs in the SAN infrastructure: VSAN88, VSAN89. This is to allow multiple SANs and LANs to share a common infrastructure when carried using the same FCoE links between the server and FIs.

**Figure 3-37 Infrastructure VSANs**

Name	ID	Fabric ID	If Type	If Role	Transport	FCoE VLAN ID	Operational State
VSANs							
VSAN default (1)	1	Dual	Virtual	Network	Fc	4048	Ok
Fabric A							
VSANs							
VSAN sys8_FCoE_Fab_A (88)	88	A	Virtual	Network	Fc	88	Ok
Fabric B							
VSANs							
VSAN sys8_FCoE_Fab_B (89)	89	B	Virtual	Network	Fc	89	Ok

Figure 3-38 shows the vHBA configuration on each server blade. vHBAs are part of a server service profile derived from a server template, consists of two vHBA adapters per server blade. Each vHBA is placed on a unique, isolated SAN network. vHBA0 of all server blades are placed in SAN-A and vHBA1 is placed in SAN-B

**Figure 3-38 Infrastructure vHBAs**

Name	WWPN	Desired Order	Actual Order	Fabric ID	Desired Placement	Actual Placement
vHBA vHBA-A	Derived	1	Unspecified	A	Any	Any
vHBA If sys8_FCoE_Fab_A						
vHBA vHBA-B	Derived	2	Unspecified	B	Any	Any
vHBA If sys8_FCoE_Fab_B						

Figure 3-39 shows the ports used between FI and the MDS switch for SAN traffic. Although 4 ports are configured, only two of the ports are physically cabled. FI-A (fc 2/1, 2/2) connects to MDS-1 (fc 1/19, 1/20) and FI-B (fc 2/1, 2/2) connects to MDS-2 (fc 1/19, 1/20).

**Figure 3-39 Ports Used between FI and MDS Switch**

Name	Fabric ID	If Type	If Role	Transport	Administrative State
FC Interface 2/1	A	Physical	Network	Fc	Enabled
FC Interface 2/2	A	Physical	Network	Fc	Enabled
FC Interface 2/3	A	Physical	Network	Fc	Enabled
FC Interface 2/4	A	Physical	Network	Fc	Enabled

Soft zoning (using World Wide Port Name (WWPN) names) is configured on the MDS to allow servers with specific identity (WWPN) to communicate with VNX filers. Each filer connection has its own WWPN name. The following configuration shows the zoning configuration SAN-A. As mentioned before, vHBA0 of all server blades are placed in the SAN-A and vHBA1 is placed in SAN-B. The WWPN of vHBAs is obtained from the UCSM. The WWPN of VNX filers is fetched using VNX FC port properties.

```

zoneset name cvf8-Fab-a vsan 88
  zone name cvf8-temp-all-vnx5500 vsan 88
    pwwn 50:00:00:25:b5:e1:81:1f
    pwwn 50:00:00:25:b5:e1:81:3e
    pwwn 50:00:00:25:b5:e1:81:3f
    pwwn 50:00:00:25:b5:e1:81:5e
    pwwn 50:00:00:25:b5:e1:81:5f
    pwwn 50:00:00:25:b5:e1:81:7e
    pwwn 50:00:00:25:b5:e1:81:7f
    pwwn 50:00:00:25:b5:e1:81:9e
    pwwn 50:00:00:25:b5:e1:81:9f
    pwwn 50:00:00:25:b5:e1:81:ae
    pwwn 50:00:00:25:b5:e1:81:af
    pwwn 50:00:00:25:b5:e1:81:be
    pwwn 50:00:00:25:b5:e1:81:bf
    pwwn 50:00:00:25:b5:e1:81:de
    pwwn 50:00:00:25:b5:e1:81:df
    pwwn 50:06:01:64:3e:a0:36:1a < VNX

```

```
cvf6-san-mds1# show flogidatabase vsan 88
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/19	88	0xb00100	20:41:54:7f:ee:12:5d:40	20:58:54:7f:ee:12:5d:41
fc1/19	88	0xb00101	50:00:00:25:b5:e1:81:7e	20:00:00:25:b5:08:01:2f
fc1/19	88	0xb00102	50:00:00:25:b5:e1:81:ae	20:00:00:25:b5:08:01:4f
fc1/19	88	0xb00104	50:00:00:25:b5:e1:81:1f	20:00:00:25:b5:08:01:9f
fc1/19	88	0xb00108	50:00:00:25:b5:e1:81:5f	20:00:00:25:b5:08:01:af
fc1/19	88	0xb0011e	50:00:00:25:b5:e1:81:9f	20:00:00:25:b5:08:01:bf
fc1/19	88	0xb0011f	50:00:00:25:b5:e1:81:af	20:00:00:25:b5:08:01:ef
fc1/19	88	0xb00123	50:00:00:25:b5:e1:81:5e	20:00:00:25:b5:08:01:3f
fc1/19	88	0xb00126	50:00:00:25:b5:e1:81:df	20:00:00:25:b5:08:01:ff
fc1/20	88	0xb00000	20:42:54:7f:ee:12:5d:40	20:58:54:7f:ee:12:5d:41
fc1/20	88	0xb00001	50:00:00:25:b5:e1:81:7f	20:00:00:25:b5:08:01:cf
fc1/20	88	0xb00002	50:00:00:25:b5:e1:81:3e	20:00:00:25:b5:08:01:0f
fc1/20	88	0xb00004	50:00:00:25:b5:e1:81:de	20:00:00:25:b5:08:01:6f
fc1/20	88	0xb00008	50:00:00:25:b5:e1:81:3f	20:00:00:25:b5:08:01:8f
fc1/20	88	0xb0000f	50:00:00:25:b5:e1:81:bf	20:00:00:25:b5:08:01:df
fc1/20	88	0xb00010	50:00:00:25:b5:e1:81:9e	20:00:00:25:b5:08:01:5f
fc1/20	88	0xb00017	50:00:00:25:b5:e1:81:be	20:00:00:25:b5:08:01:7f
fc1/34	88	0xb00300	50:06:01:64:3e:a0:36:1a	50:06:01:60:be:a0:36:1a

We had the choice of implementing a "single initiator zoning" where each zone contains only one host server vHBA and can contain multiple storage array targets in the same zone. Instead, we implemented "multi-initiator zoning" to allow the flexibility of moving hosts between ESXi clusters. Instead of masking at the MDS level, we used VNX storage group to mask specific LUNs to the ESXi Cluster. Refer to [VNX5500 Configuration Overview, page 3-45](#) for details. The FC interface on the MDS switch is used to connect to the VNX 5500 for FC connectivity and is configured below.

```
interface fc1/34
switchport description Connection to VNX5500 port-license acquire
no shutdown
```

## VNX5500 Configuration Overview

VNX has four main configuration elements:

- The physical drives
- The storage pool
- The LUN
- The tiering policy of the LUN

[Figure 3-40](#) shows a high level overview of VNX.

**Figure 3-40 High Level Overview of VNX**

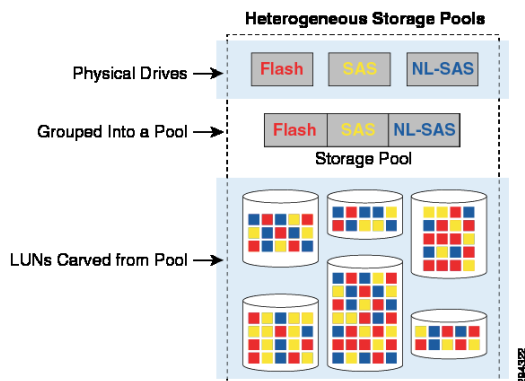


Figure 3-40 was taken from from EMC VNX Virtual Provisioning at <http://www.emc.com/collateral/hardware/white-papers/h8222-vnx-virtual-provisioning-wp.pdf>.

As discussed earlier, FAST VP is a tiering solution that can be utilized with the VNX to reduce total cost of storage ownership. FAST VP operates by continuously collecting performance statistics. Collected data is analyzed once per hour and, based on schedule, data is moved between tiers once every 24 hours during a specified relocation window. The granularity of data is 1GB. Each 1 GB block of data is referred to as a "slice." When FAST VP relocates data, it will move the entire slice to a different storage tier.

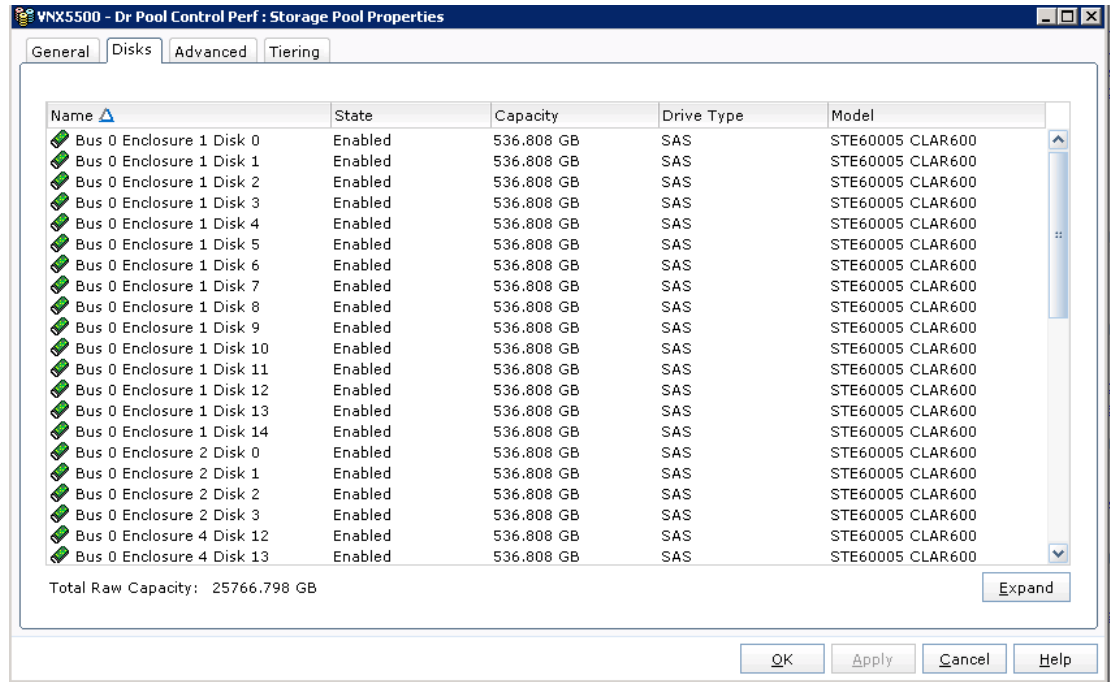
To start off, our implementation of FAST VP consists of only SAS disks: this is to provide sufficient performance to on board newly protected VMs and also to allow VNX sufficient time to identify hotter/colder slices of data. Once the desired number of tenants per storage pool is reached, NL-SAS drives can be introduced to move cold data from performance tier to capacity tier. To balance the data split between performance and capacity tier, a manual relocation at the storage pool level can

be initiated by a cloud admin through the Unisphere GUI. Both relocation rate and duration can be specified at the time of manual relocation.

In our implementation, storage pool was sized based on total IOPS to support peak VM transfer from the primary site and change rate of existing VMs under protection. 4800 IOPS was provisioned to support Journal retention and an additional 4800 IOPS to support aggregate workload change rate. In an actual deployment, IOPS requirement will vary considerably. Depending on WAN bandwidth, number of customers on boarding new VMs and change rate of existing VMs under protection, IOPS needs to be sized according to the deployment scenario.

Based on 200 IOPS per SAS disk and RAID 10 configuration, 48 SAS drives were needed to support 4800 IOPS. Figure 3-41 shows a screen capture of disk configuration.

Figure 3-41 Disk Configuration



Tenant-specific LUN is created on top of the storage pool. Each tenant is assigned a dedicated Journal LUN and workload LUN is shared between tenants. EMC FAST Cache was also utilized to provide read acceleration during the time of recovery. Total of 274GB of usable flash cache was deployed. In a real deployment scenario, the amount of flash cache should be sized based on the overall capacity of recovery workloads.

Table 3-12 Journal LUN

Journal	LUN Size (GB)
Tenant Control_1	1480
Tenant Control_2	590
Tenant Control_3	1140
Tenant Control_4	1140
Tenant Control_5	1140
Tenant Control_6	1140
Tenant Control_7	1140
Tenant Control_8	1140
Tenant Control_9	340
Tenant Control_10	340
Tenant Control_11	340
Tenant Control_12	340

**Table 3-13 Workload LUN**

Workload	LUN Size (GB)
LUN Gold-1	500
LUN Silver-1	750
LUN Silver-2	750
LUN Bronze-1	900
LUN Bronze-2	900
LUN Bronze-3	900

All of the LUNs are mapped to the corresponding storage group, based on ESXi cluster, as shown in [Table 3-14](#). As discussed in earlier sections, LUN masking is implemented at the storage array level to simplify the ability to move hosts between clusters for various test scenarios.

**Table 3-14 Storage Group DR-Bronze**

Storage Group	ESX Cluster Name	Hosts	LUNs
DR-Bronze	Bronze	cvf8-draassp-esx-4.cvfdmz.sdu	Tenant Control_1
		cvf8-draassp-esx-5.cvfdmz.sdu	Tenant Control_4
		cvf8-draassp-esx-6.cvfdmz.sdu	Tenant Control_5
		cvf8-draassp-esx-7.cvfdmz.sdu	Tenant Control_6
		cvf8-draassp-esx-8.cvfdmz.sdu	Tenant Control_7
			Tenant Control_8
			LUN Bronze-1
			LUN Bronze-2
		LUN Bronze-3	

**Table 3-15 Storage Group DR-Silver**

Storage Group	ESX Cluster Name	Hosts	LUNs
DR-Silver	Silver	cvf8-draassp-esx-2.cvfdmz.sdu	Tenant Control_2
		cvf8-draassp-esx-3.cvfdmz.sdu	Tenant Control_3
		cvf8-draassp-esx-3-1.cvfdmz.sdu	LUN Silver-1
			LUN Silver-2

**Table 3-16 Storage Group DR-Gold**

Storage Group	ESX Cluster Name	Hosts	LUNs
DR-Gold	Gold	cvf8-draassp-esx-9.cvfdmz.sdu	Tenant Control_9
		cvf8-draassp-esx-10.cvfdmz.sdu	Tenant Control_10
		cvf8-draassp-esx-11.cvfdmz.sdu	Tenant Control_11



**Table 3-16 Storage Group DR-Gold (continued)**

Storage Group	ESX Cluster Name	Hosts	LUNs
		cvf8-draassp-esx-12.cvfdmz.sdu	Tenant Control_12
			LUN Gold-1

## BMC Cloud Lifecycle Management

DRaaS 1.0 leverages existing capabilities of Cloud Orchestration for VMDC with BMC Cloud Lifecycle Management 3.1 SP1. Initial workflows of onboarding tenants, network container creation, firewall policy changes, and server load balancer updates align with VMDC 2.3 operational method and procedures. This is well documented by the SDU BMC-CLM team. Refer to [BMC Design and Implementation Guide](#) for additional details.





## CHAPTER 4

# Disaster Recovery Workflow

---

InMage CDP starts with the FX/VX agent, also known as "DataTap," which is used to monitor all writes to disk. A small amount of memory on the source machine is reserved by the DataTap (250MB). When changes are being written to primary storage, a copy of the change is stored in the reserved memory. When the memory reaches a certain size, the contents of the reserved memory are sent as a chunk to the processing server. This "write coalescing" is mainly for WAN optimization. The processing server then compresses it for WAN transmission and sends to the MT at the secondary site. Once the MT receives the compress data, it will:

1. Uncompress and store incoming data in the cache volume. One folder exists for each VMDK under protection.
2. Reverse coalescing received data into individual blocks.
3. Read/retrieve old data blocks from VMDK.
4. Write old data blocks to MT retention volume for Journal.
5. Writes individual new blocks to target VMDKs.

For every write at the source, 2 writes and 1 read exist at the destination performed by the MT.

Data protection will continue until disaster declaration. During a DR event, a predefined VM recovery plan to secondary site can be started. The recovery plan will:

1. Unmount protected disks (VMDK) attached to the MT and release the read/write lock.
2. Readdress the secondary server with new IP address and subnet.
3. Power on the secondary server (VM).

Once the primary site comes back online, a failback plan can be created. It is called a failback plan because the objective is to restore the latest data from the secondary site to the newly restored primary site. Under the covers, a failback plan works identically as a protection plan; the only difference is the direction of data protection, secondary to primary versus primary to secondary. Once failback replication reaches "Differential Sync" status, the recovery plan can be executed to bring the servers in primary site back in service.

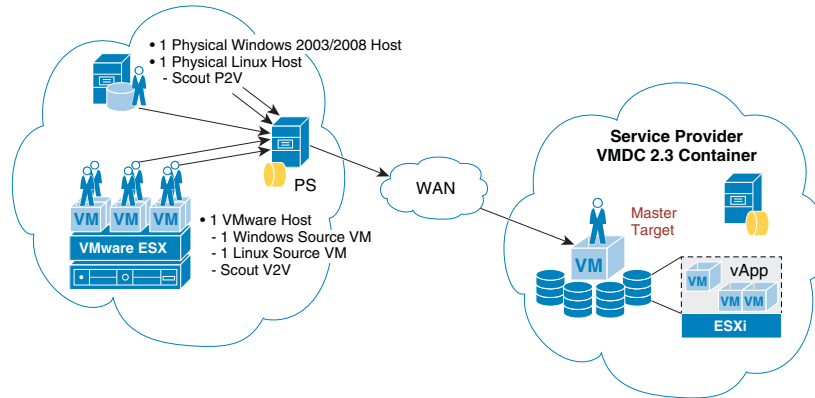
This chapter includes the following major topics:

- [Protection Workflows, page 4-2](#)
- [Recovery Workflows, page 4-37](#)
- [Failback Protection Workflows, page 4-46](#)
- [Resume Protection Workflows, page 4-78](#)
- [DR Drill Workflows, page 4-84](#)

# Protection Workflows

InMage vContinuum is a DR solution for VMware vSphere (ESX and ESXi) servers and physical servers. vContinuum protects VMs on your primary ESX/ESXi server, and your physical servers, by replicating them to a secondary ESX/ESXi server and recovering them on the secondary ESX/ESXi server when needed. vContinuum not only captures all changes, but also provides the ability to recover to any point in time during the configured retention period.

**Figure 4-1 Protection Plan Overview**



vContinuum supports two types of protection:

- **Virtual-to-Virtual (V2V):** Primary customer VMs that reside on the primary VMware ESX server are protected and recovered as secondary provider VMs on the secondary VMware ESX server.
- **Physical-to-Virtual (P2V):** Customer physical servers can be protected and recovered as a secondary provider VM on a secondary VMware ESX server.

When a V2V or P2V protection plan is first created, the initial volume replication can be bandwidth intensive. Offline Sync is a feature that allows for initial volume replication to occur offline instead of over a WAN. This technique eliminates the initial volume replication over the WAN, thereby reducing the overall WAN bandwidth required to perform the protection plan.

This section includes the following topics:

- [Setting up Virtual-to-Virtual \(V2V\) Protection Plan, page 4-2](#)
- [Setting up Physical-to-Virtual \(P2V\) Protection Plan, page 4-18](#)
- [Offline Sync, page 4-26](#)

## Setting up Virtual-to-Virtual (V2V) Protection Plan

Virtual-to-virtual protection enables the recovery of a VM, including OS partition and application data, to a prior point in time. A dedicated VM, called a "master target" or "MT," needs to be prepared on the secondary vSphere server. The MT receives all the changes of the primary server and stores them in the retention store. When a disaster event is declared, the administrator can recover a VM to a specific consistency point or point in time. InMage provides a wizard to facilitate the creation of protection plans and other actions. The vContinuum wizard will be used to create a new protection plan for five primary Windows 2008 VMs in this section.

**Note**

The following steps to configure a protection plan are based on the online Scout Help, which can be accessed from the main vContinuum page.

---

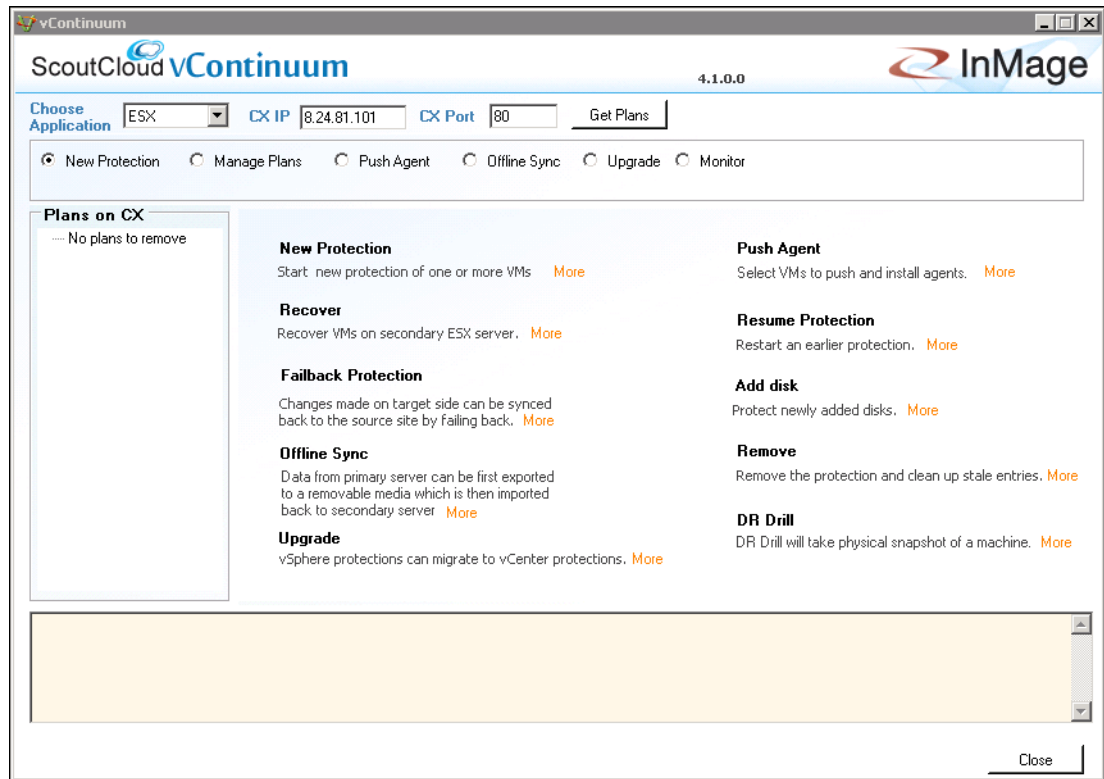
**Summary of Steps**

- 
- Step 1** Start vContinuum wizard application.
- a. Select the primary vCenter to view available source servers.
    - Select VM(s) and their volumes to protect.
  - b. Select the secondary vCenter to view available MT(s).
    - Select MT(s).
  - c. Configure replication options.
  - d. Select datastores in secondary vCenter to create recovery VM(s).
  - e. Configure recovery VM(s) configuration options (for example, network, hardware, display name, sparse retention settings).
  - f. Finalize protection plan.
- Step 2** Monitor protection plan.
- 

**Detailed Steps**

- 
- Step 1** On the Management Console, start the vContinuum wizard application using the desktop icon or Start menu shortcut **Start > Program > InMage System > VContinuum > vContinuum**:
- a. Select **ESX** in the **Choose Application** drop-down list for V2V protection plan.
  - b. Enter the CX server IP address and port number (the default is 80).
  - c. Select **New Protection** to create a new protection plan.

Figure 4-2 Creating New V2V Protection Plan in vContinuum



- d. Enter the primary vSphere/vCenter IP address, Username, Password, and Guest OS Type and then click **Get Details** to view available source servers.

Figure 4-3 Selecting Primary vSphere/vCenter

ScoutCloud vContinuum 4.1.0.0 InMage

Protect

Provide details of the vSphere/vCenter hosting your primary VM(s)

vSphere/vCenter IP	Username	Password	Guest OS Type	
6.126.104.138	root	XXXXXXXXXX	Windows	Get Details

Help

Next > Cancel

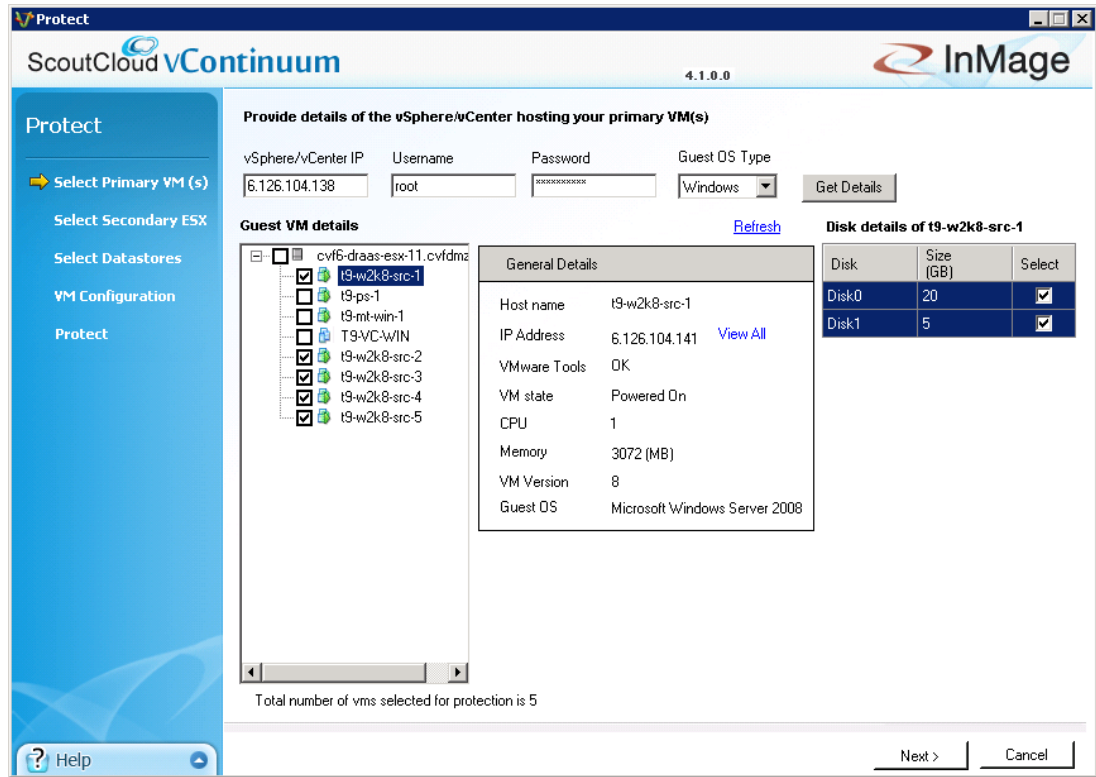
294406

- e. Select the primary VM(s) and their volumes to protect, then click **Next**.

**Note**

- When a primary VM is selected, the General Details section of the vContinuum wizard shows details about the selected VM.
- By default, all local volumes are selected for protection. Volumes can be omitted from the protection plan by deselecting the volumes, but the disk that contains the operating system must be selected or the recovered VM will not be able to start after a disaster event.

Figure 4-4 Selecting Primary VMs



- f. Enter the secondary vSphere/vCenter IP address, Username, and Password, and then click **Get Details** to view the available MT(s).

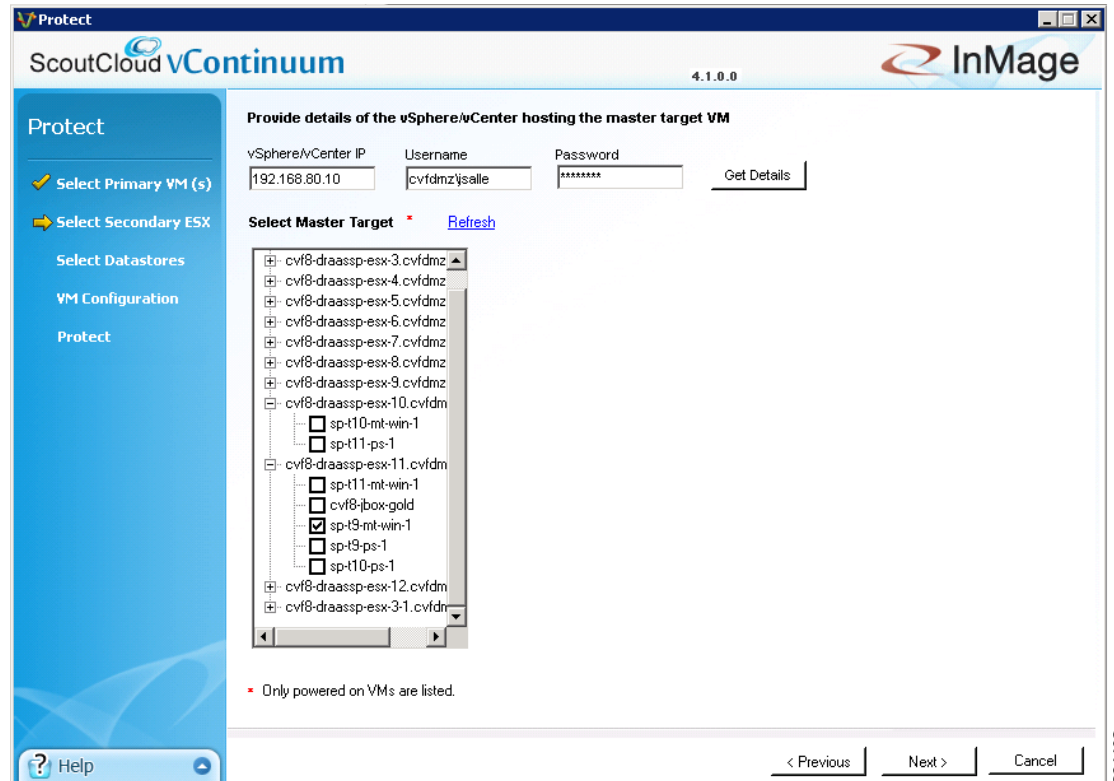
**Note**

The MT must be of the same OS family as the primary VM(s) it protects. If the primary VMs use Windows, then the MT must also be Windows. The same requirement exists for Linux servers. For more information on MT considerations, refer to [Master Target—Enterprise and Service Provider, page 3-1](#).

- g. Select the MT that will be used to protect the selected primary VM(s) and then click **Next**.



Figure 4-5 Selecting Secondary MT(s)



#### h. Configure replication options.

- In the Process server IP field, select the process server located in the primary network (e.g., Enterprise). Multiple process servers can be deployed and associated with a limited number of primary VMs for scalability.
- In the Retention size (in MB) field, enter the maximum amount of disk space to use for retention data.
- In the Retention Drive field, select the drive letter that is associated with the retention drive on the MT.
- In the Retention (in days) field, enter the maximum number of days to store retention data.



#### Note

- The amount of retention data can be limited by disk space, time, or both. For more information on retention data considerations, refer to [Retention Volume Sizing, page 3-3](#).
- The vContinuum wizard release used during testing (v.4.1.0.0) did not allow the user to configure the retention window lower than one day. If a retention window smaller than one day is desired, the retention window can be later adjusted to less than one day through the CX UI in the Protect > Volume > Settingspage.

- i. In the Consistency interval (in mins) field, enter the number of minutes between execution of the replication jobs. The replication jobs will run every x minutes generating application consistency recovery points for the primary VMs. This value determines the RPO for consistency point-based recovery.

Figure 4-6 Configuring Replication Options

ScoutCloud vContinuum 4.1.0.0 InMage

Protect

- Select Primary VM (s)
- Select Secondary ESX
- Select Datastores
- VM Configuration
- Protect

**Provide Replication options**

Primary Server Name	Total disk size(GB)	Process server IP	Retention size(in MB)	Retention drive	Retention(in days)	Consistency interval(in mins)	Select target datastore	Select
t9-w2k8-src-1	25	6.126.10...	Enter value	E	1	30		<a href="#">Not a</a>
t9-w2k8-src-2	20	6.126.10...	Enter value	E	1	30		<a href="#">Not a</a>
t9-w2k8-src-3	20	6.126.10...	Enter value	E	1	30		<a href="#">Not a</a>
t9-w2k8-src-4	15	6.126.10...	Enter value	E	1	30		<a href="#">Not a</a>
t9-w2k8-src-5	15	6.126.10...	Enter value	E	1	30		<a href="#">Not a</a>

**Datastore(s) on secondary vSphere**

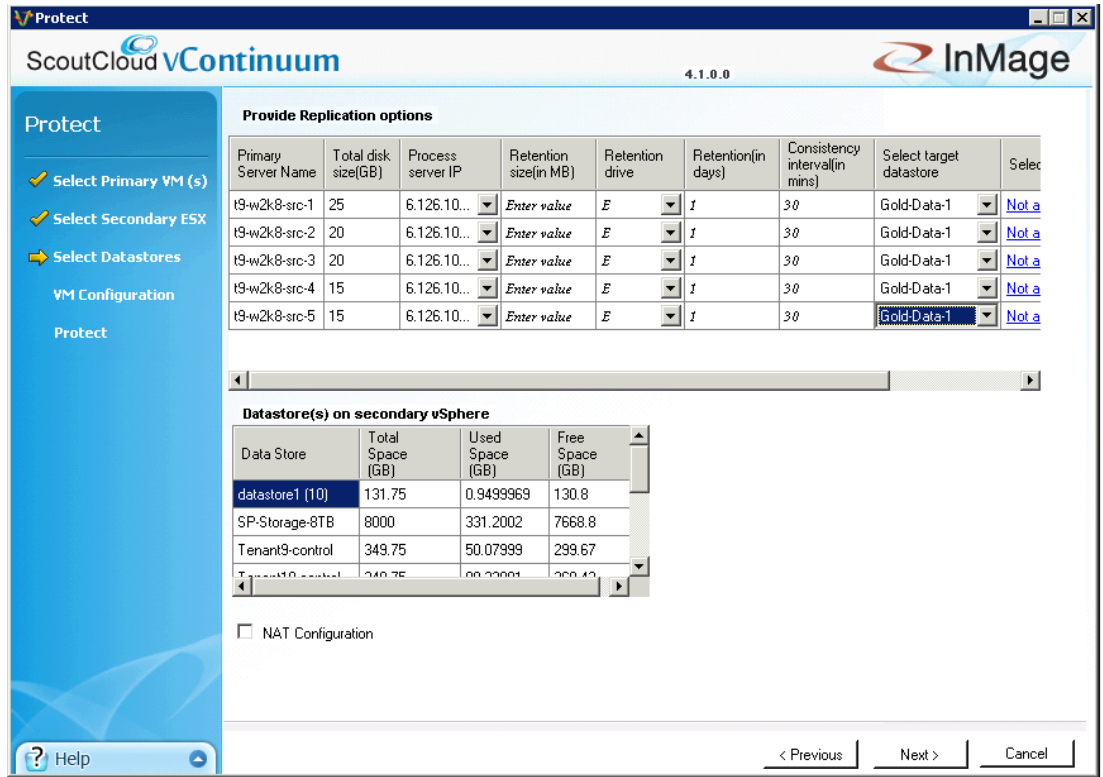
Data Store	Total Space (GB)	Used Space (GB)	Free Space (GB)
datastore1 (10)	131.75	0.9499969	130.8
SP-Storage-8TB	8000	366.5898	7633.41
Tenant9-control	349.75	50.07999	299.67
Tenant10-control	349.75	50.07999	299.67

NAT Configuration

Help < Previous Next > Cancel

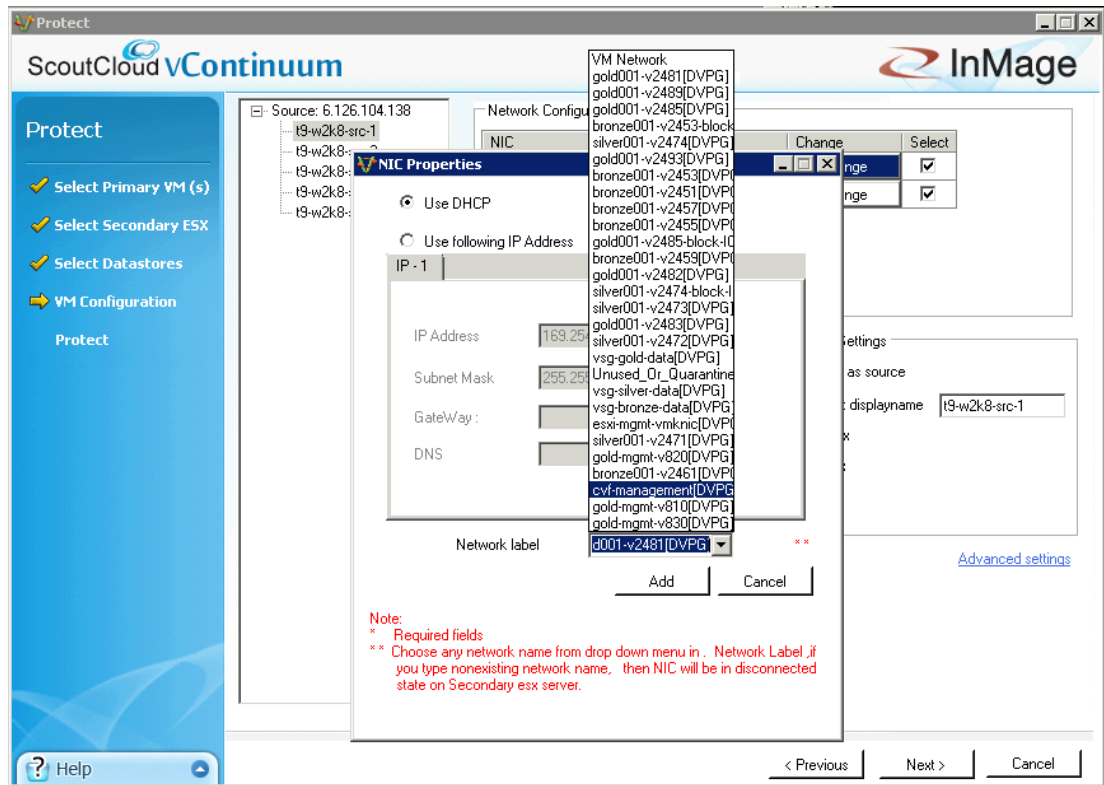
- j. Select the target datastores in secondary vCenter to create recovery VM(s) and then click **Next**.

Figure 4-7 Selecting Datastores in Secondary vCenter



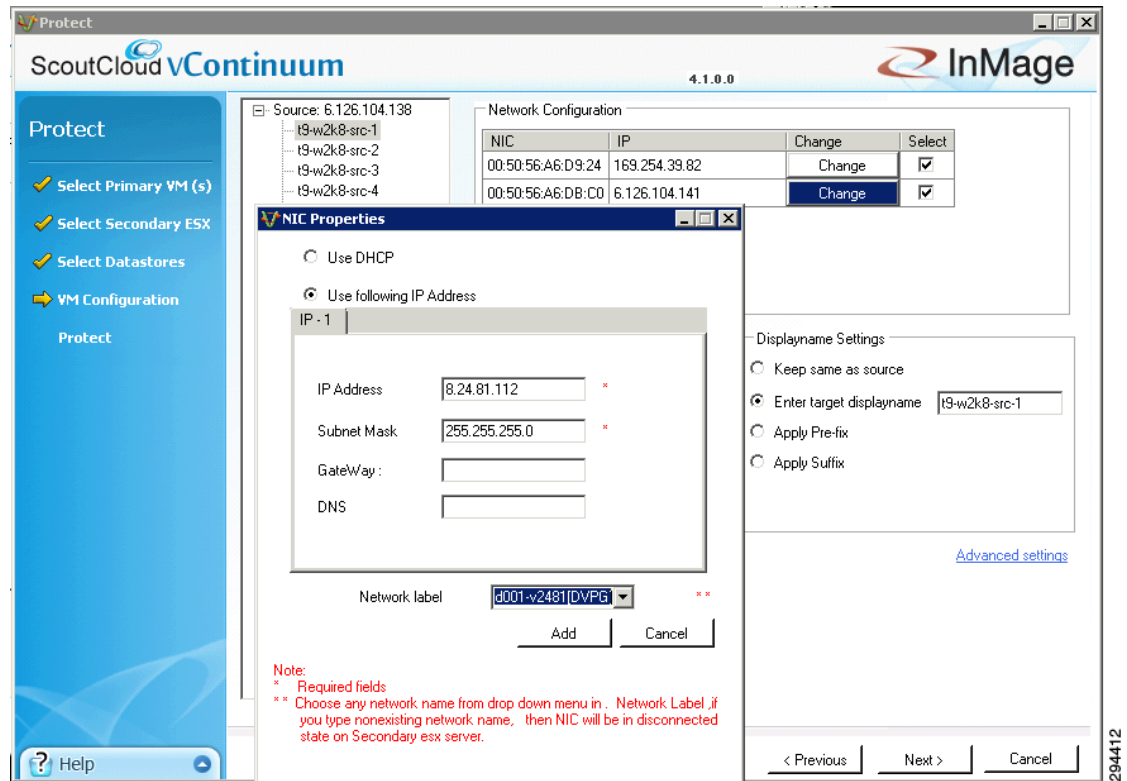
- k. Specify configuration options for recovered VM(s) in the secondary vSphere/vCenter environment (for example, network, hardware, display name, sparse retention settings).
  - Select which network interfaces to include in the recovery VM(s).
  - Configure the port group to use in the secondary vSphere/vCenter network and any new network configurations. By default, the network configuration of the primary VM(s) will be used and may need to be changed. If an interface will use DHCP in the secondary vSphere/ vCenter environment to get an IP address, then select DHCP and the appropriate port group (for example, Network Label).

Figure 4-8 Configure Network Settings for VM in Secondary vCenter (DHCP)



294411

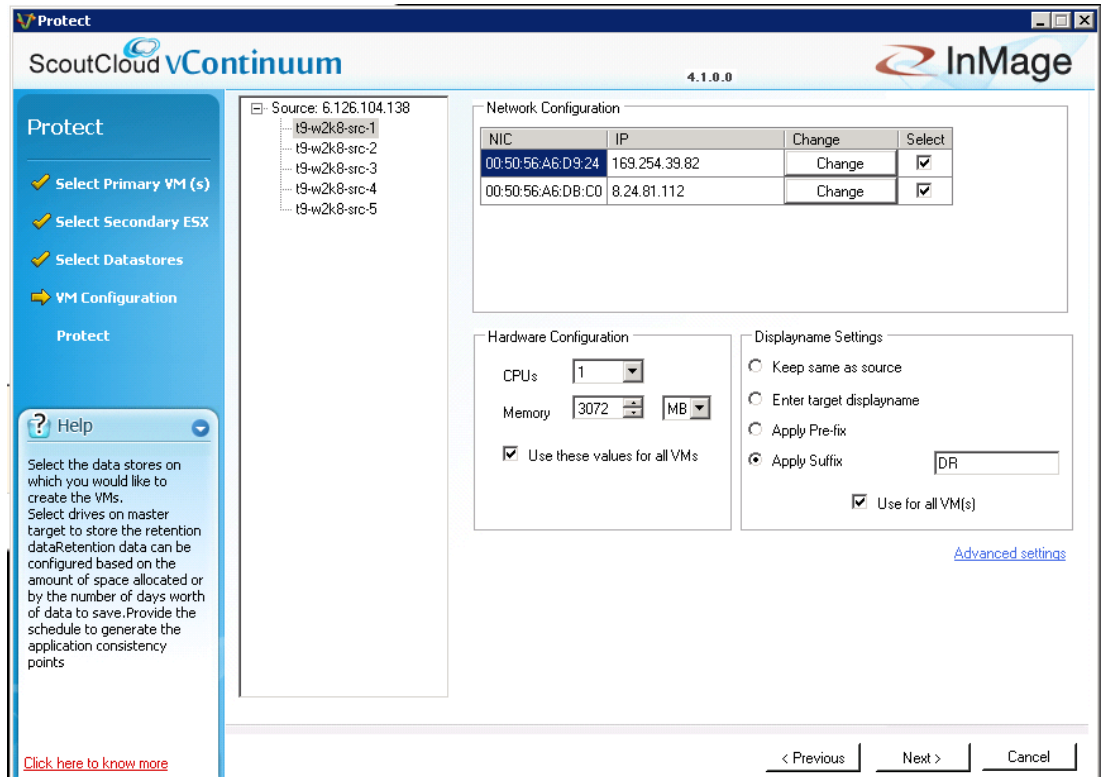
Figure 4-9 Configure Network Settings for VM in Secondary vCenter (New IP Address)



**Note** If multiple interfaces are selected, some additional configuration may be required for proper routing. Only one of the interfaces should be configured with a default route and if any static routes are required, they must be added after the recovered VM is powered up via scripts or manual configuration.

- Configure the hardware settings for the VM(s) in the secondary vSphere/vCenter. These settings can be applied to all VMs in the protection plan by selecting Use these values for all VMs.
- Configure the display name for the VM(s) in the secondary vSphere/vCenter. To apply the "Keep same as source," "Apply Pre-fix," and "Apply Suffix" options to all VMs in the protection plan, select **Use for all VM(s)**.

Figure 4-10 Configure Display Name for VM in Secondary vSphere/vCenter



- I. Click Advanced settings to access advanced settings for protection. The default configuration shown in Figure 4-11.
  - In the Sparse Retention section, advanced retention settings can be configured to have varying number of retention points based on age. A small number of retention points can be stored for weeks or months in the past, while more retention points can be stored for days in the past.
  - In the Folder Name Settings section, the directory for the VM in the datastore can be configured.
  - In the Compression section, compression type can be changed or disabled.
  - In the Encryption section, encryption can be enabled for the primary VM to the Process Server path, the Process Server to MT path, or both.
  - In the Resource pool section, a resource pool on the secondary vSphere/vCenter can be specified. Resource pools can be used to isolate tenants from each other.
  - In the Provisioning section, thin or thick provisioning can be configured.

Figure 4-11 Optional Advanced Settings for VM Protection

Advanced settings

Scout vContinuum 4.1.0.0 InMage

Advanced setting for t9-w2k8-src-1

Sparse Retention

Backups are retained for 90 days(3 months)

Provide continuous backup for latest 1 days

Advanced retention

<input type="checkbox"/>	From 1 days onwards. Provide 1 restore point per 1 hour for next days
<input type="checkbox"/>	From 1 days onwards. Provide 1 restore point per day for next weeks
<input type="checkbox"/>	From 1 days onwards. Provide 1 restore point per week for next months

Apply for all servers Total days selected 1

Folder Name Settings

Keep VM in datastore's root directory

Keep VM in datastore's sub directory

Apply for all servers

Compression

No compression

Compression

CX based  Source based

Apply for all servers

Encryption

Secure transport from Source to CX-PS

Secure transport from CX-PS to destination

Apply for all servers

Resource pool

Select resource pool on target T9-Resource-Po

Apply for all servers

Provisioning

Thin provisioning

Thick provisioning

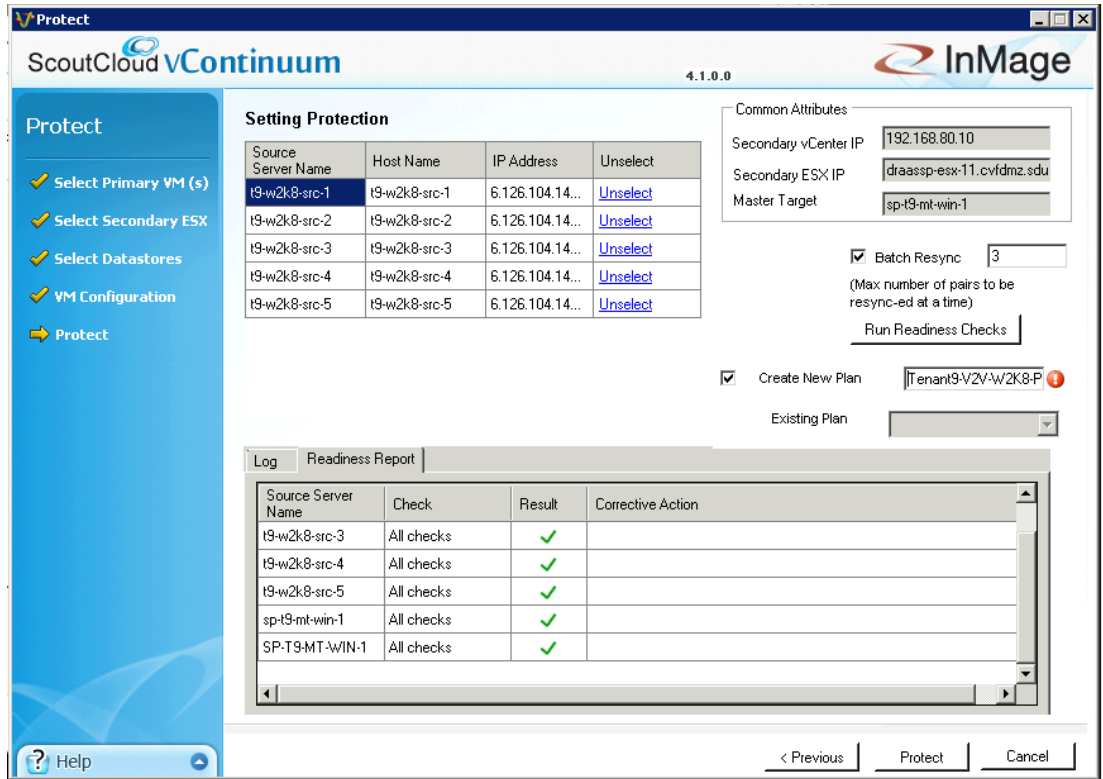
Apply for all servers

Ok Cancel

294414

- Click Next to advance to the final page.
- m. Finalize the protection plan.
  - Click **Run Readiness Checks** to perform checks.
  - Enter a name for the protection plan.
  - Click **Protect** to finalize the protection plan.

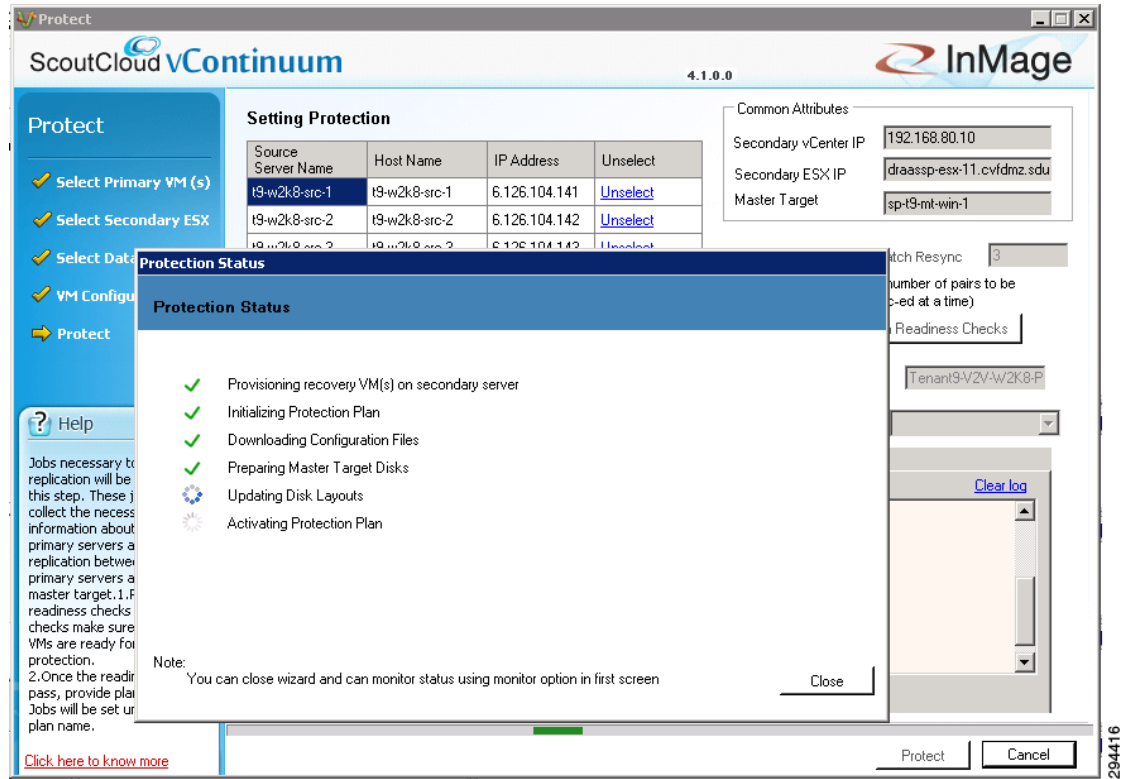
Figure 4-12 Finalize Protection Plan



**Step 2** Monitor the protection plan. After finalizing the protection plan, vContinuum goes through a number of steps to put the protection plan in place. The initialization of the protection plan can be monitored from the status window.



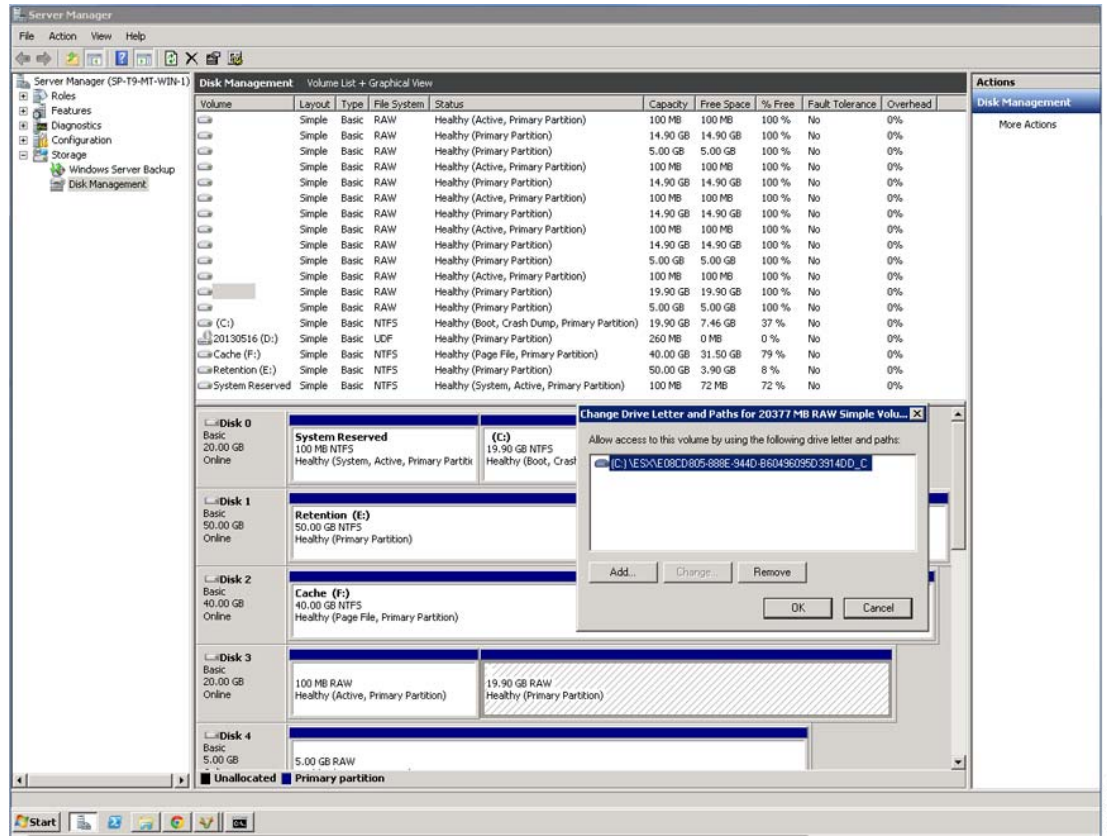
Figure 4-13 Protection Plan Initializing (vContinuum)



294416

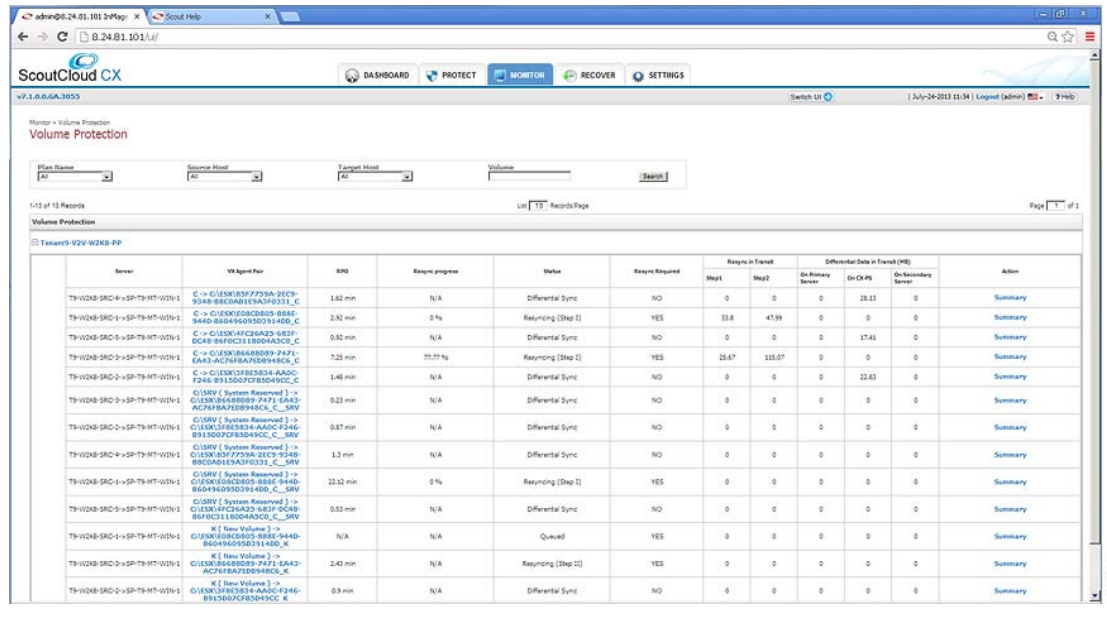


Figure 4-16 Master Target Disk Layout (Service Provider vCenter)



264419

Figure 4-17 Checking Primary VM from CX UI



264420

## Setting up Physical-to-Virtual (P2V) Protection Plan

The entire physical server, including operating system and data, can be protected using vContinuum. The physical server is replicated to VMs running on ESX servers located at the secondary site. These VMs can then be powered up at the time of disaster or whenever required. Secondary VMs can be protected back to physical servers using the steps described in the virtual-to-physical (V2P) procedure found in [“Virtual-to-Physical \(V2P\) Failback Protection” section on page 4-59](#). P2V supports both Windows and Linux operating systems.

**Note**

---

The following steps to configure a protection plan are based on the online Scout Help, which can be accessed from the main vContinuum page.

---

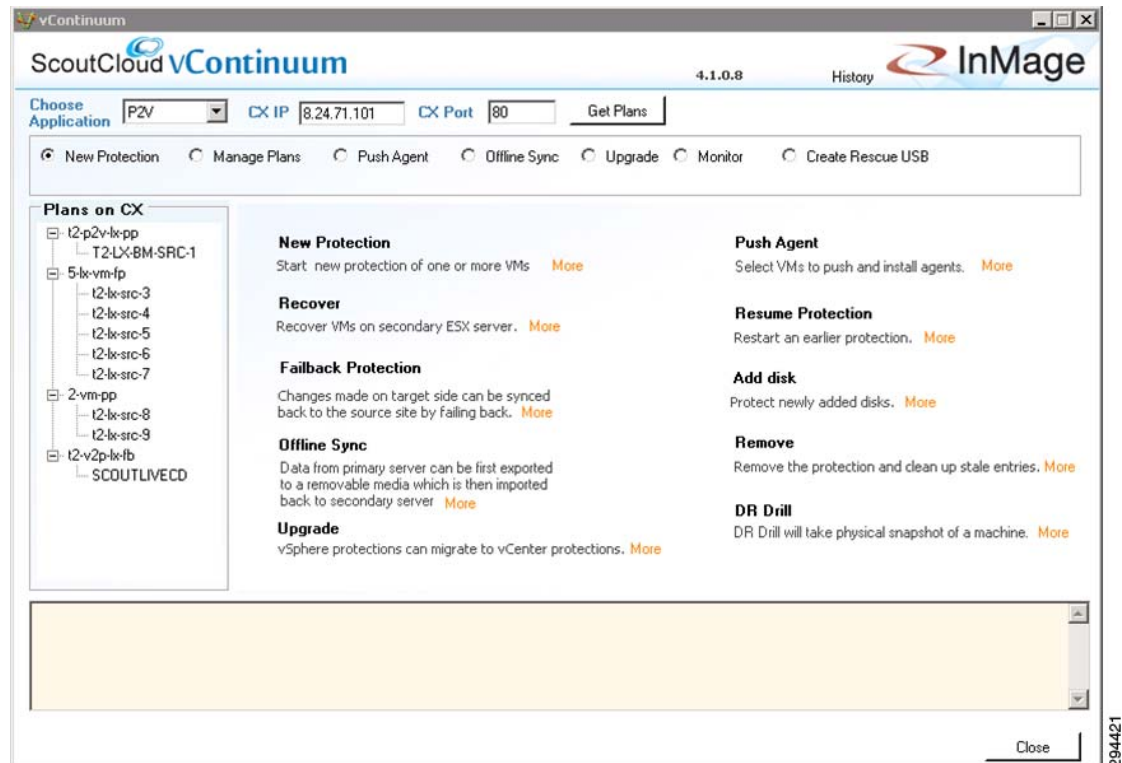
**Summary of Steps**

1. Create a new protection plan using the vContinuum wizard.
2. Select the primary physical server(s) and volume(s) to protect.
3. Select the secondary site MT.
4. Configure the replication options and select the secondary site target datastore.
5. Configure the secondary VM(s) configuration options.
6. Configure the secondary VM(s) configuration advanced settings (optional).
7. Run the readiness check, name the protection plan, and protect.
8. Monitor the protection plan status from the CX UI.

**Detailed Steps**

- 
- Step 1** Create a new protection plan using the vContinuum wizard.

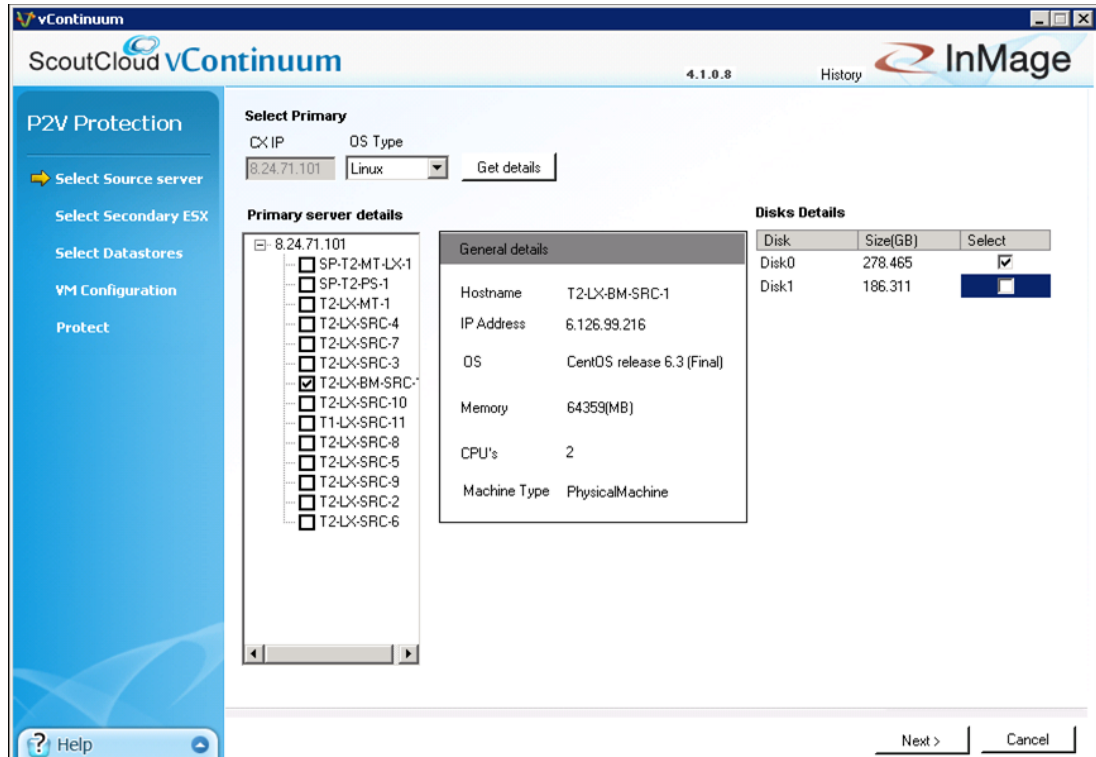
Figure 4-18 Create P2V Protection Plan using vContinuum



- On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut Start>Program>InMage System>VContinuum>vContinuum.
- Select P2V from the Choose Application drop-down list for P2V protection plan.
- Enter the CX server's IP address and port number (default is 80).
- Select New Protection to create a new protection plan.

**Step 2** Select the primary physical server(s) and volume(s) to protect.

Figure 4-19 Select Source Server



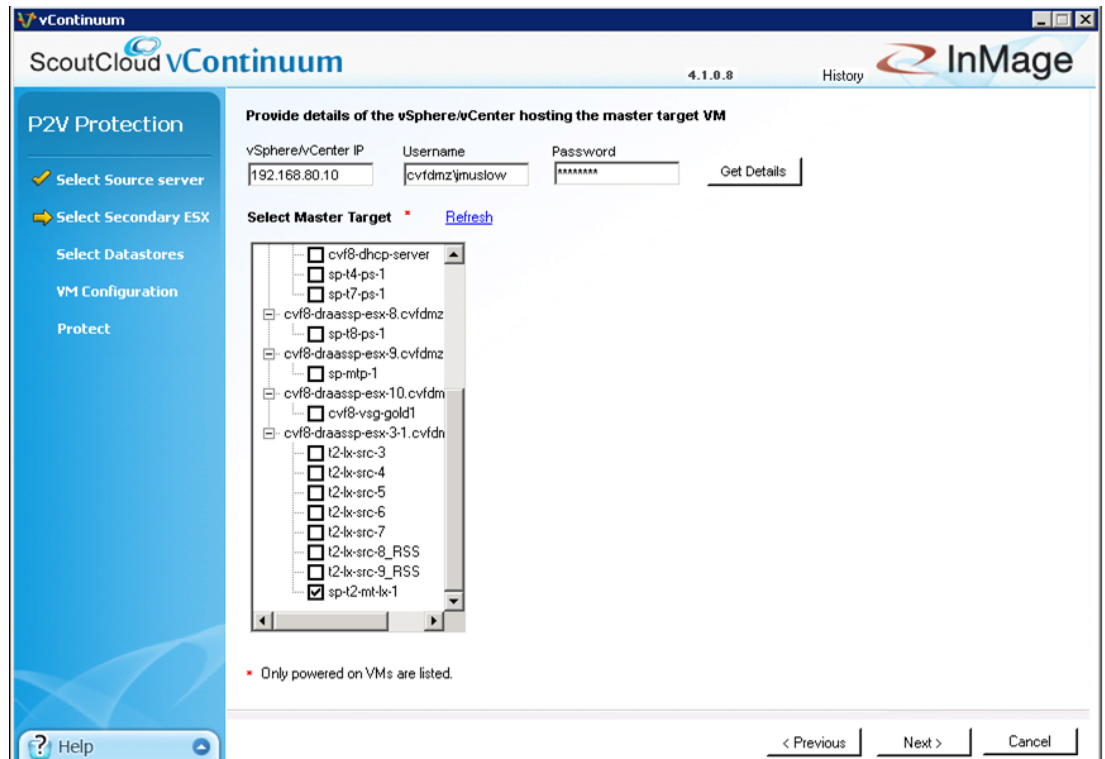
- Set the OS type to Linux and select Get Details to retrieve a list of registered Linux primary servers. The list includes both virtual and physical servers.
- Select the physical server(s) and disk(s) to protect.
- Click **Next** to continue.



**Note** By default, all local volumes are selected for protection. Volumes can be omitted from the protection plan by deselecting the volumes. When protecting Linux physical servers, disk labels displayed in vContinuum may not map to the same disk in the physical server. For example, the label disk0 displayed in vContinuum may not map to the Linux physical disk /dev/sda.

**Step 3** Select the secondary site MT.

Figure 4-20 Select Secondary ESX and Master Target



- Enter the secondary vSphere/vCenter IP address and login credentials.
- Click **Get Details** to list the available MT(s).
- Select the MT that will be used to protect the selected primary physical server(s) or VM(s).
- Click **Next**.



**Note** The MT must be of the same OS family as the primary physical server it protects. If the primary server run Linux OS, then the MT must also run Linux OS. The same requirement exists for Windows servers.

**Step 4** Configure the replication options and specify the secondary site datastore.

Figure 4-21 Configure Replication Options and Select Datastore

The screenshot shows the 'Provide Replication options' configuration screen in the vContinuum ScoutCloud interface. The interface includes a sidebar with navigation options: 'Select Source server', 'Select Secondary ESX', 'Select Datastores', 'VM Configuration', and 'Protect'. The main area contains a table for replication options and a table for secondary vSphere datastores.

Primary Server Ip	Total disk size(GB)	Process server IP	Retention size(in MB)	Retention drive	Retention(in days)	Consistency interval(in mins)	Select target datastore	Select
6.126.99.216	278	6.126.99...	Enter value	/mnt/...	1	30	Silver-Data-1	Select

Data Store	Total Space (GB)	Used Space (GB)	Free Space (GB)
datastore1	274.5	0.9500122	273.55
SP-Storage-8TB	8000	25.06006	7974.94
silver-Data-2	749.75	81.94	667.81
...	...	...	...

Below the tables, there is a checkbox for 'NAT Configuration' and navigation buttons: '< Previous', 'Next >', and 'Cancel'.

- In the Process server IP field, select the process server located in the primary network (e.g., Enterprise). The process server is used for both virtual and physical volume replication. Multiple process servers can be deployed for scale.
- In the Retention size (in MB) field, enter the maximum amount of disk space to use for retention data.
- In the Retention Drive field, select the Windows drive letter or Linux mount point that is associated with the retention drive on the MT.
- In the Retention (in days) field, enter the maximum number of days to store retention data.
- In the Consistency interval (in mins) field, enter the number of minutes between execution of the replication jobs. The replication jobs will run every x minutes generating application consistency recovery points for the primary VMs. This value determines the RPO for consistency point-based recovery.
- Click **Next** to continue.

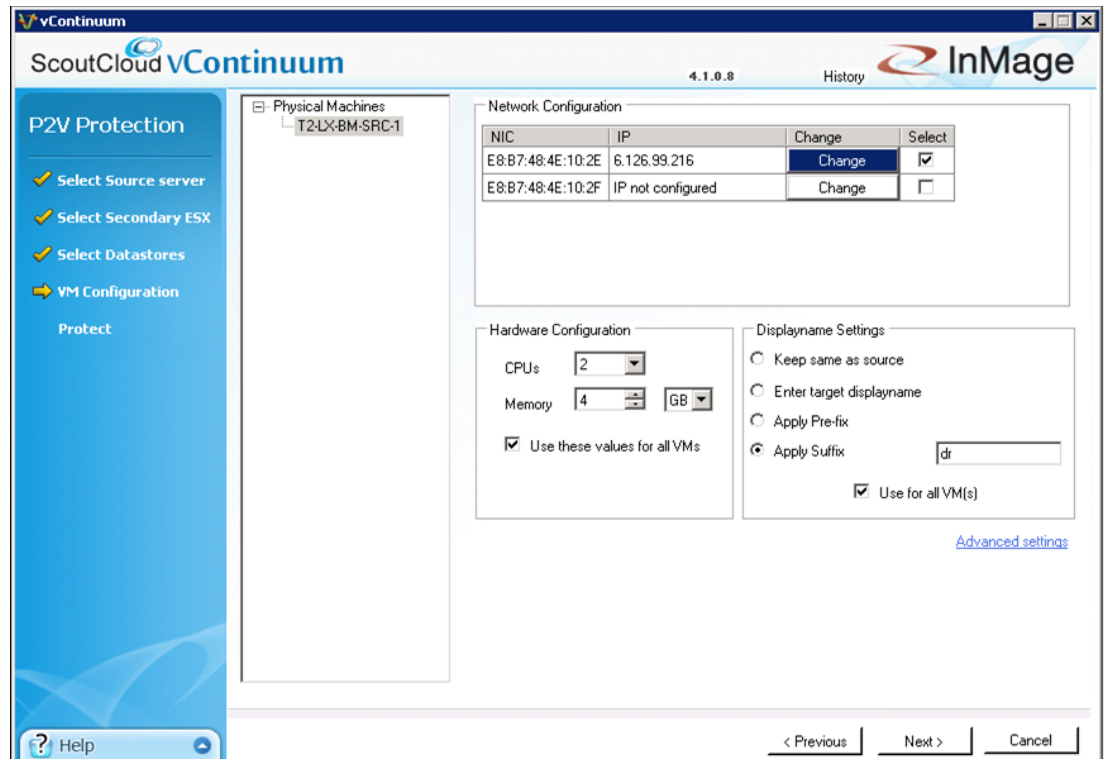
**Note**

The amount of retention data can be limited by disk space, time, or both. The vContinuum wizard release used during testing (v.4.1.0.0) did not allow the user to configure the retention window lower than one day. If a retention window smaller than one day is desired, the retention window can be later adjusted to less than one day through the CX UI in the **Protect > Volume > Settings** page.

**Step 5** Configure the secondary VM(s) configuration options.



Figure 4-22 Configure the VM Configuration



- Select the network interfaces to include in the secondary VM(s).
- Click **Change** to assign the interface IP address and port group. The address and port group can be static or assigned dynamically via DHCP.
- Configure the hardware settings for the secondary VM(s); for example, the number of CPU and size of RAM. These settings can be applied to all secondary VMs in the protection plan by selecting **Use these values for all VMs**.
- Configure the display name for the secondary VM(s) in the vSphere/vCenter. The "Keep same as source," "Apply Pre-fix," and "Apply Suffix" options can be applied to all VMs in the protection plan by selecting **Use for all VM(s)**.
- Click **Advanced settings** to access optional advanced settings for the protection plan.
- Or, click **Next** to continue.

**Note**

If multiple interfaces are selected, some additional configuration may be required for proper routing. Only one of the interfaces should be configured with a default route and if any static routes are required, they must be added after the recovered VM is powered up via scripts or manual configuration.

**Step 6** Configure the secondary VM(s) configuration optional advanced settings.

Figure 4-23 VM Configuration Advanced Settings

Advanced settings

Scout vContinuum 4.1.0.8 History InMage

Advanced setting for T2-LX-BM-SRC-1

**Sparse Retention**

Backups are retained for 90 days(3 months)

Provide continuous backup for latest 1 days

Advanced retention

<input type="checkbox"/>	From 1 days onwards. Provide 1 restore point per 1 hour for next 1 days
<input type="checkbox"/>	From 1 days onwards. Provide 1 restore point per day for next 1 weeks
<input type="checkbox"/>	From 1 days onwards. Provide 1 restore point per week for next 1 months

Apply for all servers Total days selected 1

**Folder Name Settings**

Keep VM in datastore's root directory

Keep VM in datastore's sub directory

Apply for all servers

**Compression**

No compression

Compression

CX based  Source based

Apply for all servers

**Encryption**

Secure transport from Source to CX-PS

Secure transport from CX-PS to destination

Apply for all servers

**Resource pool**

Select resource pool on target

Apply for all servers

**Provisioning**

Thin provisioning

Thick provisioning

Apply for all servers

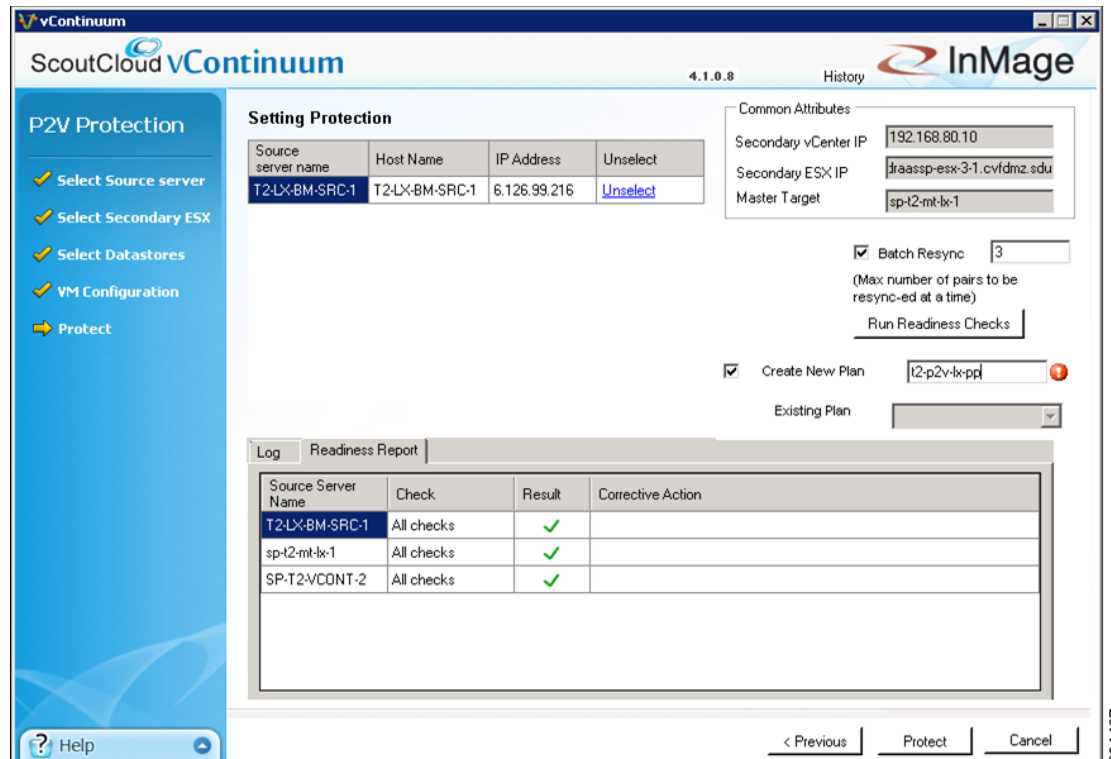
Ok Cancel

294426

- In the Sparse Retention section, advanced retention settings can be configured to have varying number of retention points based on age. Less retention points can be stored for weeks or months in the past, while more retention points can be stored for days in the past.
- In the Folder Name Settings section, the directory for the VM in the datastore can be configured.
- In the Compression section, compression type can be changed or disabled.
- In the Encryption section, encryption can be enabled for the primary VM to Process Server path, the Process Server to MT path, or both.
- In the Resource pool section, a resource pool on the secondary vSphere/vCenter can be specified. Resource pools can be used to isolate tenants from each other.
- In the Provisioning section, thin or thick provisioning can be configured.

**Step 7** Run the readiness check, name the protection plan, and protect.

Figure 4-24 Run the Protection Plan



- Click **Run Readiness Checks** to perform checks.
- Enter a name for the protection plan.
- Click **Protect** to finalize the protection plan.

**Step 8** Monitor the protection plan status from the CX UI.

Resync and differential data in transit can be monitored from the CX UI. Refer to [“RPO and Health Monitoring”](#) section on page 5-5 for details on monitoring.

Figure 4-25 Monitor the Protection Plan Status

ScoutCloud CX

v7.1.0.0.GA.3055

DASHBOARD PROTECT MONITOR RECOVER SETTINGS

Switch UI | July-25-2013 12:42 | Logout (ad)

Monitor » Application Protection » Plan Details  
Plan Details of "t2-p2v-lx-pp"

RPO Health: 1  
Protection Health: 1  
Data Consistency Health: 1  
Retention Health: 1

**Disks/Volumes/LUNs Replication**  
1-1 of 1 Records

Server	VX Agent Pair	Health	Health Issue	RPO	Resync progress	Status	Resync Required	Resync Data in Transit (MB)		Differential Data in Transit (MB)		View	
								Step1	Step2	On Primary Server	On Secondary Server		
T2-LX-8M-SRC-1->SP-T2-MT-LX-1	/dev/ada -> /dev/mapper/36000c2940678b12f7ed8a2886c5380e9	OK	N/A	1.13 min	N/A	Differential Sync	NO	0	0	0	0.05	0	Summary

**Files/Folders Replication**  
1-2 of 2 Records

Server	FX Agent Pair	Health	Status	Exit Code	Application	Job Description	Scheduled Type	Group ID	Job ID	Job Instance	View Details
SP-T2-MT-LX-1 -> SP-T2-MT-LX-1	Jar local/InMAGE/vx/falover_data/t2-p2v-lx-pp_sp-t2-mt-lx-1_218987 -> Jar local/InMAGE/vx/falover_data/t2-p2v-lx-pp_sp-t2-mt-lx-1_218987	OK	Stopping...	N/A	t2-p2v-lx-pp	Master Target - ...	Once Now	221	342	5812	Summary
T2-LX-8M-SRC-1 -> T2-LX-8M-SRC-1	Jar local/InMAGE/Fx/falover_data -> Jar local/InMAGE/Fx/falover_data	OK	Not started...	N/A	t2-p2v-lx-pp-Consistency4740	T2-LX-8M-SRC-1 - ...	Run Every	220	341	5811	Summary

294/428

## Offline Sync

The initial copy of data in a protection plan is both WAN bandwidth intensive and takes a long time to complete. The offline sync feature can be used to limit the amount of WAN bandwidth and time required for the initial protection plan sync. This is accomplished by sending this first time copy of data to the secondary site via a removable media instead of transmitting the data across the WAN. The data copied to removable media can then be shipped to the secondary site for offline import.

Primary servers are first protected to a local vSphere server in the primary site. After volume replication to the local MT is complete, the MT is shutdown and folders containing the MT VM, along with a temporary staging folder called InMAGE OfflineSync Folder, are copied to a removable media and shipped to the secondary site. Folders can then be restored to the secondary vSphere server via the Offline Sync Import feature. Once the Offline Sync Import is complete, replication resumes and data changes called differentials get sent across WAN network in a normal way.

Offline sync has the following three steps:

- 
- Step 1** Offline sync export.
  - Step 2** Transfer folders to the removable media and copy them to the secondary vSphere server.
  - Step 3** Offline sync import.
- 



### Note

This example documents two steps in the offline sync workflow: offline sync export and offline sync import. The transfer step is completely bypassed by allowing the exported data to be replicated to the secondary site MT and copied directly to a temporary staging folder of a secondary site datastore. The primary server used in this example is a physical server running CentOS.

This section presents the following topics:

- [Offline Sync Export, page 4-27](#)
- [Offline Sync Import, page 4-33](#)

## Offline Sync Export

**Note**

---

The following steps to configure offline sync export are based on the online Scout Help, which can be accessed from the main vContinuum page.

---

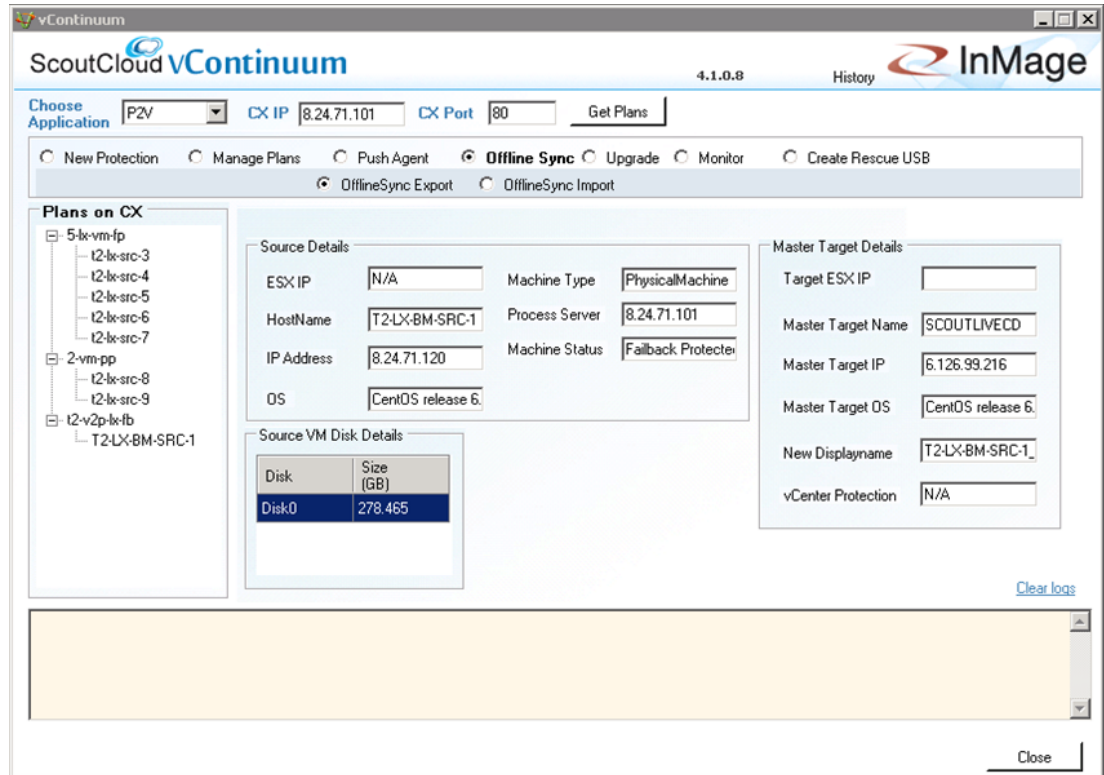
### Summary of Steps

1. Create an offline sync export plan using the vContinuum wizard.
2. Select the primary physical server(s) and volume(s) to protect.
3. Select the secondary site MT.
4. Configure the replication options and specify the secondary site datastore.
5. Run the readiness check to finalize the offline sync export plan.
6. Monitor the protection plan status from the CX UI.
7. Remove the MT from the vSphere/vCenter inventory.

### Detailed Steps

- 
- Step 1** Create an offline sync export plan using the vContinuum wizard.

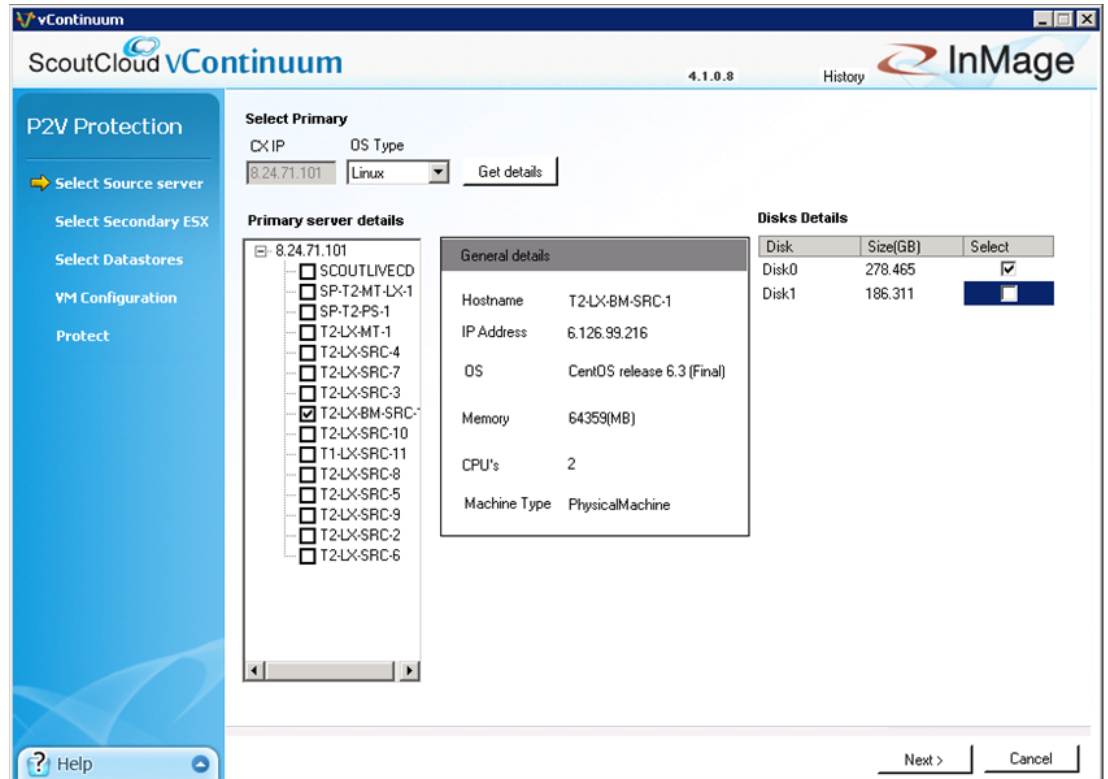
Figure 4-26 Create Offline Sync Export using vContinuum



- On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut **Start > Program > InMage System > VContinuum > vContinuum**.
- Select **P2V** from the Choose Application drop-down list for P2V protection plan.
- Enter the CX server IP address and port number (default is 80).
- Select **Offline Sync** and **OfflineSync Export** to create the offline sync export workflow.

**Step 2** Select the primary physical server(s) and disk(s) to protect.

Figure 4-27 Select Primary Physical Server to Protect



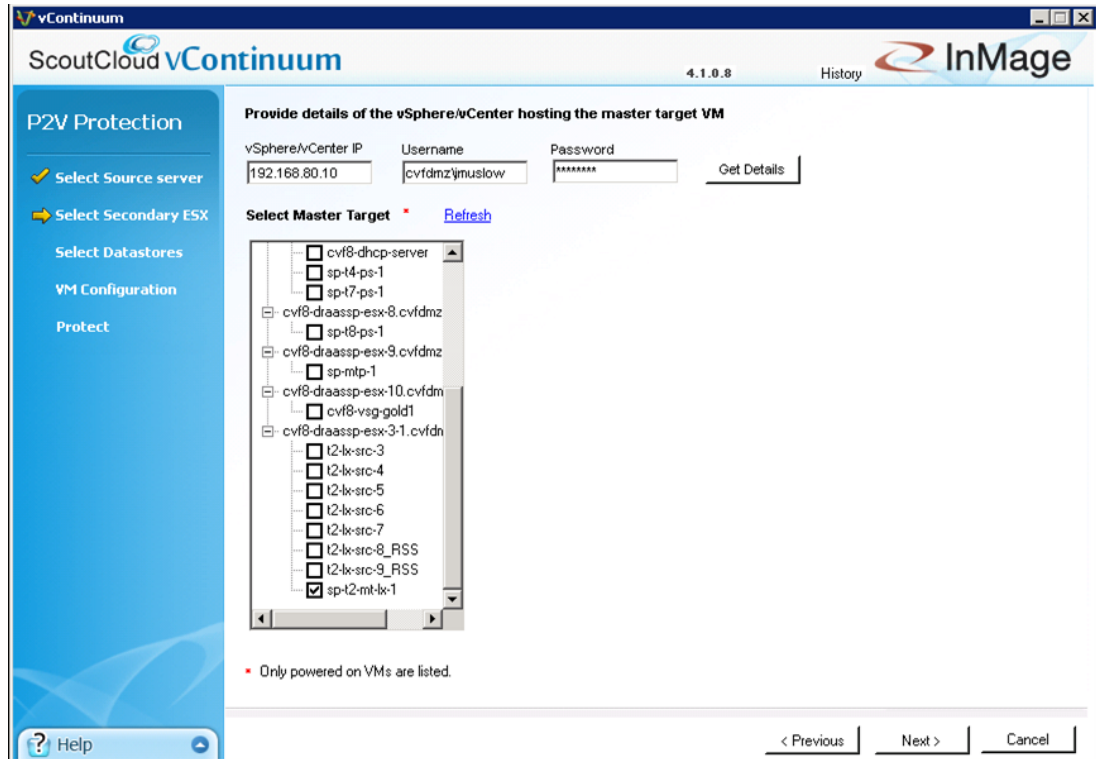
- Set the OS type to Linux and select **Get Details** to retrieve a list of available registered Linux primary servers. The list includes both virtual and physical servers.
- Select the physical server(s) and disk(s) to protect.
- Click **Next** to continue.



**Note** By default, all local volumes are selected for protection. Volumes can be omitted from the protection plan by deselecting the volumes. When protecting Linux physical servers, disk labels displayed in vContinuum may not map to the same disk in the physical server. For example, the label disk0 displayed in vContinuum may not map to the Linux physical disk /dev/sda.

**Step 3** Select the secondary site MT.

Figure 4-28 Select the Secondary ESX and Master Target



- Enter the secondary vSphere/vCenter IP address and login credentials.
- Click **Get Details** to list the available MT(s).
- Select the MT that will be used to protect the selected primary physical server(s).
- Click **Next** to continue.



**Note** The MT must be of the same OS family as the primary VM(s) it protects. If the primary VMs use Linux, then the MT must also be Linux. The same requirement exists for Windows servers.

**Step 4** Configure the replication options and specify the secondary site datastore.



Figure 4-29 Configure Replication Options and Select Datastore

ScoutCloud vContinuum 4.1.0.8 History InMage

**Provide Replication options**

Primary Server Ip	Total disk size(GB)	Process server IP	Retention size(in MB)	Retention drive	Retention(in days)	Consistency interval(in mins)	Select target datastore	Select
6.126.99.216	278	6.126.99...	Enter value	Amntre...	1	30	Silver-Data-1	Select

**Datastore(s) on secondary vSphere**

Data Store	Total Space (GB)	Used Space (GB)	Free Space (GB)
datastore1	274.5	0.9500122	273.55
SP-Storage-8TB	8000	25.06006	7974.94
silver-Data-2	749.75	81.94	667.81
63... Data-1	749.75	40.00000	709.75

NAT Configuration

< Previous Next > Cancel

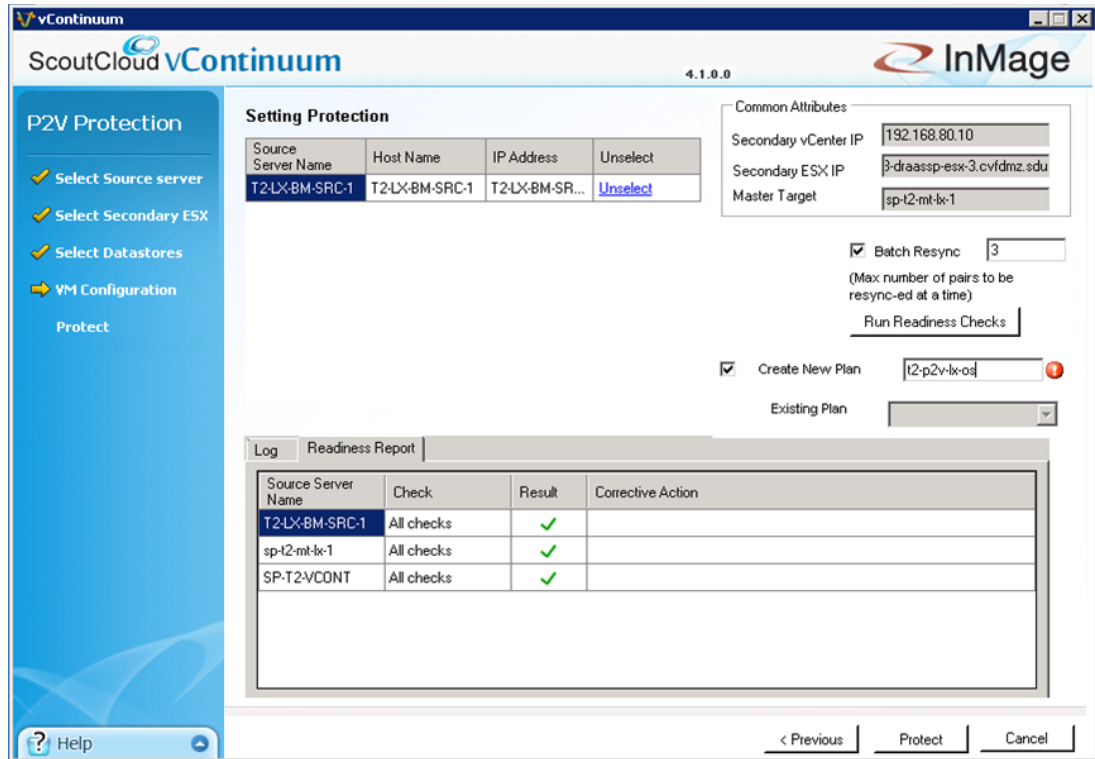
- In the Process server IP field, select the process server located in the primary network (for example, enterprise). The process server is used for both virtual and physical volume replication. Multiple process servers can be deployed for scale.
- In the Retention size (in MB) field, enter the maximum amount of disk space to use for retention data.
- In the Retention Drive field, select the Windows drive letter or Linux mount point that is associated with the retention drive on the MT.
- In the Retention (in days) field, enter the maximum number of days to store retention data.
- In the Consistency interval (in mins) field, enter the number of minutes between execution of the replication jobs. The replication jobs will run every x minutes generating application consistency recovery points for the primary VMs. This value determines the RPO for consistent point-based recovery.
- Click **Next** to continue.

**Note**

The amount of retention data can be limited by disk space, time, or both. The vContinuum wizard release used during testing (v.4.1.0.0) did not allow the user to configure the retention window lower than one day. If a retention window smaller than one day is desired, the retention window can be later adjusted to less than one day through the CX UI in the **Protect > Volume > Settings** page.

**Step 5** Run the readiness check to finalize the offline sync export plan.

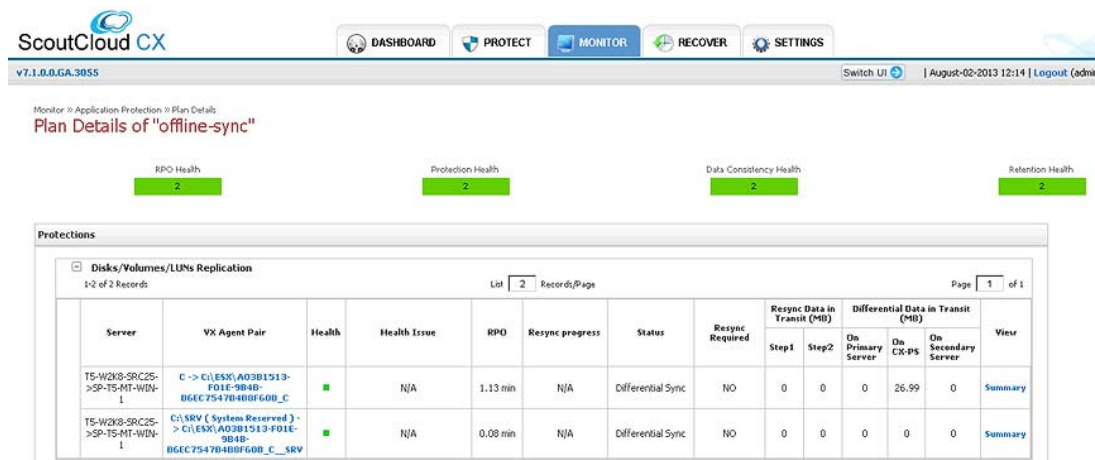
Figure 4-30 Run the Offline Sync Import Plan



- Click **Run Readiness Checks** to perform checks.
- Enter a name for the protection plan.
- Click **Protect** to finalize the protection plan.

**Step 6** Monitor the protection plan status from the CX UI.

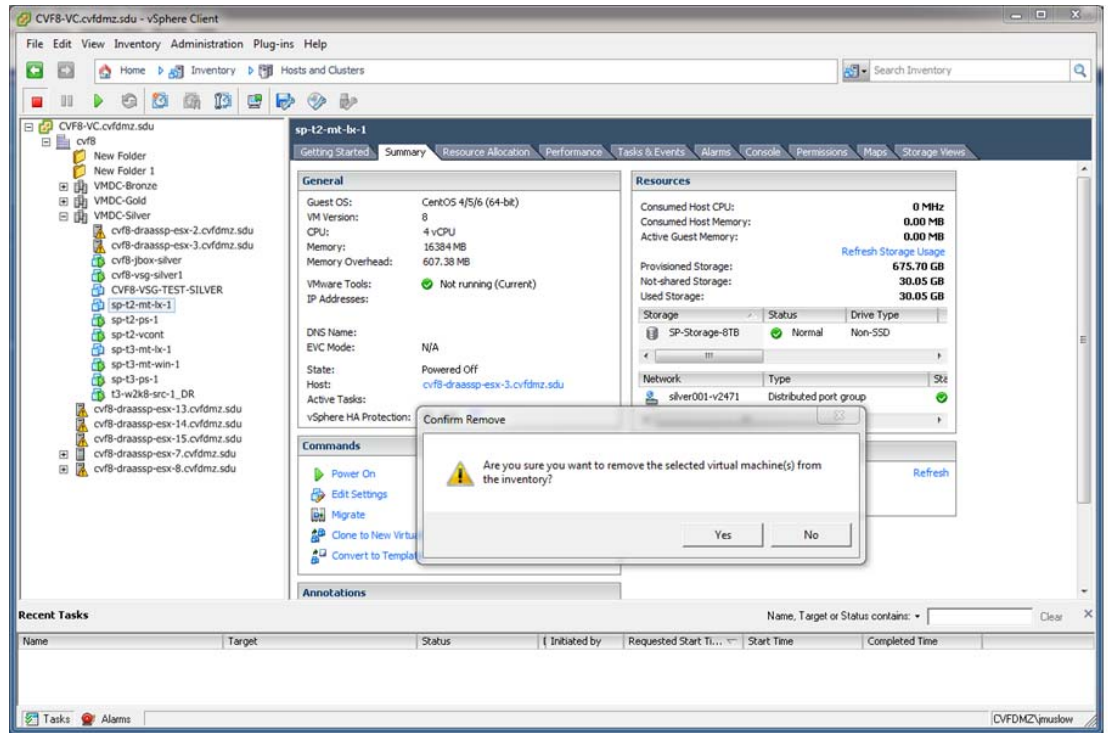
Figure 4-31 Figure 4-31. Monitor the Offline Sync Export Status using the CX UI





**Note** Volume replication is complete when all volume replications achieve Differential Sync status. Step 7 Remove the MT from the vSphere/vCenter inventory.

**Figure 4-32** *Figure 4-32. Remove the Master Target from vSphere/vCenter Inventory*



**Note** The MT should be powered off, then removed from the vSphere/vCenter inventory prior performing Offline Sync Import. Refer to [“Offline Sync Import”](#) section on page 4-33 for details on Offline Sync Import.

## Offline Sync Import



**Note** The following steps to configure offline sync import are based on the online Scout Help, which can be accessed from the main vContinuum page.

### Summary of Steps

1. Create an offline sync import plan using the vContinuum wizard.
2. Select the vSphere host and datastore.
3. Monitor the offline sync import plan in vContinuum.
4. Power on the MT.

## Detailed Steps

**Step 1** Create an offline sync import plan using the vContinuum wizard.

**Figure 4-33** Create the Offline Sync Import Workflow using vContinuum

The screenshot shows the vContinuum wizard interface. At the top, it displays 'ScoutCloud vContinuum 4.1.0.6' and the InMage logo. Below the title bar, there are input fields for 'Choose Application' (set to P2V), 'CX IP' (8.24.71.101), and 'CX Port' (80). A 'Get Plans' button is also visible. The main area contains several sections:

- Navigation:** Radio buttons for 'New Protection', 'Manage Plans', 'Push Agent', 'Offline Sync' (selected), 'Upgrade', and 'Monitor'. Under 'Offline Sync', there are radio buttons for 'OfflineSync Export' and 'OfflineSync Import' (selected).
- Plans on CX:** A tree view showing 't2-p2v-lx-os' and 'T2-LX-BM-SRC-1' (selected).
- Source Details:**
  - ESX IP: N/A
  - Machine Type: PhysicalMachine
  - HostName: T2-LX-BM-SRC-1
  - Process Server: 6.126.99.201
  - IP Address: 6.126.99.216
  - Machine Status: Protected
  - OS: CentOS release 6.
- Source VM Disk Details:**

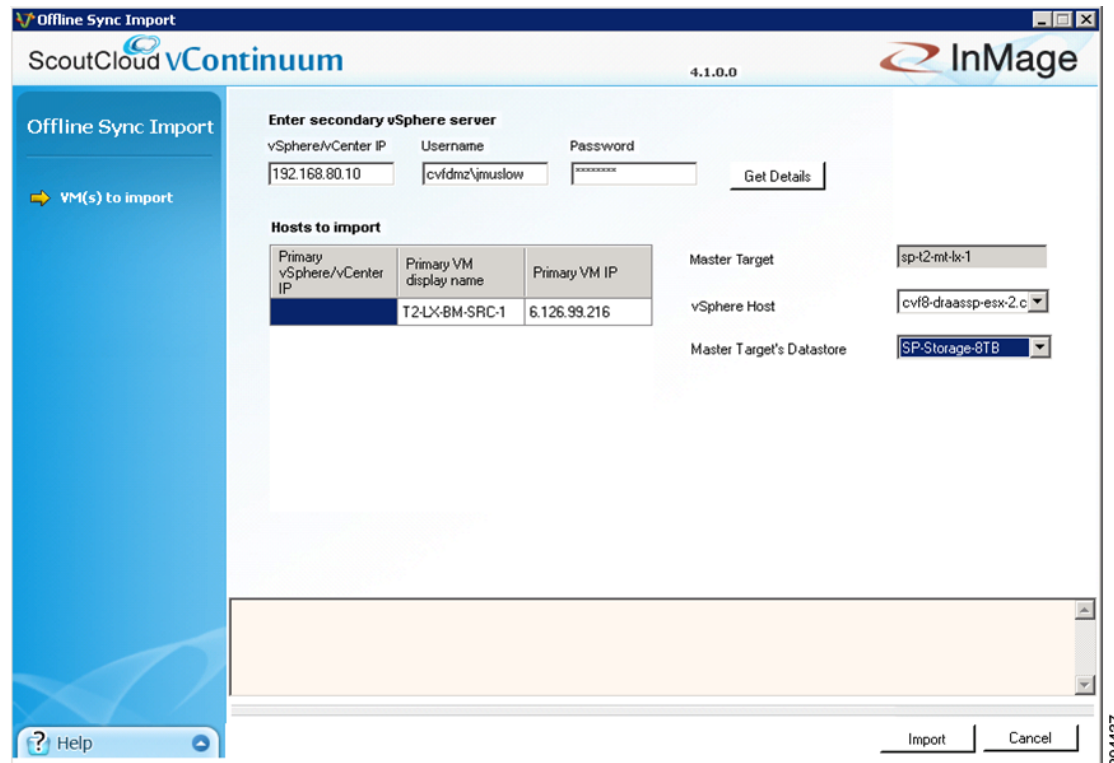
Disk	Size (GB)
Disk0	279.397
Disk1	279.397
- Master Target Details:**
  - Target ESX IP: cvf8-draassp-esx-
  - Master Target Name: sp-t2-mt-lx-1
  - Master Target IP: 8.24.71.102
  - Master Target OS: CentOS 4/5/6 (64)
  - New Displayname: T2-LX-BM-SRC-1
  - vCenter Protection: 192.168.80.10

At the bottom right, there are 'Next' and 'Close' buttons, and a 'Clear logs' link.

- On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut **Start > Program > InMage System > VContinuum > vContinuum**.
- Select **P2V** from the **Choose Application** drop-down list for P2V protection plan.
- Enter the CX server IP address and port number (default is 80).
- Select **Offline Sync** and **OfflineSync Import** to create the offline sync import workflow.
- Select the primary physical server(s) to import.
- Click **Next** to continue.

**Step 2** Select the vSphere host and datastore.

Figure 4-34 Configure the Secondary vSphere Server and Select the vSphere Host and Datastore



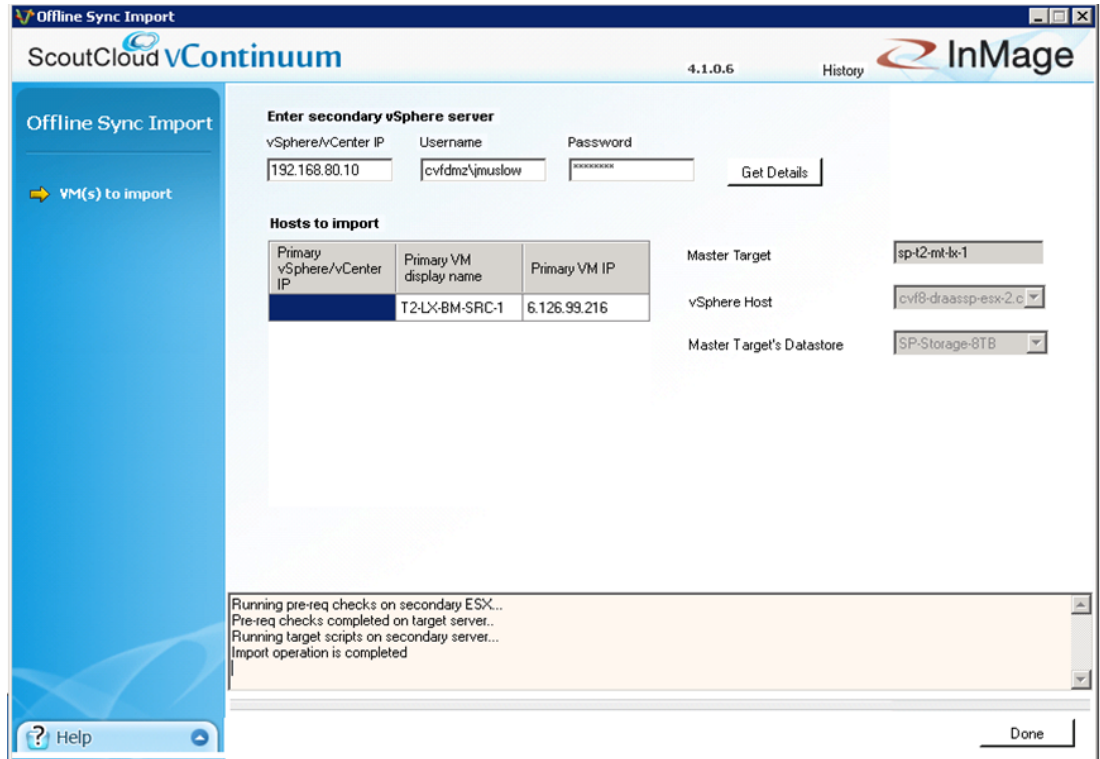
- Enter the secondary vSphere/vCenter IP address and login credentials.
- Click **Get Details** to list the available vSphere hosts and MT datastores.
- Select the vSphere ESX Host.
- Select the MT Datastore.
- Click **Import** to continue.



**Note** The exported data from the offline sync export plan was copied to the temporary staging folder on the secondary site MT and datastore. Therefore, the data transfer step via removable media is not required in this example.

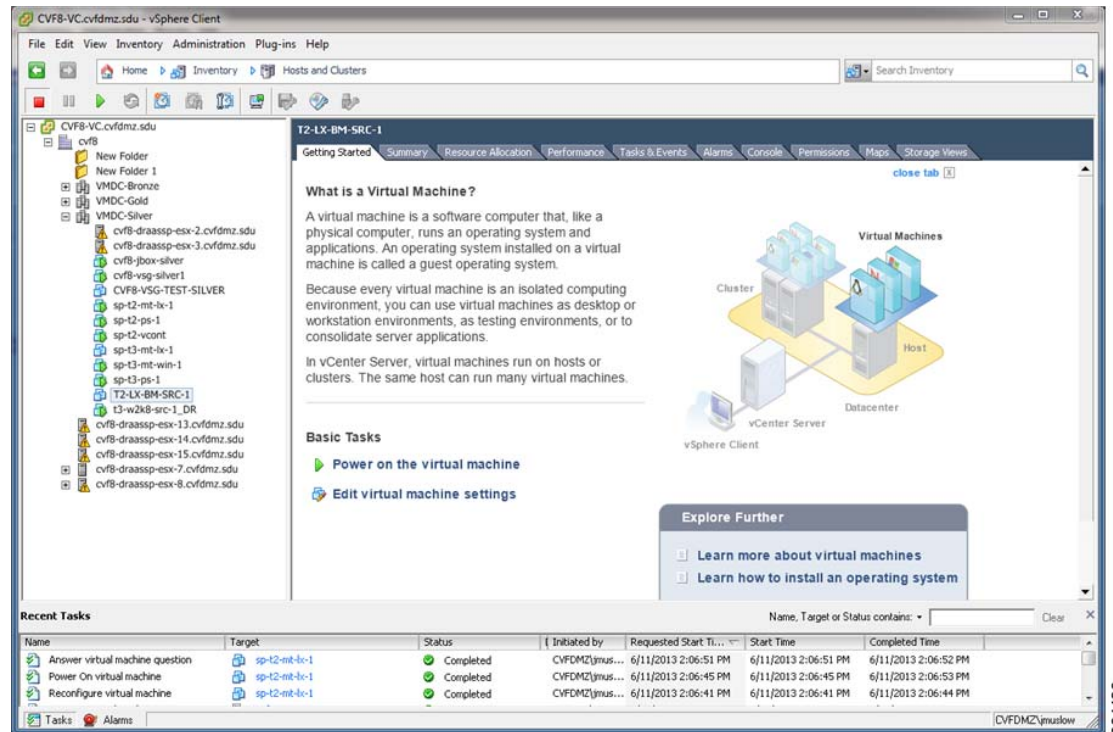
**Step 3** Monitor the offline sync import plan in vContinuum.

Figure 4-35 Run the Offline Sync Import Workflow



- Click Done once the import operation is completed. Step 4 Power on the MT.

Figure 4-36 Monitor the VMs on the Secondary vSphere vCenter



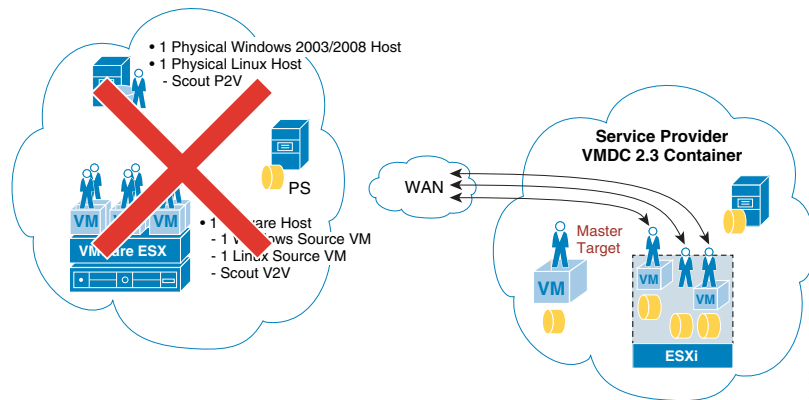
**Note** The offline sync import process is complete once the MT VM is added back to vSphere/vCenter inventory and powered on.

## Recovery Workflows

Once a protection plan is in place, the recover operation can be used to recover a primary server in a secondary vSphere environment when a disaster event occurs. This operation creates a VM on the secondary vSphere server based on a snapshot of the primary server. The snapshots can be based on the following consistency points or points in time:

- Latest application consistent point
- Latest point in time
- Consistency point near (prior to) any given time
- Specific time

Figure 4-37 Primary Server Recovery Overview

**Note**

The following steps to recover a primary server are based on the online Scout Help, which can be accessed from the main vContinuum page.

**Summary of Steps**

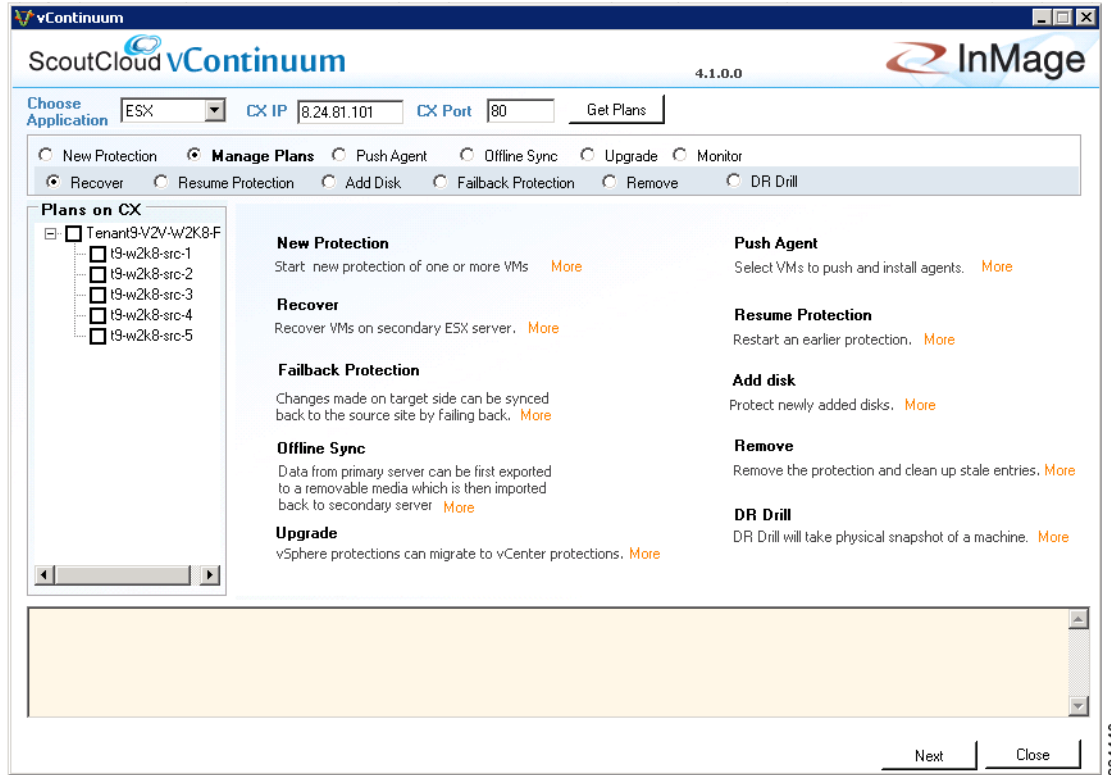
1. Start the vContinuum wizard application:
  - a. Select VM(s) to recover.
  - b. Specify snapshot to use based on time or tag.
  - c. Perform Readiness Check to make sure the VM(s) are ready for recovery.
  - d. Configure network, hardware, and display name for new VM(s).
  - e. Specify recovery job type and time to execute.
  - f. Finalize recovery and execute.
2. Monitor recovery.

**Detailed Steps**

- 
- Step 1** On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut **Start > Program > InMage System > VContinuum > vContinuum**.
- a. Select **ESX** from the **Choose Application** drop-down list to view V2V protection plans or P2V for P2V.
  - b. Enter the CX server IP address and port number (default is 80) and then click Get Plans.
  - c. Select the Manage Plans radio button and then click **Recover**.

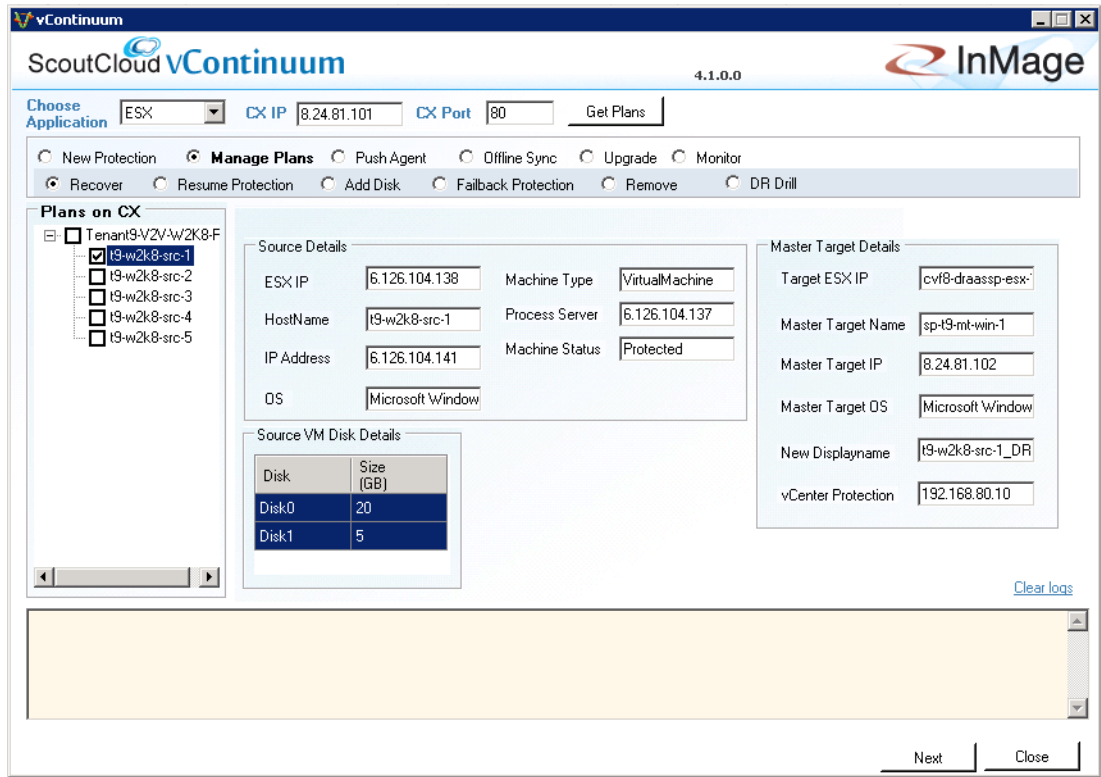


Figure 4-38 Starting Recover on vContinuum



- d. Select the primary VM(s) to recover and then click Next.

Figure 4-39 Selecting Primary VMs for Recovery



e. Specify snapshot to use based on time or tag:

- **Latest Tag:** Select this option to recover a VM to a latest tag that is common across all volumes of a VM. For example, if a VM that has three volumes (for example, C, E, and F), the latest common tag that is available across all volumes at same time point across all volumes is used.
- **Latest Time:** Select this option to recover a VM to a latest common point time among all volumes of a VM. Only common time points where volumes are in green state (data mode) are considered. For example, if a VM that has three volumes (for example, C, E, and F), the latest common time where all three volumes are in green state (data mode) is used.
- **Tag at Specified Time (Source Time Zone):** Select this option to recover a VM to a common tag prior to the specified time. For example, if a VM that has three volumes (for example, C, E, and F), the latest tag available prior to that time is used. The time provided is converted to GMT and compared against the timestamps in the retention logs. The closest consistency point prior to the time provided will be used to recover the VM.
- **Specific Time (Source Time Zone):** Select this option to recover a VM to a common point in time among all volumes of a VM. The time provided is converted to GMT and compared against the time stamps of the primary server. All recovery times are based on primary server's time stamps and not the secondary or management console times.

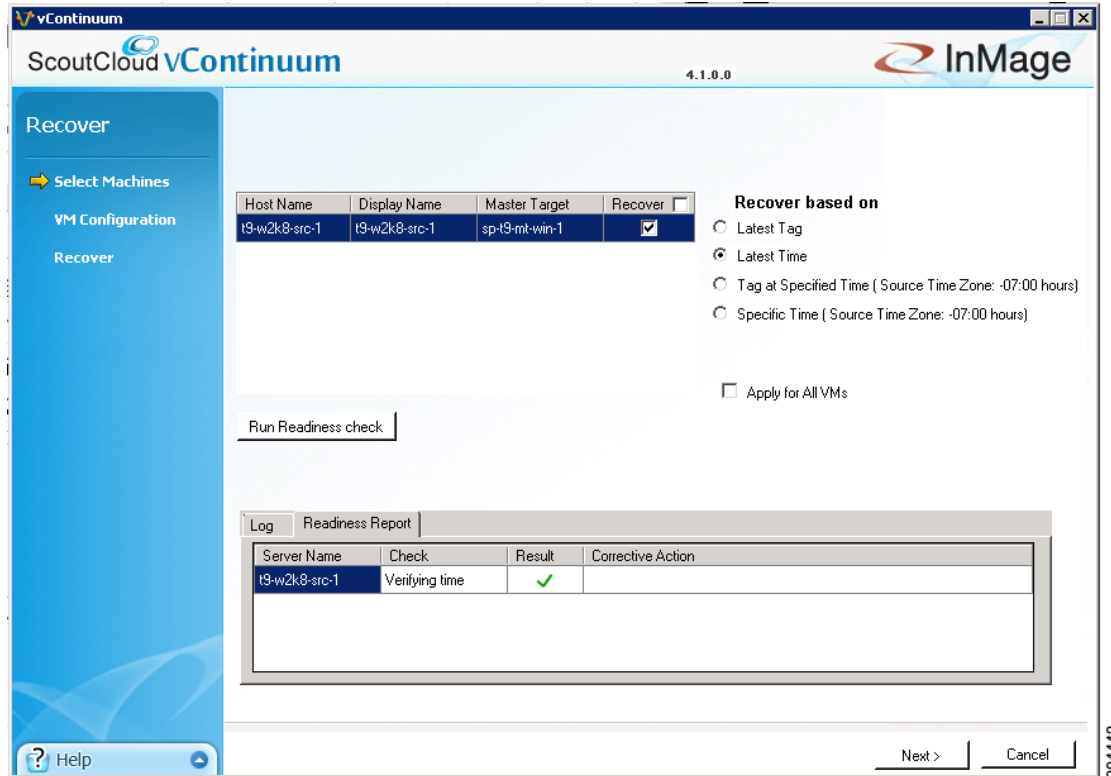


**Note**

Click Apply for all VMs to perform the recovery for all VMs at the specified snapshot type.

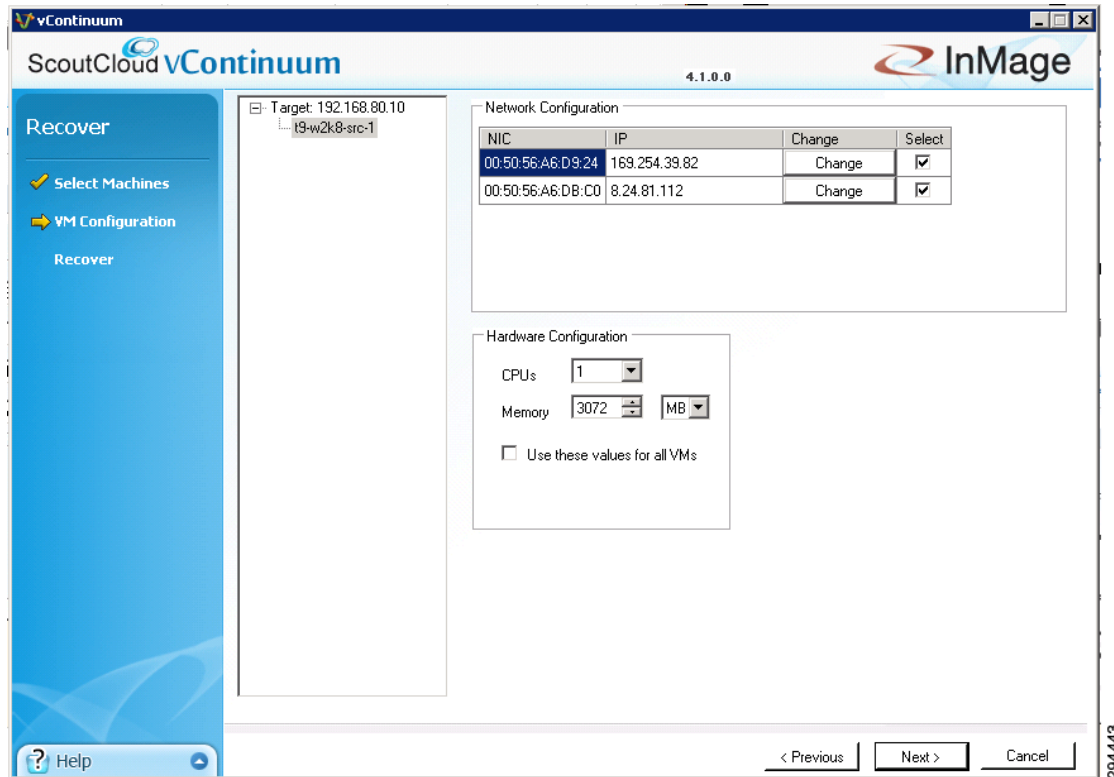
- f. Click **Run Readiness Check** to make sure the VM(s) are ready for recovery. If the check passes, click **Next** to advance to the next page.

**Figure 4-40** Running Readiness Check for Recovery



- g. Configure network and hardware settings for the new VM(s) if these configurations need to be different than what was defined in the original protection plan. In [Figure 4-41](#), no changes were made to the settings in the protection plan. Click **Next** to move to the next page.

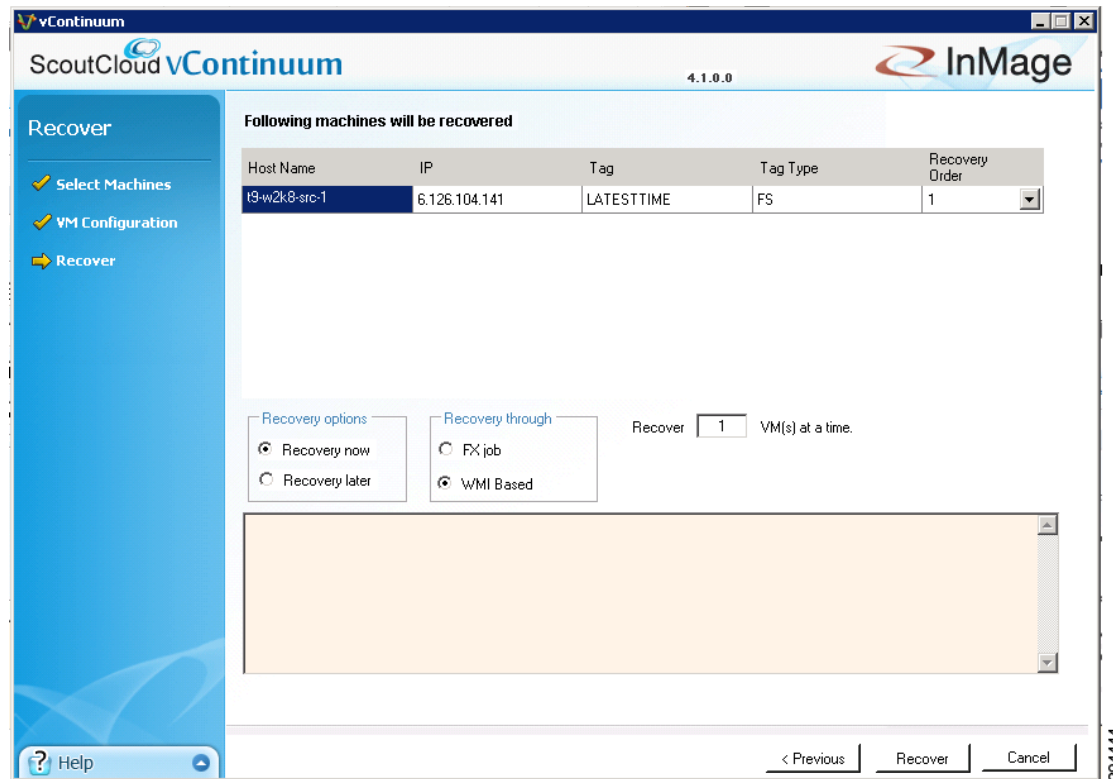
Figure 4-41 Modifying Network and Hardware Settings for Recovery



- h. Select the order in which the new VMs should be powered up. For example, if some VM(s) are dependent on another VM to up and running before they should power up, then you want to have the dependent VM(s) power up last. The default Recover Order is all "1" and all VM(s) will be powered up within seconds of each other.
- i. Specify recovery job type and time to execute based on one of the following three ways:
  - Recovery Option set to Recovery Now, Recovery Through set to FX Job, and plan name entered into Recovery Plan Name.
  - Recovery Option set to Recovery Now, Recovery Through set to WMI Based, and no plan name required.
  - Recovery Option set to Recovery Later, and plan name entered into Recovery Plan Name. If you select Recovery Later, an FX job is created, which can manually started at any time.

In Figure 4-42, the recover operation is configured for an immediate WMI-based recovery.

Figure 4-42 Finalizing Recover Operation

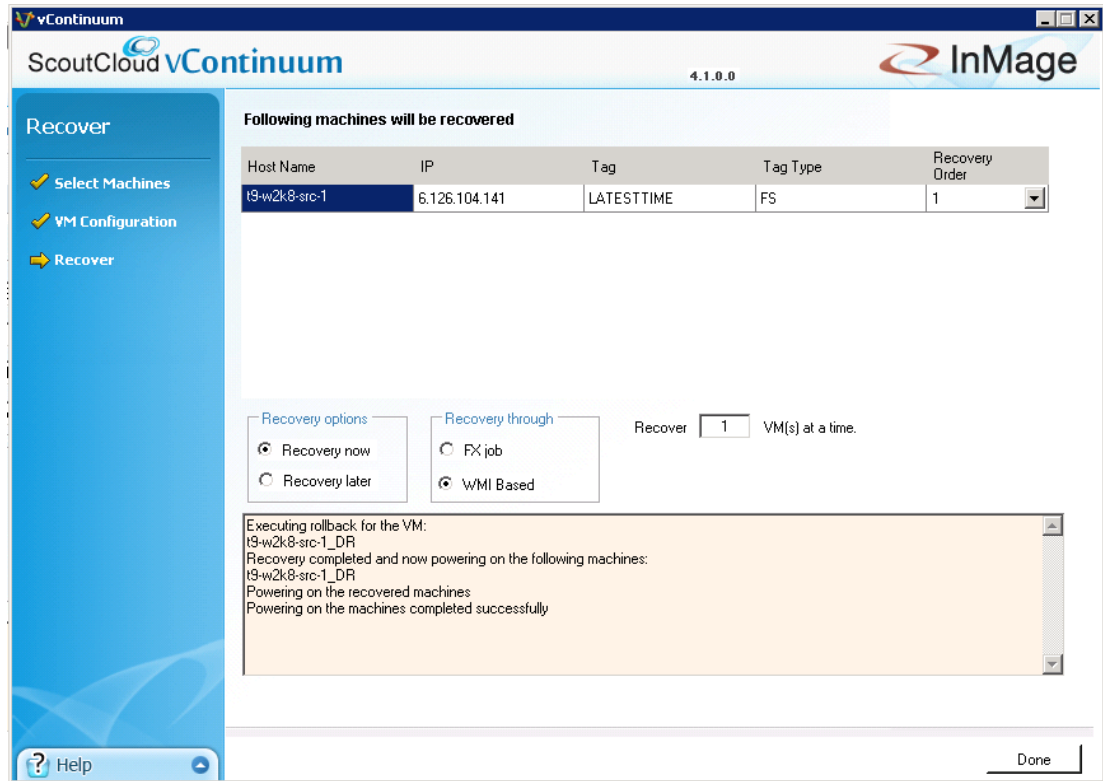


- j. Click **Recover** to start the recovery operation.
- k. Monitor recovery.

After starting the recovery, vContinuum goes through several steps to execute the recovery, which can be monitored from the status window.

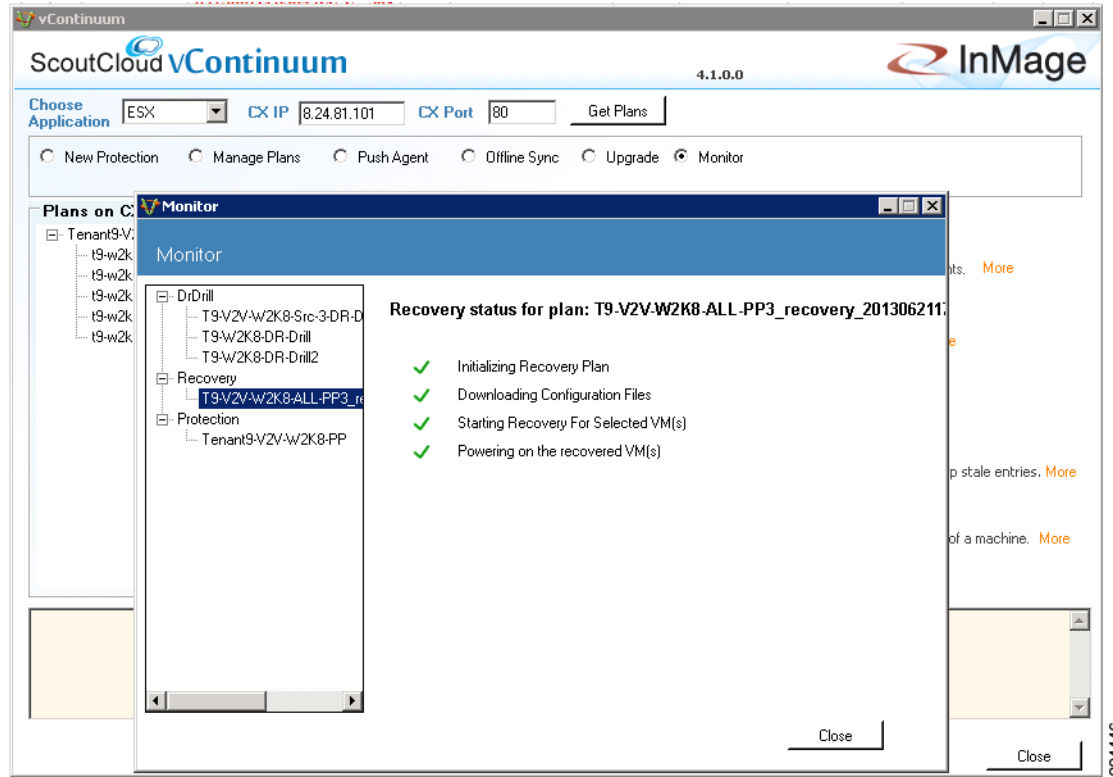
294444

Figure 4-43 Recover Starting (vContinuum)



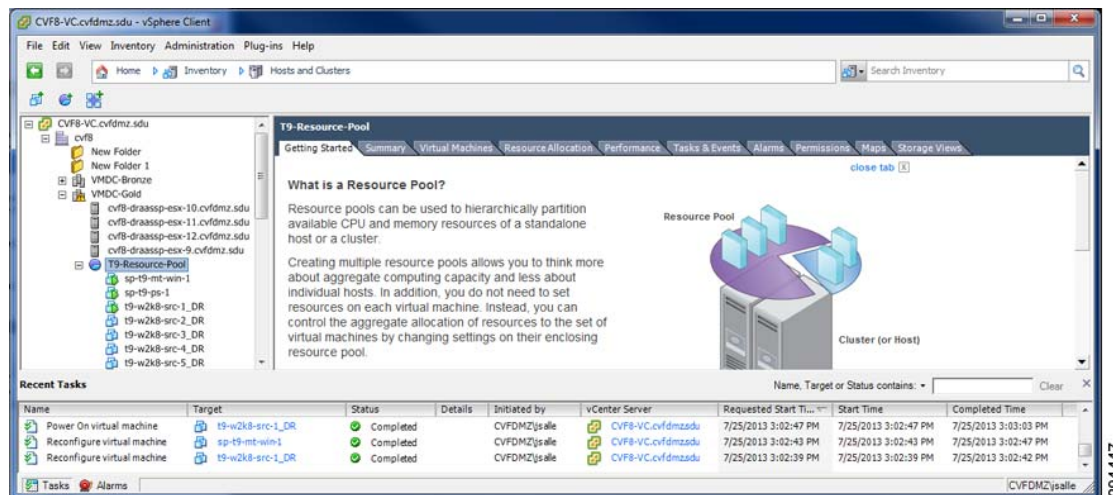
294445

Figure 4-44 Recover Monitoring (vContinuum)



294446

Figure 4-45 Recovered VM (Service Provider vCenter)



294447

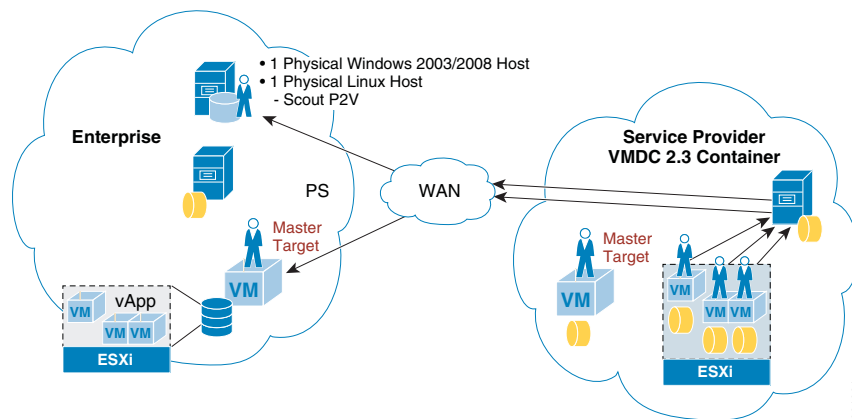
**Note**

Before executing the recovery, ensure that sufficient resources (for example, memory, CPU) are available on the target ESX/ESXi. If there are insufficient resources, the recovered VMs cannot be powered on and will have to be manually powered on.

## Failback Protection Workflows

Prepare a MT on the primary vSphere server before starting the protection. Failback operation replicates any new changes made on the secondary VMs back to the primary VM after failover. Failback can be done only on those VMs that are failed over to the secondary server. Replication pairs are set from secondary VMs running on secondary vSphere to the MT running on primary vSphere server.

**Figure 4-46** Failback Protection Overview



This section includes the following topics:

- [Virtual Failback Protection, page 4-46](#)
- [Virtual-to-Physical \(V2P\) Failback Protection, page 4-59](#)

## Virtual Failback Protection

**Note**

The following steps to configure and execute a failback recovery are based on the online Scout Help, which can be accessed from the main vContinuum page.

### Summary of Steps

1. Create failback protection plan.
  - a. Start vContinuum wizard application.
  - b. Select secondary VM(s) for failback recovery.
  - c. Select MT on primary vSphere/vCenter server.
  - d. Finalize protection plan.



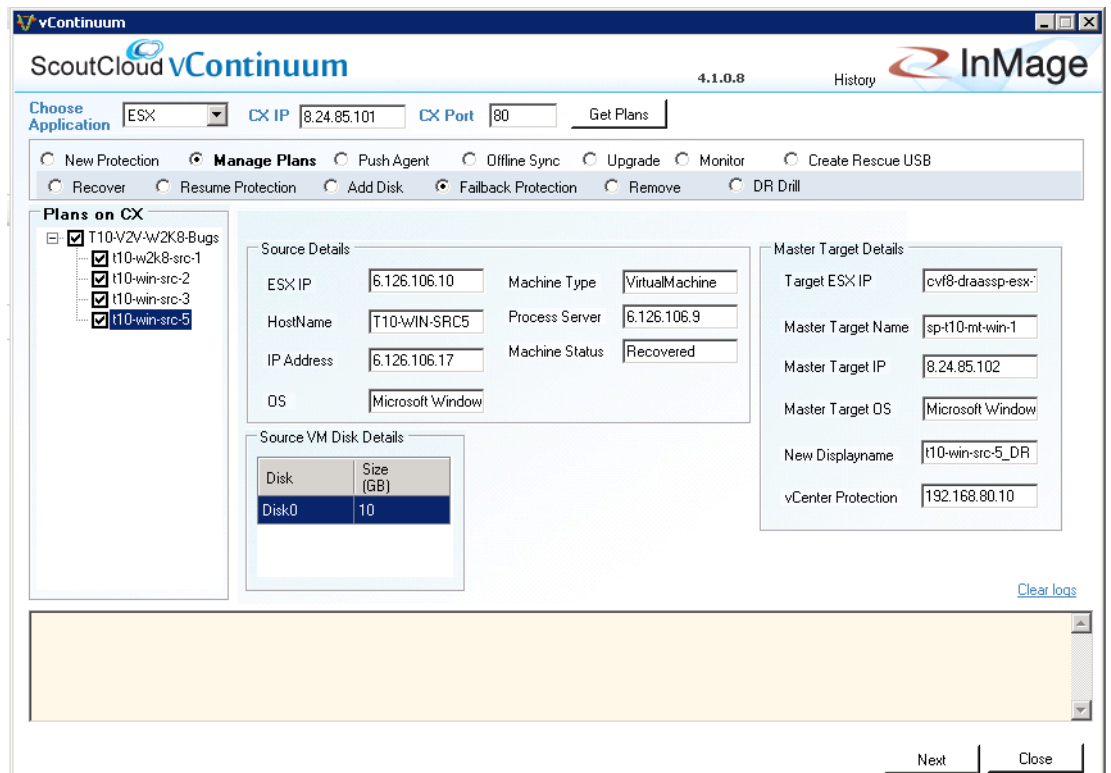
2. Monitor failback protection.
3. Execute failback recovery.
  - a. Specify snapshot to use based on time or tag.
  - b. Perform Readiness Check to make sure the VM(s) are ready for DR Drill.
  - c. Configure network, hardware, and display name for new VM(s).
  - d. Select datastore(s) for new VM(s).
  - e. Enter failback recovery plan name and initiate drill.
4. Monitor failback recovery.

### Detailed Steps

#### Step 1 Create failback protection plan.

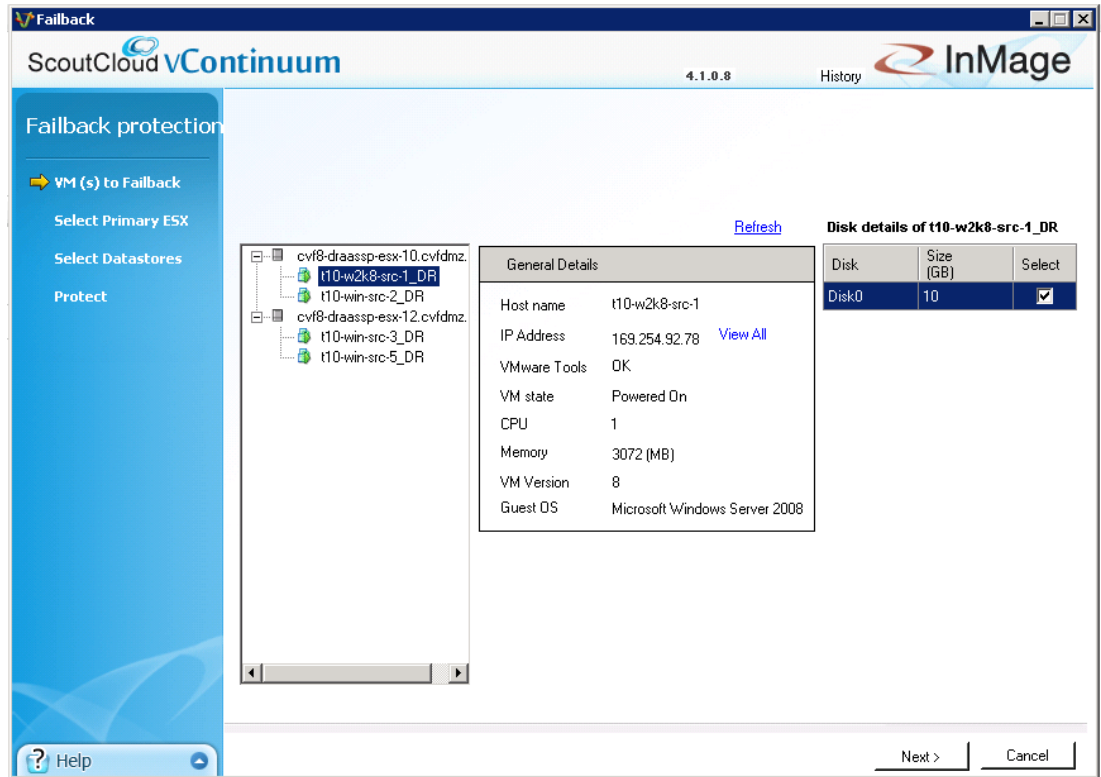
- a. On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut **Start>Program>InMage System>VContinuum>vContinuum**.
- b. Select **ESX** from the **Choose Application** drop-down list to view V2V protection plans.
- c. Enter the CX server IP address and port number (default is 80), then click **Get Plans**.
- d. Select the **Manage Plans** radio button and then click **Failback Protection**.
- e. Select the secondary VM(s) for failback recovery and click **Next**.

**Figure 4-47 Starting Failback Recovery Protection on vContinuum**



- f. Verify the selected secondary VM(s) for failback recovery and click **Next**.

Figure 4-48 Verify Secondary VM(s) Selected for Failback Protection



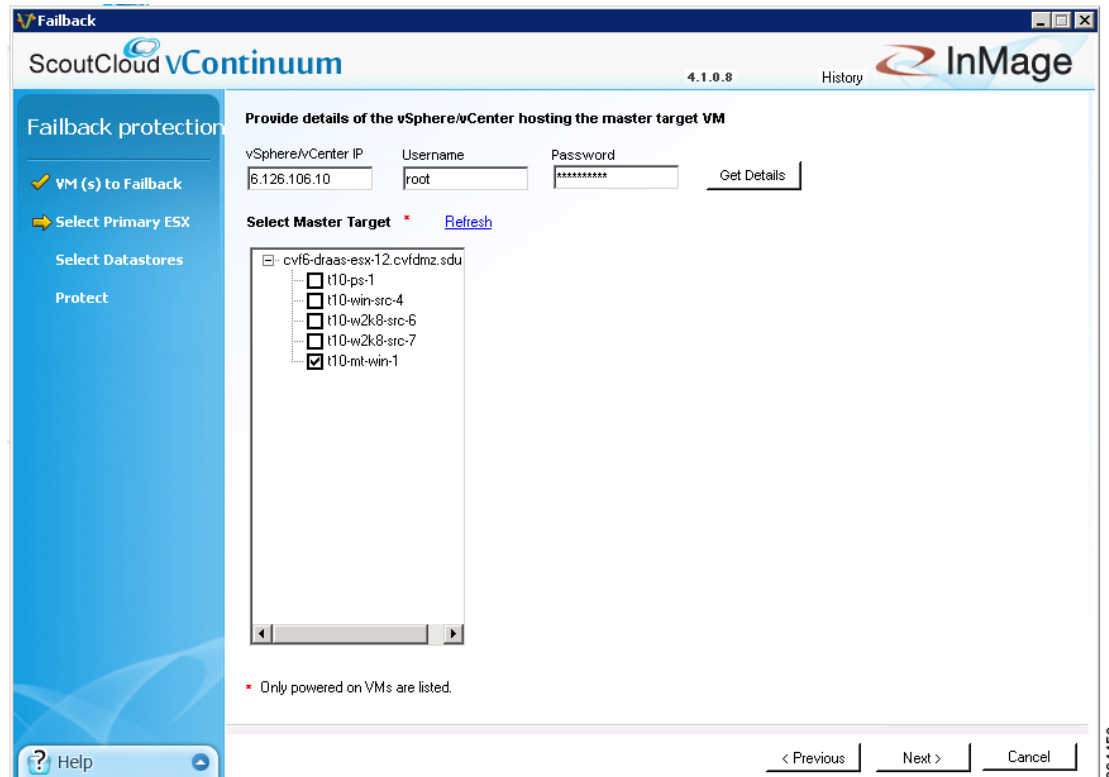
- g. Enter the primary vSphere/vCenter IP address, Username, and Password, then click **Get Details** to view the available MT(s).



**Note** The MT must be of the same OS family as the secondary VM(s) it protects. If the secondary VMs use Windows, then the MT must also be Windows. The same requirement exists for Linux servers. For more information on MT considerations, refer to [Master Target—Enterprise and Service Provider, page 3-1](#).

- Select the MT that will be used to protect the selected secondary VM(s) and click **Next**.

Figure 4-49 Selecting Primary Master Target(s)



h. Configure replication options.

- In the Process server IP field, select the process server located in the secondary network (for example, SP). Multiple process servers can be deployed and associated with a limited number of secondary VMs for scalability.
- In the Retention size (in MB) field, enter the maximum amount of disk space to use for retention data.
- In the Retention Drive field, select the drive letter that is associated with the retention drive on the MT.
- In the Retention (in days) field, enter the maximum number of days to store retention data.



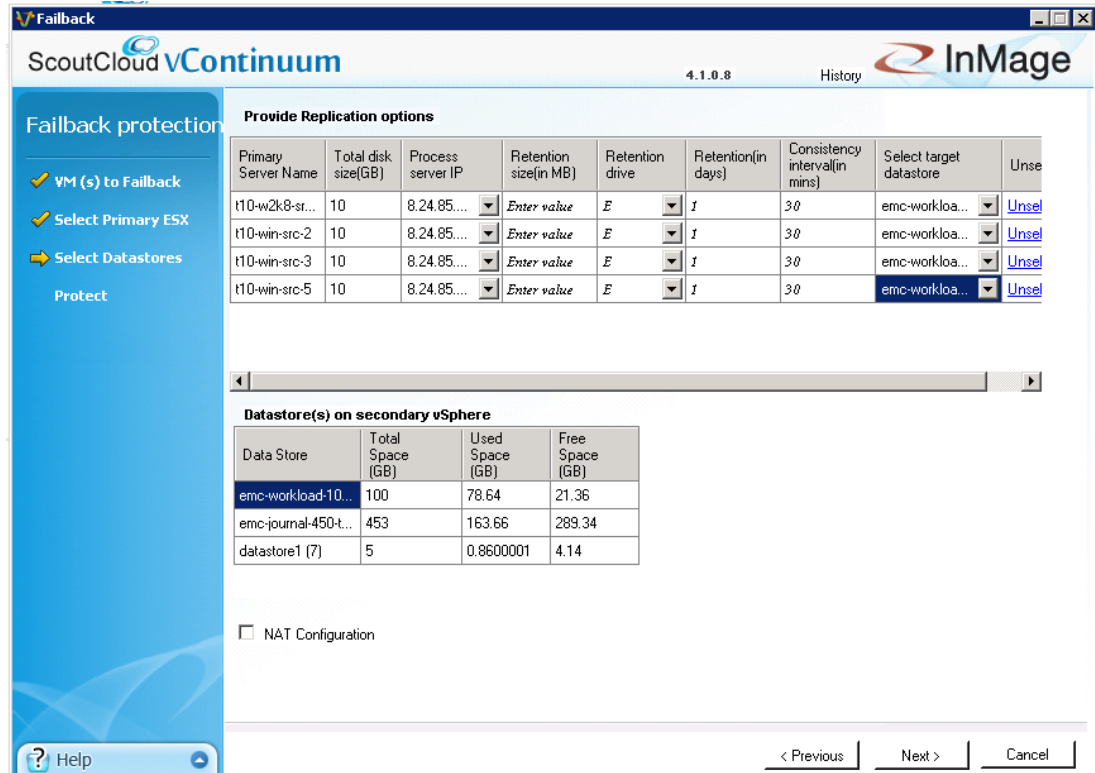
**Note**

- The amount of retention data can be limited by disk space, time, or both. For more information on retention data considerations, refer to [“Retention Volume Sizing” section on page 3-3](#).
- The vContinuum wizard release used during testing (v.4.1.0.0) did not allow the user to configure the retention window lower than one day. If a retention window smaller than one day is desired, the retention window can be later adjusted to less than one day through the CX UI in the **Protect > Volume > Settings** page.

- In the Consistency interval (in mins) field, enter the number of minutes between execution of the replication jobs. The replication jobs will run every x minutes generating application consistency recovery points for the secondary VMs. This value determines the RPO for consistency point based recovery.

- i. Select the target datastores in primary vCenter to create recovery VM(s) and click **Next**.

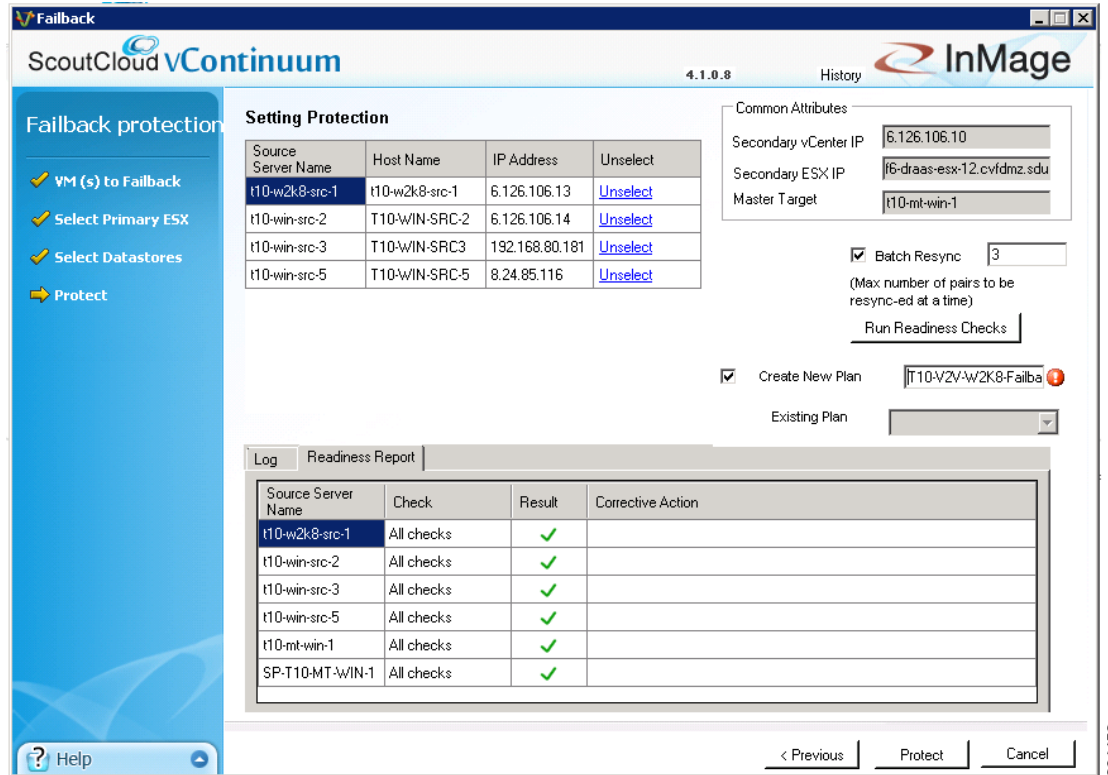
**Figure 4-50** Selecting Datastores in Secondary vCenter



- j. Finalize failback protection plan.
- Click **Run Readiness Checks** to perform checks.
  - Enter a name for the failback protection plan.
  - Click **Protect** to finalize the failback protection plan.

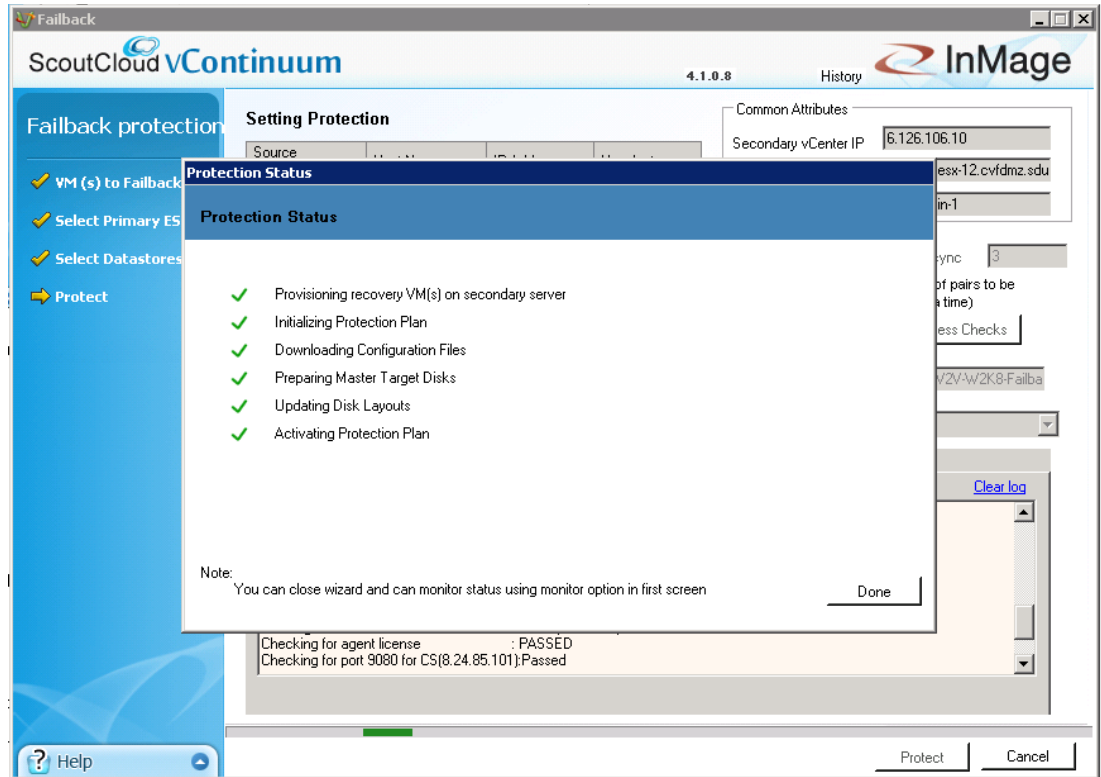
294451

Figure 4-51 Finalize Protection Plan



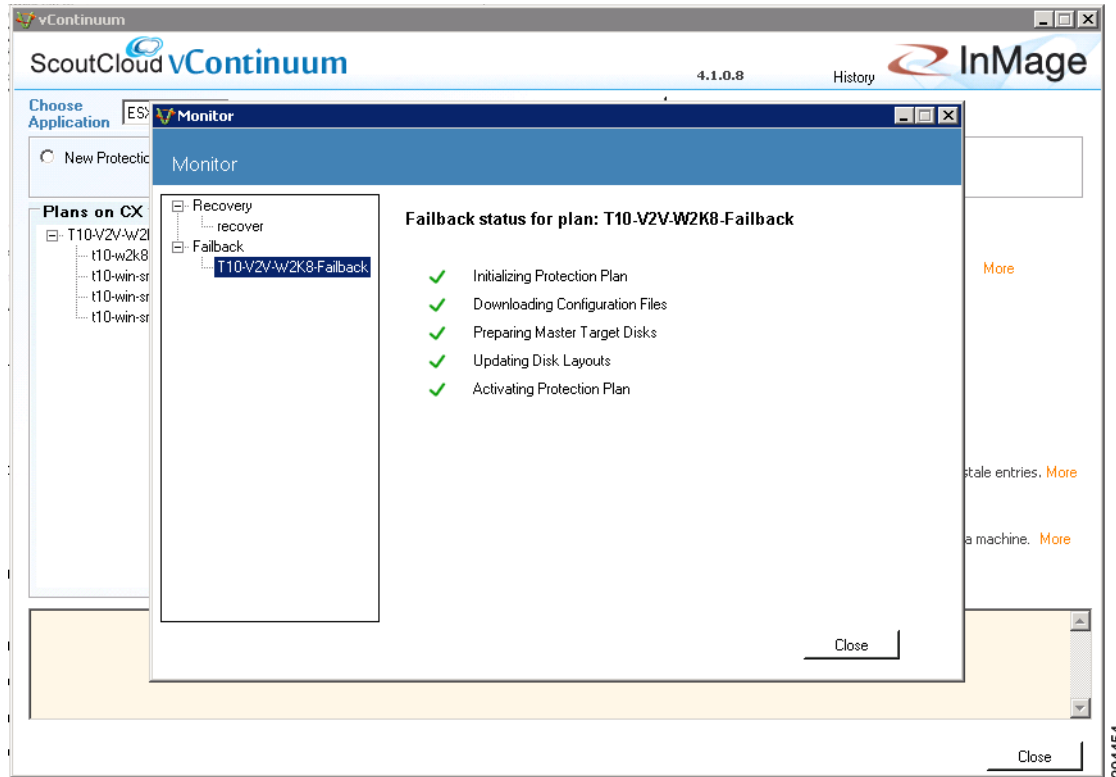
**Step 2** Monitor protection plan. After finalizing the failback protection plan, vContinuum goes through several steps to put the protection plan in place. The initialization of the protection plan can be monitored from the status window.

Figure 4-52 Failback Protection Plan Initializing (vContinuum)



294463

Figure 4-53 Recover Monitoring (vContinuum)



294454

Figure 4-54 Failback Replication Syncing (CX UI)

Monitor » Application Protection » Plan Details  
Plan Details of "T10-V2V-W2K8-Failback"

RPO Health: 8  
Protection Health: 8  
Data Consistency Health: 8

Protections

Disks/Volumes/LUNs Replication

1-8 of 8 Records

Server	VX Agent Pair	Health	Health Issue	RPO	Resync progress	Status	Resync Required	Resync Data In Transit (MB)		Differential Data In Transit (MB)			View
								Step1	Step2	On Primary Server	On Cl-PS	On Secondary Server	
T10-WIN-SRC-5->T10-MT-WIN-1	C -> C:\ESX\907E3FD9-B53A-F74A-ACECB16168E87A25_C	■	N/A	2.15 min	34.09 %	Resyncing (Step I)	YES	4.47	183.18	0	0	0	Summary
T10-WIN-SRC-3->T10-MT-WIN-1	C -> C:\ESX\83A9DEBF-FABE-7640-966EBCF7F18639AF_C	■	N/A	2.35 min	0 %	Resyncing (Step I)	YES	4.93	183.61	0	0	0	Summary
T10-W2K8-SRC-1->T10-MT-WIN-1	C -> C:\ESX\5986DC01-A9D5-964D-97D1B16E2B40B799_C	■	N/A	1.17 min	N/A	Differential Sync	NO	0	0	0	100.03	0	Summary
T10-WIN-SRC-2->T10-MT-WIN-1	C -> C:\ESX\51BA93FB-7575-B340-A1E306BA7F69C64D_C	■	N/A	1.05 min	N/A	Differential Sync	NO	0	0	0	17.06	0	Summary
T10-W2K8-SRC-1->T10-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\5986DC01-A9D5-964D-97D1B16E2B40B799_C_SRV	■	N/A	1.35 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T10-WIN-SRC-2->T10-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\51BA93FB-7575-B340-A1E306BA7F69C64D_C_SRV	■	N/A	1.47 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T10-WIN-SRC-5->T10-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\907E3FD9-B53A-F74A-ACECB16168E87A25_C_SRV ( System Reserved )	■	N/A	0.68 min	N/A	Resyncing (Step II)	YES	0	0	0	0	0	Summary
T10-WIN-SRC-3->T10-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\83A9DEBF-FABE-7640-966EBCF7F18639AF_C_SRV	■	N/A	2.63 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary

204455

**Step 3** Once the failback protection plan is completed and replication synced, a failback recovery can be executed.

- Select ESX from the Choose Application drop-down list to view V2V protection plans.
- Enter the CX server's IP address and port number (default is 80), then click Get Plans.
- Select the Manage Plans radio button and then click Recover.
- Select the secondary VM(s) for failback recovery and then click Next.



Figure 4-55 Starting Failback Recovery on vContinuum

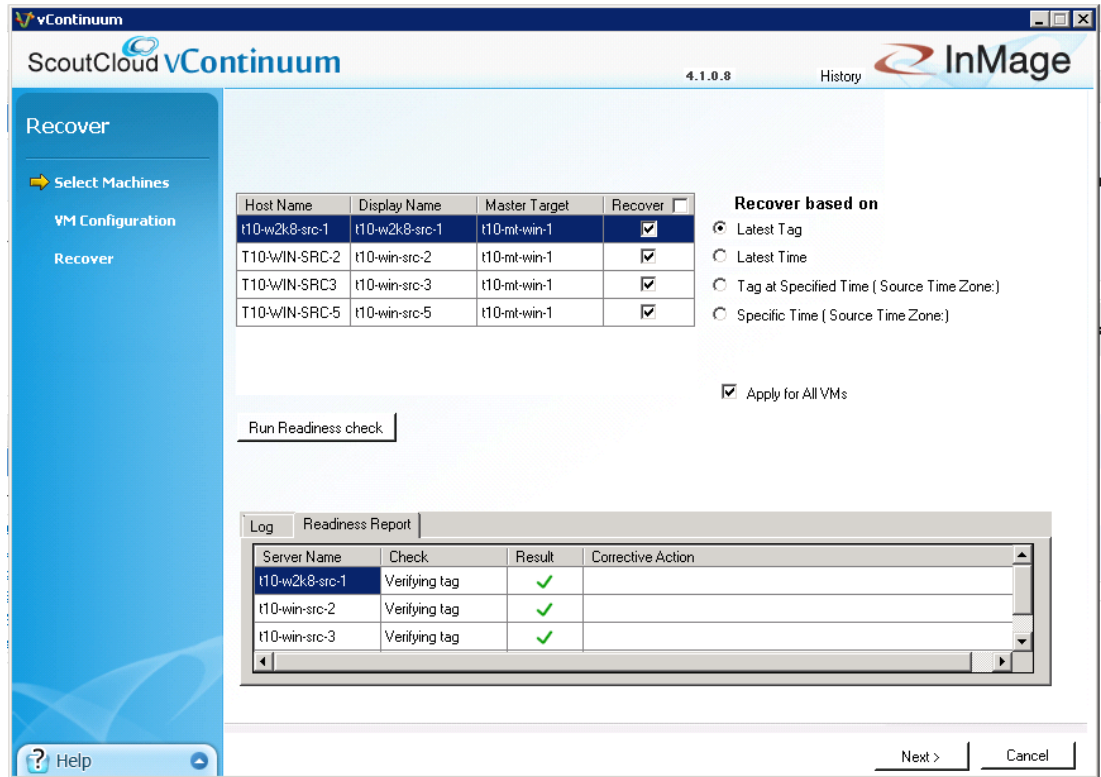
- e. Specify snapshot to use based on time or tag.
- **Latest Tag:** Select this option to recover a VM to a latest tag that is common across all volumes of a VM. For example, if a VM that has three volumes (for example, C, E, and F), the latest common tag that is available across all volumes at same time point across all volumes is used.
  - **Latest Time:** Select this option to recover a VM to a latest common point time among all volumes of a VM. Only common time points where volumes are in green state (data mode) are considered. For example, if a VM that has three volumes (for example, C, E, and F), the latest common time where all three volumes are in green state (data mode) is used.
  - **Tag at Specified Time (Source Time Zone):** Select this option to recover a VM to a common tag prior to the specified time. For example, if a VM that has three volumes (for example, C, E, and F), the latest tag available prior to that time is used. The time provided is converted to GMT and compared against the timestamps in the retention logs. The closest consistency point prior to the time provided will be used to recover the VM.
  - **Specific Time (Source Time Zone):** Select this option to recover a VM to a common point in time among all volumes of a VM. The time provided is converted to GMT and compared against the time stamps of the secondary server. All recovery times are based on secondary server's time stamps and not the primary or management console times



**Note** Click Apply for all VMs to perform the recovery for all VMs at the specified snapshot type.

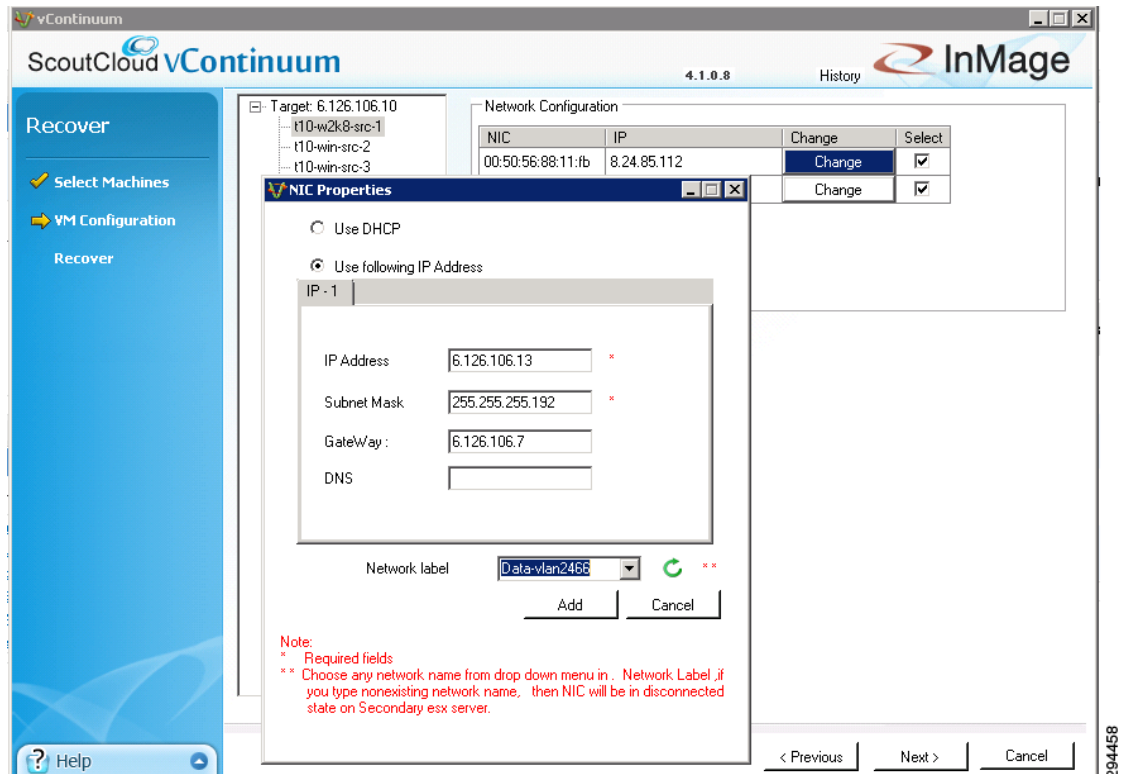
- f. Click **Run Readiness Check** to make sure the VM(s) are ready for recovery. If the check passes, click **Next** to advance to the next page.

Figure 4-56 Running Readiness Check for Failback Recovery



- g. Configure network and hardware settings for the new VM(s) if these configurations need to be different than what was defined in the original protection plan. In Figure 4-57, no changes were made to the settings in the protection plan. Click **Next** to move to the next page.

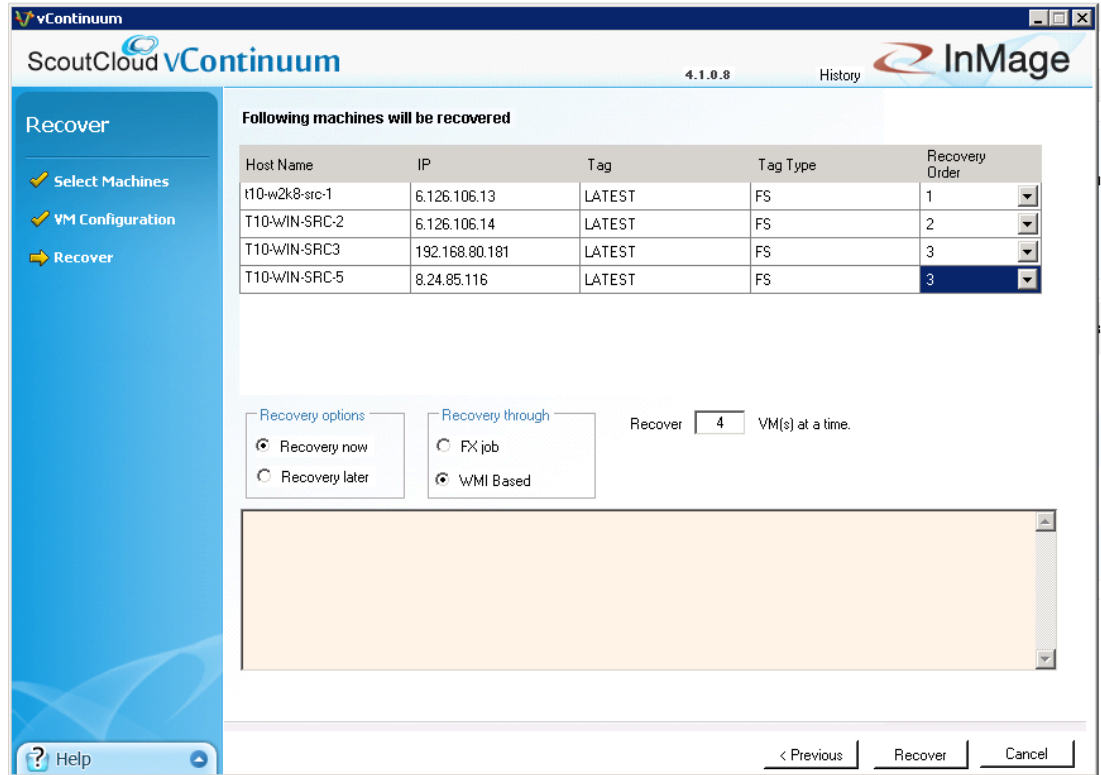
Figure 4-57 Modifying Network and Hardware Settings for Failback Recovery



- h. Select the order the new VMs should be powered up in. For example, if some VM(s) are dependent on another VM to up and running before they should power up, then you want to have the dependent VM(s) power up last. The default Recover Order is all "1" and all VM(s) will be powered up within seconds of each other.
- i. Specify recovery job type and time to execute based on one of the following three ways:
  - Recovery Option set to **Recovery Now**, Recovery Through set to **FX Job**, and plan name entered into Recovery Plan Name.
  - Recovery Option set to **Recovery Now**, Recovery Through set to **WMI Based**, and no plan name required.
  - Recovery Option set to Recovery Later, and plan name entered into Recovery Plan Name. If you select Recovery Later, an FX job is created, which can manually started at any time.

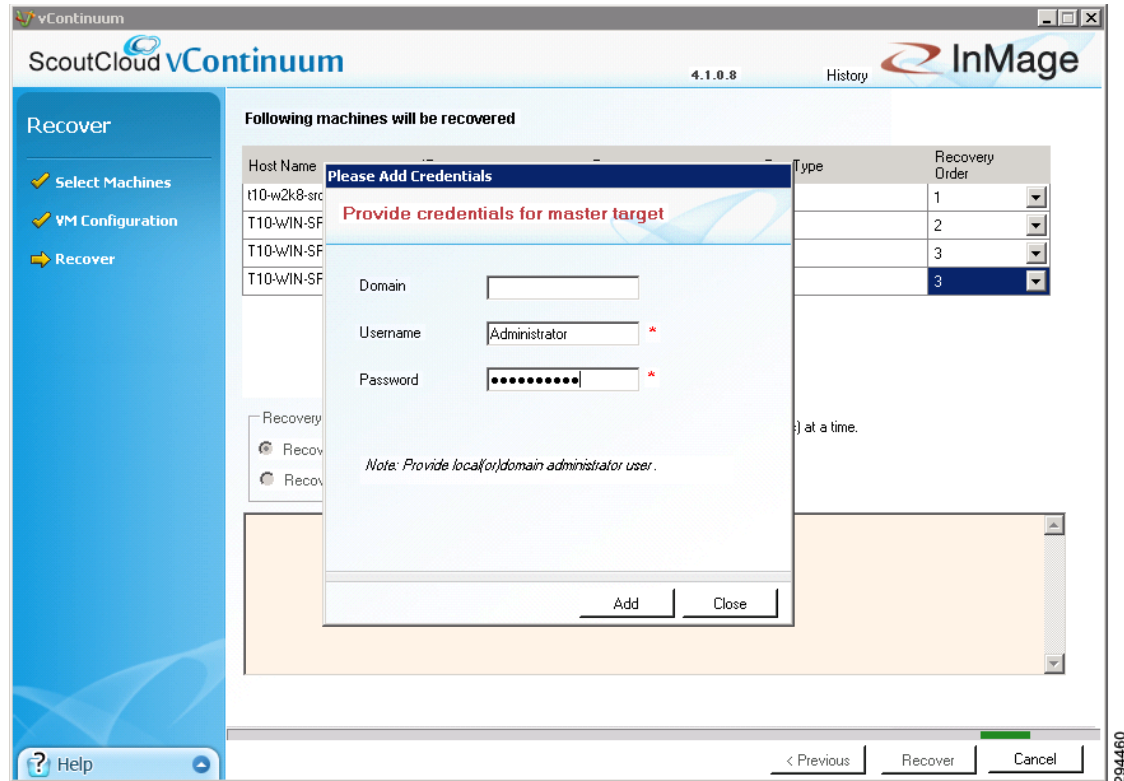
In Figure 4-58, the recover operation is configured for an immediate WMI-based recovery. The powering up of the servers will be staggered with T10-W2K8-SRC-1 first, T10-WINSRC-2 second, and both T10-WIN-SRC3 and T10-WIN-SRC-5 third.

Figure 4-58 Finalizing Recover Operation



294459

Figure 4-59 Primary Master Target Credentials Required



**Note** The protection plans for recovery and failback recovery no longer exist in vContinuum or the CX UI. A new recovery protection plan needs to be created to re-establish disaster recovery protection.

## Virtual-to-Physical (V2P) Failback Protection

Failback operation replicates any new changes made on the secondary VMs back to the primary physical server after failover. Failback can be done only on physical servers that are failed over to the secondary server. The V2P failback protection supports both Linux and Windows operating systems. A V2P failback of a physical server running CentOS is documented in this section.

Virtual to physical (V2P) failback protection has the following steps:

- Step 1** Prepare the Physical Server.
  - Connect to the server console and boot from the InMage LiveCD. In this step, edit the physical server network, DNS, and firewall configurations.
- Step 2** Prepare the USB disk.

- The minimum size of the USB disk should 8GB. The physical server will provide the MT function for the failback process. Because of this, the USB disk will be configured with two Linux partitions. The InMage Unified Agent is installed on the first partition and the MT retention drive is installed on the second partition.
- Step 3** Create the Failback protection plan.
- The V2P failback is a similar to the protection plan, but in the reverse direction. The primary VM running in the secondary site is failed back to the physical server in the enterprise. The V2P failback plan uses the secondary site process server and the physical server as the MT.
- Step 4** Run the Recovery plan.
- Run the failback recovery after the failback plan achieves differential sync status.
- 

This section includes the following topics:

- [Prepare the Physical Server, page 4-60](#)
- [Prepare the USB Flash Drive, page 4-63](#)
- [Create the V2P Failback Plan, page 4-68](#)
- [Recover the Physical Server, page 4-73](#)

## Prepare the Physical Server

Connect to the physical server console and boot from the InMage LiveCD. In this step, edit the physical server network, DNS, and firewall configurations.

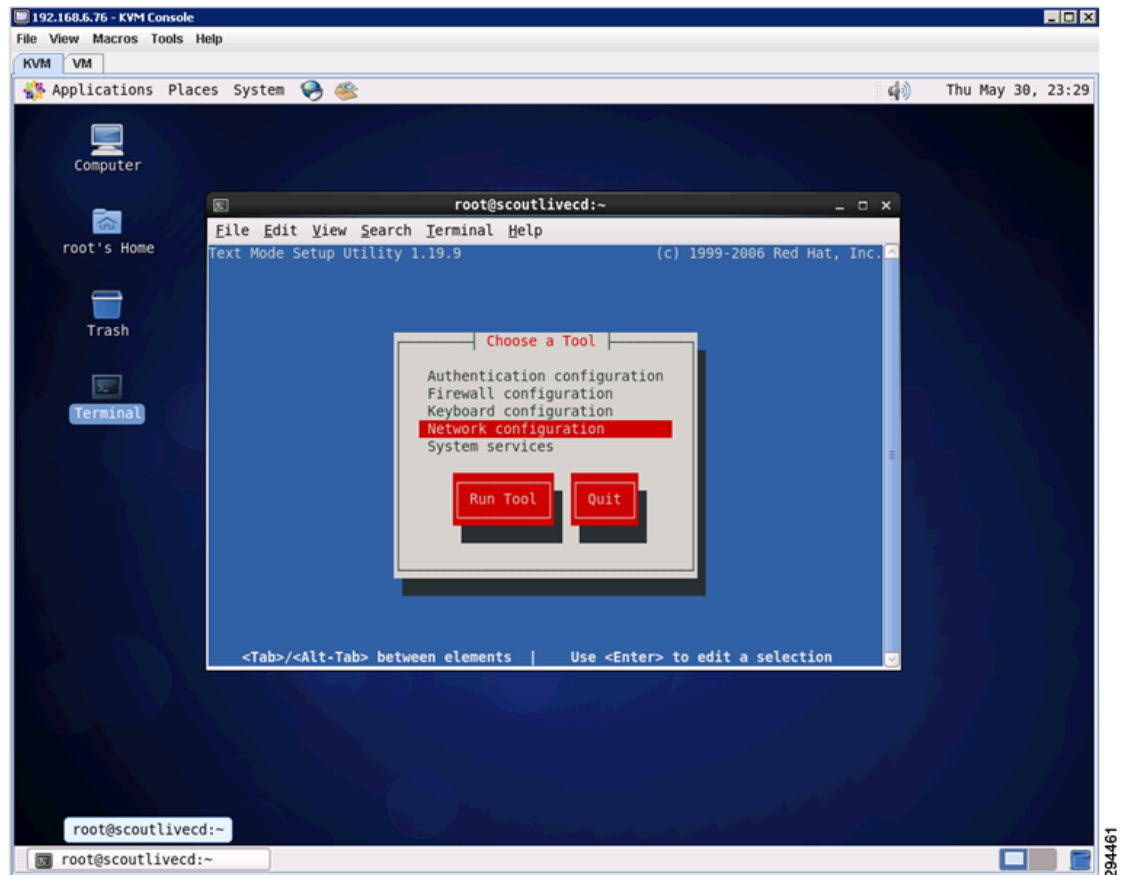
### Summary of Steps

1. Boot the InMage LiveCD and run Setup.
2. Edit the physical server network configuration.
3. Edit the physical server DNS configuration.
4. Edit the physical server firewall configuration.
5. Configure the physical server /etc/hosts file.
6. Configure the physical server /etc/sysconfig/network file.
7. Enable the physical server interface.
8. Restart the network service.

### Detailed Steps


---

- Step 1** Boot the InMage LiveCD and run setup.

**Figure 4-60** Run Setup to Configure Network and Firewall Configurations

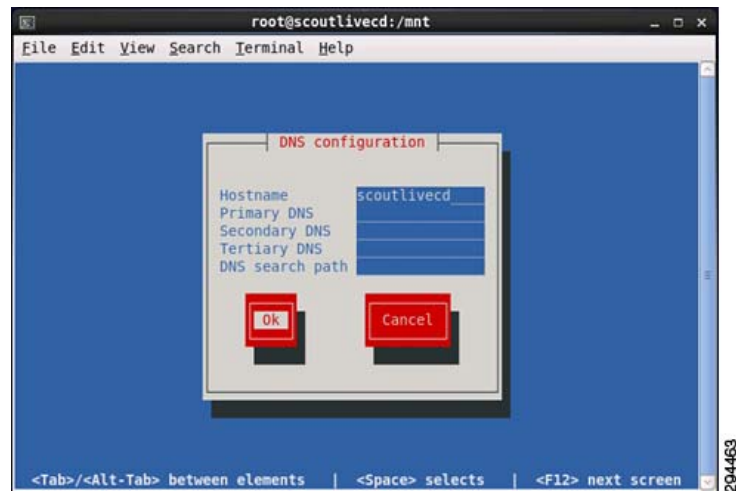
**Step 2** Edit the physical server network configuration:

- Configure the IP address.
- Configure the subnet mask.
- Configure the gateway.

**Figure 4-61** Edit the Network Configuration

**Step 3** Edit the physical server DNS configuration:

- Configure the hostname.

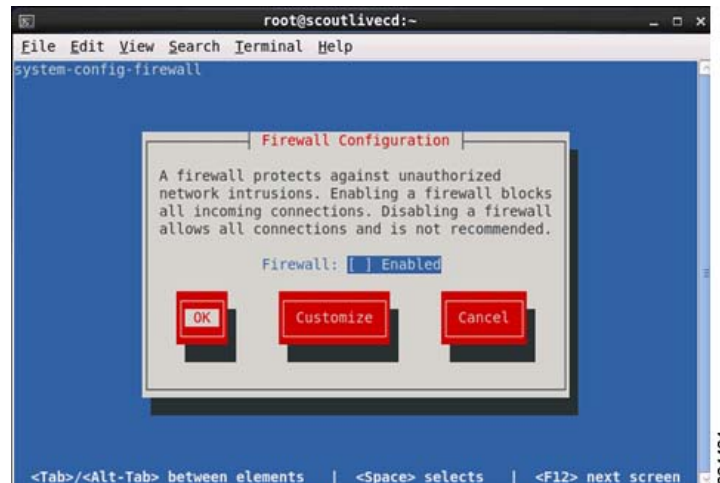
**Figure 4-62** Edit the DNS Configuration

**Step 4** Edit the physical server firewall configuration:

- Disable the firewall service.



Figure 4-63 Disable the Firewall



- Step 5** Configure the physical server `/etc/hosts` file. Add the hostname entry `scoutlivecd` to the `/etc/hosts` file.
- ```
6.126.99.216scoutlivecdscoutlivecd
127.0.0.1localhost.localdomain localhostscoutlivecd::1localhost.localdomain localhost6
localhostscoutlivecd
```
- Step 6** Configure the physical server `/etc/sysconfig/network` file. Configure the hostname `scoutlivecd` to the `/etc/sysconfig/network` file
- ```
NETWORKING=yes
NETWORKING_IPV6=no
HOSTNAME=scoutlivecd
```
- Step 7** Enable the physical server interface. Set the `ONBOOT` option to `Yes`.
- ```
DEVICE=eth0 BOOTPROTO=none
ONBOOT=yes
NETMASK=255.255.255.192
TYPE=Ethernet
HWADDR=e8:b7:48:4e:10:2e
IPADDR=6.126.99.216
GATEWAY=6.126.99.199
IPV6INIT=no USERCTL=no
```
- Step 8** Restart the network service.
- ```
service network restart
```

## Prepare the USB Flash Drive

The minimum size of the USB disk should be 8GB. The physical server will provide the MT function for the failback process. Because of this, the USB disk will be configured with two Linux partitions. The InMage Unified Agent is installed on the first partition and the MT retention folder is installed on the second partition.

### Summary of Steps

1. Insert the USB disk into the physical server and copy the InMage Unified Agent onto the USB disk.
2. List the server's available disk devices.
3. Create the disk partition for the InMage Unified Agent on the USB disk.

4. Create the disk partition for the InMage Retention Folder on the USB disk.
5. Verify the newly created USB disk partitions.
6. Create the directories to be mount points for the two USB disk partitions.
7. Format the newly created USB disk partitions.
8. Mount the newly created USB disk partitions.
9. Install the InMage Unified Agent software.
10. Run the InMage Unified Agent install script.
11. Agree to the license terms and conditions
12. Specify where to install the Unified Agent installation software.
13. Configure the primary role of Scout Agent.
14. Configure the host agent Global settings.
15. Configure the host agent Agent settings.

### Detailed Steps

**Step 1** Insert the USB disk into the physical server. Insert the USB disk and copy the InMage Unified Agent onto the USB disk.

**Step 2** List the server's available disk devices.

```
[root@scoutlived /]# fdisk -l
Disk /dev/sda: 299.0 GB, 298999349248 bytes 255 heads, 63 sectors/track, 36351
cylinders Units = cylinders of 16065 * 512 = 8225280 bytes Sector size
(logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512
bytes Disk identifier: 0x00023254
Device BootStartEndBlocksId System
/dev/sda1 *164512000 83 Linux
Partition 1 does not end on cylinder boundary.
/dev/sda264363522914785288e Linux LVM
Disk /dev/sdb: 200.0 GB, 200049647616 bytes 255 heads, 63 sectors/track, 24321
cylinders Units = cylinders of 16065 * 512 = 8225280 bytes Sector size
(logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512
bytes Disk identifier: 0x393e8c41
Device BootStartEndBlocksId System
```



#### Note

- The `fdisk -l` option is used to list the partition table(s).
- `/dev/sda` is the original disk partition, and `/dev/sdb` is the USB disk.

**Step 3** Create the disk partition for the InMage Unified Agent on the USB disk.

```
[root@scoutlived /]# fdisk /dev/sdb
WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
switch off the mode (command 'c') and change display units to sectors (command 'u').
Command (m for help): n Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-24321, default 1): Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-24321, default 24321): +3126M
Command (m for help): w
The partition table has been altered!
```

Calling `ioctl()` to re-read partition table.

**Note**

- The `fdisk <disk>` is used to change the partition table(s).
- Create a 3 GB primary partition on the USB disk that will contain the InMage Agent software.

**Step 4** Create the disk partition for the InMage Retention Folder on the USB disk.

```
[root@scoutlivecd /]# fdisk /dev/sdb
WARNING: DOS-compatible mode is deprecated. It's strongly recommended to switch off
the mode (command 'c') and change display units to sectors (command 'u').
Command (m for help): n
Command action
e extended
p primary partition (1-4)
p
Partition number (1-4): 2
First cylinder (401-24321, default 401):
Using default value 401
Last cylinder, +cylinders or +size{K,M,G} (401-24321, default 24321):
Using default value 24321
Command (m for help): w
The partition table has been altered!
Calling ioctl() to re-read partition table. Syncing disks.
```

**Note**

- The `fdisk <disk>` is used to change the partition table(s).
- Create the second primary partition on the remaining free space of the USB disk. This partition will contain the MT retention folder.

**Step 5** Verify the newly created USB disk partitions.

```
[root@scoutlivecd /]# fdisk -l /dev/sdb
Disk /dev/sdb: 200.0 GB, 200049647616 bytes 255 heads, 63 sectors/track, 24321
cylinders Units = cylinders of 16065 * 512 = 8225280 bytes Sector size
(logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512
bytes Disk identifier: 0x393e8c41
Device BootStartEndBlocksId System
/dev/sdb114003212968+ 83 Linux
/dev/sdb240124321 192145432+ 83 Linux
```

**Note**

- The `fdisk -l` option is used to list the partition table(s).
- `/dev/sdb1` is used for the InMage Unified Agent software
- `/dev/sdb2` is used for the InMage MT retention drive

**Step 6** Create the directories to be mount points for the two USB disk partitions.

```
[root@scoutlivecd /]# mkdir /mnt/InMageAgent
[root@scoutlivecd /]# mkdir /mnt/InMageCDP
```

**Note**

- `/mnt/InMageAgent` is used for the InMage Unified Agent Software
- `/mnt/InMageCDP` is used for the MT retention folder

**Step 7** Format the newly created USB disk partitions.

```
[root@scoutlivecd /]# mkfs.ext3 /dev/sdb1
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
201200 inodes, 803242 blocks
40162 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=826277888
25 block groups
32768 blocks per group, 32768 fragments per group
8048 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912
Writing inode tables: done
Creating journal (16384 blocks): done
Writing superblocks and filesystem accounting information: done
  This filesystem will be automatically checked every 38 mounts or 180 days,
  whichever comes first. Use tune2fs -c or -i to override.
[root@scoutlivecd /]# mkfs.ext3 /dev/sdb2
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
12009472 inodes, 48036358 blocks
2401817 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
1466 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000, 23887872
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
  This filesystem will be automatically checked every 38 mounts or 180 days,
  whichever comes first. Use tune2fs -c or -i to override.
```

**Note**

- The `mkfs.ext3` command is used to format the partitions
- `/dev/sdb1` contains the InMage Unified Agent software
- `/dev/sdb2` contains the InMage MT retention folder

**Step 8** Mount the newly created USB disk partitions.

```
[root@scoutlivecd /]# mount /dev/sdb1 /mnt/InMageAgent
[root@scoutlivecd /]# mount /dev/sdb2 /mnt/InMageCDP
```

**Note**

- The `mount` command is used to mount the partitions
- `/dev/sdb1` contains the InMage Unified Agent software
- `/dev/sdb2` contains the InMage MT retention folder

**Step 9** Install the InMage Unified Agent software.

- The tar command is used to extract the Unified Agent binary.
- Extract the files the newly created USB disk partition mounted to the /mnt/InMage/Agent directory.

```
[root@scoutlivecd InMageAgent]# tar xvzf
InMage_UA_7.1.0.0_RHEL6-64_GA_14May2013_release.tar.gz
InMageFx-7.1.0.0-1.x86_64.rpm
InMageVx-7.1.0.0-1.x86_64.rpm
.fx_build_manifest
install
uninstall.sh
install_fx
install_vx
EULA.txt
conf_file
.vx_version
.fx_version
OS_details.sh
[root@scoutlivecd InMageAgent]#
```

**Step 10** Run the InMage Unified Agent install script. Select option 3 to install both File and Volume replication agents.

```
[root@scoutlivecd InMageAgent]# ./install

You can install the following :
File Replication Agent
Volume Replication Agent
Both
Please make your choice (1 or 2 or 3) here . Default [3]: 3
```

**Step 11** Agree to the license terms and conditions. Enter Y to agree to the license terms and conditions.

```
Please press (Y/y) if you agree to the license terms and conditions: y
```

**Step 12** Specify where to install the Unified Agent installation software. The agent software will be installed on the newly created USB disk partition.

```
Where do you want to install the InMage UA agent (default /usr/local/InMage) : /mnt/
InMageAgent
```

**Step 13** Configure the primary role of Scout Agent. Select option 2 to set the primary role of this agent to MT.

```
What is the Primary Role of this Agent ?
```

```
Scout Agent
```

```
Select 'Scout Agent' for installation on servers that need to be protected, or for
servers that act as targets in a failover/failback situation.
```

```
Master Target
```

```
Select 'Master Target' for installation on a VMWare VM that acts as the protection
target for other protected physical or virtual servers.
```

```
Please make your choice ? (1/2) [Default: 1] 2
```

**Step 14** Configure the host agent Global settings. Configure the IP address and port number of the CX server, the default port number is 80.

```
+ +
| Host Config Interface|
|Pick the command you wish to run. |
| Press ? for help. |
```

```

+ +
| Global Agent NAT Logging Quit|
+ +
CX Server settings
+ +
| Enter IP Address|
|IP: 8.24.71.101|
+ +
+ +
|Enter Port number|
|Port: 80 |
+ +

```

**Step 15** Configure the host agent Agent settings. Set the application cache directory to /mnt/InMageAgent.

```

+ +
| Host Config Interface|
|Pick the command you wish to run. |
| Press ? for help.|
+ +
| Global Agent NAT Logging Quit|
+ +
VX Agent
+ +
| Application Cache Directory:|
|/mnt/InMageAgent|
+ +
+
|Note: Changing cache directory requires the following steps.
| Not following these steps can result into data loss.
|
Stop svagents
Wait for svagent and child processes to stop completely
Create the new cache directory
Move contents from old cache directory to new cache directory
change cache directory using hostconfigcli/hostconfigui
start svagents
+

```

## Create the V2P Failback Plan

### Summary of Steps

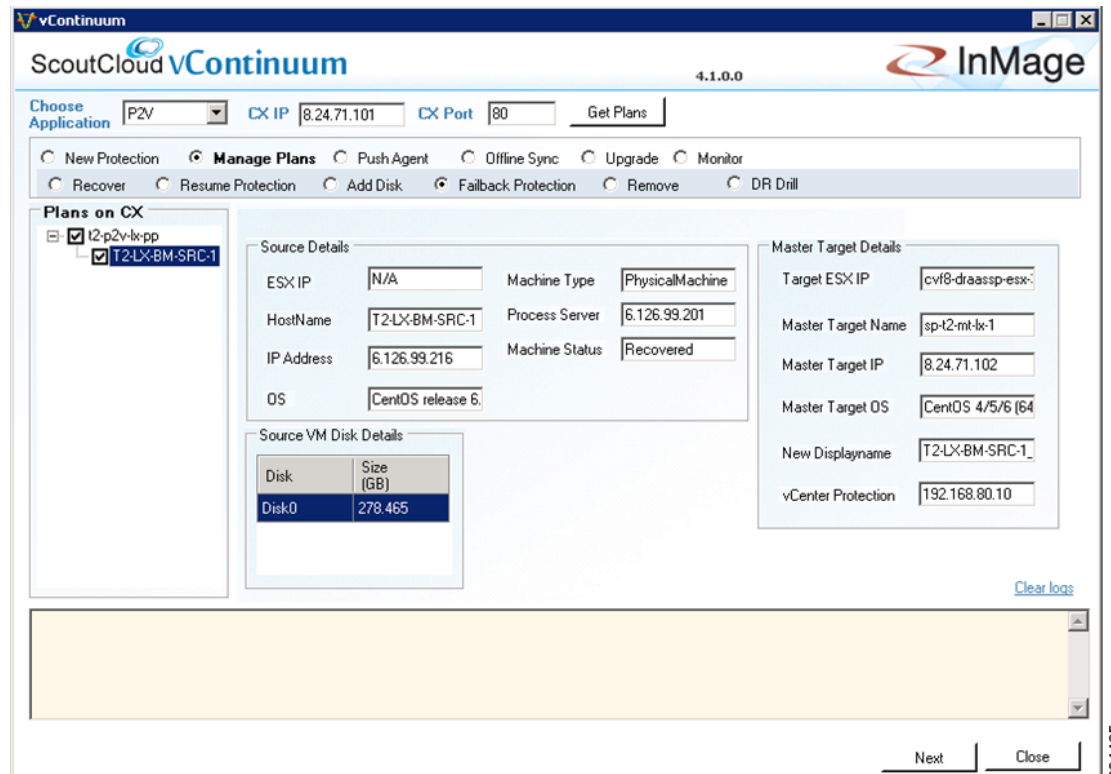
1. Create failback protection plan.
2. Select secondary VM(s) disk to include in the failback protection plan.
3. Select the physical server.
4. Configure V2P failback replication options and finalize the failback protection plan.
5. Monitor failback protection plan.

### Detailed Steps

---

**Step 1** Create the failback protection plan.

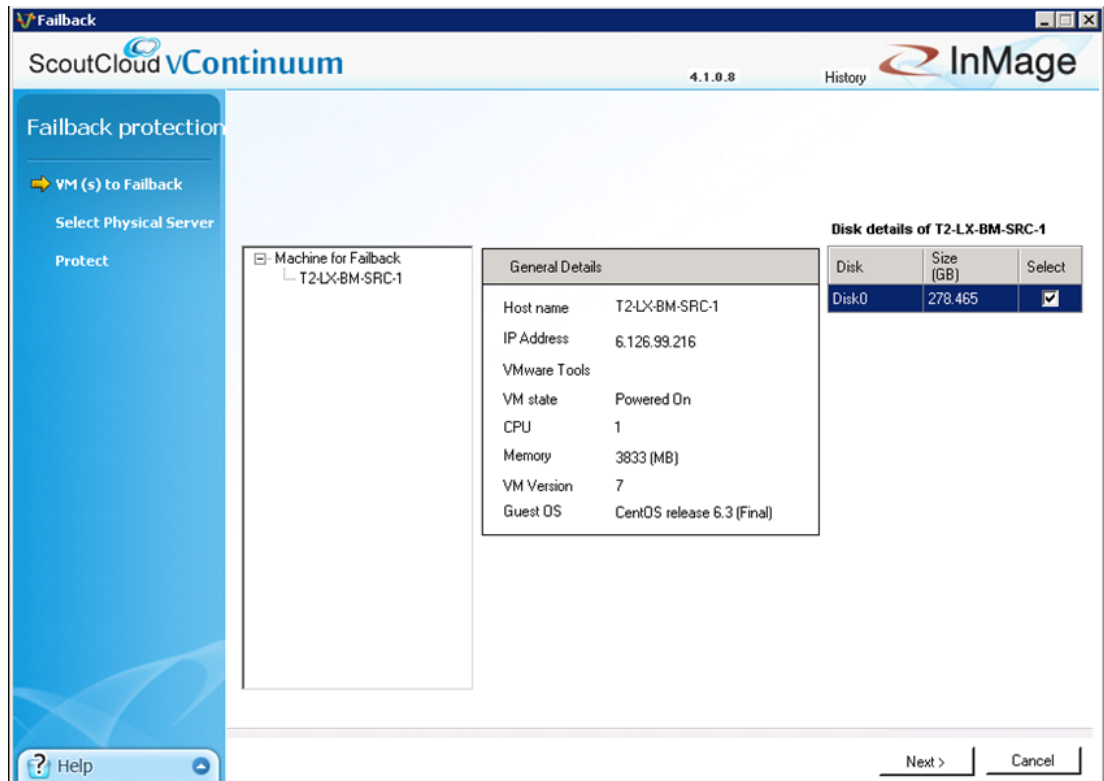
Figure 4-64 Create the V2P Failback Protection Plan in vContinuum



- On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut **Start>Program>InMage System>VContinuum>vContinuum**.
- Select **P2V** from the **Choose Application** drop-down list to view V2P protection plans.
- Enter the CX server IP address and port number (default is 80), then click **Get Plans**.
- Select the **Manage Plans** radio button and click **Failback Protection**.
- Select the secondary VM(s) for failback recovery.
- Click **Next** to continue.

**Step 2** Select secondary VM(s) disk to include in the failback protection plan.

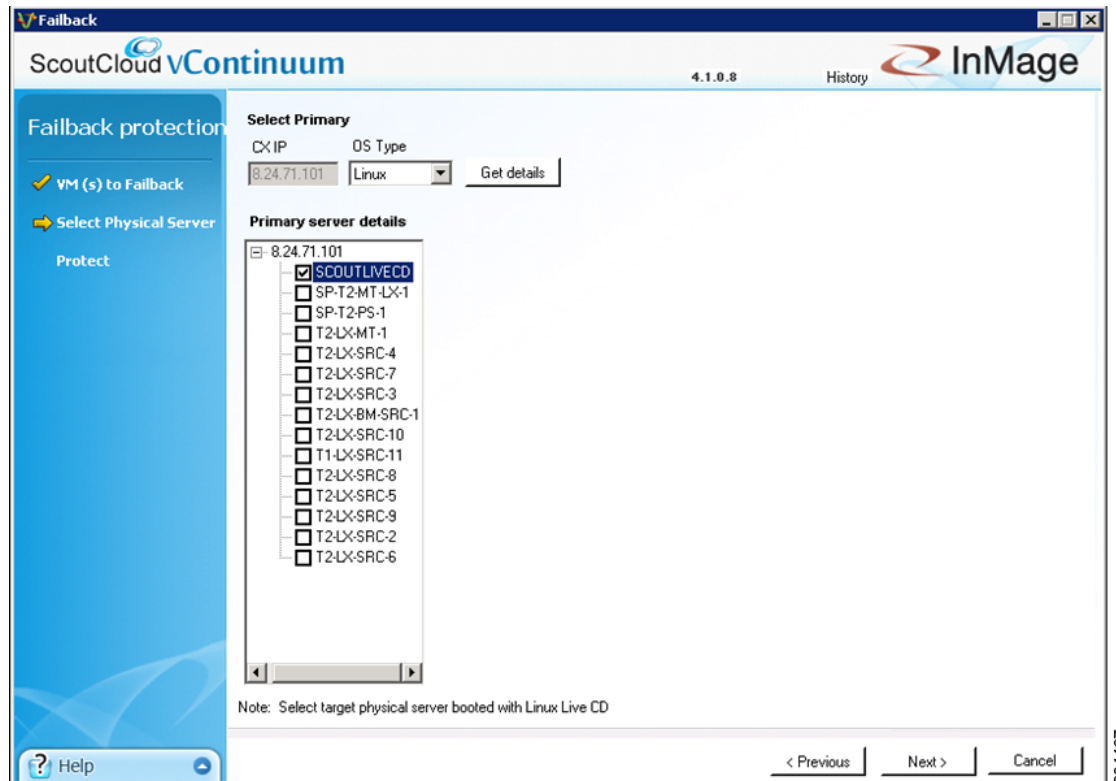
Figure 4-65 Select the VM and VM Disk to Failback



**Step 3** Select the physical server.



Figure 4-66 Select the Target Physical Server



- Set the OS Type to Linux.
- Click **Get details** to list the primary physical servers available for failback.
- Select the physical server booted with the Linux LiveCD.
- Click **Next** to continue.

**Step 4** Configure V2P failback replication options and finalize the failback protection plan.

Figure 4-67 Select Physical Server Disk and Retention Policy

Source VM Name: T2-LX-BM-SRC-1  
Source VM IP: 6.126.99.216  
Target Physical Server Name: SCOUTLIVECD  
Target Physical Server IP: 6.126.99.216

Source Disk Name	Size (in KBs)	Select Target Disk (in KBs)
/dev/sda	291991552	/dev/sda(291991552)

Retention Policy

Retention Drive: /mnt/InMageCDP

Retention in Days: 1

Consistency interval (mins): 30

Process Server: 8.24.71.101

Enter plan name: t2-v2p-lx-fb

Readiness Check

Source Server Name	Check	Result	Corrective Action
T2-LX-BM-SRC-1	All checks	✓	
SCOUTLIVECD	All checks	✓	

Note: Please don't select target USB disk where Unified agent is installed as target

< Previous Protect Cancel

- In the Process server IP field, select the process server located in the secondary site (for example, service provider).
- In the Retention Drive field, select the mount point /mnt/InMageAgent that is associated with the retention drive on the master target.
- In the Retention (in days) field, enter the maximum number of days to store retention data.
- In the Consistency Interval field, enter the Consistency Interval in mins. Jobs will run every x minutes generating application consistency recovery points in primary VMs using which you can recover at the time of recovery. This value determines the RPO in case of consistency point-based recovery.
- In the Select Target Disk field, select the physical server's target disk(s).
- Click **Next** to continue.

**Step 5** Monitor the failback protection plan, wait for differential sync status.

Figure 4-68 Monitor the Failback Status

ScoutCloud CX v7.1.0.0.GA.3055

Monitor > Application Protection > Plan Details  
Plan Details of "t2-v2p-lx-fb"

RPO Health: 1  
Protection Health: 1  
Data Consistency Health: 1

**Disks/Volumes/LUNs Replication**  
1-1 of 1 Records

Server	VX Agent Pair	Health	Health Issue	RPO	Resync progress	Status	Resync Required	Resync Data in Transit (MB)		Differential Data in Transit (MB)			View
								Step1	Step2	On Primary Server	On CX-PS	On Secondary Server	
T2-LX-8M-SRC-1->SCOUTLIVECD	/dev/sda -> /dev/sda	■	N/A	0.15 min	N/A	Differential Sync	NO	0	0	0	0.02	0.39	Summary

**Files/Folders Replication**  
1-2 of 2 Records

Server	FX Agent Pair	Health	Status	Exit Code	Application	Job Description	Scheduled Type	Group ID	Job ID	Job Instance	View Details
SCOUTLIVECD -> SCOUTLIVECD	/mnt/InMageAgent/vx/falover_data/t2-v2p-lx-fb_SCOURLIVECD_262982 -> /mnt/InMageAgent/vx/falover_data/t2-v2p-lx-fb_SCOURLIVECD_262982	■	Stopping...	N/A	t2-v2p-lx-fb	Master Target - ...	Once Now	224	345	5815	Summary
T2-LX-8M-SRC-1 -> T2-LX-8M-SRC-1	/usr/local/InMage/Fx/falover_data -> /usr/local/InMage/Fx/falover_data	■	Not started...	N/A	t2-v2p-lx-fb-Consistency1555	T2-LX-8M-SRC-1-...	Run Every	223	344	5814	Summary

294469



**Note** The failback protection is complete once all volume replications reach differential sync status.

## Recover the Physical Server

Note The following steps to configure and execute a failback recovery are based on the online Scout Help, which can be accessed from the main vContinuum page.

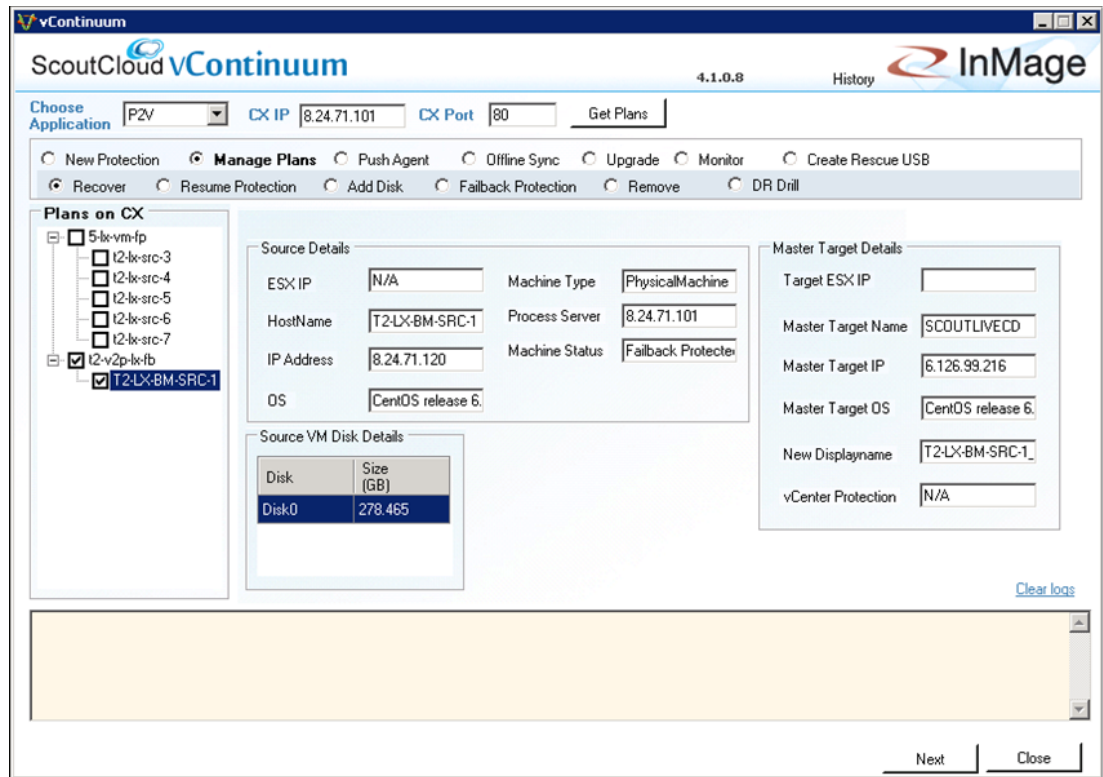
### Summary of Steps

1. Create an failback recovery plan using the vContinuum wizard.
2. Select the VM(s) to recover.
3. Configure the network settings for the physical server.
4. Configure the NIC properties.
5. Run the recovery plan.

### Detailed Steps

- Step 1** Create an failback recovery plan using the vContinuum wizard.

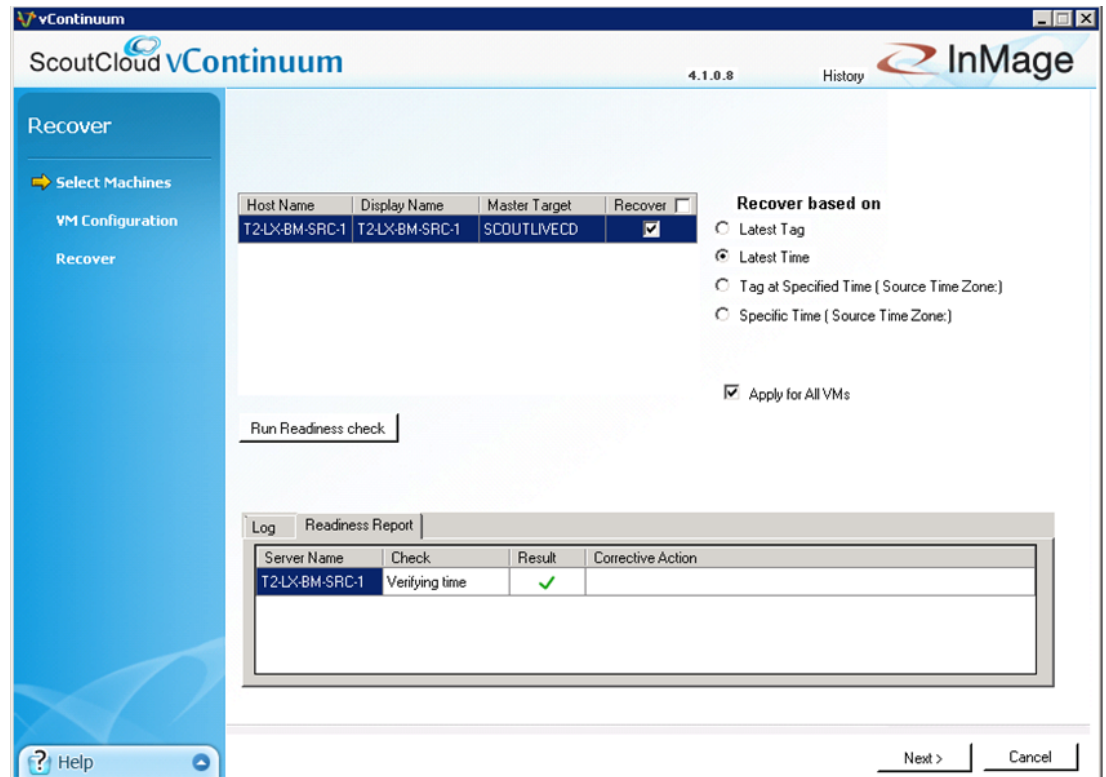
Figure 4-69 Create the V2P Failback Recovery Plan in vContinuum



- Select **P2V** from the **Choose Application** drop-down list to view V2P failback plans.
- Enter the CX server IP address and port number (default is 80) and then click **Get Plans**.
- Select the **Manage Plans** radio button and then click **Recover**.
- Select the secondary VM(s) for failback recovery.
- Click **Next** to continue.

**Step 2** Select the VM(s) to recover and specify whether to base the recovery on time or tag.

Figure 4-70 Select Machines to Recover



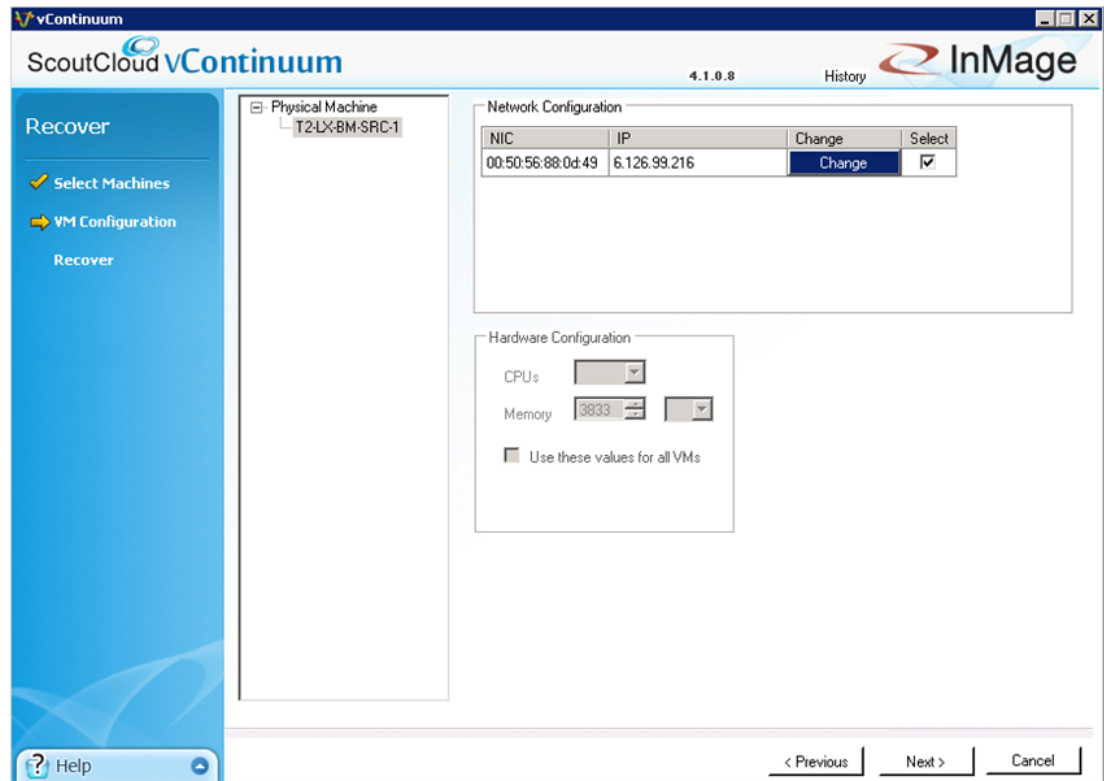
- **Latest Tag:** Selecting this option recovers to a latest tag that is common across all protected disks of a server. For example, consider a server that has three disks: /dev/sda,/dev/sdb,/dev/sdc. By using this option, you would recover to a latest common tag that is available across all disks at same point.
- **Latest Time:** Selecting this option recovers to a latest common point in time among all disks protected from a server. Only the latest common point in time where disks are in green state (data mode) is considered. For example, consider a server that has three disks: /dev/sda,/dev/sdb,/dev/sdc. By using this option, you would recover to a latest common point in time where all three disks are in green state (data mode).
- **Tag at Specified Time:** Selecting this option recovers to a common tag prior to the specified time. For example, consider a server that has three disks: /dev/sda,/dev/sdb,/dev/sdc. By using this option, you would recover to a latest tag available prior to that time.
- **Specified time:** Selecting this option recovers to common point in time among all the disks protected by you.



**Note** Click Apply for all VMs to perform the recovery for all VMs at the specified snapshot type.

**Step 3** Configure the network settings for the physical server.

Figure 4-71 Configure the Network and Hardware Configurations



- For each VM being recovered, click **Change** to configure the network configuration.
- Click **Next** after all selected machines are configured.



**Note** The network configuration can be assigned statically or dynamically via DHCP. When the physical server is recovered and powered on it will contain these new values provided. If no new network settings are provided, original settings will be retained.

**Step 4** Configure the NIC properties.

Figure 4-72 Configure the NIC Properties

**NIC Properties**

Use DHCP

Use following IP Address

IP - 1

IP Address: 6.126.99.216 \*

Subnet Mask: 255.255.255.192 \*

GateWay: 6.126.99.199

DNS:

Network label: [ ] \*\*

Add Cancel

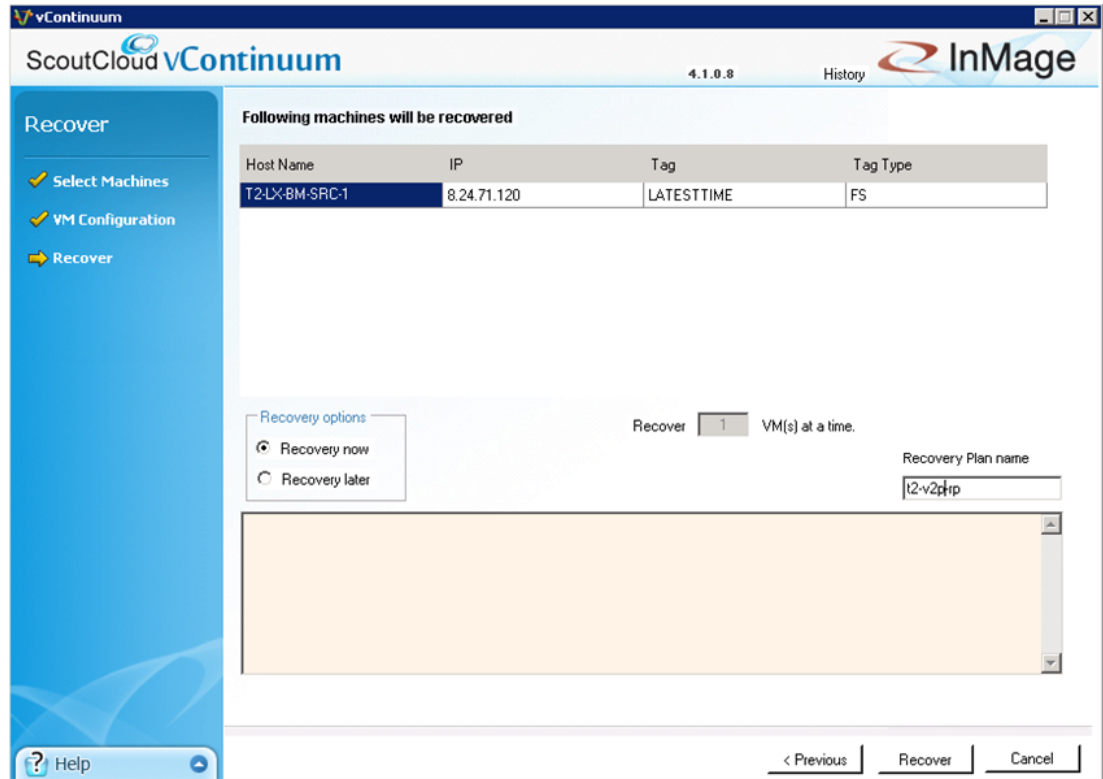
Note:  
 \* Required fields  
 \*\* Choose any network name from drop down menu in . Network Label ,if you type nonexisting network name, then NIC will be in disconnected state on Secondary esx server.

204473

- Select the NIC properties mode, static or dynamic (DHCP).
- If NIC properties are set to static, configure the IP address, Subnet Mask, Gateway, and DNS server.
- Configure the interface VLAN.
- Click **Add** to continue.

**Step 5** Run the recovery plan.

Figure 4-73 Recover the Physical Server



- Specify the recovery job type and time to execute based on one of the following two ways:
- Recovery Option set to Recovery Now, and plan name entered into Recovery Plan Name.
- Recovery Option set to Recovery Later, and plan name entered into Recovery Plan Name, a Recovery Later can be manually started at any time.

## Resume Protection Workflows

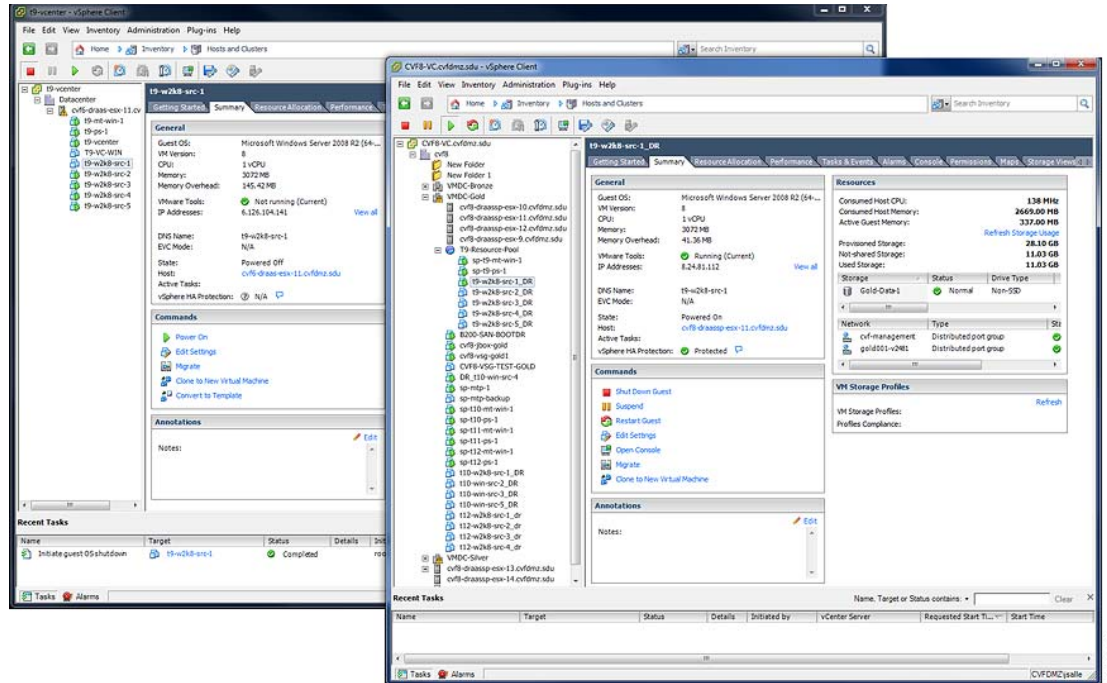
When a primary server is recovered at a secondary vCenter, two ways to switch back to the primary server exist:

- The first is a resume operation and it will discard any changes that occurred while the secondary server was been online and bring the primary server back online and into the existing protection plan. The volume replication will continue based on the last recovery.
- The second is a failback operation and similar to the recover operation, the primary server will be created from a snapshot of the secondary server based on a consistency point or point in time.

In “[Recovery Workflows](#)” section on page 4-37, a single primary VM in a protection plan was recovered to the secondary vCenter. This can be seen in [Figure 4-74](#) and will be the starting state for the resume protection discussion in this section.



Figure 4-74 Single VM Recovered to Secondary vCenter



**Note**

The following steps to recover a primary server are based on the online Scout Help, which can be accessed from the main vContinuum page.

**Summary of Steps**

1. Start vContinuum wizard application.
2. Select VM(s) for which to resume protection.
3. Finalize resume and execute.
4. Monitor resume.

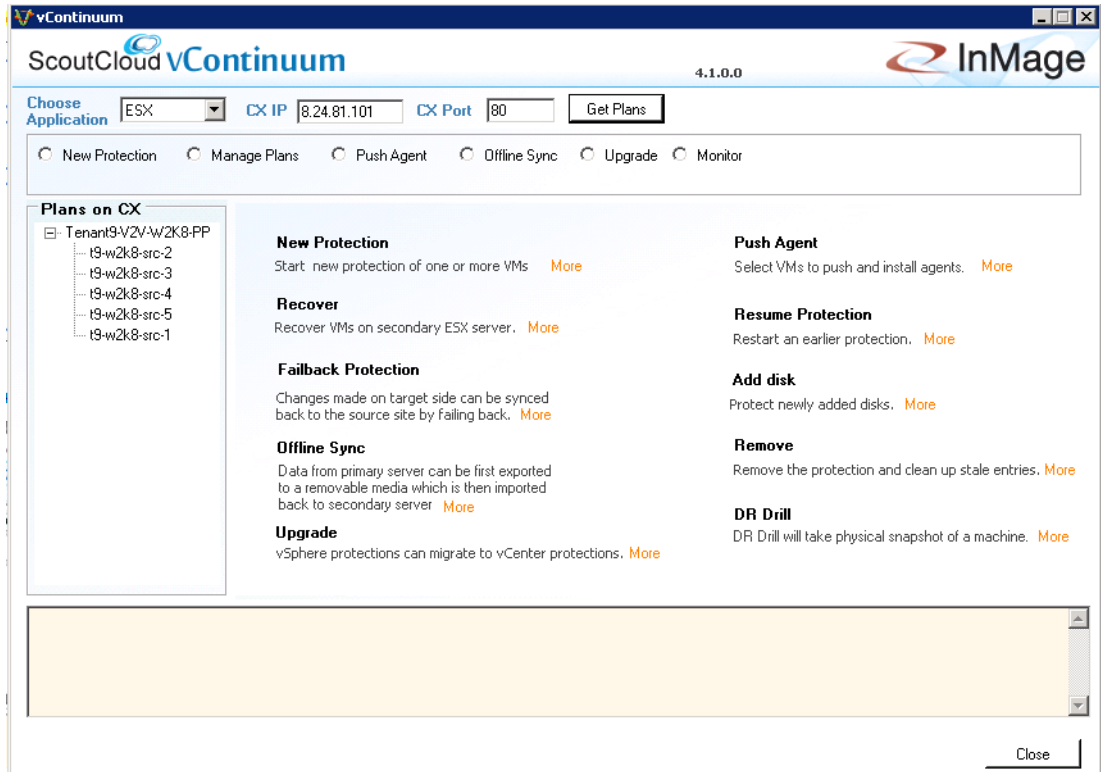
**Detailed Steps**

**Step 1**

On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut **Start>Program>InMage System>VContinuum>vContinuum**.

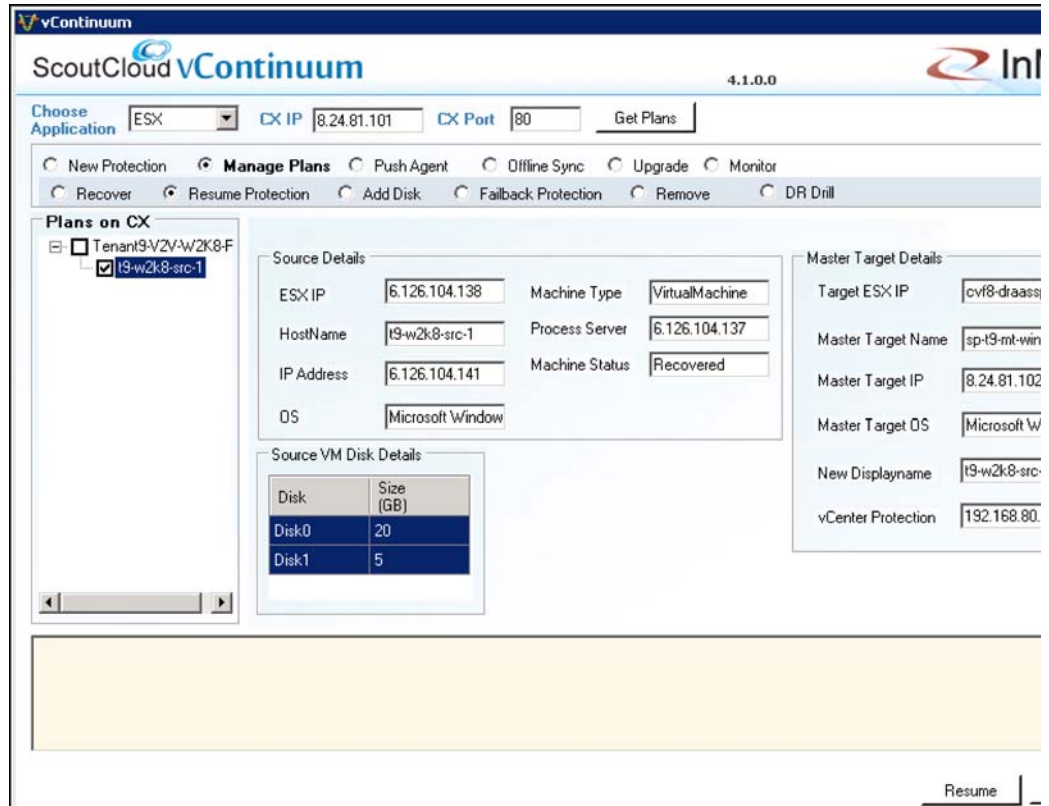
- a. Select ESX from the Choose Application drop-down list to view V2V protection plans or P2V for P2V.
- b. Enter the CX server IP address and port number (default is 80), then click **Get Plans**. Looking at Figure 4-75, there are five primary servers in the Tenant9-V2V-W2K8-PP protection plan. Actually, the last server in the list was recently recovered to the secondary vCenter.

Figure 4-75 Existing Plans on vContinuum



- c. Select the Manage Plans radio button and then click Resume.
- d. Select the primary VM t9-w2k8-src-1.

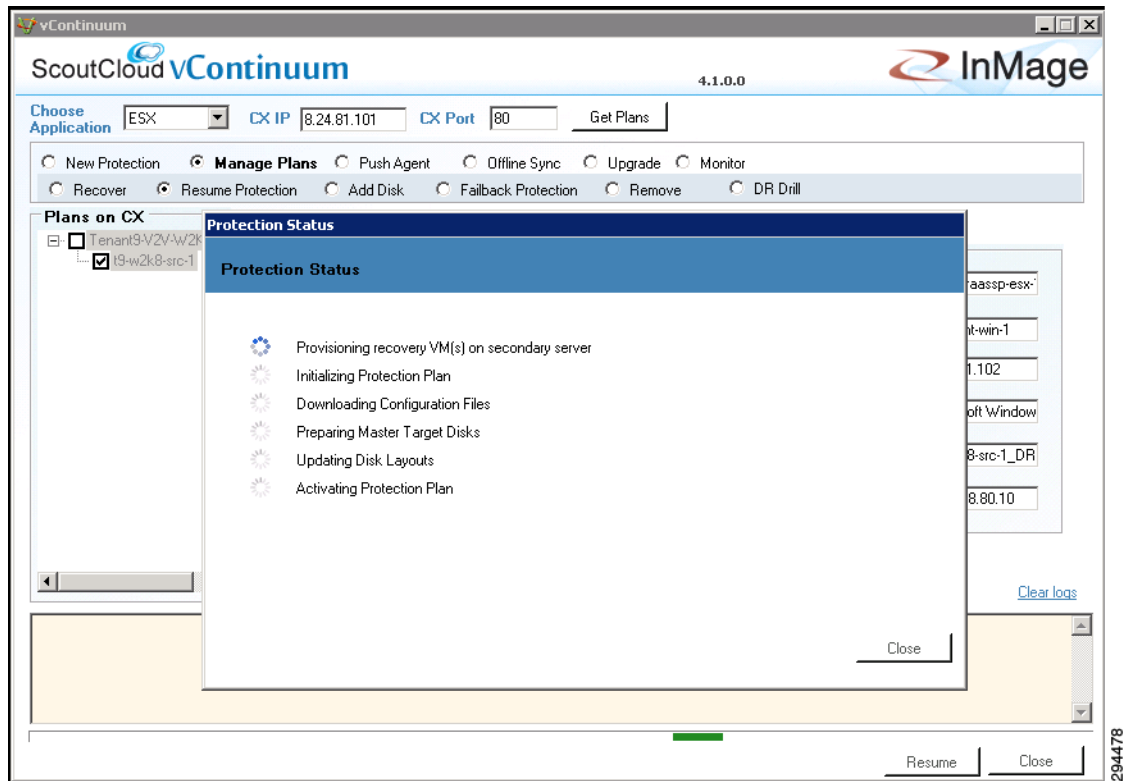
Figure 4-76 Selecting VM for Resume



- e. Click Resume to start the resume operation.

**Step 2** Monitor recovery. After starting the resume, vContinuum goes through a number of steps to execute the recovery, which can be monitored from the status window.

Figure 4-77 Resume Operation Started



294478

Figure 4-78 Secondary VM Shutdown Confirmation

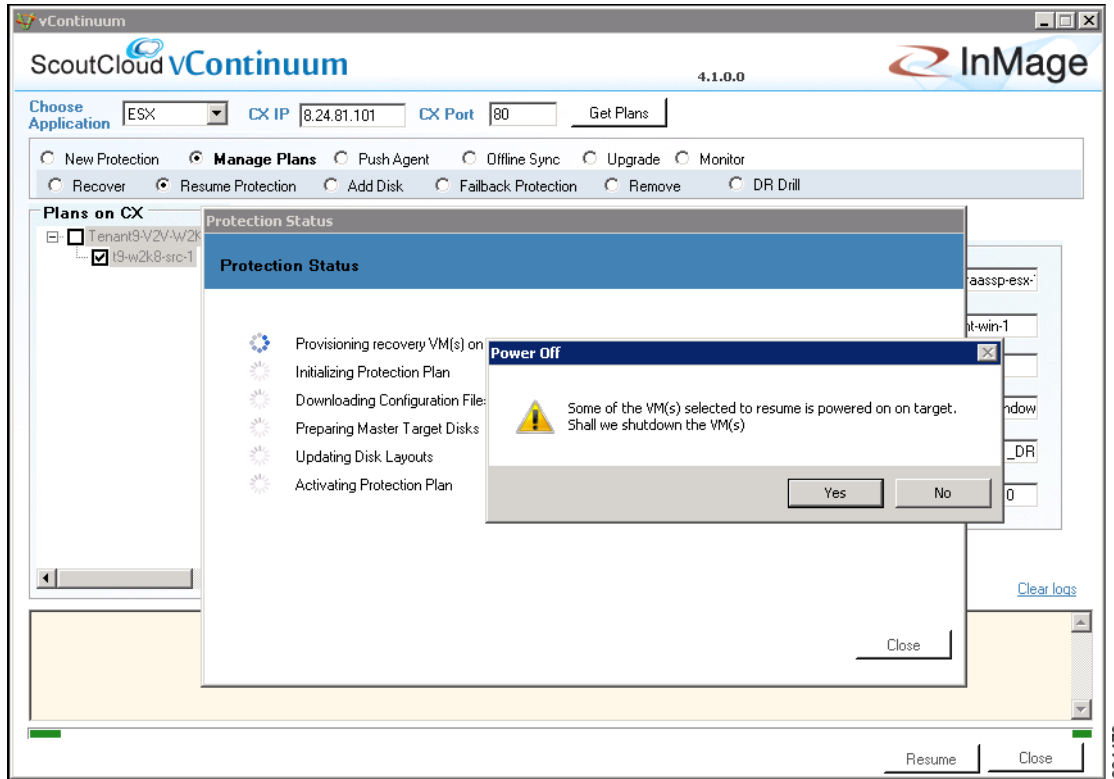


Figure 4-79 Secondary VM Shutdown in Secondary vCenter

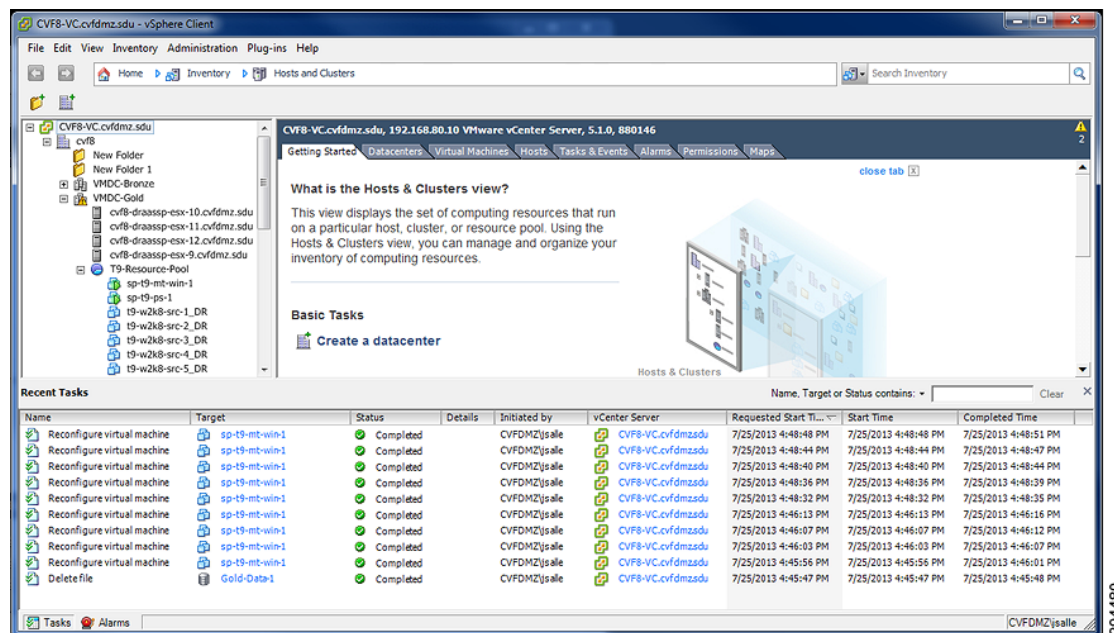


Figure 4-80 Volume Resyncing After Resume

Server	VxAgent Pair	Health	Health Issue	RPO	Resync progress	Status	Resync Required	Resync Data in Transit (MB)		Differential Data in Transit (MB)			View
								Step1	Step2	On Primary Server	On CP-PS	On Secondary Server	
T9-W2XB-SRC-3->SP-T9-MT-WIN-1	C -> C:\ESX\B6688D89-7471-EA43-AC76BA7EDB948C6_C	■	N/A	1.48 min	N/A	Differential Sync	NO	0	0	0	23.12	0	Summary
T9-W2XB-SRC-2->SP-T9-MT-WIN-1	C -> C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C	■	N/A	0.9 min	N/A	Differential Sync	NO	0	0	0	17.5	0	Summary
T9-W2XB-SRC-4->SP-T9-MT-WIN-1	C -> C:\ESX\85F7759A-2EC9-9348-B8C0AD1E9A3F0331_C	■	N/A	1.4 min	N/A	Differential Sync	NO	0	0	0	23.38	0	Summary
T9-W2XB-SRC-1->SP-T9-MT-WIN-1	C -> C:\ESX\E08CD805-888E-944D-B60496095D3914DD_C	■	N/A	2.58 min	0 %	Resyncing (Step 1)	YES	0.12	21.27	0	0	0	Summary
T9-W2XB-SRC-5->SP-T9-MT-WIN-1	C -> C:\ESX\4FC26A25-683F-DC48-86F0C31180D4A5C0_C	■	N/A	1.18 min	N/A	Differential Sync	NO	0	0	0	23.15	0	Summary
T9-W2XB-SRC-5->SP-T9-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\4FC26A25-683F-DC48-86F0C31180D4A5C0_C_SRV	■	N/A	1.15 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T9-W2XB-SRC-3->SP-T9-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\B6688D89-7471-EA43-AC76BA7EDB948C6_C_SRV	■	N/A	0.27 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T9-W2XB-SRC-2->SP-T9-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C_SRV	■	N/A	0.92 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T9-W2XB-SRC-4->SP-T9-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\85F7759A-2EC9-9348-B8C0AD1E9A3F0331_C_SRV	■	N/A	0.27 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T9-W2XB-SRC-1->SP-T9-MT-WIN-1	C:\SRV ( System Reserved ) -> C:\ESX\E08CD805-888E-944D-B60496095D3914DD_C_SRV	■	N/A	2.63 min	N/A	Resyncing (Step 1)	YES	0	0.03	0	0	0	Summary
T9-W2XB-SRC-3->SP-T9-MT-WIN-1	K ( New Volume ) -> C:\ESX\B6688D89-7471-EA43-AC76BA7EDB948C6_K	■	N/A	1.38 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T9-W2XB-SRC-2->SP-T9-MT-WIN-1	K ( New Volume ) -> C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_K	■	N/A	0.95 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary
T9-W2XB-SRC-1->SP-T9-MT-WIN-1	K ( New Volume ) -> C:\ESX\E08CD805-888E-944D-B60496095D3914DD_K	■	N/A	1.58 min	N/A	Resyncing (Step 1)	YES	0	0.02	0	0	0	Summary

294481

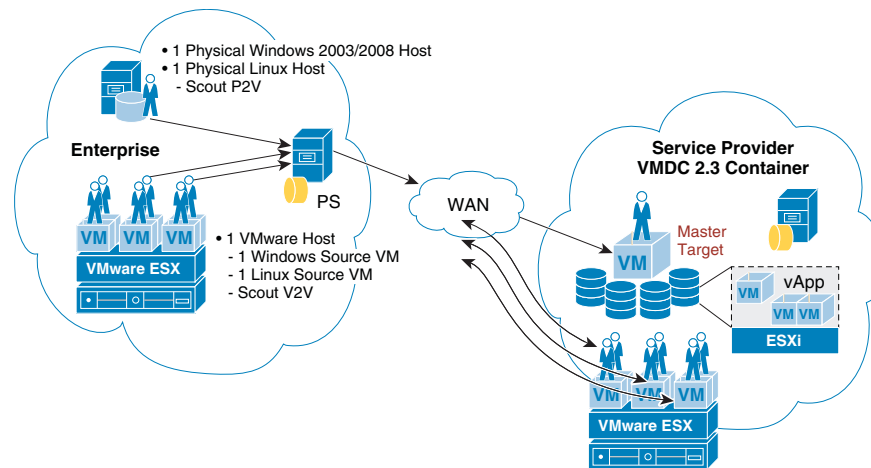
## DR Drill Workflows

The Disaster Recovery (DR) Drill allows the administrator the ability to verify primary VM protection by creating a secondary VM using a physical snapshot from any of the following recovery points:

- Latest application consistent point
- Latest point in time
- Consistency point near (prior to) any given time
- Specific time

DR Drill does not impact protection pairs, so all VMs currently under protection will remain protected while the DR Drill is being performed.

Figure 4-81 DR Drill Overview

**Note**

The following steps to configure and execute a DR Drill are based on the online Scout Help, which can be accessed from the main vContinuum page.

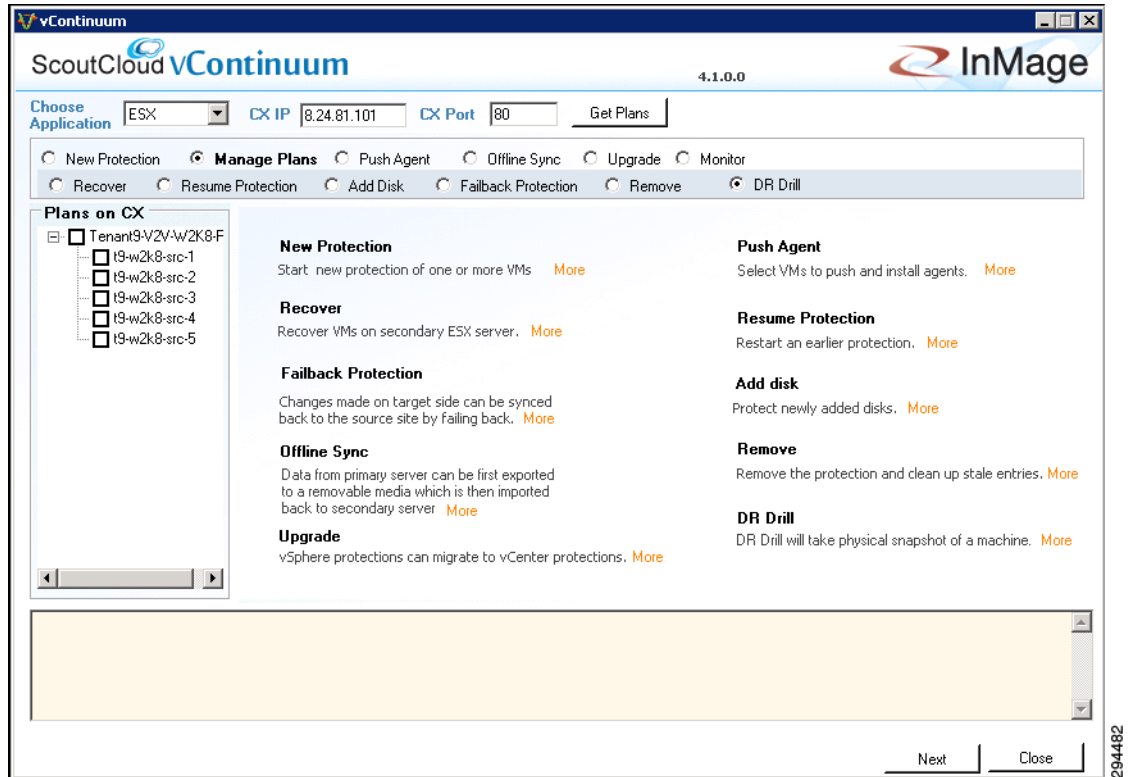
**Summary of Steps**

1. Start vContinuum wizard application.
  - a. Select VM(s) for DR Drill.
  - b. Specify snapshot to use based on time or tag.
  - c. Perform Readiness Check to make sure the VM(s) are ready for DR Drill.
  - d. Configure network, hardware, and display name for new VM(s).
  - e. Select datastore(s) for new VM(s).
2. Monitor DR Drill progress from CX UI or vContinuum wizard application.

**Detailed Steps**

- Step 1** On the Management Console, start the vContinuum wizard application via the desktop icon or Start menu shortcut **Start>Program>InMage System>VContinuum>vContinuum**.
- a. Select **ESX** from the **Choose Application** drop-down list to view V2V protection plans or P2V for P2V.
  - b. Enter the CX server IP address and port number (default is 80) and click **Get Plans**.
  - c. Select the **Manage Plans** radio button and click **DR Drill**.

Figure 4-82 Starting DR Drill on vContinuum



- d. Select the primary VM(s) for the DR Drill and then click Next.

294482



Figure 4-83 Selecting Primary VMs for DR Drill

ScoutCloud vContinuum 4.1.0.0 InMage

Choose Application: ESX CX IP: 8.24.81.101 CX Port: 80 Get Plans

New Protection
  Manage Plans
  Push Agent
  Offline Sync
  Upgrade
  Monitor
  Recover
  Resume Protection
  Add Disk
  Failback Protection
  Remove
  DR Drill

Plans on CX

- [-] Tenant9-V2V-W2K8-F
  - t9-w2k8-src-1
  - t9-w2k8-src-2
  - t9-w2k8-src-3
  - t9-w2k8-src-4
  - t9-w2k8-src-5

Source Details

ESX IP: 6.126.104.138 Machine Type: VirtualMachine  
 HostName: t9-w2k8-src-5 Process Server: 6.126.104.137  
 IP Address: 6.126.104.145 Machine Status: Protected  
 OS: Microsoft Window

Source VM Disk Details

Disk	Size (GB)
Disk0	15

Master Target Details

Target ESX IP: cvf8-draassp-esx:  
 Master Target Name: sp-t9-ml-win-1  
 Master Target IP: 8.24.81.102  
 Master Target OS: Microsoft Window  
 New Displayname: t9-w2k8-src-5\_DR  
 vCenter Protection: 192.168.80.10

Clear logs

Next Close

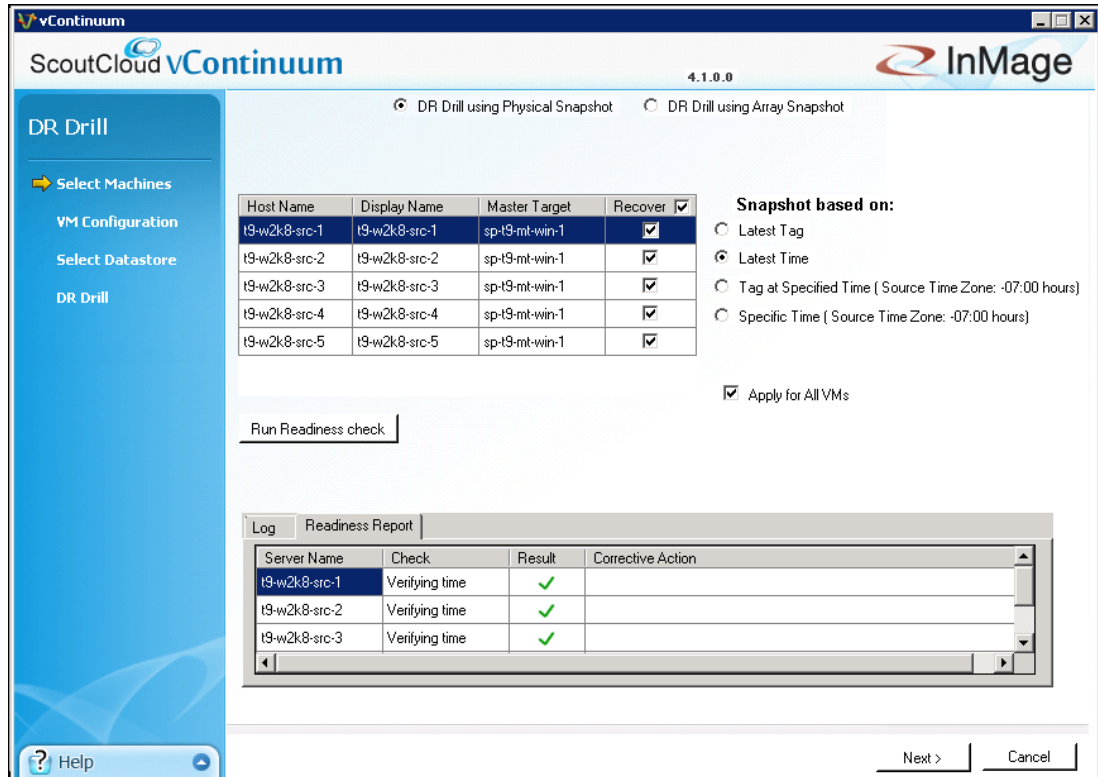
- e. Specify snapshot to use based on time or tag.
- **Latest Tag:** Select this option to recover a VM to a latest tag that is common across all volumes of a VM. For example, if a VM that has three volumes (e.g., C, E, and F), the latest common tag that is available across all volumes at same time point across all volumes is used.
  - **Latest Time:** Select this option to recover a VM to a latest common point time among all volumes of a VM. Only common time points where volumes are in green state (data mode) are considered. For example, if a VM that has three volumes (e.g., C, E, and F), the latest common time where all three volumes are in green state (data mode) is used.
  - **Tag at Specified Time (Source Time Zone):** Select this option to recover a VM to a common tag prior to the specified time. For example, if a VM that has three volumes (e.g., C, E, and F), the latest tag available prior to that time is used. The time provided is converted to GMT and compared against the timestamps in the retention logs. The closest consistency point prior to the time provided will be used to recover the VM.
  - **Specific Time (Source Time Zone):** Select this option to recover a VM to a common point in time among all volumes of a VM. The time provided is converted to GMT and compared against the time stamps of the primary server. All recovery times are based on primary server's time stamps and not the secondary or management console times



**Note** Click Apply for all VMs to perform the DR Drill for all VMs at the specified snapshot type.

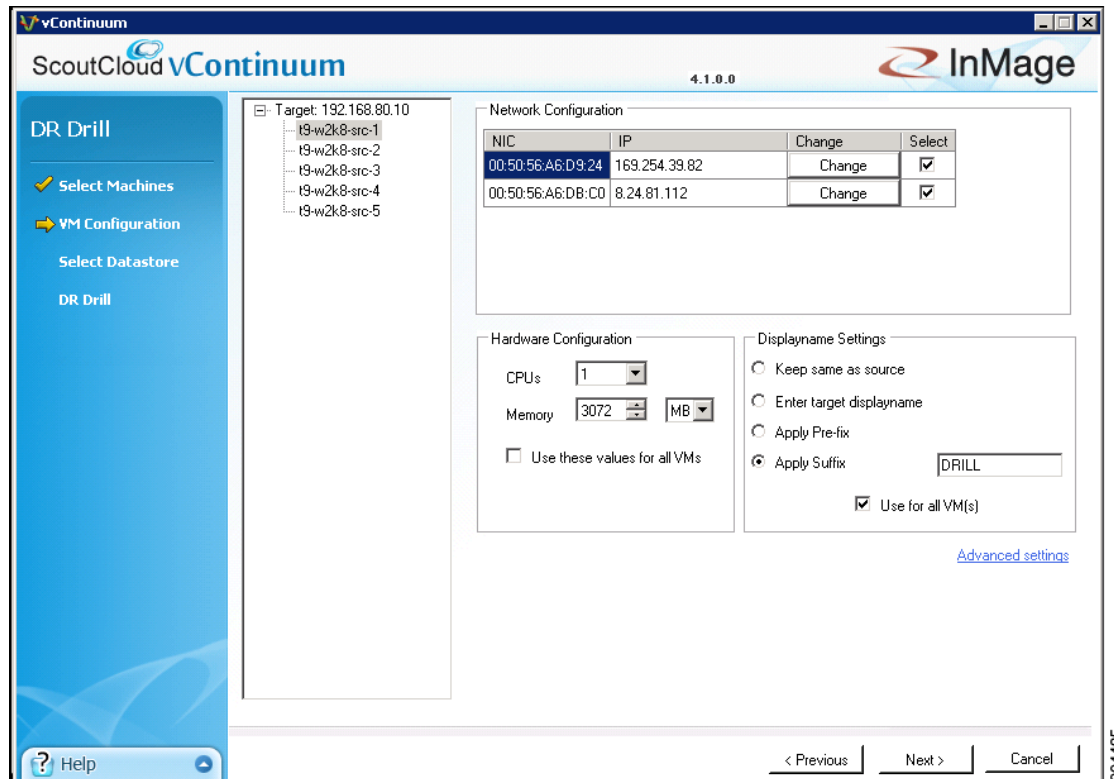
- f. Click Run Readiness Check to make sure the VM(s) are ready for DR Drill. If the check passes, then click Next to advance to the next page.

Figure 4-84 Running Readiness Check for DR Drill



- g. Configure network, hardware, and display name for new VM(s) if these configurations need to be different than what was defined in the original protection plan. In Figure 4-85, we are just appending "DRILL" to the end of each recovery VM display name, no other changes were made.

Figure 4-85 Modifying Network, Hardware, Display Name Settings for DR Drill



- Click Advanced settings to access advanced settings for DR Drill. The default configuration shown in Figure 4-86. The options are greyed out except for the folder name section.
- In the **Folder Name Settings** section, the directory for the VM in the datastore can be configured.

Figure 4-86 Optional Advanced Settings for DR Drill

**Advanced settings**

Scout vContinuum 4.1.0.0 InMage

**Advanced setting for t9-w2k8-src-1**

Sparse Retention

Backups are retained for 90 days(3 months) N/A

Provide continuous backup for latest

Advanced retention

<input type="checkbox"/>	From 0 day onwards. Provide <input type="text" value="1"/> restore point per <input type="text" value=""/> hour for next <input type="text" value=""/> days
<input type="checkbox"/>	From 0 day onwards. Provide <input type="text" value="1"/> restore point per day for next <input type="text" value=""/> weeks
<input type="checkbox"/>	From 0 day onwards. Provide <input type="text" value="1"/> restore point per week for next <input type="text" value=""/> months

Apply for all servers Total days selected

Folder Name Settings

Keep VM in datastore's root directory

Keep VM in datastore's sub directory

Apply for all servers

Compression

No compression

Compression

CX based  Source based

Apply for all servers N/A

Encryption

Secure transport from Source to CX-PS N/A

Secure transport from CX-PS to destination

Apply for all servers

Resource pool

Select resource pool on target

Apply for all servers N/A

Provisioning

Thin provisioning

Thick provisioning N/A

Apply for all servers

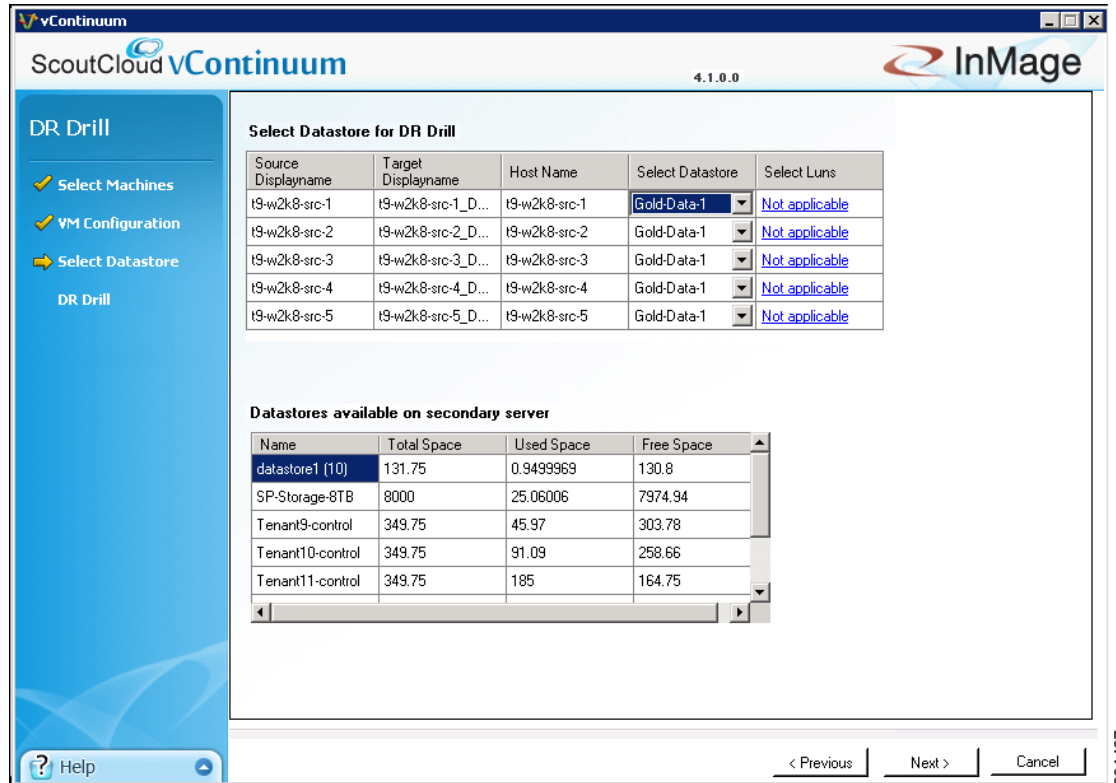
Ok Cancel

294486

– Click **Next** to advance to the final page.

- h. Select the datastore(s) for the new VM(s) that will be created for the DR Drill, then click Next.

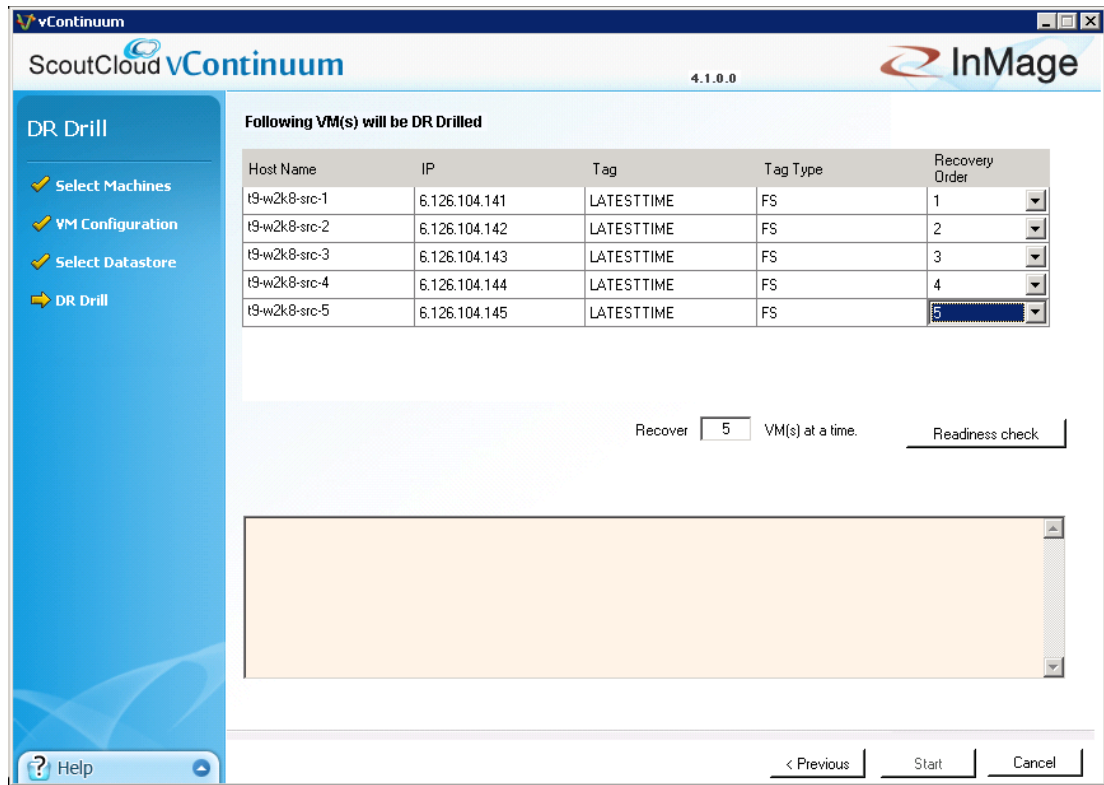
Figure 4-87 Selecting Datastores for DR Drill



- i. Select the order the new VM(s) should be powered up in. For example, if some VM(s) are dependent on another VM to up and running before they should power up, then you want to have the dependent VM(s) power up last. The default Recover Order is all "1" and all VM(s) will be powered up within seconds of each other. In Figure 4-88, the Recovery Order was changed to 1 through 5, so each DR Drill VM will be powered up about 90 seconds after the previous VM.

294487

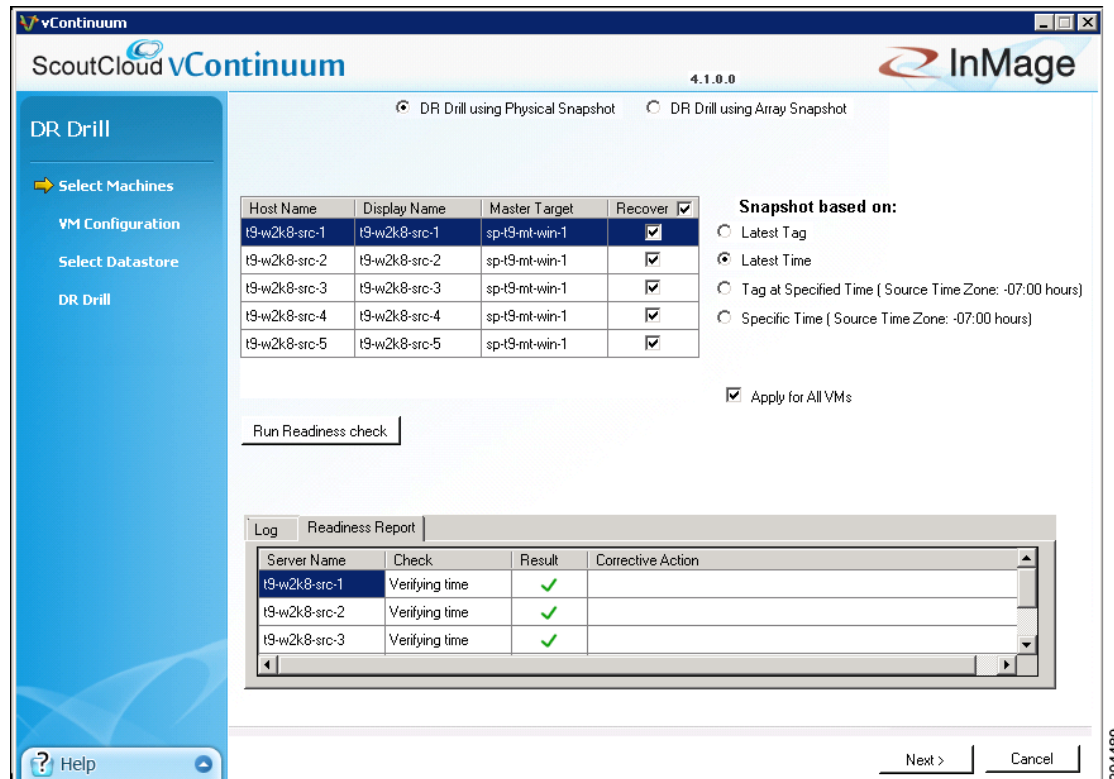
Figure 4-88 .Recovery Order for DR Drill



- j. Click **Readiness check** to verify the Master Target is ready to perform a DR Drill.

294488

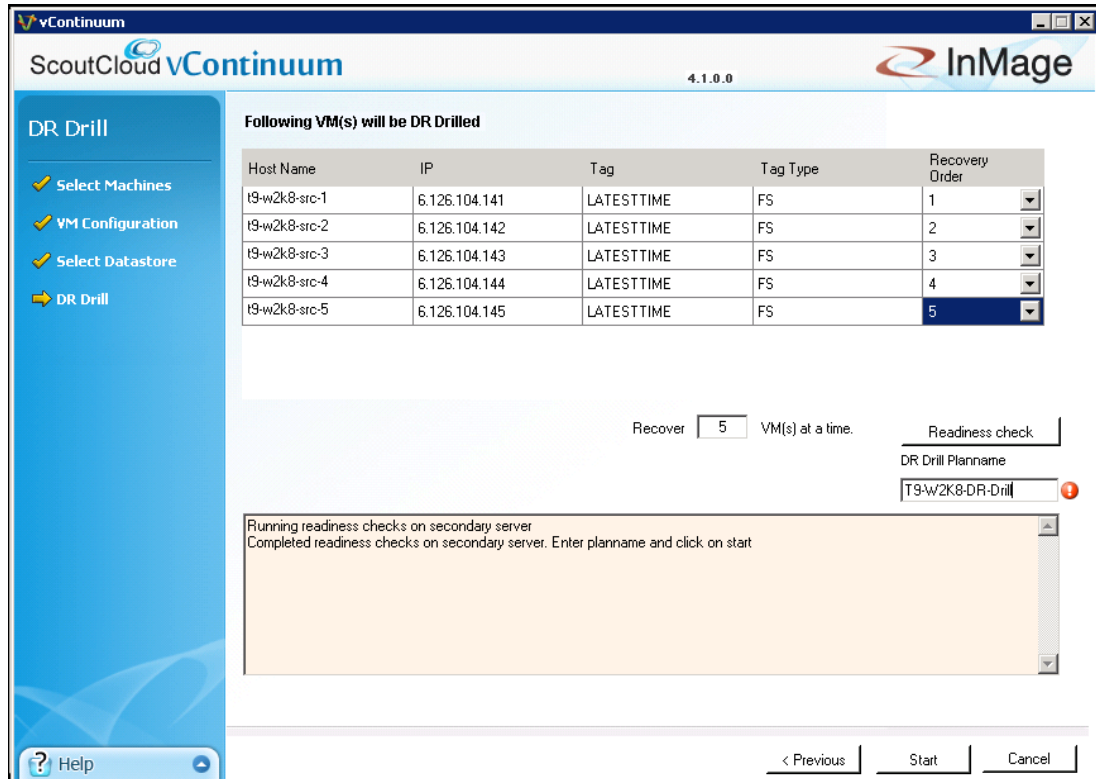
Figure 4-89 Readiness Check for DR Drill



- k. Enter DR Drill plan name and initiate drill.

294489

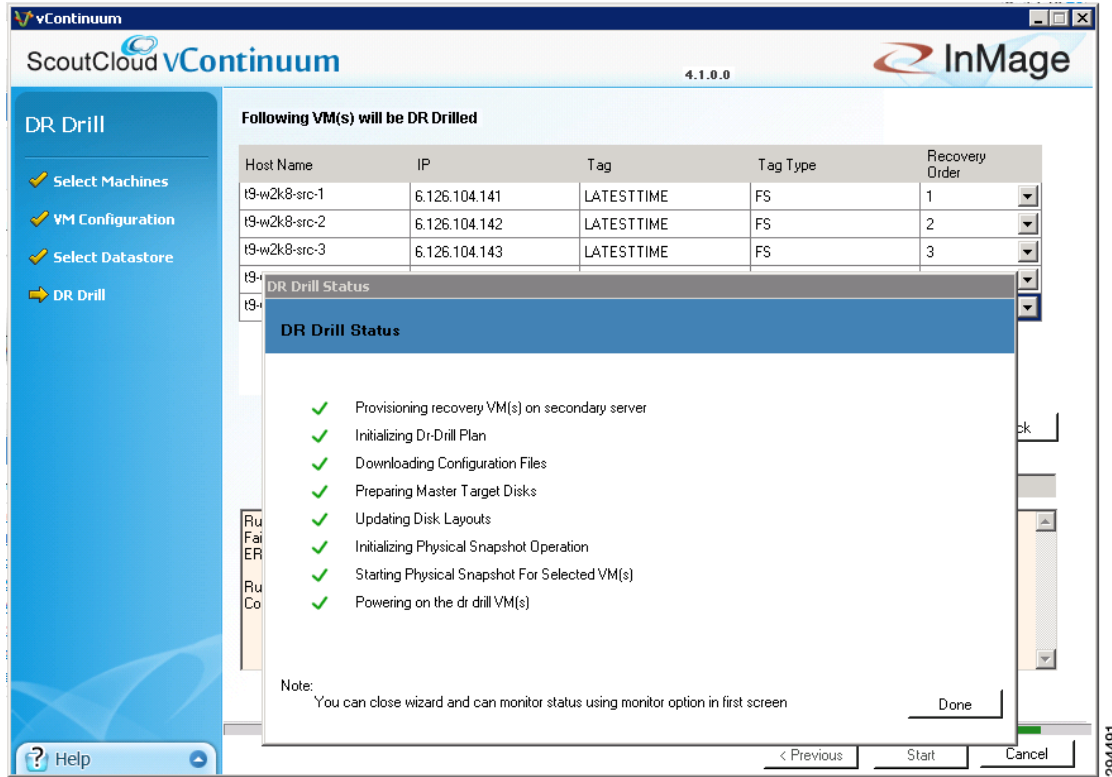
Figure 4-90 Naming and Starting DR Drill



**Step 2** Monitor DR Drill progress. After starting the DR Drill, vContinuum goes through several steps to execute the drill, which can be monitored from the status window.

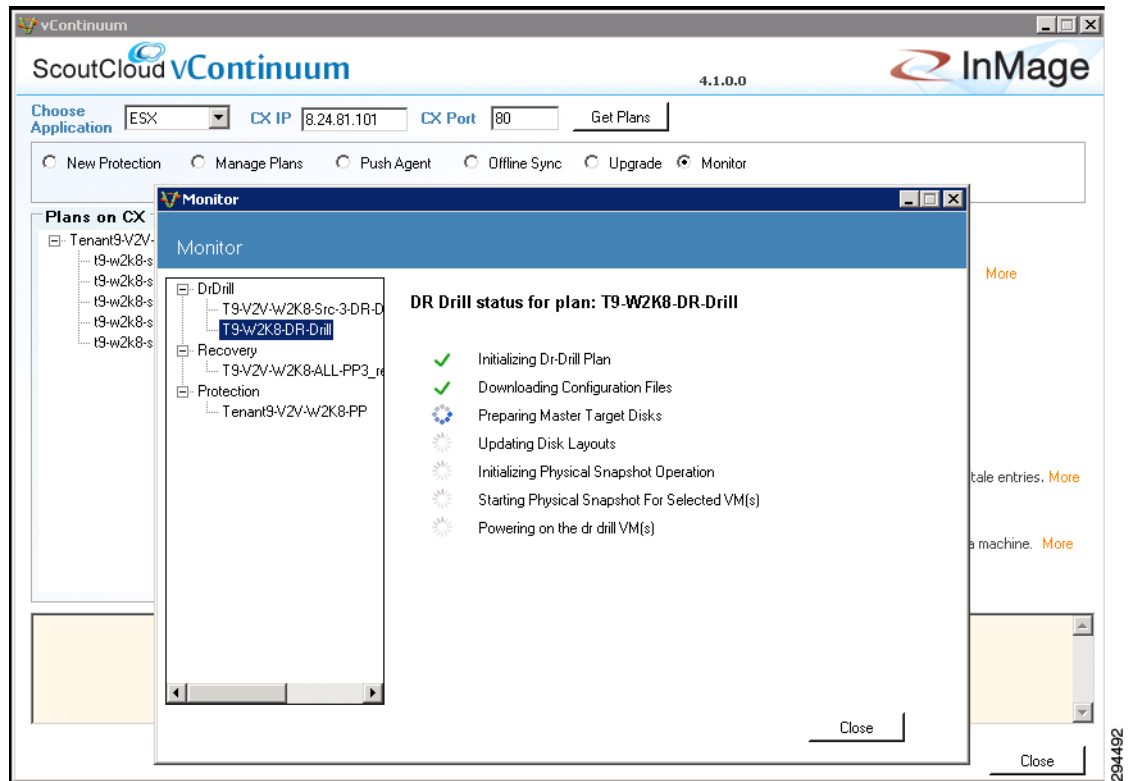


Figure 4-91 DR Drill Starting (vContinuum)



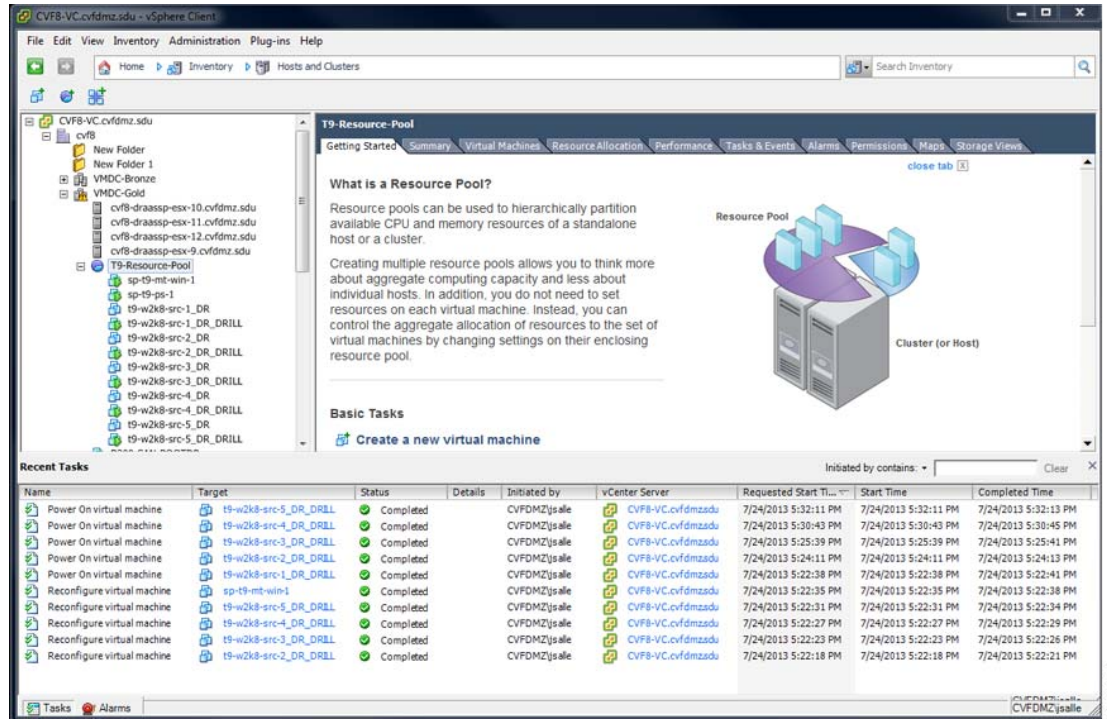
294491

Figure 4-92 DR Drill Monitoring (vContinuum)



294492

Figure 4-93 DR Drill VMs (Service Provider vCenter)



294493

Figure 4-94 DR Drill Status

The screenshot shows the ScoutCloud CX Monitor interface. The main heading is "Plan Details of 'T9-W2K8-DR-Drill2'". Under the "Protections" section, there are two sub-sections: "Disks/Volumes/LUNs Replication" and "Files/Folders Replication".

The "Files/Folders Replication" section contains a table with the following data:

Server	FX Agent Pair	Health	Status	Exit Code	Application	Job Description	Scheduled Type	Group ID	Job ID	Job Instance	View Details
SP-T9-MT-W2N-1 -> SP-T9-MT-W2N-1	C:\Program Files (x86)\InMage Systems\Falover\Data\T9-W2K8-DR-Drill2_20086 -> C:\Program Files (x86)\InMage Systems\Falover\Data\T9-W2K8-DR-Drill2_20086	Green	Completed	0	T9-W2K8-DR-Drill2	Master target - ...	Once Now	105	153	5686	Summary

An "FX Agent Pair Details" pop-up window is open, showing the following information:

More Details	<a href="#">Log: Trending</a>
Start Time	2013-07-24 17:58:26
End Time	2013-07-24 18:26:57
Last Update Time	2013-07-24 18:26:57
Data Compression	N/A
Sync Compression	100.00
Bytes Changed	0

At the bottom of the page, there is a "Start" button and a "Auto refresh this page in every 60 seconds" checkbox which is checked. A vertical ID "294494" is visible on the right side of the screenshot.



## CHAPTER 5

# Monitoring, Best Practices, Caveats, and Troubleshooting

---

This chapter includes the following major topics:

- [InMage Resource Monitoring, page 5-1](#)
- [Implementation Best Practices, page 5-8](#)
- [Caveats, page 5-12](#)
- [Troubleshooting, page 5-15](#)

## InMage Resource Monitoring

InMage components are deployed at both the SP and enterprise. Depending on where the components reside, resource utilization can be monitored with a combination of various tools. As the VMDC CLSA team has done extensive work in the area of service assurance in a SP VMDC cloud, interested readers should refer to the VMDC CLSA for additional details. On the Enterprise side, if an Enterprise already deploys a comprehensive and well-managed infrastructure and systems monitoring program to enable proactive and make better infrastructure planning decisions based on historical trending and detailed usage analysis, we recommend those Enterprises to simply incorporate InMage components into their existing monitoring framework. This section does not intend to repeat previous CLSA recommendations or provide guidance on Enterprise end-to-end monitoring; instead we are focusing on specific metrics that an Enterprise or SP can gather based on our lab implementation.

Metrics such as storage IOPS, IO size, WAN bandwidth utilization, CPU usage on the primary server, CPU usage on the processing server, and RPO are all important metrics to monitor for performance and capacity planning. Most of those statistics are available directly from InMage or can be accessed through the environment that InMage connects to:

- WAN Bandwidth: InMage CX-CS server, vCenter statistics (V2V), Netflow (P2V)
- LAN Bandwidth Per Virtual/Physical Machine: InMage CX-CS server, vCenter statistics (V2V), NetFlow (P2V)
- CPU (PS): InMage CX-CS server, vCenter statistics (V2V), SNMP/SSH
- CPU (Agent): vCenter (V2V), Windows PerfMon, SNMP/SSH
- RPO: InMage
- IO: perfmon, iostat and drvutil utility from InMage

This section will focus on statistics monitoring using InMage. vCenter monitoring is well documented; refer to vSphere Performance Monitoring for monitoring details and operations. NetFlow monitoring can be deployed at key observation points such as the server access layer, fabric path domains, and WAN to gain visibility into LAN / WAN bandwidth and application performance. The Cisco NetFlow Generation Appliance (NGA) 3240 introduces a highly scalable, cost-effective architecture for cross-device flow generation in today's high-performance data centers. Refer to the [3240 Datasheet](#) for details.


- [Bandwidth Monitoring, page 5-2](#)
- [Scout Server Health Monitoring, page 5-3](#)
- [RPO and Health Monitoring, page 5-5](#)
- [I/O Monitoring, page 5-7](#)

## Bandwidth Monitoring

Although in different formats, LAN/WAN bandwidth reporting can be generated directly from the CXCS server or from the RX server for a particular customer/tenant. From the CX-CS UI statistics are grouped based on roles:

Network traffic statistics for all ScoutOS-based devices, the PS and CX server, are available from the CX by accessing Monitor > Network Traffic Report. Statistics are stored in RRD databases maintained by the RRDtool, available in 24 hours, week, month and year intervals. As in most RRD implementations, statistics are more granular for the 24 hour interval and less granular for older statistics. [Figure 5-1](#) is an example of network traffic rate for an PS server.

**Figure 5-1** Sample Network Traffic Rate for a PS Server



Network traffic statistics for Unified Agent-based devices, the source and MT server, are available from the CX by accessing Monitor > Bandwidth Report. Similar to ScoutOS statistics, data are stored in RRD databases maintained by the RRDtool. Daily aggregate statistics are also available. [Figure 5-2](#) is an sample daily bandwidth report.

Figure 5-2 Sample Daily Bandwidth Report

Bandwidth Report		Custom Report														
Bandwidth Report for T11-W2K8-SRC-8 ( 6.126.103.84 )																
Select Host <span>T11-W2K8-SRC-8</span> <span>Last Day Last Week Last Month Last Year</span>																
Month: 2013 Jul																
	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	0 B
Out	82.7 GB	83 GB	83.65 GB	83.57 GB	83.03 GB	82.91 GB	83.16 GB	83.07 GB	81.91 GB	81.52 GB	80.11 GB	80.68 GB	80.59 GB	80.27 GB	83.56 GB	85.7 GB
Max	82.7 GB	83 GB	83.65 GB	83.57 GB	83.03 GB	82.91 GB	83.16 GB	83.07 GB	81.91 GB	81.52 GB	80.11 GB	80.68 GB	80.59 GB	80.27 GB	83.56 GB	85.7 GB
Sum	82.7 GB	83 GB	83.65 GB	83.57 GB	83.03 GB	82.91 GB	83.16 GB	83.07 GB	81.91 GB	81.52 GB	80.11 GB	80.68 GB	80.59 GB	80.27 GB	83.56 GB	85.7 GB
	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Total
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	0 B
Out	86.02 GB	85.72 GB	84.9 GB	83.59 GB	83.91 GB	83.41 GB	75.12 GB	76 GB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	1.93 TB
Max	86.02 GB	85.72 GB	84.9 GB	83.59 GB	83.91 GB	83.41 GB	75.12 GB	76 GB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	1.93 TB
Sum	86.02 GB	85.72 GB	84.9 GB	83.59 GB	83.91 GB	83.41 GB	75.12 GB	76 GB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	1.93 TB

The RX UI reports only aggregate network traffic from Unified Agents based on a time interval. Traffic report for the PS and CX are not available from the RX.

Although the majority of the performance stats are not directly exportable from the GUI, fetching data directly from the RRD database is fairly simple and straightforward. The following is an example of fetching bandwidth data from Jul 24 2013 07:46:18 to 08:03:48:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# pwd
/home/svsystems/052E4A5E-8195-1341-90A49C18364A0532
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# rrdtool fetch bandwidth.rrd
AVERAGE --start 1374651978 --end 1374653028
in out
1374652200: 0.0000000000e+00 8.1258306408e+08 1374652500: 0.0000000000e+00
2.2297785464e+08 1374652800: 0.0000000000e+00 7.3103023488e+08 1374653100:
0.0000000000e+00 4.4805078912e+08
```

Customized graphs can be generated easily as well using the RRDtool:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# rrdtool graph xgao.png --start
1374651978 --end 1374653028 DEF:myxgao=bandwidth.rrd:out:AVERAGE LINE2:myxgao#FF0000
```

## Scout Server Health Monitoring

Scout Server statistics for CPU, memory, disk and free space are directly available from the CX-CS portal. Statistics are displayed at near real time at the CX-CS UI dashboard. Historical performance data are kept in round-robin databases (RRD) and maintained by the RRDtool similar to the bandwidth reports.

**Table 5-1 Scout Server Health Statistics**

Resource	Process Server	CX-CS Server
System Load	Yes	Yes
CPU Load	Yes	Yes
Memory Usage	Yes	Yes
Free Space	Yes	Yes

**Table 5-1 Scout Server Health Statistics (continued)**


Resource	Process Server	CX-CS Server
Disk Activity	Yes	Yes
PS Services	Yes	NA
Web Server	NA	Yes
Database Server	NA	Yes
CS Services	NA	Yes

In a single core system, your System Load/Load Average should always be below 1.0, meaning that when a process asks for CPU time it gets it without having to wait. It is important to keep in mind that PS/CS server are typically deployed with multiple cores. When monitoring such system, the rule of thumb is max load should not exceed number of cores. Use the following command to figure out the number of cores on your system:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# cat /proc/cpuinfo|grep processor
processor : 0
processor : 1
processor : 2
processor : 3
```

CPU load or CPU percent is the amount of time in an interval that the system's processes were found to be active on the CPU. While it can be a good indicator of overall utilization when used in conjunction with system load, it is important to remember that CPU percent is only a snapshot of usage at the time of the measurement, this statistics alone is not a good indication of overall utilization.

Monitoring memory usage on a linux system can be tricky because RAM is used to not only store user application data, but also kernel data as well as cache/mirror data stored on the disk for fast access (Page Cache). Page Cache can consume large amount of memory in general, anytime a file is read, file data goes into memory in forms of page cache. Inode and Buffer cache are kernel data cached in memory. In a typical system it is perfectly normal to see memory usage increases linearly over time as in [Figure 5-3](#):

**Figure 5-3 Memory Usage**

When memory usage reaches some watermark, the kernel starts to reclaim memory from the different cache described above. Swap statistics can be a good indication of overall memory health. `vmstat` can be used to monitor swap usage:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# vmstat 5
procs memoryswapio---- --systemcpu
r b swpd free buff cache si sobibo in cs us sy id wa st
```



```
00 0 3726088 769220 717836 0 0 0 20 0 0 1 0 99 0 0
00 0 3725972 769220 717836 0 0 0 33 697 746 0 0 100 0 0
00 0 3726004 769220 717840 0 0 1 35 891 988 0 0 99 0 0
```

"si" and "so" are abbreviations for "swap in" and "swap out", the size of the swap can indicate the health of the system:

1. small "si" and "so" is normal, available memory is sufficient to deal with new allocation requests.
2. large "si" and small "so" is normal as well.
3. small "si" and large "so" indicates that when additional memory is needed, system is able to reclaim / allocate sufficient memory.
4. large "si" and "so" indicates system is "thrashing" and we are running out on memory Large "so" and "si" should be avoided. Small swap in and large swap out considered to be normal.

Dirty write is another metric to indicate if system is running low on memory. This information can be obtained from meminfo file:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# more /proc/meminfo | grep Dirty
Dirty:704 kB
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]#
```

Disk activity should align with the underline storage configuration. Refer to [“Implementation Best Practices”](#) section on page 5-8.

## RPO and Health Monitoring

Recovery Point Objective (RPO) is the maximum amount of data loss tolerated during disaster recovery. Depending on data change rate, WAN, and storage, RPO values can fluctuate. InMage, by default, will attempt to maintain close to zero RPO through CDP. In reality, achievable RPO is a function of available resources both at the Primary and secondary data center. Monitoring real time RPO is fairly simple using InMage CX-CS or RX. Achievable RPO is reported per volume for each virtual/physical server under protection. In addition to real time reporting, InMage also provides historical trending as part of the Health Report.

Real time RPO can be monitored directly from the CX-CS Dashboard under protection details on the **Plan Summary > Protection Details** page as shown in [Figure 5-4](#).

**Figure 5-4 RPO Protection Details**

Server	VX Agent Pair	Health	Health Issue	RPO	Resync progress	Status	Resync Required	Resync Data in Transit (MB)		Differential Data in Transit (MB)			View
								Step1	Step2	On Primary Server	On CX-PS	On Secondary Server	
T2-LX-SRC-5->T2-LX-MT-1	/dev/sda -> /dev/mapper/36000c293c8fe36199f21e1144ea473b	OK	N/A	0.57 min	N/A	Differential Sync	NO	0	0	0	0.03	0	Summary
T2-LX-SRC-4->T2-LX-MT-1	/dev/sda -> /dev/mapper/36000c29edbee19152675651240d667a	OK	N/A	1.27 min	N/A	Differential Sync	NO	0	0	0	0.04	0	Summary
T2-LX-SRC-7->T2-LX-MT-1	/dev/sda -> /dev/mapper/36000c2953dc97d22899764e1a64e4f06	OK	N/A	1.12 min	N/A	Differential Sync	NO	0	0	0	0.04	0	Summary
T2-LX-SRC-3->T2-LX-MT-1	/dev/sda -> /dev/mapper/36000c29d8d02ac650ca186350fe5300	OK	N/A	1.1 min	N/A	Differential Sync	NO	0	0	0	0.06	0	Summary
T2-LX-SRC-6->T2-LX-MT-1	/dev/sda -> /dev/mapper/36000c29b93e7bfc6bf901ab68b369fe	OK	N/A	1.2 min	N/A	Differential Sync	NO	0	0	0	0.05	0	Summary
T2-LX-SRC-7->T2-LX-MT-1	/dev/sdb -> /dev/mapper/36000c29d830c8a4ac6ca42cd91084aac	OK	N/A	1.22 min	N/A	Differential Sync	NO	0	0	0	0	0	Summary

If the reported RPO is worse than the target RPO, it is important to find out where the bottleneck may reside. The "Differential Data in Transit" column in Figure 5-4 is designed to quickly identify potential bottlenecks:

- If the majority of the differential data resides on the Primary Server, the bottleneck is most likely within the primary site. Health of the Primary server, campus LAN connectivity, and Processing server health should be investigated.
- If the majority of the differential data resides on the CX-PS server, this could be an indication of WAN congestion or CX-PS server processing delay.
- If the majority of the differential data resides on the secondary server, then the health of the MT needs to be investigated. It is possible the MT couldn't process the incoming data fast enough or recovery site storage isn't sufficient to keep up with the change rate.
- If the majority of the differential data resides on both the CX-PS and secondary server, then the issue may be due to small I/O size, low compression ratio, or slow data draining at the target. Refer to "I/O Monitoring" section on page 5-7 for details.

Health Report provides a historical view of all volumes under protection. Data are sorted by date and by server. It displays health details such as data changes with and without compression, available retention window in days, RPO threshold and violation, available consistency points, and so on. Refer to Figure 5-5 for a complete list of options.

**Figure 5-5 RPO Health Report**

Date	Data changes (in HBytes)		Retention Window (Days)		RPO		No. of hours RPO not met	Data Flow Controlled (Hours)	Retention log reset?	Available Consistency Points	Protection Coverage
	With Compression	Without Compression	Policy	Available	Threshold	Max.					
Jul 24, 2013	17.26	229.18	1	0.12	30 min	2.77 min	0.03	0	NO	22	100%
Jul 25, 2013	35.18	468.3	1	0.87	30 min	1.88 min	0	0	NO	40	100%
Jul 26, 2013	50.11	537.96	1	1.02	30 min	5.63 min	0	0	NO	0	100%
Jul 27, 2013	8.74	114.48	1	1	30 min	1.34 min	0	0	NO	0	100%
<b>Total:</b>	111.29	1349.92	N/A	N/A	N/A	N/A	0.03	0	N/A	N/A	100%

Key parameters to monitor are:

- **Retention Window Available:** If available value is less than the configured value, this may be an indication actual change rate exceeds the allocated retention journal space. Refer to [Retention Volume Sizing, page 3-3](#) for additional details. For time-based retention policy, increase the size of the MT retention volume may provide a temporary relief.
- **Maximum RPO:** If Maximum RPO reached exceeds the RPO threshold, this could be an indication of WAN and storage delay / bottleneck. Refer to the previous discussion on real time RPO monitoring for details.
- **Data Flow Controlled:** This occurs when incoming change rate from the primary server exceeds the 8 GB buffer allocated by the PS server. When in data flow-controlled state, instead of sending all change blocks to the process server, the primary server will cache all changes locally in memory (up to 250 MB). Once the PS buffer utilization drops below 8 GB, the normal exchange of data between the primary and PS sever will resume. Flow controlled state is also known as thrashing state; it is intended to give the PS server sufficient time to process data it already has by backing off the primary server. This condition occurs if there's lack of sufficient WAN bandwidth or storage resources at CX-CS / MT. Refer to "Scout Server Storage and Compute Implementation" section on page 3-6 for details.

- **Available Consistency Points:** The available consistency point should align with retention window available. If the retention window has issues, there will be similar issues for available consistency points. Under rare circumstances where the retention window is normal, but available consistency points are low or does not exist, confirm if the proper version of VSS is installed with all required OS patches, remove third party backup software, if any, and finally confirm if sufficient disk space is available on the primary server. VSS requires at least 650MB of free disk. CX-CS log also keeps track of failure reasons, refer to “[Configuration and Processing Server Logging](#)” section on [page 5-21](#) for additional details.

The Health Report is available from both the CX-CS UI and the RX.

## I/O Monitoring

Applications will generally vary greatly in their I/O, and can change over time as the workload changes. Understanding the I/O is critical guarantee application performance. The first thing to understand is how much I/O the application under protection is doing, total I/Os per Second (IOPs), and the consistency of the I/O load.

- Is it constant, or bursty?
- If the load is bursty, what are the sizes of the largest spike, and duration / frequency of the spike?
- Does the spike occur predictably? or random?
- What is the growth trend of existing workloads under protection?
- Does IOPs growth correlate with capacity growth, or is performance growth out-pacing capacity growth?

Understanding the I/O size is also critical: some applications do small block I/O (less than 64K), and while others do large streaming I/Os, sequential write of MBs of data to disk. There can be differences even within the same application. Take MS-SQL for example: online transaction processing (OLTP) workloads tend to select a small number of rows at a time. These transfers are fairly small in size. Data warehouse applications tend to access large portions of the data at a time. These operations result in larger I/O sizes than OLTP workloads do. I/O size is important because it impacts storage bandwidth and latency. 1K IOPS with 4K IO size results in 4 MB/s storage throughput, while 64K I/O size will drive 16x the bandwidth. As I/O size increases, amount of time to write data to disk (latency) also increases.

The information in [Table 5-2](#) needs to be obtained before onboarding a new customer and should be constantly monitored for existing customers.

**Table 5-2** I/O Monitoring Information

Counter	Description
Disk Reads/sec Disk Writes/sec	Measures the number of IOPs.
Average Disk sec/ Average Disk sec/Write	Measures disk latency. Numbers vary, but here are the optimal values for averages over time: <ul style="list-style-type: none"> <li>• 1 - 5 milliseconds (ms) for Log (ideally 1 ms or less on average)</li> <li>• 5 - 20 ms for Database Files (OLTP) (Ideally 10 ms or less on average)</li> <li>• Less than or equal to 25-30 ms for Data (decision support or data warehouse)</li> </ul>

**Table 5-2** I/O Monitoring Information (*continued*)

Counter	Description
Average Disk Bytes/Read Average Disk Bytes/Write	Measures the size of I/Os being issued.
Current Disk Queue Length	Displays the number of outstanding I/Os waiting to be read or written from the disk.
Disk Read Bytes/sec Disk Write Bytes/sec	Measures total disk throughput.

In a Windows-based system, the counters above are available in Perfmon. For more information about Performance Monitoring on your specific version of Windows, refer to Microsoft support sites. Linux hosts can use `iostat` to gather similar performance statistics. Refer to [Monitoring Storage with Iostat](#) for additional details.

## Implementation Best Practices

### Master Target Total VMDK

A maximum of 60 VMDKs/RDMs can be attached to a MT, of which at least three are already in use (MT OS, cache, and retention VMDK). That leaves 57 slots free. The following factors should be considered when determining the maximum number of VMDKs a single MT should protect:

- **DR Drill:** When performing a drill, InMage attaches all the disk snapshots of a VM to the MT. That is, if one of the VMs being protected to the MT has three disks, then you will need at least three SCSI IDs open to be able to perform a drill for that VM.
- **Future Growth:** Application growth; move from a medium share point deployment to a large deployment. Total allocated SCSI slots should not exceed 25 - 30 (MT OS and retention included). Therefore, thirty more are open for DR drill, if the desire is to perform a DR drill for all VMs mapped to a MT all at once, as opposed to few at a time. As a best practice, InMage recommends to not exceed forty SCSI slots on a MT. The remaining slots are reserved for disks/ VM addition or DR drill.

### Master Target Cache

As a design guideline, 500MB of disk space per VMDK under protection should be reserved on the cache volume. The total size of the cache volume is a function of total number of volumes under protection:

**Size of Cache Volume = (Total number of volumes under protection) \* (500MB per volume)**

### Recovery Plan

There is one recovery plan per protection plan. To achieve RTO SLAs for large environments, you can have pre-defined recovery plans created ahead of time with the "Recover later" option and trigger them all at the same time from within vContinuum or the Scout server UI.

### Scout Server

To properly tune the system, monitor the data in transit per volume and hourly data change history. If a particular server is generating a large change rate or if there's a WAN/Storage bottleneck, it is recommended to proactively increase the cache disk on the processing server. The main intent behind increasing these thresholds is to cache data on the process server instead of the sources.

### Retention Policy

At the time of protection you can provide space or time based policy or both for retention. You could change this policy whenever you need from CX-GUI. Best practice for consistency interval and retention length depends on the following factors:

- Disk space available for retention.
- Number of servers and change rate from those servers assigned to retention drive.
- Nearest point you can go back and recover data using book mark or tags for application servers.

**Table 5-3 InMage Storage Configuration**

Change Rate	Disk Size	Type of Disk	Number of Disk	RAID
300GB	390GB	10K /15K	8	RAID 1+0
700GB	790GB	10K /15K	12	RAID 1+0
1TB	790GB	10K /15K	24	RAID 1+0

### Storage Array Sizing

Profiling on the I/O characteristics of all customer work loads to determine the type of storage configuration is required at the SP's cloud. Classify the workload into the following:

- Steady State
- Customer Onboard
- DR Recovery / Drill

For each use case, characterize the worst case number of VM, read / write ratio, and average IO size. As an example:

- Workload 1: Normal Operation
  - Maximum of 250 VMs
  - Each VM requires 96 IOPS on SP-provided storage average
  - Storage is characterized as 33% read / 33% random write / 34% sequential write
  - Average read/write size is 20KB
  - 75GB of space required per VM
- Workload 2: Onboarding
  - Maximum of 3 VMs at any given time
  - Each VM requires 80 IOPS average
  - 100% sequential write
  - Average read/write size is greater than 16KB
  - 75GB of space required per VM
- Workload 3: DR Operation / Recovery
  - Maximum of 250 VMs
  - Each VM requires 96 IOPS on SP-provided storage average
  - Storage is characterized as 67% read / 33% random write
  - Average block size is 20KB
  - 75GB space required per VM

- Workload 4: Mixed Normal Operation and Recovery
  - Maximum of 250 VMs
  - Each VM requires 96 IOPS on SP-provided storage average
  - Storage is characterized as 50% read / 25% random write / 25% random read
  - Average block size is 20KB
  - 75GB space required per VM

Based on each workload characteristic and storage vendor, determine if combinations of SSD / SAS / SATA could fit into your environment. Both FlexPod and vBlock offer auto tiering, but there are some major differences in terms of implementation.

- VMAX:
  - Performance Time Window: 24/7/365 for continuous analysis
  - Data Movement Window: 24/7/365 for continuous data movement
  - Workload Analysis Period: 24/7/365 for continuous analysis
  - Initial Analysis Period: Can be configured to be between 2 hours and 4 weeks, The default is 8 hours.
  - FAST-VP Relocation Rate: 1 to 10, 786KB chunks
  - Promotion to Cache: Immediate
- VNX:
  - Performance Time Window: 24/7/365 for continuous analysis
  - Data Movement Window: Once every 24 hours
  - Workload Analysis Period: Once an hour
  - Initial Analysis Period: Once an hour
  - FAST-VP Relocation Rate: Low/Medium/High, 1GB chunk
  - Promotion to Cache: Immediate
- NetApp:
  - Performance Time Window: 24/7/365 for continuous analysis
  - Data Movement Window: 24/7/365 for continuous data movement
  - Workload Analysis Period: 24/7/365 for continuous analysis
  - Promotion to Cache: Immediate
  - Random Write < 16K IO size - SSD
  - Random Write > 16K IO Size - SAS /SATA
  - Sequential Writes - SAS / SATA

Depending on storage platform and vendor, I/O size can influence if write cache can be optimized. Data movement window can influence whether a special on-boarding strategy needs to be implemented. Check with the storage vendor to determine the optimal storage configuration.

### **InMage Interaction with MS SQL**

Replication and backup products truncate application logs through VSS. However, the VSS writer implementation for MS SQL does not expose the ability to truncate logs. This is different from the behavior of MS Exchange VSS Writer, for example, which exposes the API to truncate logs. Due to this, InMage does not have a way to truncate MS SQL logs. For DB type of apps, the recommendation is to continue native application backup on a regular basis to maintain log sizing.

### **Application Consistency**

In situations where Windows 2003 (Base, SP1) source machines have applications that require application quiesce, it is strongly suggested to upgrade to Windows 2003 SP2 to overcome the VSS-related errors.

### **vSphere Version**

To provide failover from enterprise to SP, the secondary vSphere (SP) version should be either the same or higher than the source (enterprise) vSphere server. To perform a failback from SP to

Enterprise, enterprise vSphere version should be either the same or higher than the SP vSphere. vSphere server may need to be upgraded if failback is required.

### **OS Recommendations**

For new installations, InMage recommends:

- Secondary ESXi Platform: ESXi 5.1
- MT Platform for Windows: Windows 2008 R2 Enterprise Edition
- MT Platform for Linux: CentOS 6.2 64-bit
- CX Scout OS: CentOS 6.2 64-bit
- vContinuum: Windows 2008 R2

Protect Windows 2012 VM with ReFS Filesystem. This requires matching the 2012 MT.

### **SSH Client**

Linux bulk agent install requires SSH access to originate from the primary server. Ensure SSH client is installed on all Linux primary servers.

### **Tenant Self Service**

One way to achieve complete self service capability is to allow tenants to have access to the tenant-specific vContinuum Server. For security reasons you may prefer to create a non-administrator role on the vCenter for each tenant vContinuum user. The following privileges should be selected:

- Datacenter
- Datastore
- Folder
- Host
- Network
- Resource
- Storage views
- Virtual machine
- vSphere Distributed Switch

Using vSphere RBAC, assign tenant-specific vSphere resources to the newly created tenant user/role.

### Upgrade Sequence

General upgrade sequence follows:

- Upgrade the CX-CS.
- Upgrade the processing server.
- Upgrade the MT.
- Upgrade the agent on source physical or virtual server. Use the CX-CS UI instead of vContinuum to upgrade the agents.

Always refer to the release notes for the exact sequence.

## Caveats

Refer to the following InMage documents:

- InMage RX Release Notes and Compatibility Matrix Documents  
[http://support.inmage.net/partner/poc\\_blds/14\\_May\\_2013/Docs/RX/](http://support.inmage.net/partner/poc_blds/14_May_2013/Docs/RX/)
- InMage CX Release Notes and Compatibility Matrix Documents  
[http://support.inmage.net/partner/poc\\_blds/14\\_May\\_2013/Docs/Scout/](http://support.inmage.net/partner/poc_blds/14_May_2013/Docs/Scout/)
- InMage vContinuum Release Notes and Compatibility Matrix Documents  
[http://support.inmage.net/partner/poc\\_blds/14\\_May\\_2013/Docs/vContinuum/](http://support.inmage.net/partner/poc_blds/14_May_2013/Docs/vContinuum/)

### Disk Removal

vContinuum does not support disk removal. To remove a disk, first remove the VM from the protection plan and then add VM without the removed disk back into the protection plan. This operation requires a complete resync between enterprise and service provider.

### Storage Over Allocation

vContinuum does not check for available capacity on the MT retention drive at the first protection. It is possible to reserve capacity beyond what is available. No alarms are generated for this condition; use the following procedure to confirm if you are running into this condition:

1. Create a new protection plan from the vContinuum.
2. Select any random Primary VM.
3. Select secondary ESX host.
4. When selecting data stores, confirm if the retention drive is available as an option.

If the retention drive is available, then the available capacity has not been over-subscribed. If the OS volume is the only available option, it is strongly recommended to manually reduce the retention size on protected VMs from the CX UI.

### Protection Plan

A protection plan can only map to a single MT. No mechanisms exist to migrate a protection plan between MTs.

### Continuous Disk Error

Linux MT continuously generates the following error:



```
Feb 26 11:08:32 mtarget-lx-22 multipathd: 36000c29f5decda4811ca2c34f29a9fdc: sdf -
directio checker reports path is down
Feb 26 11:08:32 mtarget-lx-22 kernel: sd 4:0:0:0: [sdf] Unhandled error code
Feb 26 11:08:32 mtarget-lx-22 kernel: sd 4:0:0:0: [sdf] Result:
hostbyte=DID_NO_CONNECT driverbyte=DRIVER_OK
Feb 26 11:08:32 mtarget-lx-22 kernel: sd 4:0:0:0: [sdf] CDB: Read(10): 28 00 00 00 00
00 00 00 08 00
```

- **Root Cause:** Each MT can have up to four SCSI controllers; each controller can have up to 15 disks. When adding a new SCSI controller, VMware requires a disk to be associated with the controller. Since disks are only added during protection time, the InMage workaround for this is to associate a dummy disk to the SCSI controller and delete the dummy disk once the controller is added. When OS attempts to acquire locks to disks that no longer exist, this causes the continuous error log. The error above is cosmetic in nature and can safely be ignored.

### CX Integration with RX

If a CX-CS server is configured with dual NICS (one to communicate with the primary server and the second to communicate with the RX), use the push method instead of pull when adding the CX-CS to the RX. This is a known limitation.

### VMware Tools

For virtual-to-virtual protection, updated VMware tools are required for InMage. vContinuum will not add a server to a protection plan if VM tools are not started or out of date.

### Statistics Export

The network traffic rates cannot be exported from the CX-CS UI; it is only available from the dashboard as a daily, monthly and yearly graph. However, it is possible to access the raw data (RRD format) from the CS/PS Server.

### Agent Status not reflected in CX Alerts and Notifications Window

CX-CS server does not raise any alarms in the CX Alerts and Notifications window when a source machine fails into bitmap mode. The machine's bit map mode cannot be recovered in the SP VPC.

### Recovery Plan

1. Tenant-created recovery plans are not visible from vContinuum until execution. Once executed, recovery status can be monitored from vContinuum.
2. SP-created recovery plans (recover later) from vContinuum are intended to be SP managed and not visible to the tenant via multi-tenant portal.
3. Protect > Manage protected Files/Folders is a single source of truth. Both tenant-created and SP-created recovery plan can be viewed.

For recover later plans, outside of the Multi-Tenant Portal, there is no mechanism to view VMs included in a recovery plan. The assumption is the recovery plan should map exactly to a protection plan. If protection plan changes, a new recovery plan should be created.

### Protection Plan Span Across Multiple Master Targets

All protected VMs disk have to reside in a single MT and cannot be spanned over multiple MTs.

### Vmotion

Vmotion of MT from one ESX host to another is supported. Vmotion of MT storage will not work due to UUID changes.

**Multi-tenant Portal**

Multi-tenant portal does not support protection plan modifications. The following are InMage roadmap items:

- Add disk to a protected VM (Windows / Linux).
- Add VM (Windows / Linux) to an existing plan.
- Add physical servers (Windows / Linux) to an existing protection plan.

**Portal Update**

The default sync interval between the CX-CS and RX is 10 minutes. The tenant portal can lag behind the CX-CS UI by up to 10 minutes.

**vContinuum Agent Install**

vContinuum agent install wizard has the following limitations:

- Indicates if a server already has agents installed.
- Ability to select which network interface to use to push agent.
- NAT configuration for servers with multiple IPs.

vContinuum agent install may fail on servers with multiple NICs. Workaround is to manually install agent.

**VM Name**

During Windows failback, VM name configured in the primary vCenter may be modified to match the VM name in the secondary site. This issue is not observed on Linux machines.

**Recovery Plan**

If a disaster recovery is launched from the CX portal, depending on sequence of execution, successful recovery job maybe reported as failed. This is an UI reporting issue. To ensure accurate recovery job reporting, use the vContinuum to initiate the recover job.

**Failback Readiness Check**

When performing bulk failback recovery from secondary site to primary site using tag consistency, it is possible that readiness check may pass even when portion of the VM(s) have not reached tag consistent status. Workaround is to visually inspect the vContinuum log, ensure there are no errors before proceeding with live failback.

**Recovery Readiness Check from the RX Portal**

RX does not validate vCenter access as part of recovery readiness check.

**Documented Procedure for Extending a Linux Volume/Disk**

The procedure for extending a Linux volume/disk is missing in the online documentation. The procedure is documented in Appendix B Extending a Linux Volume.

**vContinuum V2P Recover Does Not Complete**

The vContinuum wizard V2P recovery plan fails to reboot the physical server after the recovery is complete. At this point in the recovery plan, the recovery is complete, yet the recovery script never completes. Reboot the physical server from the console to complete the V2P recover plan.

**Linux P2V and V2P Recover Fails to Bring Up Ethernet Interface**

The P2V and V2P recover plan for a Linux physical server fails to restore the network configuration. The interface network configuration can be recovered by entering the service network restart command. However, this is work-around is not persistent and requires reentering the service network restart command after each reboot.

**V2P Failback Plan Requires Resync**

The V2P failback for a Linux physical server requires a Resync immediately after volume replication is complete. Once the volume replication achieves differential sync status, use the CX UI to manually Resync the volume replications.

**Recovery Plan Readiness Check**

The vContinuum wizard does not fail the recovery plan when any primary VM in the recovery plan fails the readiness check. Instead, the vContium wizard removes any primary VMs from the recovery plan that fail the readiness check and allows the user to proceed.

In situations where a large number of primary VMs are being recovered, it is possible that the user may not be aware that some of primary VMs in the recovery plan failed the readiness check. This can happen when the primary VM that fails the readiness check is not visible within the current vContinuum window and the user must scroll down to see the error.

## Troubleshooting

Troubleshooting of DR events, both successful and failed, is critical to maintain business continuity. In an environment where business service level agreements (SLAs) are tied to speed to recovery, effective troubleshooting is an integral part of accomplishing those goals. Successful troubleshooting starts with the ability to trace an event from beginning to end, by starting from the highest level and drilling down towards each component. The ability to correlate multiple events and presentation of those events is a fundamental requirement.

Configuring and setting up protection plan is a multi-step process, it starts from tenant-side discovery and reaches steady state with Differential Sync. [Table 5-4](#) provides a brief workflow description, software components involved, and relevant log files to review when failure occurs.

**Table 5-4 Troubleshooting DR Events**

Sequence	Workflow Description	Software Components	Relevant Logs
1	Primary Site Discovery	vContinuum, vCenter/ESXi/ Physical Server - Primary Site	<ol style="list-style-type: none"> <li>1. vContinuum               <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs\vContinuum.log</li> <li>b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest\MasterConfigFile.xml</li> </ol> </li> <li>2. vCenter Status (tenant)               <ol style="list-style-type: none"> <li>a. VMware Tools Status</li> <li>b. vCenter Security</li> </ol> </li> </ol>
2	Secondary Site Discovery	vContinuum, vCenter/ESXi - Service Provider Secondary Site	<ol style="list-style-type: none"> <li>1. vContinuum               <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs\vContinuum.log</li> <li>b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest\MasterConfigFile.xml</li> </ol> </li> <li>2. vCenter Status (tenant)               <ol style="list-style-type: none"> <li>a. VMware Tools Status</li> <li>b. vCenter Security</li> </ol> </li> </ol>
3	Datastore Selection	vContinuum, vCenter/ESXi - Service Provider Secondary Site	<ol style="list-style-type: none"> <li>1. vContinuum               <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs\vContinuum.log</li> <li>b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest\MasterConfigFile.xml</li> </ol> </li> <li>2. Retention drive size. Refer to "Storage Over Allocation" in Caveats.</li> <li>3. vCenter Security</li> </ol>
4	Target VM and Network Configuration	vContinuum, vCenter/ESXi - Service Provider Secondary Site	<ol style="list-style-type: none"> <li>1. vContinuum               <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs\vContinuum.log</li> <li>b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest\MasterConfigFile.xml</li> </ol> </li> <li>2. vCenter Security</li> </ol>
5	Readiness Check vContinuum, vCenter/ESXi/ Physical Server (Primary and Secondary Site)	vContinuum	<ol style="list-style-type: none"> <li>1. vContinuum               <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs\vContinuum.log</li> </ol> </li> </ol>

Table 5-4 Troubleshooting DR Events (continued)

Sequence	Workflow Description	Software Components	Relevant Logs
6	Activate Plan <ul style="list-style-type: none"> <li>• Export and Import VMX</li> <li>• Create VMDKs</li> <li>• Attach VMDKs to MT</li> <li>• Create protection pairs</li> </ul>	vContinuum, vCenter, ESXi/Physical Server (Primary and Secondary Site), CS, MT	<ol style="list-style-type: none"> <li>1. vContinuum               <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs</li> </ol> </li> <li>2. ESXi/vCenter (SP)               <ol style="list-style-type: none"> <li>a. vCenter Log</li> </ol> </li> <li>3. CX-CS               <ol style="list-style-type: none"> <li>a. Monitor - CX logs - job_log_XXXX</li> </ol> </li> </ol>
7	Unified Agent Pulling config	CX-CS, Unified Agent	<ol style="list-style-type: none"> <li>1. Unified Agent               <ol style="list-style-type: none"> <li>a. configuratorapitestBed.exe. Refer to InMage Tools</li> </ol> </li> <li>2. CX-CS               <ol style="list-style-type: none"> <li>a. Monitor - CX logs - configurator_register_host_static_info</li> </ol> </li> </ol>
8	Resync Step 1	UA (source & MT), CS, PS	<ol style="list-style-type: none"> <li>1. UA host logs (source &amp; MT)               <ol style="list-style-type: none"> <li>a. Monitor - Host Logs</li> </ol> </li> <li>2. CS/PS logs               <ol style="list-style-type: none"> <li>a. /home/svsystems/target host id/target volume/resync. Look for the oldest files that are not being processed</li> <li>If completed_hcd files are not processed, look for transfer errors at the source (Monitor - Host Logs)</li> <li>If completed_sync files are not processed, look for transfer errors at the destination</li> <li>b. Check for WAN / bpm policies (Setting - Manage Bandwidth Usage)</li> <li>c. /home/svsystems/transport/log/cxps.err.log</li> <li>d. mysql</li> <li>e. tmanager( Monitor -&gt; CX logs -&gt; volsync)</li> </ol> </li> </ol>

Table 5-4 Troubleshooting DR Events (continued)

Sequence	Workflow Description	Software Components	Relevant Logs
9	Resync Step 2	UA (source & MT), CS, PS	<ol style="list-style-type: none"> <li>1. UA host logs (source &amp; MT) <ol style="list-style-type: none"> <li>a. Monitor - Host Logs</li> </ol> </li> <li>2. CS/PS logs <ol style="list-style-type: none"> <li>a. /home/svsystems/&lt;target host id&gt;/&lt;target volume&gt;/diffs <pre>[root@ent1-scout-1 diffs]# ls completed_diff_P130197922886501968_13708331 1_E130197922887125969_13708350_WE1.dat.gz completed_diff_P130197922887125969_13708350 1_E130197922887593970_13708428_WE1.dat.gz completed_diff_P130197922887593970_13708428 1_E130197922888061971_13708513_WE1.dat.gz monitor.txt pre_completed_diff_P130197922888061971_1370 1_E130197922888997972_13708670_WE1.dat</pre> </li> <li>b. Check for WAN / bpm policies (Setting - Manage Bandwidth Usage)</li> <li>c. /home/svsystems/transport/log/cxps.err.log</li> <li>d. mysql</li> <li>e. tmanager ( Monitor - CX logs - volsync)</li> </ol> </li> </ol>
10	Differential Sync	UA (source & MT), CS, PS	<ol style="list-style-type: none"> <li>1. UA host logs (source &amp; MT) <ol style="list-style-type: none"> <li>a. Monitor -&gt; Host Logs</li> </ol> </li> <li>2. CS/PS logs <ol style="list-style-type: none"> <li>a. /home/svsystems/&lt;target host id&gt;/&lt;target volume&gt;/diffs</li> <li>b. Check for WAN / bpm policies (Setting -&gt; Manage Bandwidth Usage)</li> <li>c. /home/svsystems/transport/log/cxps.err.log</li> <li>d. mysql</li> <li>e. tmanager ( Monitor - CX logs - volsync)</li> </ol> </li> </ol>

**Table 5-5 Troubleshooting Recovery Plan**

Sequence	Workflow Step	Software Components Involved	Logs
1	VM, Recovery point (Tag) selection and readiness check	vContinuum, vCenter / ESXi (Secondary Site), MT (Secondary Site)	<ol style="list-style-type: none"> <li>1. vContinuum <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs</li> <li>b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest\Recovery.xml</li> </ol> </li> <li>2. MT <ol style="list-style-type: none"> <li>a. cdpli for retention information. Refer to InMage Tools</li> </ol> </li> </ol>
2	Target VM network configuration and recovery sequencing	vContinuum, vCenter/ESXi (SP)	<ol style="list-style-type: none"> <li>1. vContinuum <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs</li> </ol> </li> </ol>
3	Activation (Rollback disks, Apply network configuration, Detach disks from MT, Power on VMs)	vContinuum, vCenter/ESXi (SP), MT, CS	<ol style="list-style-type: none"> <li>1. vContinuum <ol style="list-style-type: none"> <li>a. C:\Program Files (x86)\InMage Systems\vContinuum\logs</li> </ol> </li> <li>2. ESXi/vCenter (SP) <ol style="list-style-type: none"> <li>a. vCenter Log</li> </ol> </li> <li>3. MT <ol style="list-style-type: none"> <li>a. Monitor -&gt; Hosts</li> </ol> </li> <li>4. CX-CS <ol style="list-style-type: none"> <li>a. Monitor -&gt; CX logs -&gt; job_log_xxxx</li> </ol> </li> </ol>

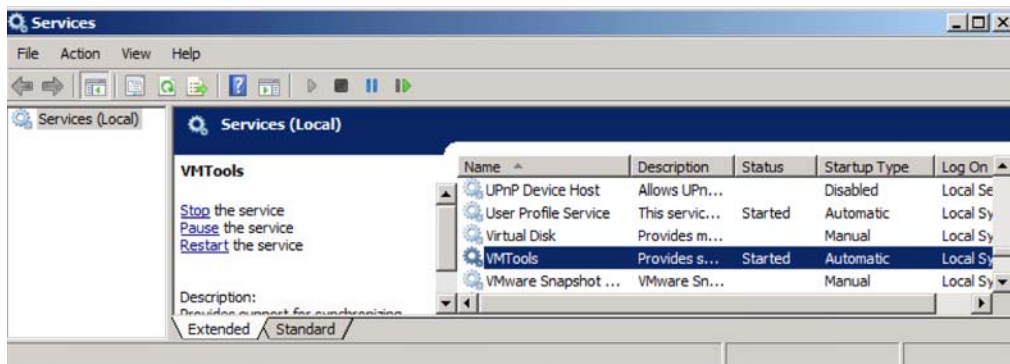
This section will introduce important log files and dependencies for InMage components:

- [VMware Tools, page 5-19](#)
- [vContinuum Logging, page 5-20](#)
- [Configuration and Processing Server Logging, page 5-21](#)
- [InMage Tools, page 5-22](#)

## VMware Tools

For V2V protection, VMware Tools are required for source-side discovery. To confirm if VMware Tools is running on Windows, log onto the source server and under Services confirm that the VMtool is started, as shown in [Figure 5-6](#).

Figure 5-6 VMware Tools



On Linux, issue:

```
ps -ef | grep vmttoolsd
```

## vContinuum Logging

vContinuum is a stateless application that communicates with the CX-CS server to create and manage various protection and recovery plans. Based on user inputs, it pulls/pushes configuration changes to/from the CX-CS server. Detailed exchange between vContinuum and CX-CS server log are all logged on the vContinuum server under:

```
C:\Program Files (x86)\InMage Systems\vContinuum\logs
C:\Program Files (x86)\InMage Systems\vContinuum\Latest:
```

Important vContinuum log files are:

1. vContinuum.log - Use this log file to follow detailed events for jobs triggered from the vContinuum. The following is an example of Recovery Plan:

```
7/30/2013 10:58:53 AM Parameter grp id Task1 Initializing Recovery Plan: 7/30/2013
10:58:53 AM Name
This will initialize the Recovery Plan.It starts the EsxUtil.exe
binary for Recovery: 7/30/2013 10:58:53 AM Description
Completed: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/
EsxUtil.log
7/30/2013 10:58:53 AM Parameter grp id Task2
Downloading Configuration Files: 7/30/2013 10:58:53 AM Name
The files which are going to download from CX are1.
Recovery.xml: 7/30/2013 10:58:53 AM Description
Completed: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/
EsxUtil.log
7/30/2013 10:58:53 AM Parameter grp id Task3
Starting Recovery For Selected VM(s): 7/30/2013 10:58:53 AM Name
The following operations going to perform in this task:1.
Remove pairs for all the selected VMs2. Completes network
related changes for all VMs3. Deploys the source disk layout
on respective target disk(in case of windows): 7/30/2013 10:58:53 AM Description
Completed: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/ EsxUtil.log
7/30/2013 10:58:53 AM Parameter grp id Task4
Powering on the recovered VM(s): 7/30/2013 10:58:53 AM Name
This will power-on all the recovered VMs1. It will detach
```



```


all the recovered disks from MT2. Power-on the recovered VMs:
7/30/2013 10:58:53 AM Description
InProgress: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/ EsxUtil.log
7/30/2013 10:59:18 AM paragroup <FunctionRequest Name="MonitorESXProtectionStatus"
Id=" " include="No"><Parameter Name="HostIdentification"
Value="0451173A-C182-8D46-9C18A7E2E844E42D" / ><Parameter Name="StepName"
Value="Recovery"/><Parameter Name="PlanId" Value="9"/></FunctionRequest>
7/30/2013 10:59:18 AM Count of parametergroup 2

```

Files in the "C:\Program Files (x86)\InMage Systems\vContinuum\Latest" are XML files that includes inventory information regarding:

- Primary Server under protection.
- Secondary Service Provider ESXi Environment.
- Detailed Storage information in Primary and Secondary site.
- Detailed network profile information in Primary and Secondary site. [Figure 5-7](#) is an example of MasterConfigFile (detailed protection plan).

**Figure 5-7** MasterConfigFile Example



## Configuration and Processing Server Logging

Logs required to troubleshoot the processing server and CS-CX can be accessed directly from the CS-CX UI. Summary of available logs can be found by navigating to Monitor > CX Logs. [Table 5-6](#) shows the logs that are most relevant for troubleshooting:

**Table 5-6** Configuration and Processing Server Logs

Log	Description
volsync	tmanager logs
Jobs_log	File replication logs. Logs associated with protection and recovery plan.
gentrends	Logs responsible for RRD trending graph.
perf_fr_job_cfg	Performance data gathering and management for gentrend.

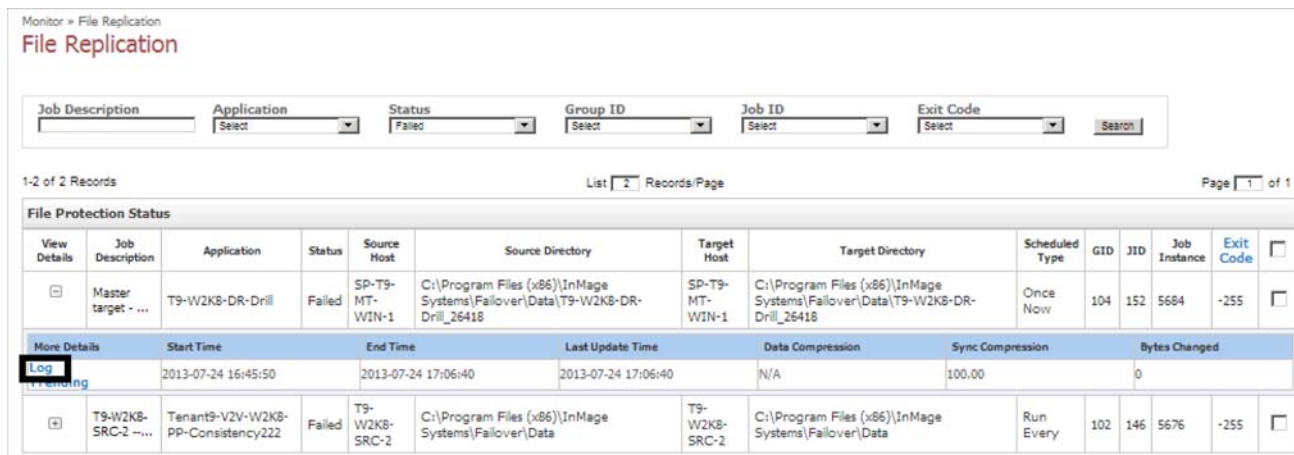
**Table 5-6** Configuration and Processing Server Logs

Log	Description
bmptrace	Bandwidth shaping module.
TrapLog	SNMP messages.

Troubleshooting directly from Monitor > CX logs can be overwhelming since it is an aggregation point for all logs. The best way to retrieve and view logs is to have some context around the failure. Having the right log file corresponding to the failure event simplifies the root cause isolation and avoids going down the wrong path due to conflicting error messages across different log files. To simplify the debugging process:

1. Always start from Monitor > File Replications.
2. Filter replication logs based on Status. Status > Failed > Search.
3. Find the corresponding failed job based on application name and then click on view details.
4. From the details windows, click on log to open the corresponding log file relating to the failure event. [Figure 5-8](#) is a screen capture:

**Figure 5-8** View Log



## InMage Tools

The following InMage tools can be used to for troubleshooting:

1. **DrvUtil - Primary Server.** DrvUtil utility can be used to inspect and modify InMage DataTap driver settings on the primary server. For monitoring and troubleshooting purposes, Drvutil --ps and --pas options can be useful to monitor IO Size, number of dirty blocks in queue, size of the change block and so on. Other features of drvutil should not be used unless instructed by InMage support team.
2. **ConfiguratorAPITestBed - Primary Server.** This tool is most useful when troubleshooting initial data protection issues (Resync). First confirm if "dataprotection" process is running; if it's not, use configuratorapitestbed.exe to check if svagent is receiving configuration settings from the CX-CS correctly.

```
C:\Program Files (x86)\InMage Systems>ConfiguratorAPITestBed.exe --default
C:\Program Files (x86)\InMage Systems>ConfiguratorAPITestBed.exe --custom --ip
```

```
8.24.81.101 --port 80 --hostid 3F8E5834-AA0C-F246-B915D07CFB5D49CC
```

This is a Windows-based utility; there isn't a equivalent Linux version. To find out configuration settings for a Linux primary servers, use any available Windows DataTap agent and run the ConfiguratorAPITestBed command using the --hostid option where the hostid is the id of the Linux host.

- 3. cdpcli - Master Target.** Use this utility on the MT to gather information regarding replication statistics, IO pattern, and protected volume.

```
c:\Program Files (x86)\InMage Systems>cdpcli.exe --listtargetvolumes
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C
C:\ESX\4FC26A25-683F-DC48-86F0C31180D4A5C0_C
C:\ESX\85F7759A-2EC9-9348-B8C0AD1E9A3F0331_C
C:\ESX\B6688D89-7471-EA43-AC76FBA7EDB948C6_C
C:\ESX\E08CD805-888E-944D-B60496095D3914DD_C
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C__SRV
C:\ESX\4FC26A25-683F-DC48-86F0C31180D4A5C0_C__SRV
C:\ESX\85F7759A-2EC9-9348-B8C0AD1E9A3F0331_C__SRV
C:\ESX\B6688D89-7471-EA43-AC76FBA7EDB948C6_C__SRV
C:\ESX\E08CD805-888E-944D-B60496095D3914DD_C__SRV
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_K
C:\ESX\B6688D89-7471-EA43-AC76FBA7EDB948C6_K
C:\ESX\E08CD805-888E-944D-B60496095D3914DD_K
c:\Program Files (x86)\InMage Systems>cdpcli.exe --displaystatistics -- vol="C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C"
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C\ is a symbolic link to C:\ESX \3F8E5834-AA0C-F246-B915D07CFB5D49CC_C
```

```
##### REPLICATION
```

```
Target Volume Name:          STATISTICS #####
                               C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C
Diffs pending in CX:         30391423
Diffs pending in Target:     0
Current RPO (secs):          96
Apply rate (Bytes/sec):      12191886
Apply time (secs):           0
```

```
c:\Program Files (x86)\InMage Systems>cdpcli.exe --showsummary --vol="C:\ESX \3F8E5834-AA0C-F246-B915D07CFB5D49CC_C"
Database:E:\Retention_Logs\catalogue
\2460F4D5-7C71-5745-9804B2F
FB039366A\C\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C\ef118abbc9\cdpv3.db
Version:3
Revision:2
Log Type:Roll-Backward
Disk Space (app):235020800 bytes
Total Data Files:10
Recovery Time Range(GMT): 2013/7/31 13:37:8:730:417:7 to
2013/7/31 13:49:13:5:533:1
c:\Program Files (x86)\InMage Systems>
c:\Program Files (x86)\InMage Systems>cdpcli.exe --iopattern --vol="C:\ESX \3F8E5834-AA0C-F246-B915D07CFB5D49CC_C"
Io Profile:
```

size	%Access	%Read	%Random	Delay	Burst	Alignment	Reply
512B		0100		1000		lsector	none
512B		00		00		lsector	none
512B		0100		00		lsector	none
512B		20		1000		lsector	none
4KB		12100		1000		lsector	none
4KB		120		00		lsector	none
4KB		12100		00		lsector	none
4KB		140		1000		lsector	none

8KB	2100	1000	lsector none
8KB	20	0	lsector none
8KB	2100	0	lsector none
8KB	50	1000	lsector none
16KB	2100	1000	lsector none
16KB	20	0	lsector none
16KB	2100	0	lsector none
16KB	40	1000	lsector none
64KB	3100	1000	lsector none
64KB	30	0	lsector none
64KB	3100	0	lsector none
64KB	60	1000	lsector none
256KB	2100	1000	lsector none
256KB	20	0	lsector none
256KB	2100	0	lsector none
256KB	40	1000	lsector none
1MB	0100	1000	lsector none
1MB	00	0	lsector none
1MB	0100	0	lsector none
1MB	10	1000	lsector none
4MB	0100	1000	lsector none
4MB	00	0	lsector none
4MB	0100	0	lsector none
4MB	10	1000	lsector none

[c:\Program Files \(x86\)\InMage Systems](#)>

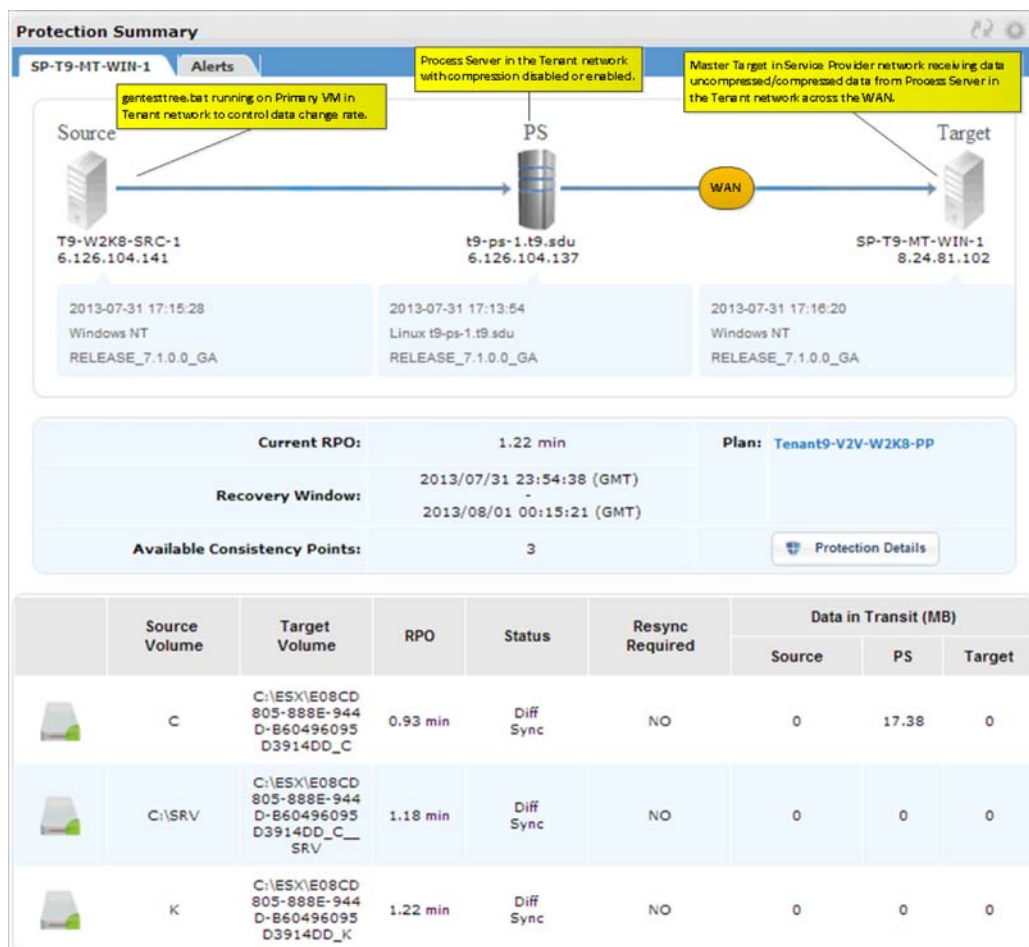


# APPENDIX A

## Characterization of Replication Process

As part of the testing effort for this DRaaS solution, the impact of the primary VM's data change rate on CPU utilization, network bandwidth, and disk I/O were investigated. The tests were executed on a single tenant with one Windows 2008 virtual server being protected by a MT running Windows 2008 in the SP's network. The process server, which was located in the tenant network, sent data to the MT across a simulated WAN link. [Figure A-1](#) shows this setup along with some details about the protection plan being used.

**Figure A-1** Protection Summary for Characterization Test Setup



**Note**

The results presented in this section should be used to understand the general effects of data change rates and compression on the system and should be not used for planning purposes by themselves. InMage provides recommendations and analysis tools to assist with network and resource planning that should be used.

InMage has a script to generate data changes on the primary server to put a load on the disaster recovery system. The script calls an executable that writes data to a temporary folder, waits for a short duration, removes the written data, and then loops back to the start. The data writes are captured by the InMage agent on the primary server and sent to the local processing server. The processing server then compresses those changes, if compression is enabled, and the MT pulls those changes to the protection VMDK in the SP network.

In the following script, the GenerateTestTree.exe executable is called with three parameters; only the second parameter (e.g., size of the data to write each time) was changed for each iteration. This value was stepped from 1 up to 250.

**gentesttree.bat Test Script**

```
:loop1
@echo on
"c:\scripts\GenerateTestTree.exe" 0 10 8 C:\temp1
ping -n 5 localhost
"c:\Program Files (x86)\InMage Systems\rm.exe" -rf C:\temp1
goto loop1
```

The GenerateTestTree.exe executable has the following syntax:

```
GenerateTestTree.exe <mode=0(write)|1(verify)> <size MB> <random seed int> <dest dir>
```

At each interval of the test script, the system was allowed to settle for an hour or more and then a number of metrics was collected. Between some iterations, we observed data being cached at the process server so we forced a restart of the replication process to flush that data.

- Metrics captured from the Tenant vCenter:
  - **Primary VM Disk Write Rate (WR)**—Measured in MBps. This is the actual amount of data changes resulting from the gentesttree.bat test script.
  - **Process Server BW Input and Output**—Measured in MBps (bytes), but converted to Mbps (bits). The difference between these two values will be the amount of compression the process server can accomplish.
  - **Primary VM (Agent) CPU Utilization**—Measured in MHz. This is the amount of CPU required by the agent.
  - **Process VM (PS) CPU Utilization**—Measured in MHz. This is the amount of CPU required by the process server.
- Metrics captured from the SP vCenter:
  - **Master Target Disk Read Rate (RR)**—Measured in MBps. This is the read rate required by the MT to compare blocks.
  - **Master Target Disk Write Rate (WR)**—Measured in MBps. This is the write rate by the MT to write old block to log and new block to VMDK.
  - **Master Target CPU Utilization**—Measured in MHz, This is the amount of CPU required by the MT.

This section presents the following topics:

- [Replication with Compression Disabled, page A-3](#)
- [Replication with Compression Enabled, page A-4](#)
- [Comparison of Compression Enabled and Disabled, page A-5](#)

## Replication with Compression Disabled


Figure A-2 shows the summary results for each of the seven (7) iterations when compression was disabled on the process server in the tenant network.

**Figure A-2 Summary of Results with Compression Disabled**

GenerateTestTree Params			Enterprise (Primary VM, PS)					Service Provider (MT)		
Mode	Size	Seed	Pri VM Disk WR	PS BW (In)	PS BW (Out)	Pri VM CPU	PS CPU	Disk RR	Disk WR	CPU
0	1	8	0.375	3.39968	3.883008	96	167	0.395	1.095	365
0	10	8	1.91	16.46592	16.1792	200	230	1.956	5.622	505
0	50	8	3.909	33.83296	32.768	293	315	4.027	11.784	739
0	100	8	4.33	37.35552	35.55328	338	330	6.65	12.728	788
0	150	8	5.336	41.598976	41.20576	345	338	5.05	14.96	865
0	200	8	6.922	46.465024	43.4176	369	362	5.34	15.745	890
0	250	8	6.968	46.61248	44.212224	372	367	5.399	15.991	910


In Figure A-3, the CPU utilization for the primary VM, process server, and MT is plotted against the primary VM disk write rate. The utilization on the primary VM and process server are very close consuming 100-200 MHz for the lowest disk write load and approaches 400 MHz for the upper limit of the load script (7 MBps). On the SP side, the MT consumes more cycles for the same workload. The MT starts at a little under 400 MHz and climbs up to around 900 MHz at the highest loads.

**Figure A-3 Chart of CPU Utilization vs Primary VM Disk Write Rate (Compression Disabled)**



In Figure A-4, the input and output data rates for the network interface of the process server are plotted against the primary VM disk write rate. The data rate coming from the primary VM and going out to the MT are very close, which is expected since the process server is not performing any compression on the data.

Figure A-4 Chart of Bandwidth vs Primary VM Disk Write Rate (Compression Disabled)



## Replication with Compression Enabled

Figure A-5 shows the summary results for each of the seven (7) iterations when compression was enabled on the process server in the tenant network.

Figure A-5 Summary of Results with Compression Enabled


GenerateTestTree Params			Enterprise (Primary VM, PS)					Service Provider (MT)		
Mode	Size	Seed	Pri VM Disk WR	PS BW (In)	PS BW (Out)	Pri VM CPU	PS CPU	Disk RR	Disk WR	CPU
0	1	8	0.375	3.35872	2.58048	92	212	0.394	0.981	369
0	10	8	1.934	16.760832	12.509184	202	459	2.039	5.284	580
0	50	8	3.951	33.923072	27.189248	296	787	4.04	11.065	893
0	100	8	4.325	37.126144	28.844032	383	843	4.425	12.06	944
0	150	8	6.712	45.62944	38.985728	386	1006	5.683	16.891	1202
0	200	8	6.943	45.703168	38.928384	391	983	5.623	17.011	1148
0	250	8	7.796	59.523072	36.53632	423	1296	5.158	14.589	1116

In Figure A-6, the CPU utilization for the primary VM, process server, and MT is plotted against the primary VM disk write rate. The utilization on the primary VM consumes around 100-200 MHz for the lowest disk write load and approaches 400 MHz for the upper limit of the load script (7 MBps). As expected, this is similar to the no compression test results, since nothing has changed on the primary VM.

Looking at the CPU utilization for the process server, we see that the starting consumption is close to the no compression results (e.g., around 200 MHz), but almost triples in consumption near the upper limit of the load script. The highest consumption is around 1000 MHz, while the no compression results yielded around 350 MHz. The CPU utilization of the MT is about the same as the compression disabled case for the smallest loads, but consumes significantly more CPU resources at the higher loads.

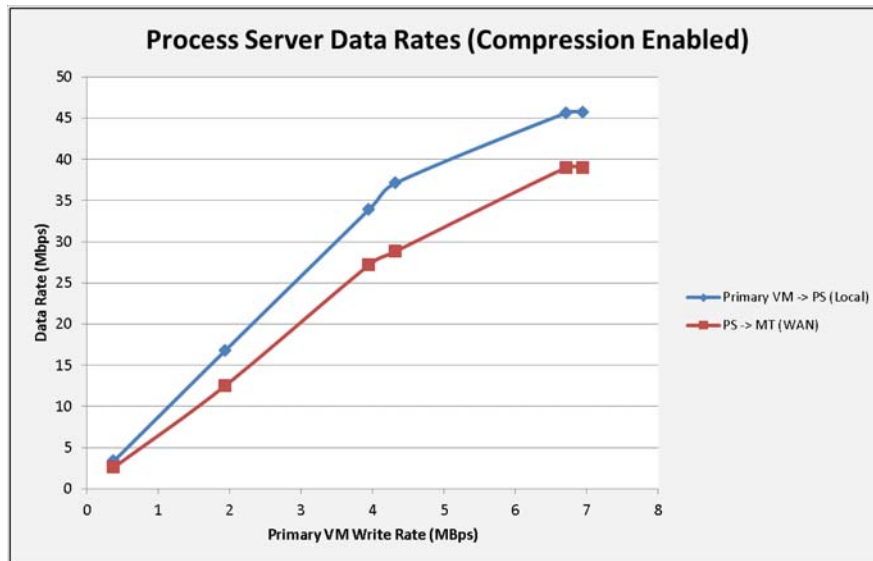


Figure A-6 Chart of CPU Utilization vs Primary VM Disk Write Rate (Compression Enabled)



In Figure A-7, the input and output data rates for the network interface of the process server is plotted against the primary VM disk write rate. The data difference between the two lines is the bandwidth savings on the WAN link due to the compression being applied by the process server.

Figure A-7 Chart of Bandwidth vs Primary VM Disk Write Rate (Compression Enabled)



## Comparison of Compression Enabled and Disabled

In this section, results from the compression disabled and compression enabled test cases will be directly compared. In Table A-1, we see that at the lower data change rates, there is a 26.95% CPU resource cost on the process server when compression is enabled. The cost approaches 200% at the higher data change rates.

**Table A-1 Comparison of Process Server CPU Costs with Compression Enabled**

<b>PS CPU MHz (Compression Disabled)</b>	<b>PS CPU MHz (Compression Enabled)</b>	<b>% CPU Cost to Enable Compression on PS</b>
167	212	26.95
230	459	99.57
315	787	149.84
330	843	155.45
338	1006	197.63
362	983	171.55

In [Table A-2](#), the CPU utilization of the MT on the SP side for both the compression disabled and compression enabled test cases is compared. The right-most column shows the % CPU resource cost when compression is used by the process server for data it sends to the MT. The MT consumes about the same amount of cycles at the lower data change rates (e.g., only about 1% CPU resource cost) when compression is enabled, but consumes 30-40% more CPU resources at the higher data rates. The maximum consumption was around 1200 MHz when compression was enabled and only around 890 MHz when compression was disabled.

**Table A-2 Comparison of Master Target CPU Costs with Compression Enabled**

<b>MT CPU MHz (Compression Disabled)</b>	<b>MT CPU MHz (Compression Enabled)</b>	<b>% CPU Cost to Enable Compression on PS</b>
365	369	1.10
505	580	14.85
739	893	20.84
788	944	19.80
865	1202	38.96
890	1148	28.99



# APPENDIX **B**

## Extending a Linux Volume

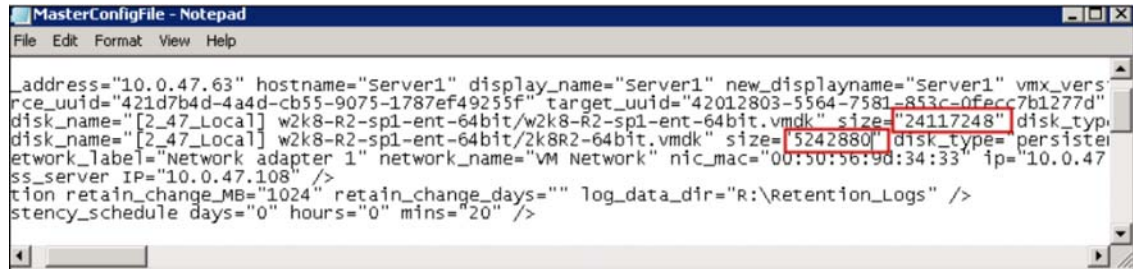
When you resize a protected volume (extend the volume beyond its original size) in the production site, the replication pair status changes to a Paused/VolumeResized State. The following procedure demonstrates how to bring the replication pair to diff sync state.

- Step 1** After extending disk on source VM, you should see that the replication pair status is Paused/Volume Resized with Resync required marked as YES and Alerts in CX UI.
- Step 2** Stop the VX Agent service on the MT server.
- Step 3** Note the SCSI IDs of the disks on the MT and detach the corresponding disk from the MT and extend the disk on DR-VM (make sure that extended size should not be less than extended size at production).
- Step 4** Attach the disk back to the MT after extension to the same SCSI ID that was before detaching from the MT.
- Step 5** Start the VX Agent service on the MT server.
- Step 6** Log in to vContinuum server and launch vContinuum. Click New Protection and do discovery for this VM to get new sizes in info.xml (Path: C:\Program Files (x86)\InMage Systems\vContinuum\Latest)
- Step 7** Check the disk sizes for extended disks: in Info.xml search source VM name and check size.

**Figure B-1** Check Disk Size in Info.xml File

```
MasterConfigFile - Notepad
File Edit Format View Help
_address="10.0.47.63" hostname="Server1" display_name="Server1" new_displayname="Server1" vmx_vers
rce_uuid="421d7b4d-4a4d-cb55-9075-1787ef49255f" target_uuid="42012803-5564-7581-853c-0fecc7b1277d"
disk_name="[2_47_Local] w2k8-R2-sp1-ent-64bit/w2k8-R2-sp1-ent-64bit.vmdk" size="23068672" disk_typ
disk_name="[2_47_Local] w2k8-R2-sp1-ent-64bit/2k8R2-64bit.vmdk" size="4194304" disk_type= persiste
etwork_label="Network adapter 1" network_name="VM Network" nic_mac="00:50:56:9d:34:33" ip="10.0.47
ss_server IP="10.0.47.108" />
tion retain_change_MB="1024" retain_change_days="" log_data_dir="R:\Retention_Logs" />
stency_schedule days="0" hours="0" mins="20" />
```

- Step 8** Open the MasterConfigFile from C:\Program Files (x86)\InMage Systems\vContinuum\Latest in the MT and search for the source VM name.

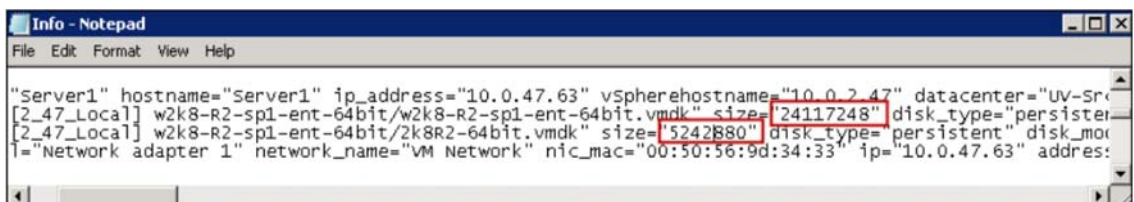
**Figure B-2 Find Primary VM Entry in MasterConfigFile**


```

_address="10.0.47.63" hostname="Server1" display_name="Server1" new_displayname="Server1" vmx_ver=
rce_uuid="421d7b4d-4a4d-cb55-9075-1787ef49255f" target_uuid="42012803-5564-7581-853c-0f6cc7b1277d"
disk_name="[2_47_Local] w2k8-R2-sp1-ent-64bit/w2k8-R2-sp1-ent-64bit.vmdk" size="24117248" disk_typ
disk_name="[2_47_Local] w2k8-R2-sp1-ent-64bit/2k8R2-64bit.vmdk" size="5242880" disk_type="persiste
etwork_label="Network adapter 1" network_name="VM Network" nic_mac="00:50:56:9d:34:33" ip="10.0.47
ss_server IP="10.0.47.108" />
tion retain_change_MB="1024" retain_change_days="" log_data_dir="R:\Retention_Logs" />
stency_schedule days="0" hours="0" mins="20" />

```

**Step 9** Modify the source VM disk size to the new size.

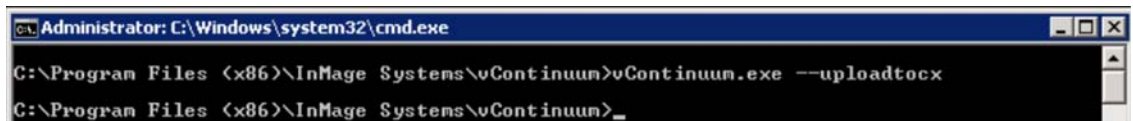
**Figure B-3 Modify Primary VM Volume Size**


```

"server1" hostname="Server1" ip_address="10.0.47.63" vspherehostname="10.0.2.47" datacenter="uv-sr
[2_47_Local] w2k8-R2-sp1-ent-64bit/w2k8-R2-sp1-ent-64bit.vmdk" size="24117248" disk_type="persiste
[2_47_Local] w2k8-R2-sp1-ent-64bit/2k8R2-64bit.vmdk" size="5242880" disk_type="persistent" disk_mor
l="Network adapter 1" network_name="VM Network" nic_mac="00:50:56:9d:34:33" ip="10.0.47.63" address:

```

**Step 10** Upload this file to CX by running the following command in MT: C:\Program Files (x86)\InMage Systems\vContinuum>vContinuum.exe -uploadtocx.

**Figure B-4 Upload File to CX Server**


```

C:\Program Files (x86)\InMage Systems\vContinuum>vContinuum.exe --uploadtocx
C:\Program Files (x86)\InMage Systems\vContinuum>_

```

**Step 11** On the primary Linux VM:

- Stop the service ( command: service vxagent stop).
- Issue the stop filtering on the protected device for which volume resize is required.
 

```
./inm_dmit --op=stopflt --src_vol=<device name>
```
- Do the resize of the device and the filesystem, if any.
- Issue the start filtering on this device.
 

```
./inm_dmit --op=startflt --src_vol=<device name>
```
- Start the service (service vxagent start). Step 12 Run the above command for all the protected disks that are extended. Step 13 Resume the replication through the CX UI and make sure to restart Resync for the resized volumes.



# APPENDIX **C**

## Acknowledgements

---

The following individuals have authored content in this CVD:

- Bharani Ramaswamy
- Xiao Hu Gao
- Sreenivasa Edula
- Jason Salle
- Jonathan Muslow
- Sunil Cherukuri
- Alex Foster





# APPENDIX **D**

## Glossary

---

### A

AMI	Amazon Machine Images
ASASM	Adaptive Security Appliance Services Module
ASR	(Cisco) Aggregation Services Router
AWS	Amazon Web Services

### B

BCO	BMC Capacity Optimization
BSA	BMC Server Automation
BNA	BMC Network Automation
BPPM	BMC ProactiveNet Performance Management

### C

CDP	Continuous Data Protection
CLM	BMC Cloud Lifecycle Management
CSA	ZenOSS Cloud Service Assurance
CX-CS	Configuration Server
CX-PS	Process Server

### D

DC	data center
DML	Definitive Media Library
DRaaS	Disaster Recovery as a Service
DRE	Data Redundancy Elimination

### E

EH	end host
F	
FC	Fibre Channel
FCoE	Fibre Channel over Ethernet

<b>G</b>	
GSLB	Global Site Load Balancing
GSS	Global Site Selector
<b>H</b>	
HBA	host bus adapter
<b>I</b>	
IaaS	Infrastructure as a Service
ICS	Integrated Compute Stack
IOPS	Input/Output Operations Per Second
<b>L</b>	
LC	linecard
<b>M</b>	
MDS	Multilayer Director Switch
MSP	Managed Services Provider
MT	master target
<b>N</b>	
NAS	Network-Attached Storage
NIC	network interface card
<b>O</b>	
OLTP	online transaction processing
OOB	out of band
OTV	Overlay Transport Virtualization
<b>P</b>	
P2V	Physical to Virtual Server
PIN	Place in the Network
<b>R</b>	
RAID	random array of integrated disks
RBAC	role-based access control
RD	route distinguisher
RDD	round-robin database
RESTful	Representational state transfer
RHI	Route Health Injection
RM	Resource Manager
RPO	Recovery Point Objective
RR	route reflector



**S**

SaaS	Software as a Service
SAN	Storage Area Network
si	swap in
SLB	Server Load Balancing
so	swap out
SOI	Service Offering Instance
SRM	VMware Site Recovery Manager

**T**

TFO	Transport Flow Optimization
-----	-----------------------------

**U**

UDP	User Datagram Protocol
-----	------------------------

**V**

V2V	(Zerto) Virtual to Virtual Server
VM	virtual machine
VMDC	Virtualized Multi-Tenant Data Center
VMDK	Virtual Machine Disk
VMNIC	VM Network Interface Card
VRF	virtual routing and forwarding
VSG	Virtual Security Gateway
vSLB	Virtual Server Load Balancer

**W**

WAAS	Wide Area Application Services
WAE	Wide Area Application Engines
WAFL	Write Anywhere File Layout
WLM	Workload Manager
WOC	WAN Optimization Controller
WWPN	World Wide Port Name

