Cisco Cloud Architecture for the Microsoft Cloud Platform: Backup as a Service Implementation Guide

Solutions

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.

Service Provider Segment
Cloud and Network Solutions
Cisco Cloud Architecture for the Microsoft Cloud Platform Solution

Cisco Cloud Architecture for the Microsoft Cloud Platform: Backup as a Service Implementation Guide

Part: CCAMCP1-1

© 2015 Cisco Systems, Inc. All rights reserved.

CONTENTS

Preface

This document details the overall proven architecture, design guidelines, and recommendations behind the Cisco Backup-as-a-Service (BaaS) architecture solution leveraging Cisco Cloud Architecture for Microsoft Cloud Platform (CCA-MCP) and Commvault® Simpana® software. The document describes the BaaS solution scope, approach, and architectural resources within the Commvault and CCA-MCP data center environments.

Purpose

This document provides design recommendations and describes the architecture of the Commvault BaaS solution within the Cisco CCA-MCP architecture. This document is based on the in-lab verification of the BaaS reference architecture solution using Commvault in a Cisco CCA-MCP test environment.

Scope

This solution design guide discusses multiple Cisco technologies and products that are part of the CCA-MCP architecture and also Commvault Simpana software and technologies.

Audience

This guide is intended for, but not limited to, system architects, network/compute/storage design engineers, systems engineers, field consultants, advanced services specialists, and customers wanting to understand how to design, operate, manage or consume a BaaS architecture leveraging Cisco and Commvault technologies. This guide assumes that the reader is familiar with the basic concepts of backup, restore, IP protocols, Quality of Service (QoS), High Availability (HA), Layer 4 (L4) - Layer 7 (L7) services, DC platforms and technologies, SAN, as well as Microsoft Hyper-V, VMware ESXi hypervisors, and OpenStack cloud virtualization. This guide also assumes that the reader is aware of general system requirements and has knowledge of Enterprise or Service Provider network and DC architectures, platforms, and virtualization technologies.

CHAPTER 1

Introduction

Data is the lifeblood of every organization today. Yet, before that data can become an asset, it must first be efficiently protected and managed. Cloud Providers delivering a comprehensive data protection strategy can help end customers avert hardware, logical and physical failures. With explosive data growth rates and mounting infrastructure costs, protecting data assets can be complex and difficult for enterprises to navigate, opening great opportunity for those Cloud Providers that are ready.

According to the ESG 2015 Spending intentions survey[++], "improving data backup and recovery" and "managing data growth" are two of the top three priorities for IT organizations within enterprises. Also, according to the survey, more than 66% of the survey respondents believe that there will be significant amount of increase in spending with Cloud Computing Services.

IDC agrees that offering Backup-as-a-Service (BaaS) to enterprises is a significant market opportunity for Cloud Providers. In fact, IDC estimates that the aggregate market for backup and recovery as a service to be $673.2 million in 2015, growing to $1.02 billion in 2018. This demonstrates high growth opportunities for Cloud Providers that can quickly deliver comprehensive data protection solutions and services with the rapid time-to-value enterprise customers demand. Refer to Examining Commvault's FY16 Key Initiatives.

This document encapsulates proven architecture, design guidelines, and recommendations for enabling Cloud Service Providers (CSPs) to launch BaaS solutions powered by Commvault Simpana software. Simpana software is the industry-leading solution for backup and recovery delivering a single platform for integrated data and information management. The solution delivers a truly holistic approach to protecting, managing and accessing data.

With BaaS, Service Providers can enhance their total addressable market with a proven architecture that increases productivity, reduce hardware and software costs, and mitigate risks. Ultimately, it will also reduce time-to-market and time-to-revenue while fueling revenue growth with a distinctive solution that offers a competitive, differentiated edge to capitalize on these dramatic Service Provider growth opportunities.

Business & Technology Use Cases

Cisco's BaaS reference architecture enables Cloud Service Providers (CSPs) to offer backup and recovery services to workloads outside of the CSP's management domain that are either customer premises environments or collocated environments. In addition, CSPs can offer data protection and data survivability services on workloads within the provider's Virtual Private Cloud (VPC) environment and management domain. The Backup & Recovery solutions described in this document are designed to provide a new set of related capabilities allowing CCA-MCP providers to enhance their addressable market, financial performance, and differentiation versus offering commoditized cloud solutions.

Use Case Overview

This document for Backup as a Service (BaaS) powered by Commvault covers the following three uses cases:

- **In-Cloud BaaS**—The In-Cloud BaaS offerings enable customers to leverage workloads within CSP's VPC environment and management domain. The Service Providers will enable customers to have backup capabilities for their IaaS work loads with various policy offerings. With this use case, the CSP shall have the ability to provide:
  1. Backup and recovery services of the workloads within the primary VPC in the customer environment.
  2. Replication of backup data to CSP's remote VPC to protect against site failures. Commvault software is deployed at each CCA CSP site and supports replication, providing local backup/restore and site survivability.
- **Remote BaaS**—The Remote BaaS offerings enable customers to perform backups at their data centers and replicate the backup data to CSP clouds remotely. The CSPs own or manage the remote site for providing recovery services to the customer. With this use case, the CSP shall have the ability to provide Backup and Recovery service for production virtual servers from a customer data center to the CSP VPC along with local (i.e. customer data center) backup/recovery capabilities. The CSP deploys the Commvault solution components at the customer site. Backups occur from the production servers to the Commvault servers at the customer site and gets replicated to Cisco CCA-MCP based CSP remote site.
- **Remote BaaS without Local Data Retention**—The Remote BaaS offerings without local retention enable customers to perform backups to Service Providers clouds located remotely. The CSPs own or manage the remote site for providing recovery services to the customer. With this use case, the CSP shall have the ability to provide Backup and Recovery service for production virtual servers from a customer data center to a CSP VPC. This service does not require local storage capacity at the enterprise data centers to host the backup data. The backup data gets directly copied over to the CCA-MCP based CSP's remote data center by leveraging Commvault software.

Cisco Solution Powered by Commvault Overview & Benefits

This section describes a high level overview of the benefits of using a Cisco Solution Powered by Commvault.

Cisco Cloud Architecture for Microsoft Cloud platform Overview

The CCA-MCP infrastructure is the foundation on which a variety of cloud services are offered. The base infrastructure consists of a set of data center devices that are setup and connected and configured prior to adding tenant services.

Service Providers build data centers using physical components to implement compute, storage and data center networking to create a pool of resources that are then used to offer services to tenants. Tenant services are offered using these physical resources, and provisioned and managed using automation software to enable consumption of these services. When tenants are on boarded, cloud containers are created from the pool of resources, to provide a slice of resources that include compute, storage and networking. This container is securely isolated from other tenants that are consuming similar services, thereby providing isolation for multi-tenant services.

The cloud services that are enabled in the CCA-MCP Solution are the Infrastructure as a service and Platform/Software as a service. Each of these services are described in a service configuration guide, and require the data center physical infrastructure to be built and the resource pools created and ready to onboard these services.

Figure 1-1 Cisco Cloud Architecture for the Microsoft Platform

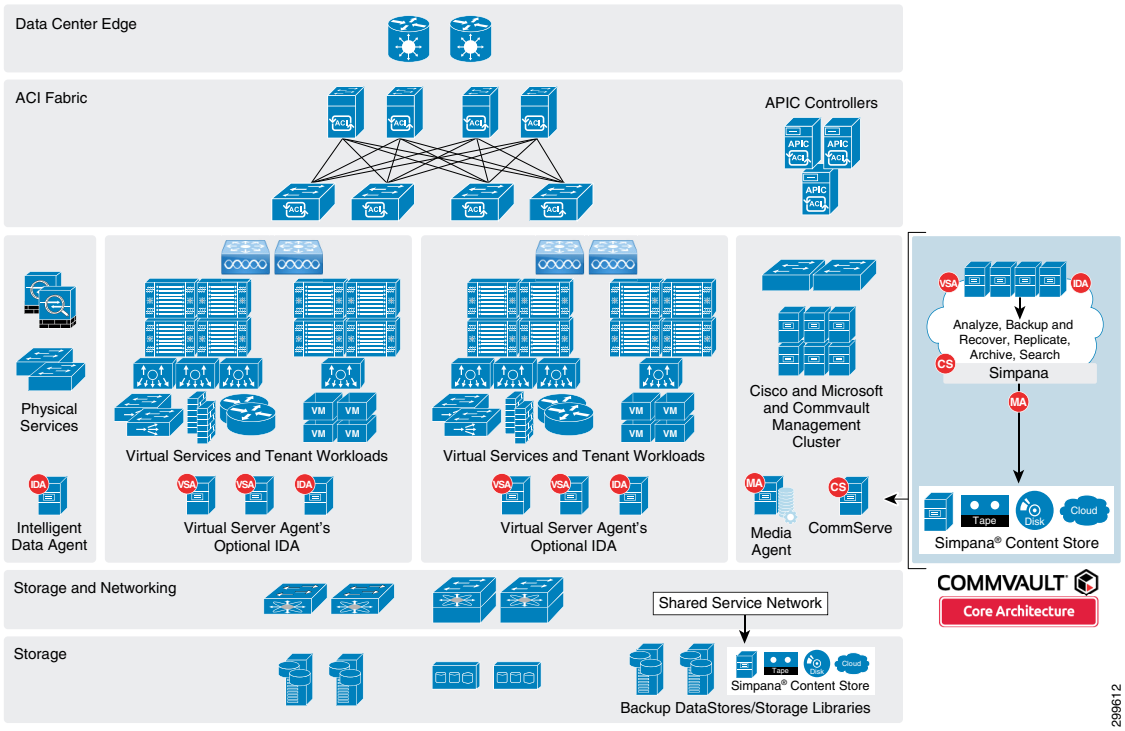

The architecture of this solution is built using a layered approach enabling a modular design. This enables one to deploy a scalable solution with expansion capability being added in modular units.

1. Data Center Network
2. Compute for Tenant workloads
3. Storage and SAN
4. Service Tiers and differentiated services
5. Cloud Management

Commvault Simpana Overview

Commvault Simpana software is built from the ground up on a single platform and unifying code base for integrated data and information management. All functions share the same DNA and back-end technologies to deliver the unparalleled advantages and benefits of a truly holistic approach to protecting, managing, and accessing data.

Simpana software offers investment protection with a core software platform that is flexible, modular, and ready to conquer new challenges as they emerge. Simpana software does the jobs of many point-level products, only better, more cost effectively, and much more simply.

It all starts with the single platform to power highly efficient cloud infrastructures. The Simpana platform contains individually licensable module to Analyze, Replicate, Protect, Archive, and Search your data. And because these modules share a common set of back-end services and advanced capabilities, they effortlessly talk to one another through the platform to solve a myriad of problems related to the storage and access of customer's data and information.

With Commvault Simpana software, CSPs can deliver "cloud scale" shared services to end customers with multi-tenancy capabilities—all from a single platform. With built-in automation and customized reporting, CSPs and customers can spend less time on routine administration, and more time delivering value to the business. Commvault cloud-ready software enables CSPs to accelerate their time to market, expand revenue opportunities and boost profitability of cloud and managed services.

CHAPTER 2

Use Cases

This section describes business use cases of Backup-as-a-Service (BaaS) for both Enterprise and Service Provider customers.

Business Use Cases

Continuing rapid data growth is driving enterprises to embrace new approaches for their backup and recovery needs. Cloud adoption offers many benefits technically and economically. This also provides opportunity for Cloud Service Providers to meet customer demands and further monetize their cloud investments with BaaS offerings.

For Enterprises

In most large organizations, the reality is that the level of IT complexity has funneled the majority of spending towards “keeping the lights on” versus funding innovation-driven efforts. The exponential growth of data and the underlying infrastructure to support that data is one of the primary sources of cost, operational complexity, and inflexibility.

Today, there are certain strategic drivers that must be tackled such as:

- How can customers leverage the cloud?
- How can customers harness data for greater insight or agility?
- How can customers secure and govern an ever-increasing data volume?

Cisco’s BaaS solution is a highly scalable yet operationally streamlined offering—built on Cisco Cloud Architecture for Microsoft Cloud Platform and Commvault’s Simpana software. Designed to deliver a flexible, SLA-driven approach to enterprise data management, this BaaS solution encompasses the Commvault Simpana suite of products. In addition, this solution also includes globally delivered services that ensure customer success in the deployment, ongoing management, and expansion of the platform.

For Cloud Service Providers

The business of delivering IT as a service—applications, infrastructure, or platform—is undergoing a significant transformation. Innovation cycles are faster and more compressed. Pricing and the ability to manage a large environment profitably is more challenging. And differentiating between large-scale, born-in-the-cloud providers and new entrants proves difficult, even for established brands.

How do you achieve revenue leverage when each new service requires incremental investment of infrastructure, tools, and operations staff? How do you compete in the market when the deployment, integration, training, and ongoing service delivery is measured in months or quarters?

Commvault Simpana delivers data efficiencies likened unto virtualization benefits.

Cisco's CCA-MCP architecture complements Commvault Simpana by providing a superior foundation for cloud computing by unifying computing, networking, storage, and management in a common platform designed to automate deployment and management across physical and virtual resources. CCA-MCP enables highly secure, multitenant deployments by embedding security at each layer of the data center.

Cisco's CCA-MCP architecture is also closely integrated with the service orchestration that provide configuration and provisioning automation.

This solution is a fully-featured CSP offering—built on Cisco CCA-MCP and Commvault Simpana platform—that enables rapid, profitable introduction of differentiated data management services.

As a storage infrastructure-agnostic platform designed to leverage existing storage investments, Commvault Simpana encompasses the complete Commvault product portfolio, covering a wide range of data protection and data archiving capabilities. In addition, the solution also includes globally delivered services that assist in the deployment and integration of Commvault into a service provider's back-office systems, including portals, billing systems, ticketing and service management.

Commvault Technology Use Cases

This solution enables the following use cases to provide backup and recovery services for enterprise workloads running at the customer's data centers and for the IaaS workloads running with in the service provider's cloud.

In-Cloud BaaS

As enterprises embrace cloud to deploy their critical applications, CSPs are under more pressure to offer enterprise-grade services in the cloud in an efficient manner. Backup and recovery is becoming a common value-added service for IaaS workloads.

This use case provides backup and recovery services to workloads within the CSP's VPC environment and management domain. BaaS will enable customers to have backup and recovery capabilities for their IaaS workloads with various policy offerings.

In-Cloud BaaS enables following functionality (Figure 2-2):

- Provides the ability to backup and recover the workloads within the primary VPC.
- Backup and Recovery functionality at the application, file, and VM level.
- Provides replication of backup data to a remote VPC for recovery against site failures.

In this solution for cloud backup of IaaS workloads, the CSP deploys the backup service consisting of multiple components from Commvault running on Cisco hardware, which will be used by all the cloud tenants enabling backup and recovery functionality. The service also provides backup data replication to a remote CSP data center offering site survivability as an option.

The solution provides tenants with full control by providing self-service capabilities to backup, restore, and monitor data.

Figure 2-1 In-Cloud BaaS Logical Diagram

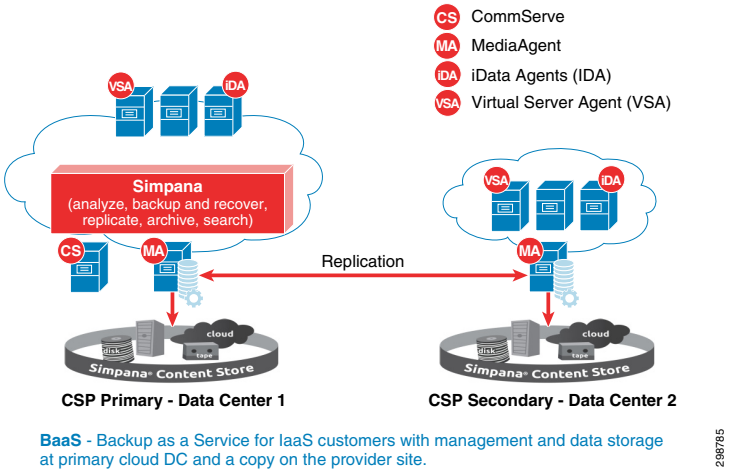

Remote BaaS

This use case enables customers to perform backups at their local data centers and also to replicate the backup data to the remote CSP's cloud without owning, managing, or incurring the expense of a remote site for recovery purposes. The local backup can be used for faster recovery, when needed.

The Remote BaaS service enables the following functionality (Figure 2-2):

- Backup and Recovery service for production physical and virtual servers from a customer data center to an CSP VPC along with local recovery capabilities.
- Backup and Recovery functionality at the application, file, and VM level.

Figure 2-2 Remote BaaS Logical Diagram

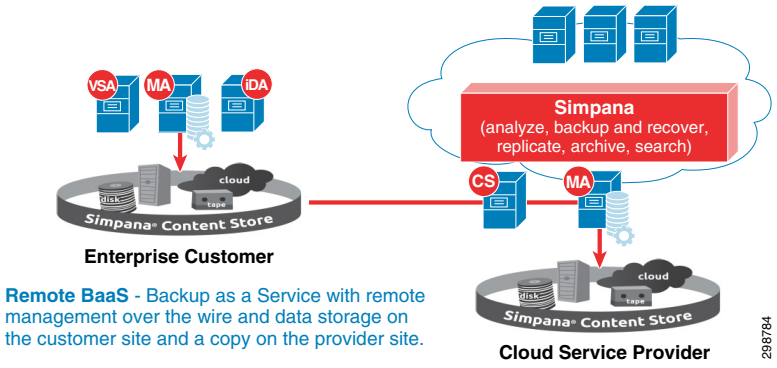

Remote BaaS Without Local Data Retention

This use case enables customers to backup their on-premises data to a remote CSP's cloud with zero capital investment and a low operational expense, without owning, managing, or incurring the expense of storage or a remote site for recovery purposes.

The Remote BaaS without Local Data Retention enables the following functionality (Figure 2-3):

- Secure backup storage capacity in the Cloud.
- Backup and Recovery service for production servers from a customer data center to an CSP VPC.
- Backup and Recovery functionality at the application, file, and VM level.

Figure 2-3 Remote BaaS without Local Retention

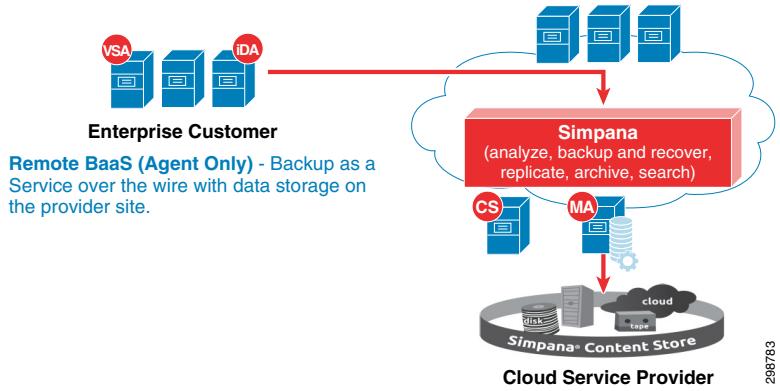

CHAPTER 3

Design Overview

The CCA-MCP BaaS solution enables CSPs to offer backup and recovery services to customers to protect their physical and virtual servers. These service offerings are enabled when a CSP deploys the CCA-MCP based infrastructure and then overlays the backup solutions from Cisco's partner Commvault.

Besides BaaS, Commvault technology enables several other Data management services including ediscovery, archiving, deduplication etc.

Commvault supports the most common hypervisors including HyperV, VMware and Openstack KVM. Besides BaaS, Commvault could be utilized for tenant onboarding or VM migration between hypervisor technologies.

BaaS Design/Architecture

This solution architecture utilizes the Commvault Simpana backup solution to enable BaaS capabilities on a CCA-MCP based cloud.

CCA-MCP solution provides the DC infrastructure architecture for cloud data centers to host and offer Infrastructure as a service, Platform as a service and Software as a service to customers.

CCA-MCP BaaS solution addresses the following design principles and architectural goals:

- Secure multi-tenancy
- Secure, modular, and highly available cloud
- Self Service
- Efficient data protection with deduplication and encryption
- Scalability

Figure 3-1 provides an overview on the joint CCA-MCP and Commvault BaaS architecture.

Figure 3-1 BaaS CCA-MCP Architecture

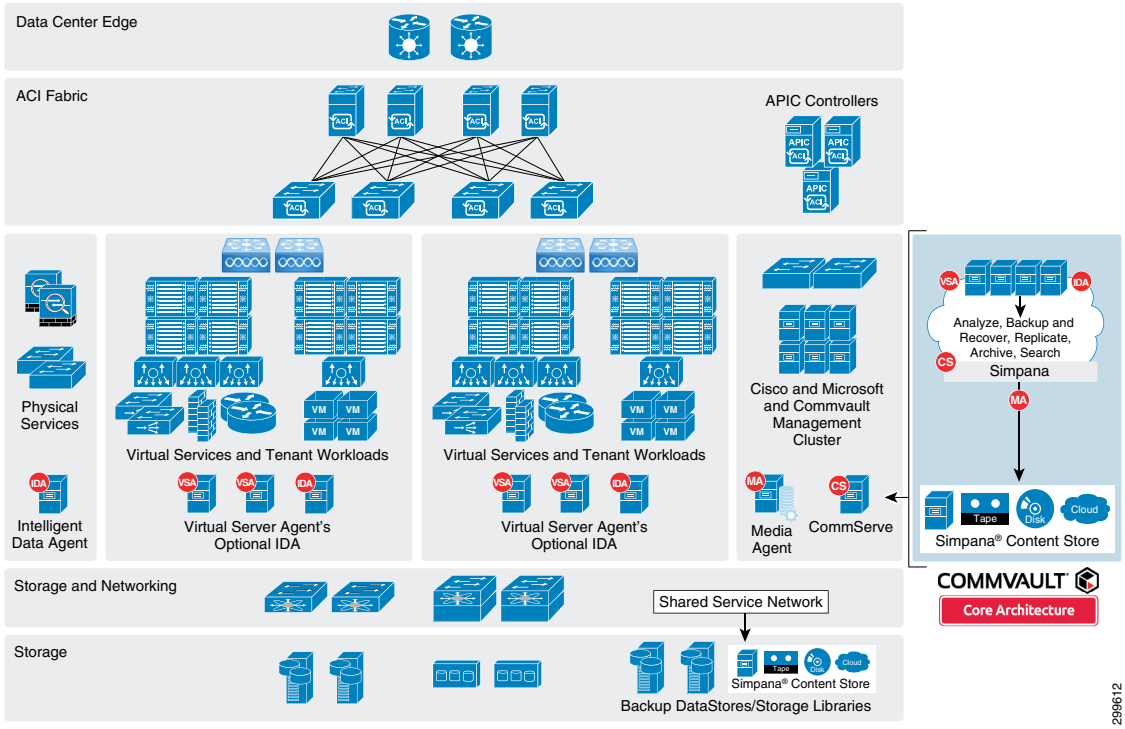

Commvault Simpana Solution Components

Commvault has developed a modular approach to sizing and building infrastructure components called Building Blocks (Figure 3-2). Each BaaS environment, called a CommCell, consists of the following Simpana® component roles:

- **CommServe (CS)**—The scheduling, job history, media management, and data management orchestrator.
- **MediaAgent (MA)**—The workhorse of the environment that manages deduplication database and the data transmission between clients and storage media.
- **Client**—The client owns the data to be managed, protected, and where the Intelligent Data Agent (iDA) is installed.
- **ContentStore**—All Simpana-managed data resides within the ContentStore - a secure, deduplicated virtual repository. Data is automatically stored and tiered according to user-defined policies, while a shared, intelligent index catalogs data versions and locations across snapshot, backup and archive copies to find data when users need it.
- **Intelligent Data Agent (iDA)**—Provides unified protection and recovery for most common operating systems, databases, and applications. This is installed on the client server or VM.
- **Virtual Server iDataAgent (VSA)**—Provides a unified protection and recovery vehicle for all virtual machine data in your virtual environments. In addition to complete protection of entire virtual machines for disaster recovery, the Virtual Server Agent provides more granular backup and recovery options. This is installed on the Hyper-V servers.

- **Web Proxy (PXY)**—The Simpana® Web Proxy role is typically deployed in a DMZ and serves the web console interface as well as being enabled as a communications proxy. It provides a security separation layer preventing direct connectivity and the core service infrastructure from clients and web console users on public or untrusted networks.
- **Web Console (WC)**—A web-based application that allows end-users to manage their file data. The console behaves as a self-service application allowing you to perform backup, restore, download and other operations.
- **Reporting Server (RS)**—An automated reporting system that helps you to monitor all of the CommCell computers in your organization on a central reporting Web site. Reports include metrics such as the number of CommCell computers installed with a particular software version. Reports also contain information about individual CommCells, such as SLA performance, job errors, and deduplication rates.

Figure 3-2 Simpana Components

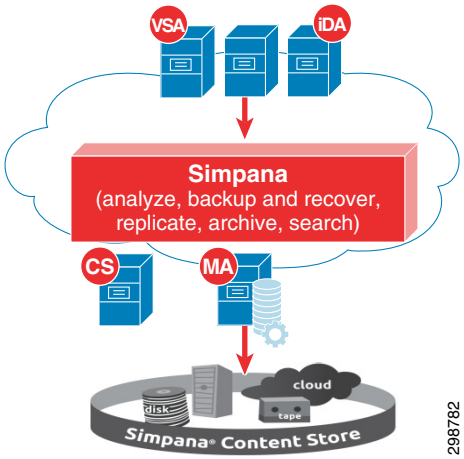

A logical representation of the Cisco BaaS solution/architecture resources for a single Cloud Service Provider is shown in Figure 3-3.

Figure 3-3 Cisco BaaS Solution Logical Topology

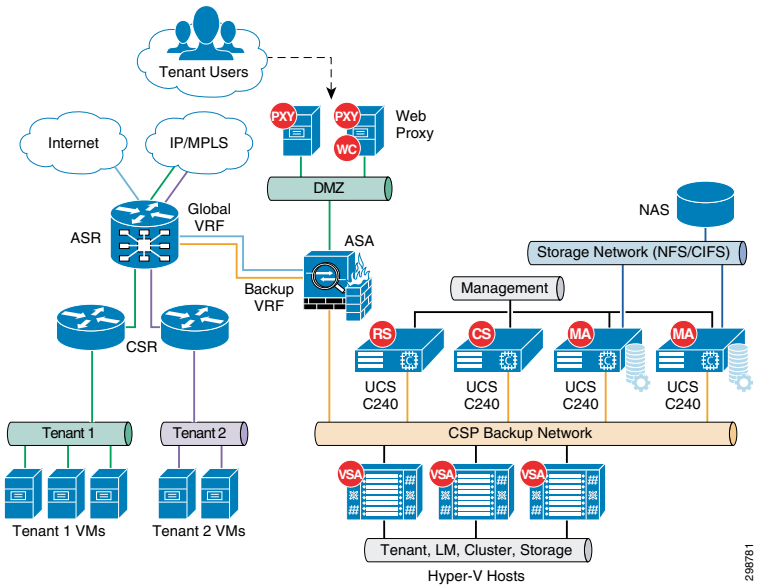

A typical CCA-MCP CSP environment will have tenants running applications—for example, Oracle, Exchange, SQL—in physical or virtual environments. Those apps are typically accessed by clients via an IP network, and leveraging production storage (either SAN, NAS, or even DAS).

Commvault Simpana platform is deployed non-disruptively, side-by-side with production in a dedicated and secure backup network as shown in Figure 3-3. This network needs to be made available to the tenants for them to consume the shared multitenant BaaS. The architecture also includes a DMZ network to host the Commvault web proxy servers, the web proxy servers are used by the CSP's who do not want to expose any of the Commvault components to the end customers directly. With the deployment of these proxies, tenants within the cloud and the remote enterprise customers can access the proxies and still have all the Self Service backup and recovery functionality without having direct access to the CS and MA servers.

Commvault allows customers to use any kind of storage as the content store for the backup data. Within the solution we have included the Cisco C3160 servers as the MA servers with built-in storage capacity and also have included traditional NAS storage with Cisco C240 servers deployed as the MA servers accessing the shared NAS storage.

This solution supports three use cases:

- In Cloud BaaS for workloads running in the cloud
- Remote BaaS for workloads running on the customer premises (local retention)
- Remote BaaS for workloads running in the customer premises without local retention

In-Cloud BaaS

The In-Cloud BaaS use case uses the CCA-MCP Architecture as the CSP cloud. The CSP can offer backup as a service to the tenant workloads running within the CCA-MCP based CSP cloud. Figure 3-4 shows the high-level architecture of In-Cloud BaaS for IaaS workloads between two CSP cloud data centers.

Figure 3-4 In-Cloud BaaS Architecture

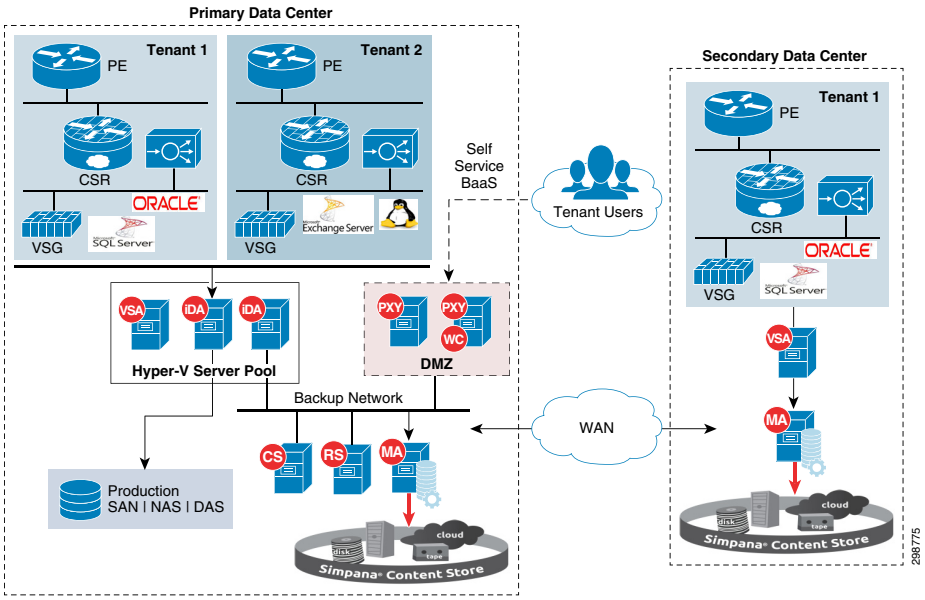

In a multitenant environment, each customer is mapped as a separate CCA-MCP tenant where the necessary network security is provided and traffic segregation is maintained.

All the tenants within the CCA cloud share the multitenant enabled Simpana deployment, including backup storage attached to the Commvault Media Agent Servers to store the backup data. The tenants will have the ability to also replicate the backup data to remote Media Agent Servers in a Secondary CSP DC within the cloud, when they choose service offering which includes data replication (Figure 3-5).

This provides multitenant-enabled In-Cloud BaaS to the tenants and also enables site survivability by providing offsite data backup.

The IaaS workloads within the CCA-MCP cloud are deployed on a shared Hyper-V cluster. The Hyper-V servers hosting the tenant workloads have a Commvault VSA installed, which is used to back up the tenant workloads with the help of Hyper-V snapshots.

The tenant separation and isolation is provided according to CCA-MCP best practices using network containers. The CCA-MCP Storage Architecture remains unchanged in the BaaS solution. There are a few key additions to the CCA-MCP architecture, including a dedicated backup network used for deploying Commvault components which helps segregate the backup traffic from the tenant production traffic.

Another addition to the CCA-MCP architecture is the extension of the Backup network across the Service Provider data centers. This network will be used for carrying the replication traffic between the Commvault Media Agents, as well as communication between the Commvault management components.

Figure 3-5 Architecture Across CSP Data Centers

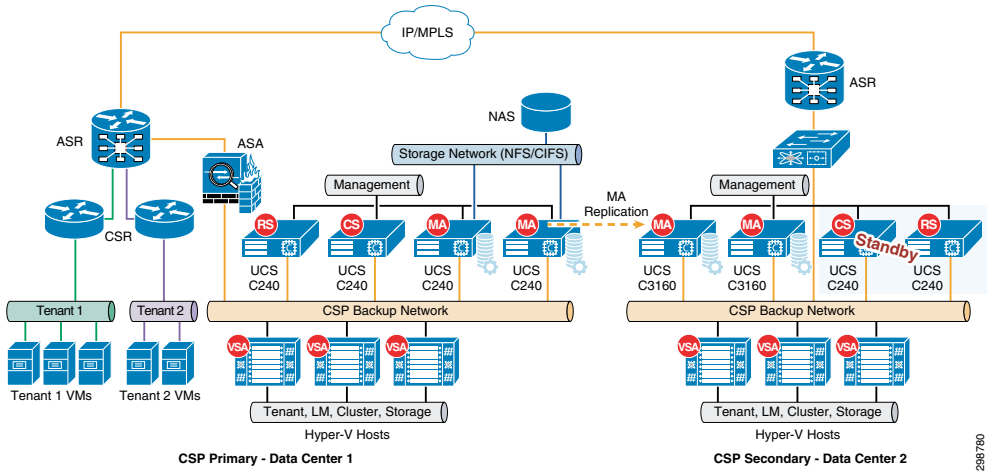

The data replicated to the secondary CSP data center can be used to support scenarios such as primary data center failure, spinning up the VMs along with production data to support additional use cases such as Test/Dev and analytics, etc.

The Secondary data center in the architecture will be used to host the standby Commvault components such as the CommServe Management and Reporting servers, which can be used if the primary servers are unavailable.

Remote BaaS

The Remote BaaS use case allows Enterprises running applications at their local data centers to backup data on-site and to also leverage a CCA-MCP CSP cloud to save their backup data remotely.

This Remote BaaS architecture includes the Cisco UCS C240 servers used as the Media Agents, these servers will be deployed by the CSP at the Enterprise customers' data centers. The MA servers can be built based on the amount of production data that needs to be backed up by including the capacity and the compute resources accordingly.

Any other UCS C-Series servers can also be used as the MA server based on the requirement. The WAN connectivity from the Enterprise data centers is being provided by a Cisco Cloud Services Router (CSR 1000V) within the solution, it allows the enterprises to extend a WAN to off-premises clouds and cloud service providers to offer enterprise-class networking services to their tenants.

The CSP can include the CSR 1000V as a VM on the MA server to provide network connectivity and the backup software in one single device to the Enterprise customers, which eases solution deployment and new customer onboarding for CSPs.

Solution management is provided by Commvault components deployed within the Cloud. The Enterprise customers can access the portal running from the cloud to discover the productions servers and initiate a backup or perform a recovery of the workloads either on-premise or within the Cloud VPC.

The on-premises MA Servers provides local backup capabilities to the Enterprise for faster recovery with a local copy and minimizes the backup window leveraging LAN throughput.

The replicated traffic between the MA servers from the enterprise data center to the cloud CSP’s cloud is deduplicated unique data, minimizing WAN bandwidth and cloud storage requirements. This offsite copy can be used to restore the data back to the original customer's premises or can be recovered within the CCA cloud anywhere the customer might need it.

Data-in-transit encryption is necessary to keep the backup data secure while in transit. Establishing an IPsec tunnel between the customer's data center and the CSP's data center using the CSR 1000V routers as the tunnel endpoints can enable this data encryption. Commvault is also capable of encrypting the replicated data between the source and destination locations, that can be optionally implemented if customers prefer.

Figure 3-6 covers the high-level architecture of Remote BaaS for workloads hosted at customer data centers.

Figure 3-6 Remote BaaS Architecture

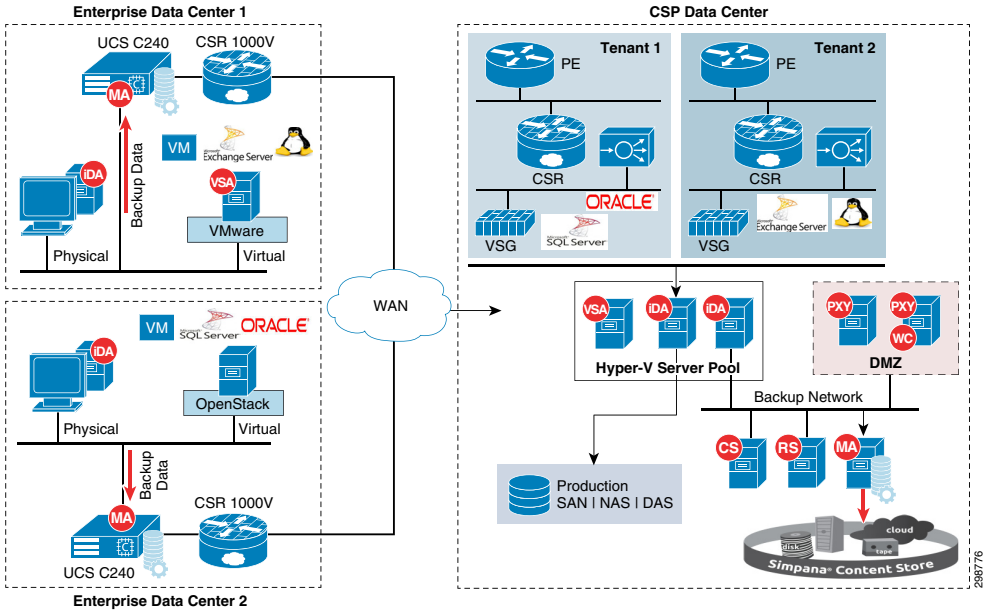

Remote BaaS without Local Retention

The Remote BaaS without Local Retention use case allows Enterprises running applications at their local data centers to remotely backup data to a CCA CSP cloud.

This is similar to the Remote BaaS use case, but will not offer the local backup and recovery functionality. Customers will have an option of recovering the data from the Cloud back to their data centers or recover the workloads in the Cloud VPC.

The MA servers with in the cloud will be used to host the backup data; there is no requirement to have the MA servers deployed locally at the Enterprise data centers.

Figure 3-7 shows the high-level architecture of Remote BaaS without local retention for workloads hosted at customer's data centers.

Figure 3-7 Remote BaaS without Local Retention

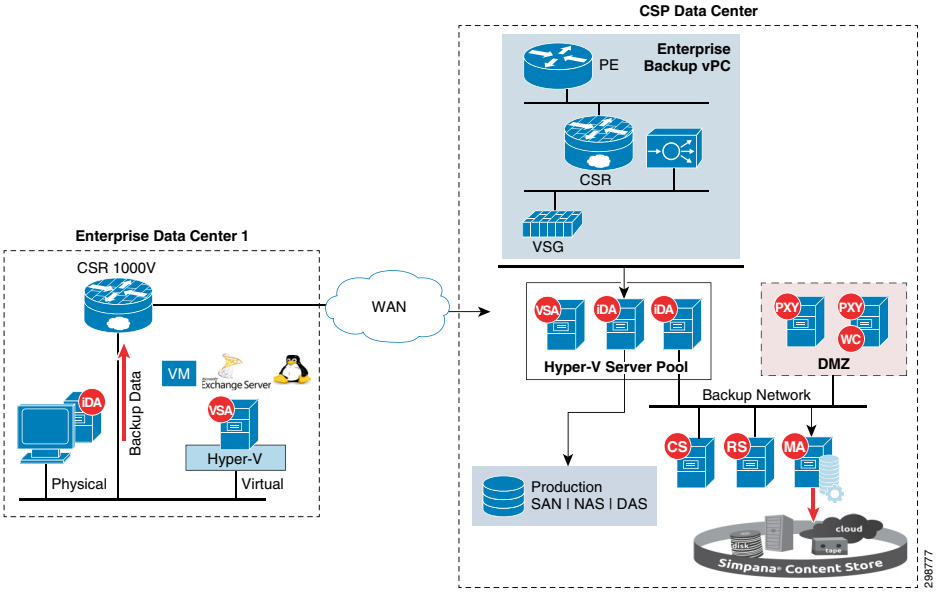

WAN Connectivity

There are multiple connectivity options for tenants and end-users to connect to their in-cloud resources. Some of these mechanisms include:

- L3 Connectivity
  - L3VPN (MPLS) based, where the tenant sites connect to the Cloud DC through MPLS-VPN services
  - IP (Internet) based, where clients access cloud resources directly across the Internet
- L2 Connectivity
  - Layer-2 (VLAN-extension) based, where the tenant sites connect to the cloud DC through L2VPN services like VPLS and EoMPLS

The BaaS solution will support any of these interconnect mechanisms for connecting enterprise DC to the CCA-MCP based provider cloud.

This release of BaaS implements L3-based connectivity with L3VPN and Internet-based connectivity (Figure 3-8).

Connectivity to the CSP cloud can be enabled by the CSP by deploying a Cisco CSR 1000V or Cisco ISR G2. The ISR G2 is the second generation Integrated Services router that is designed to meet the application demands of today's medium-sized branches and to evolve to cloud-based services. They deliver virtualized applications and highly secure collaboration through widest array of WAN connectivity at high performance that offers concurrent services.

The Cisco CSR 1000V Cloud Services Router provides a cloud-based virtual router that is deployed on a virtual machine (VM) instance on x86 server hardware. The Cisco CSR 1000V router is a virtual platform that provides selected Cisco IOS XE security and switching features on a virtualization platform.

When the Cisco CSR 1000V virtual IOS XE software is deployed on a VM, the Cisco IOS software functions just as if it were deployed on a traditional Cisco hardware platform. You can configure different features depending on the supported Cisco IOS XE software image. The Cisco CSR 1000V supports a subset of Cisco IOS XE software features and technologies.

The Cisco CSR 1000V provides secure connectivity from the enterprise premise (such as a branch office or data center) to the public or private cloud.

The intent of the CCA-MCP BaaS solution is to keep the Enterprise DC architecture generic so as to provide the greatest coverage for the CSP customer base.

The Service Provider cloud data center is based on the Cisco's CCA-MCP solution. The infrastructures built using this architecture hosts tenant IaaS workloads within the network containers, which provide security and enable multitenancy.

Customer data centers can be connected to the CSP's cloud via MPLS-VPN or Internet to the respective tenant container. In the case of MPLS-based connectivity, the ASR 9000 or ASR1000 Series Routers are used as MPLS Provider Edge (PE) routers in the data center, providing L3VPN connectivity to the provider IP/MPLS WAN network. Tenants within the CCA-MCP Architecture get their own virtual service appliances as part of the network container for their IaaS workloads. VLANs are used for connecting the tenant routing instance (CSR 1000V) to the tenant Virtual Routing and Forwarding (VRF) instances on the ASR 9000 WAN router.

Figure 3-8 L3VPN WAN Connectivity

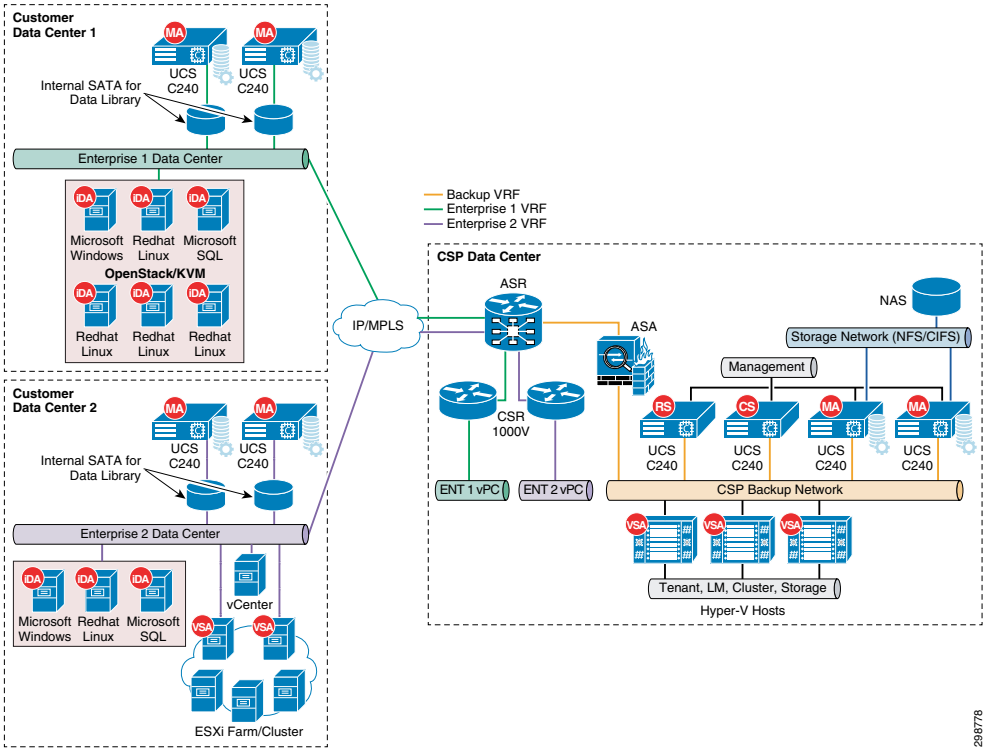

The BaaS Architecture includes a dedicated Backup VRF on the ASR 9000 WAN router and have all the Commvault components deployed behind an ASA firewall connected to this VRF. In the CCA-MCP Solution, there is a ASA firewall cluster, and used in multi-context mode – one context can be used for BaaS or an existing service context may be used, depending on the security and operational needs of the deployment.

The solution requires bi-directional communication between the customer VRF's and the Backup Services VRF, with various components communicating with each other across the data centers that needs to carry management and replication traffic. To enable this communication the Service provider has to Import/Export the BaaS customer VRF routes into the Backup VRF and vice versa.

Figure 3-9 shows the Internet based connectivity scenario were the Enterprise customers are connected to the CCA-MCP CSP's data center via Internet. The ASR 9000 PE WAN router is also connected to the Internet, via either global table or a Internet VRF. A shared VLAN is used for access to the global/Internet routing space of the ASR 9000, the ASA firewall gets connected to the ASR 9000 on this shared VLAN for access to the global/Internet routing space of the ASR 9000.

In this scenario, customers from remote data centers will have access only to the Commvault proxy servers that are placed in the backup DMZ network, this helps service providers to protect the Commvault components and not expose them with public IP addresses. The proxy servers accept incoming tunnel connections from internal servers (CS, MA) in the cloud and from customers' sites, these connections are authenticated and encrypted. Knowing source and destination client names for every tunneled control/data connection, proxy works as a PBX forwarding this control/data traffic between established tunnels.

Static NAT is used to dynamically translate the private IP addresses of the Commvault proxies to public IP addresses, translating the private addresses in the internal DMZ private network into legal, routable addresses that can be used on the public Internet.

Figure 3-9 Enterprise to Cloud Service Provider via Internet

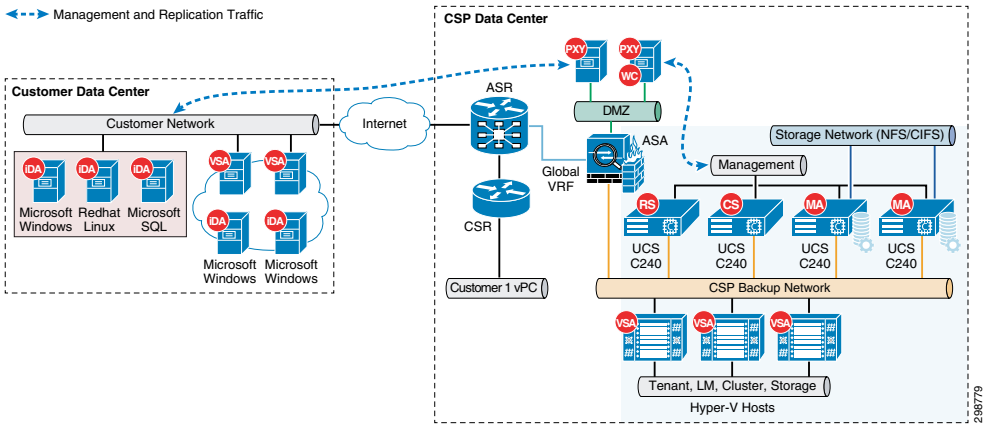

Cisco Storage Server as Converged MA

The Cisco UCS Storage Rack Server is an advanced, modular, high-storage-density rack server targeted at storage-driven use cases. Combining industry-leading performance and scalability, the UCS C3160 directly targets environments deploying any software-defined and distributed storage environments. The rack server offers the highest levels of drive density.

The Cisco UCS Storage Server offers following features and capabilities:

- Enterprise-class redundancy with full featured Redundant Array of Independent Disks (RAID) plus Just a Bunch of Disks (JBOD)
- Standalone management interface (Cisco Integrated Management Controller)
- No data migration required when replacing or upgrading server nodes
- No need for extended depth racks

The following are the specifications at a glance:

- High-density, bare-metal, x86-based enterprise storage server
- Supports up to 360 TB of modular storage capacity
- Optimized for high throughput performance, high capacity, and small footprint
- Enterprise-class redundancy with full featured RAID plus JBOD
- Standalone management interface (CIMC)
- Up to 256 GB of memory
- Up to 62 drive bays
- Up to 4 GB of RAID cache

Cisco Storage Server is a 4U chassis, designed to operate both in standalone environments and as part of the Cisco Unified Computing System. The system is targeted at the service provider, storage server, and big data markets.

The chassis can accommodate 1 or 2 Network IO Modules, 1 or 2 server modules, 56 3.5” drives, and 4 PSUs. One of the server slots can be used by a storage expansion module for an additional 4 3.5” drives. The server modules can also accommodate 2 SSDs for internal storage dedicated to that module. SAS expanders are configurable to assign the 3.5” drives to individual server modules.

Cisco Storage Server will be delivered in two SKUs, referred to as C3160 and C3260.

C3160 is an accelerated, TTM-driven program. It has the following characteristics:

- Support for a single server only
- All storage is assigned to the single server
- Operates in standalone mode only (no UCSM support)
- Uses mLOM-based NIOMs
- Functionally behaves like a traditional C-Series server

C3260 is a feature driven program with the following characteristics:

- Supports single or dual server
- Individual drives are assignable to either server
- Can operate in standalone or UCSM mode
- Uses Cisco 3rd Gen VIC 1300 with 40Gbps support
- Adds chassis-level functionality in standalone mode
- All shared components (for example, storage, fans, PSUs) are configured at a chassis-level scope
- Server-specific components (for example, boot order, KVM) are managed at a server-level scope

Figure 3-10 shows the modular Architecture of the UCS C3160 Server.

Figure 3-10 Architecture of UCS C3160 Server

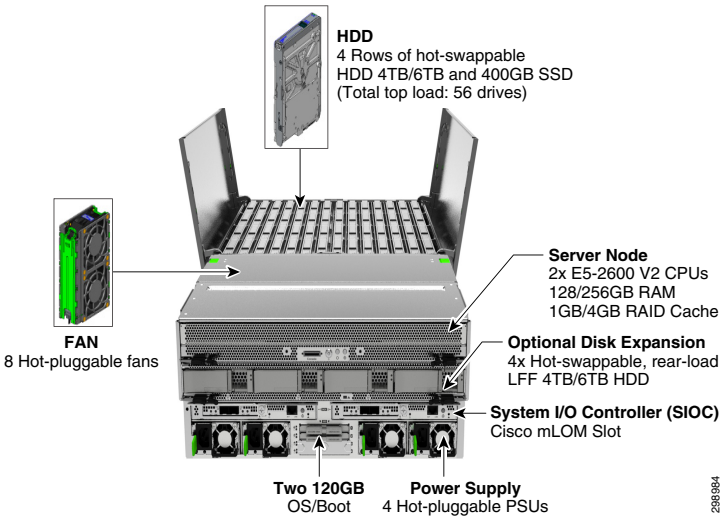

Commvault Architecture/Design Considerations

Commvault has developed easy to consume BaaS Sizing Guidelines based on use cases articulated in this document. The BaaS Sizing Guidelines can be leveraged in a highly repeatable fashion as capacity and performance thresholds are achieved within an environment.

The initial and predicted growth of the use case and service will dictate which scale model to use to meet capacity and demand. The BaaS Sizing Guidelines offerings are segmented into three types depending on projected size of capacity (i.e. size of data that needs to be protected) service uptake ranges over 12 months:

1. **Small**—50-150TB
2. **Medium**—151-500TB
3. **Large**—501-1PB+

Within each of the BaaS Sizing segments, Commvault has defined (3) scale points to define Simpana® role requirements to providing a cost efficient roadmap to service deployment and growth. For each defined scale point, Commvault defines the total Simpana® role requirement at that scale point. As an example, a Small configuration will require initially at 50TB (2) Large MediaAgents. Furthermore, the Small will require a total of (4) Large MediaAgents at 150TB. Therefore, two (2) additional Large MediaAgents are required to scale from 50TB to 150TB.

A detailed depiction of the BaaS-Sizing Guidelines are shown in Table 3-1.

Table 3-1 Backup as a Service—Sizing Guide

Commvault SP Level	Small	Medium	Large
Forecasted FETB Size 12mo	50-150TB	151-500TB	501-1PB+
CommServe	Datacenter Size	Enterprise Size	Enterprise Size
Clients	Up to 2,500 Servers	Up to 10,000 Servers	Up to 10,000 Servers
Jobs/24hr	100,000	200,000	200,000
Concurrent Jobs	101-300	301 to 1,000	301 to 1,000

Table 3-1 Backup as a Service—Sizing Guide (continued)

Commvault SP Level	Small			Medium			Large		
Concurrent Throughput	4-16TB/hr			8-40TB/hr			20-72TB/hr		
Java Console Connections	Up to 30 Concurrent			Up to 50 Concurrent			Up to 50 Concurrent		
Web Console Connections	Up to 1,000 Concurrent			Up to 1,600 Concurrent			Up to 2,800 Concurrent		
Scaling Points	**50TB**	**100TB**	**150TB**	**151**	**350TB**	**500TB**	**501TB**	**750TB**	**1PB**
Concurrent Streams	100	200	200	200	300	500	500	700	900
Media Agents	(2) L	(4) L	(4) L	(4) L	(6) L	(10) L	(10) L	(14) L	(18) L
Proxy Node (optional)	1	1	3	2	3	6	6	12	18
Web Console Nodes	1	1	1	1	1	2	1	1	2
Reporting Server	Enterprise Size			Enterprise Size			Enterprise Size		

Table Key

- **L**—Large MediaAgent Size

Assumptions

- Configuration uses Commvault Building Blocks and meets Best Practices.
- FusionIO card (or equivalent) is being used and hosting multiple Deduplication Databases per MediaAgent.
- Deduplication Databases are configured in pairs for scale and redundancy.
- When reaching 3,600 clients receives warning and requires Commvault review before exceeding 4,000 clients.
- Client data profile is on average 100-250GB.
- Micro & Big Capacity Small Data profiles will impact number of jobs.
- Service providers schedule jobs.
- Metrics Enterprise Reporting server will be running on a separate server.
- Proxy nodes optional to suit network topology requirement or increase availability or minimum 2 of each role type.

Note Commvault Architecture and design guidelines represent current Commvault views on this topic as of the date of publication and is subject to change at any time without notice.

Commvault Multi-Tenancy

This section will describe how Commvault achieves secure multi-tenancy within a single CommCell environment.

Taxonomy

When speaking with Cloud Service Providers, multi-tenancy is an extremely important and sought after feature. Simply put, Commvault defines multi-tenancy as the secure separation and management of shared resources between defined entities. When dissecting multi-tenancy for data management, Commvault believes there are eight areas that make a solution multi-tenant:

- Management Server
- User Management
- Policies
- Data Mover
- Network (Proxies, Firewall, & Bandwidth)
- Security
- Reporting
- Graphical User Interface (GUI)

Commvault Simpana® is the only data management software that provides multi-tenancy for each area in a single platform. The following sections detail Simpana multi-tenancy features specific for CSPs.

Management Server

In Commvault Simpana® software the CommServe is the central management server. Simpana® can isolate and logically manage tenants separately within the same CommServe regardless whether the configuration of underlying components are shared or dedicated. For example some tenants may require having dedicated data movers (known as MediaAgents) or storage, whereas other tenants it may be perfectly acceptable to utilize a shared environment. Simpana CommServe can manage any of the examples referenced above within a single CommServe. Meaning a service provider does not have to deploy and manage multiple CommServes to satisfy most tenants' needs. Service providers will only have to install multiple CommServes if the tenant requires a completely physically isolated data management instance or has to manage more than 20,000 clients.

User Management

At Simpana's core multi-tenancy is enabled through its robust implementation of Role Based Access Control (RBAC) as part of Simpana's overall security framework. Simply put, Simpana® can have multiple users accessing the platform without any knowledge of each other or access to their data. Managing individual user permissions may be acceptable for some individual enterprises. However, at the service provider level this would quickly become unmanageable. Therefore, Simpana® has created the concept of roles with a common set of attributes and permissions. Service providers will create two categories of roles within Simpana, which are described as follows:

- **Cloud Service Provider Roles**—Reserved for service provider administrative staff and created to manage the overall service across all customers.
- **Customer or Entity Roles**—Designated to consumers of the service with common local data permission, however restricted to their own data.

Typical roles restrict functional tasks such as backup and restore (including locations), as well as who can access report or delete protected data. For a full list capability and permitted actions (otherwise known as permissions) descriptions, refer to Simpana® User Capabilities and Permitted Actions or Simpana® Capabilities and Permitted Actions by Feature.

Clients

The end-user controlled laptops, servers, or virtual machines that require protection are designated as clients within Simpana®. Agents are modules installed on clients to protect a specific type of data such as the file system, database, or application. During agent installation, each agent is issued a SSL certificate by the CommServe. This provides secure authentication and agent identification to prevent possible data breaches through spoofing. It does this by using more common username-password agent authentication techniques by competitive solutions.

Client Computer Groups

The power of Client Computer Groups provides the service provider administrator the flexibility to group resources by a multitude of parameters. Groups can be automatically updated as new or existing clients meet the designated criteria (known as Smart Client Computer Groups). Typical Client Computer Group use cases for service providers are:

- Customers
- Service Plan
- Waiting Room for new, but unauthorized client
- Hostname
- Operating System
- Network configuration (IP address or firewall rules)
- Installed Application or Agent

Simpana® can reduce the administration of Client Groups through a rule based automatic assignment approach called Smart Client Computer Groups. To view a full listing of rules that can be set for Smart Computer Group, refer to Simpana® Smart Client Groups.

Policies

Managing your data management environment at the individual user or single tenant level would quickly become unmanageable, therefore using a policy-based approach is critical for scaling. Simpana® has two types of policies that can be applied with fine granularity or broadly for rapid changes:

- **Storage Policies**—define where data should be protected, how many copies and for how long
- **Schedule Policies**—when data should be protected

Storage Policy

Storage Policy directs data and its secondary copies to a specified storage target, level of protection, and defines the retention period. The power of Storage Policies can group or segment data in a public or private categories, which provides flexibility depending on Service Offering defined to tenants. Through the use of Storage Policies some tenants can share a storage target to optimize service cost, whereas some tenants may have a dedicated data target per tenant for privacy requirements. Both examples can be provided within a single instance of the Simpana® data management platform. Commvault significantly differentiates itself as a multi-tenant leader because of the granularity that a Storage Policy can be associated:

- Tenant
- Sub-tenant

- Service plan
- Application group
- Data type

Each of the Storage Policy association examples can be specified and applied at the Client Computer level (usually a tenant), which reduces the overall administration. Storage Policies can even be associated to a sub-client (more commonly known as a partial set of data) covering those "one-off" customer requests.

Schedule Policy

Maximizing resource utilization is important to service providers and Commvault can intelligently schedule jobs to keep resource at top utilization to achieve data protection goals. Commvault provides the ability to set the timing of a job to start, which in most cases is a data protection job (such as backup or archiving). Similar to Storage Policies, Schedule Policies can be associated at a very granular level depending on the service provider's offerings and tenant's demands:

- Tenant
- Sub-tenant
- Service plan
- Application group
- Data type

Highlighted below are some common service provider examples of schedule policies:

- **Time Slot**—a specified window of time when a job must start
- **Start Time**—an exact time for the job to begin

Commonly, tenants will request a specific start time (or window) when jobs should start. Schedule policies provide the facility for service providers to offer that option to their tenants, which can be a service uplift (ie. chargeable) or value add to a tenant.

Data Mover (a.k.a media agent)

Within Simpana® software the data mover is known as the Media Agent, where clients send their data and the Media Agent moves it to the storage target. The Storage Policies direct the Media Agent to which storage target should be used per job, which can be shared among many tenants or dedicated to a single tenant. To provide the service provider with the highest level granularity and flexibility, Media Agents can have multiple Storage Polices running simultaneously with almost any variety of configurations. Simpana Media Agents can be configured many ways for multi-tenancy and the following are the most common:

- **Private**—Dedicated hardware with the Media Agent dedicated to a single tenant, which can have a dedicated or shared CommServe managing it.
- **Multi-instance**—Single physical hardware with multiple images of the Media Agent software running at once. The service provider can satisfy private requirements and drive up hardware utilization.
- **Public—**Shared among multiple tenants.

Note Simpana deduplication database (DDB) can be isolated to a single tenant or can be shared among multiple tenants in a Public configuration.

Networking

Simpana® has extensive networking configuration options to best meet a service provider's needs as shown in Figure 3-11.

Figure 3-11 Simpana Network Configuration Options

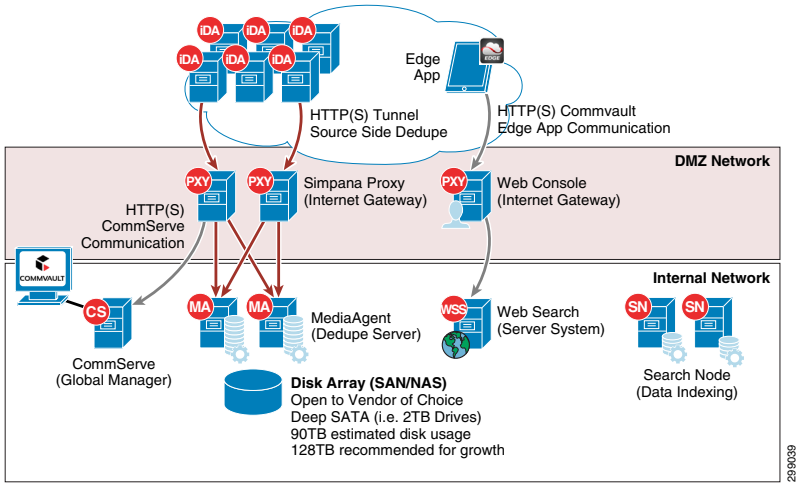

First, from a security perspective Simpana® utilizes certificate based authentication between Simpana® components and client computers. This protects against a variety of networking attacks such as spoofing. Secondly, Simpana® provides the ability to have a dedicated interfaces or shared networking interfaces among networking configurations with Data Interface Pairs (DIP). Refer to Commvault Books Online for Simpana® Data Interface Pairs details.

Firewalls

Firewalls provide security by blocking unauthorized access to networked computing and communications resources. Internet Protocol (IP) ports are configured in firewalls, permitting specific kinds of information to flow to and from opened IP address:port combinations, in specific directions (in, out or both). Firewall functionality is most often provided by either a stand-alone network appliance, or firewall software running on a general-purpose computer.

Simpana® can insert firewall rules per client allowing for tenant segregation and custom network configuration. This firewall feature provides the ability offer multiple network configurations per CommCell instance. CommCell components separated by a firewall must be configured to reach each other through the firewall using connection routes. Once configured, they can communicate to perform data management operations like backup, browse, and restore. CommCell components can be configured to operate across:

- Direct Connections using port tunnels - Direct connection with port restrictions is a setup where at least one of any two communicating computers can establish a one-to-one connection towards the other on specific ports. Three different types of direct connections, Client to CommCell, CommCell to Client, or Two way, as show in Figure 3-12.

Figure 3-12 Direct Connection—Two Way

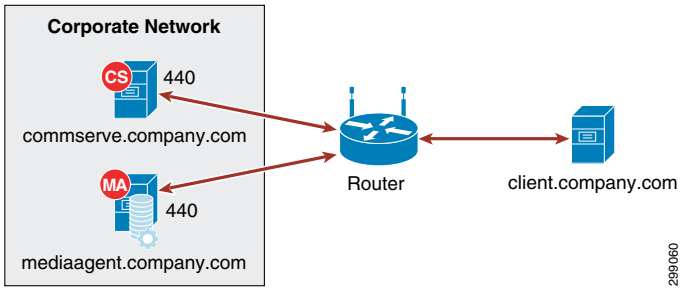

- Port-forwarding gateways - There are cases where direct connectivity setups do not work. Consider the case of the CommServe and MediaAgent being located inside a company's internal network, with the entire network being exposed to the outside world through a single IP address. Typically, this IP address belongs to a firewall or gateway that works as a NAT device for connections from the internal network to the outside. In scenarios like this, you can establish port forwarding at the gateway to forward connections coming in to specific ports to machines on the internal network that are mapped to those ports. You can then configure the client to open a direct connection to the port-forwarder's IP address on a specific port to reach a particular internal server. This creates a custom route from the client towards the internal servers. Figure 3-13 shows a client connecting to the CommServe and MediaAgent computer through a port-forwarding gateway setup.

Figure 3-13 Port-Forwarding Gateway

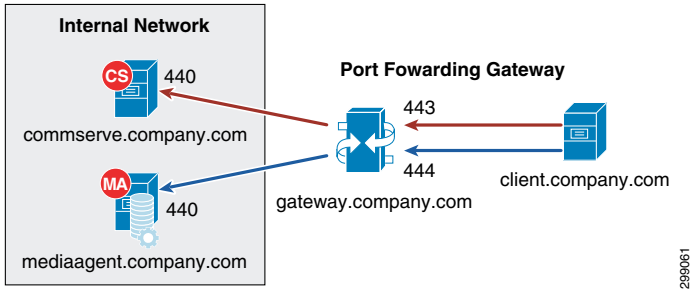

- The perimeter network (also known as a DMZ) using a Simpana proxy - Simpana proxy is a special proxy configuration where a dedicated iDataAgent is placed in a perimeter network and the firewalls are configured to allow connections (from inside and outside networks) into the perimeter network. The proxy, which is the agent running in the perimeter network authenticates, encrypts, and proxies accepted tunnel connections to connect the clients operating outside to clients operating inside. The Simpana proxy acts like a Private Branch Exchange (PBX) that sets up secure conferences between dial-in client calls. With this setup, firewalls can be configured to disallow straight connections between inside and outside networks. Figure 3-14 shows a perimeter network setup where a client from outside communicates to the CommServe and MediaAgent operating in an internal network through the Simpana proxy.

Figure 3-14 Commvault Simpana Web Proxy

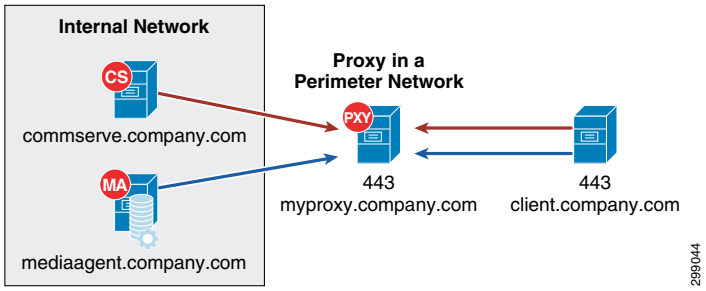

- HTTP proxies (including WiFi connections)—Consider the scenario where you are in a public location like a coffee shop, airport, hotel, or other such remote locations where Internet access is using public WiFi through a HTTP proxy. If you are a roaming user who travels frequently, you might operate the software in this scenario.
- Any Combinations of the methods listed above.

The firewall service is not restricted by a specific network configuration and can be tuned as an example per:

- Tenant
- Sub-tenant
- Client

Refer to Books Online for more information on Commvault firewall configuration.

Proxy

Proxies are an important component of service providers network security configuration to reduce the number of ports opened and provide secure data transfer between service provider and tenant. Simpana® offers two proxy configurations and within a single CommCell deployment both configurations can be used:

Private

- Dedicated proxy to the tenant
- Located at the customer or service provider's site
- Prevents the tenant's infrastructure from being Internet facing
- CommServe and Media Agent are Internet facing

Shared Proxy

- Single proxy with multiple tenants pooled together
- Located in the service provider's DMZ
- Prevents the service provider's infrastructure from being Internet facing

Network Bandwidth

Oversubscription of network resources is common place among service providers and the ability to throttle is crucial for network management. Simpana® has two available options to perform network throttling:

- **Relative—**% of available send or receive
- **Absolute—**fixed amount send or receive

More interesting for service providers is the ability to assign or even schedule network throttling through a policy based approach:

- Tenant
- Client or Client Group
- MediaAgent
- Copy jobs local or remote
- Based on IP range

For more information on Network Bandwidth information refer to Books Online.

Encryption

For a networking perspective, data can be encrypted from end-to-end from at the source as well as in-transit. Simpana® allows service providers to define encryption keys per tenant, which is discussed in more detail in the Data Level Security section.

For more information regarding Commvault encryption configuration options refer to Books Online.

Reporting

Simpana® has a robust reporting facility to show real-time and historical trending reports depending on the service provider and tenant needs. Simpana® extends user and group attributes to reporting by embedding filtering by permission set. For example, a tenant could run a capacity report, however the report view would be limited to resources assigned to that tenant. Assigning and grouping tenant resources can be accomplished in many ways and for more information refer to the user management section of this report.

Service providers can assign permissions at a report level basis. For example, a service provider could have a whole portfolio of reports and only publish certain reports subscribed to by tenants or even users.

Commvault has a service to build custom reports that are multi-tenant enabled through the Personalization Service.

For more information on the Personalization Service refer to Books Online.

Graphic User Interface (GUI)

Simpana® offers two distinctly different GUI's from a service provider perspective:

- **Administration—**Creating policies, assigning duties user/groups, & associations to a permissions, and other tasks
- **Consumption**—Viewing and executing tasks that have been delegated to a user, group, or tenant

The two Simpana® GUI interface are:

1. CommCell Console
   - Advanced administration
   - Advanced recovery requirements
2. Web Console

- View only what you own (client owner)
- View only what has been assigned (group privileges)
- End-user self-service for basic recovery options

Security

Commvault Simpana® has many security features included in the software, which have evolved and been refined over the past 20 plus years. Throughout the document there have been several discussions of security related topics. There are three specific security features relating to multi-tenancy not discussed previously:

- **Client Owner**—special permission set enabling administrator like privileges restricted to a specific client object
- **Enabling Privacy (Client side)**—restricts the administrators abilities to perform tasks on a specific client without a passphrase
- **Data Level Security**—various levels of data security from client, target, and in-transit

Client owner

Client owner is a special permission for a user limited to a particular object, usually a single or group of clients. For example, a tenant has been assigned Client Owner permissions to a server where the tenant would have administrative like privileges which would be limited in scope to that server. Included in the Client Owner permissions is access to the Web Console GUI, where the tenant would only view resource where Client owner was assigned.

Enabling Privacy

Some tenant may require additional security and assurances their privacy is being appropriately controlled in a multi-tenant environment. Simpana® has an additional privacy feature that can be enabled where a password will be required to certain tasks such as:

- View or browse data
- Restore data

The tenant would create and manage password, which would essential lock-out the service provider from performing certain tasks or viewing data. This feature is not enabled by default in Simpana® and the service provider would have to configure the options in Simpana® before being available to tenants.

For more information on Enabling Privacy refer to Books Online.

Data Level Security

As described in the Clients section (under Management Server), the CommServe generates an SSL certificate when new clients join the environment to provide an extra level of security ensuring no spoofing or rogue access to data. Simpana® provides three levels of encryption:

- **Source Side**—Encrypt at the agent
- **Target Side**—Encrypt it before you write it to storage (ie media agent)
- **Transit—**Encrypt at source, decrypts before written to storage

Service providers can enable or disable the three types of encryption at:

- Tenant

- Client
- Storage policy
- Storage array
- Off-site copy

For more information on Simpana standard ciphers and FIPS certifications refer to Books Online.

Operational Workflows

The goal of a BaaS is to provide protection to the customer's data while removing the IT overhead and providing the functionality customers are looking for, whether that is simple file protection across the Internet or full application aware backups with local copies and copies in the cloud.

Commvault Platform

Commvault Simpana is a unique, comprehensive data management platform that allows Service Provider to offer a number of different data related service offerings, (Protect/Restore, Archive/Retrieve, Replicate/Recover/Sync, Index/Search) on a number of different platforms, (Physical, Virtual, Private and Public Cloud), while providing the operational metrics required to offer reports to customers and valid capacity planning, as well as open REST API to allow for Service Provider portal integration (Figure 3-15).

Figure 3-15 Comprehensive Data Management Platform

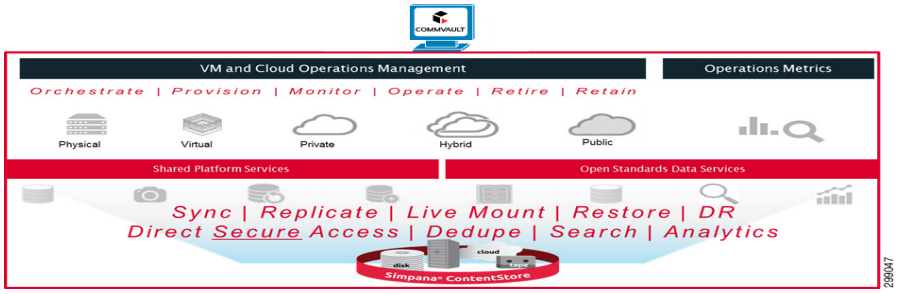

Commvault Simpana's policy based methodology allows a Service Provider to manage multiple different types of data from multiple different platforms for multiple different customers efficiently and securely. The Commvault Simpana Storage Policies act as a channel for backup and restore operations. Its chief function is to map data from its original location to a physical media, in one or more locations. The other function it servers is to determines how long the data will be retained at each given location.

Commvault Simpana allows for each Storage Policy to be configured with any number of Storage Policy Copies. There are three different types of Storage Policy Copies.

1. **Primary Copy**—First copy Simpana receives from the client.
2. **Snap Copy**—Snapshot that still resides on the disk subsystem.
3. **Secondary Copy**—Another copy of the data generated from the Primary Copy already within the Content Store.

   There are two different types of Secondary Copies:

   1. **Synchronous Copy**—Copy that contains all backup jobs (full, incremental, differential, transaction log or archive job) are written to the primary copy

2. **Selective Copy—**Allows for a specific full backup job to be copied from a source copy (either the Primary or another Synchronous Copy) to another target copy.

It is the Secondary Copies that allow Commvault Simpana to distribute data to multiple locations (logically or physically). Any MediaAgents that have connectivity between each other can pass copies between themselves.

Commvault Deduplication

Commvault Simpana Deduplication provides an efficient method to transmit and store data by identifying and eliminating duplicate blocks of data during backups. All data types from Windows, Linux, UNIX operating systems and multiple platforms can be deduplicated when data is copied to secondary storage. Deduplication allows the optimizes use of storage media by eliminating duplicate blocks of data and reduces network traffic by sending only unique data during backup operations.

Deduplication works as follows:

1. A block of data is read from the source and a signature for the block of data is generated using hash algorithm. Signatures are unique for each data block.
2. The signature is compared against a database of existing signatures for data blocks that are already on the destination storage. The database that contains the signatures is called the Deduplication Database (DDB).
3. If the signature already exists, the DDB records that an existing data block is used again on the destination storage. The associated MediaAgent writes the index information and the duplicate data block is discarded.
4. If the signature does not exist, the new signature is added to the DDB. The associated MediaAgent writes both the index information and the data block to the destination storage.

During the deduplication process:

1. Two different MediaAgents roles are used. These roles can be hosted by the same MediaAgent or different MediaAgents.
    1. **Data Mover Role**—The MediaAgent has write access to disk libraries where the data blocks are stored.
    2. **Deduplication Database Role**—The MediaAgent has access to the DDB that stores the data block signatures.
3. Data blocks can be compressed (default) and/or encrypted (optional).
4. Data block compression, signature generation, and encryption are performed in that order on the source or destination host.
5. Signature comparison is done on a MediaAgent. For performance benefits, a locally cached set of signatures on the source host can be used for the comparison. If a signature does not exist in the local cache set, it will be sent on to the MediaAgent for comparison.
6. An object (file, message, document, and so on) written to the destination storage may contain one or many data blocks. These blocks might be distributed on the destination storage. An index that is maintained by a MediaAgent tracks the location of the data blocks. This index allows the blocks to be reassembled so that the object can be restored or copied to other locations. The DDB is not involved in the restore process.

Deduplication Uses

MediaAgent-side (Storage-Side) deduplication can be used when the MediaAgent and the clients are in a fast network environment like a LAN. If used with the signature generation selected on the MediaAgent computer, it will reduce the CPU usage on the client computers by moving the processing to the MediaAgent.

Source-side (Client-Side) deduplication can be used when the MediaAgent and the clients are in a delayed or low bandwidth network environment like a WAN. It reduces the amount of data that is transferred across the network and can be used for Remote Office backup solutions. For example, Laptop Backup (DLO).

Global deduplication provides greater flexibility in defining retention policies when protecting the data. Use global deduplication storage policies to consolidate Remote Office backup data in one location or use this feature when data types like file system data and virtual machine data need to be managed by different storage policies but in the same disk library.

The DASH Full (Accelerated Synthetic Full) Backup operations can be used to increase performance and reduce network usage for full backups. The DASH Full is a Synthetic Full operation that updates the DDB and index files for existing data rather than physically copying data like a normal Synthetic Full backup.

DASH (Deduplication Accelerate Streaming Hash) Copy is a deduplication enabled storage policy copy option used by an Auxiliary Copy job to send only unique data to that copy. DASH Copy uses network bandwidth efficiently and minimizes the use of storage resources. DASH Copy transmits only unique data blocks, which reduces Auxiliary Copy job volume and time by up to 90%. Use DASH Copy when remote secondary copies can only be reachable on low bandwidth connections.

Commvault Client Protection

- The Commvault Simpana iDataAgents are the interface to the File Systems, Virtual Servers, Applications and Databases that are protected in most environments today.
- File System iDataAgent

The File System iDataAgent provides unified data protection and recovery for file systems on any number of currently available operating systems. The File System iDA is installed on each server containing files that are requiring protection, allowing for the CommCell to schedule data protection jobs. Point-in-time Recovery is available in the event of a serious disaster, as well as some number of versions back for each file. Full System recovery capabilities are available via Commvault Simpana 1-Touch and single pass backup and archive job via Commvault Simpana OnePass.

Virtual Server iDataAgent for VMware

The Virtual Server iDataAgent (VSA) for VMware is used to integrate with the vStorage API for Data Storage (VADP) to provide hypervisor image level backups from a single point of administration while still being able to isolate customer data and report on their activities (Figure 3-16). The VSA for VMware can be installed on a physical server outside of the VMware environment or it can be virtualized within the VMware environment. Either way the CommServe will communicate with the vCenter server and VADP to perform online full or incremental image backups, from which the full image can be recovered or virtual disk or specific Guest Files can be restored.

Figure 3-16 Virtual Server iDataAgent for VMware

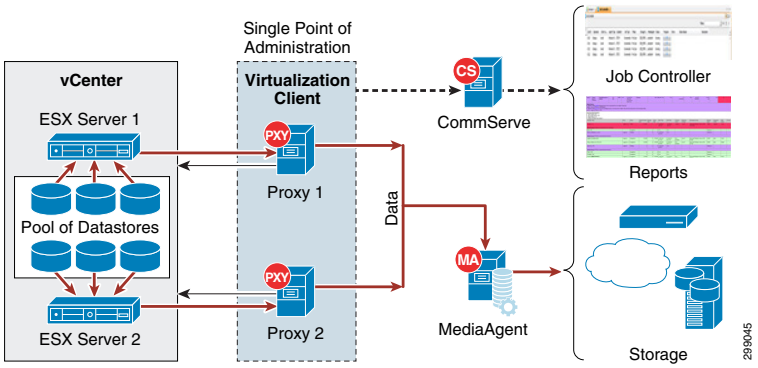

Virtual Server iDataAgent for Hyper-V

The Virtual Server iDataAgent (VSA) for Hyper-V is use to allow a Service Provider to provide hypervisor image level backups from a single point of administration while still being able to isolate customers data and report on their activities (Figure 3-17). The VSA for Hyper-V is a small agent that is installed on one or more of the Hyper-V servers within the cluster. This agent allows for the CommServe to work with Hyper-V and Volume Shadow Services (VSS) to perform online full or incremental image backup, from which the full image can be recovered or virtual disks or specific Guest Files can be restored.

Figure 3-17 Virtual Server iDataAgent for Hyper-V

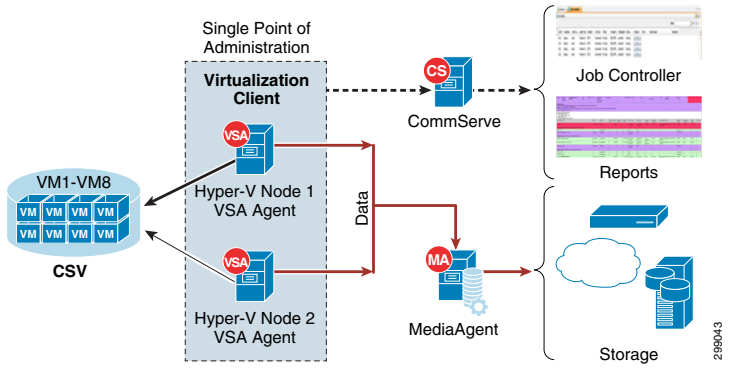

Application iDataAgent

There are multiple different Application iDataAgents available to assist with the data protection requirements of virtually all of the most commonly used application and databases, such as Microsoft Exchange, SQL Server, and SharePoint, Oracle RAC, DB2 or SAP. Each individual Application iDataAgent is created to interface with the application's or database's APIs to utilize its own data protection procedures, just directing the data to the Commvault MediaAgents and the storage attached to those. Utilizing each application or database own data protection functionality allows the data protection jobs to be application consistent, meaning the application can be paused and writes redirected during the backup process. It also allows for the most recovery options for each application or database, such as point-in-time recoveries with log replays for databases, entire SharePoint Farms or single documents, or MSSQL Databases or a single table.

CHAPTER 4

Implementation and Configuration

This chapter covers the implementation and design details configured during the testing phase of this project. The testing process involved a number of functional, operational, and negative test cases applied to each use case. This chapter does not include explicit details, like those found in a configuration guide for deployment; instead guidance is given for deployment considerations and best practices, along with links to documentation for installation, configuration, and troubleshooting.

Validation Environment

Figure 4-1 shows the scope of the lab environment used to cover all of the project's use cases and deployment scenarios. This lab environment is a subset of the CCA-MCP full architecture.

The validation environment consisted of three sites. Two of these sites were built to represent Cloud Service Providers, Service Provider 1 (SP1) and 2 (SP2). The third site was built to represent the Enterprise-class tenants. Each of these sites was connected over an IP cloud. The following sections give details of the sites' architectures and their interconnectivity. The arrows in Figure 4-1 indicate the direction of traffic for the three use cases, as follows:

- **In-Cloud BaaS—** is represented by the blue arrow from SP1 to SP2. The protection for Tenants 1 and 2 uses local backup retention in the SP1 Media Agents with remote backup data sent to the Media Agents in SP2.
- **Remote BaaS—** is represented by the red arrow from Enterprise to SP1. The protection for Tenants 6 and 7 used local backup retention in the Enterprise Media Agents with remote backup data sent to the Media Agents in SP1.
- **Remote BaaS without Local Retention—** is represented by the green arrow from Enterprise to SP1. Without local retention the protection policy for Tenant 5 transmits all backup data the Media Agents in SP1.

CCA-MCP architecture provides a shared services network segment which is accessible by all tenant VMs. CCA-MCP Infrastructure Foundation Guide Shared Services Access configuration section provides the configuration details to setup shared services access. BaaS Commvault central infrastructure components are placed in this shared services network, so it becomes accessible for tenants.

Figure 4-1 Lab Environment

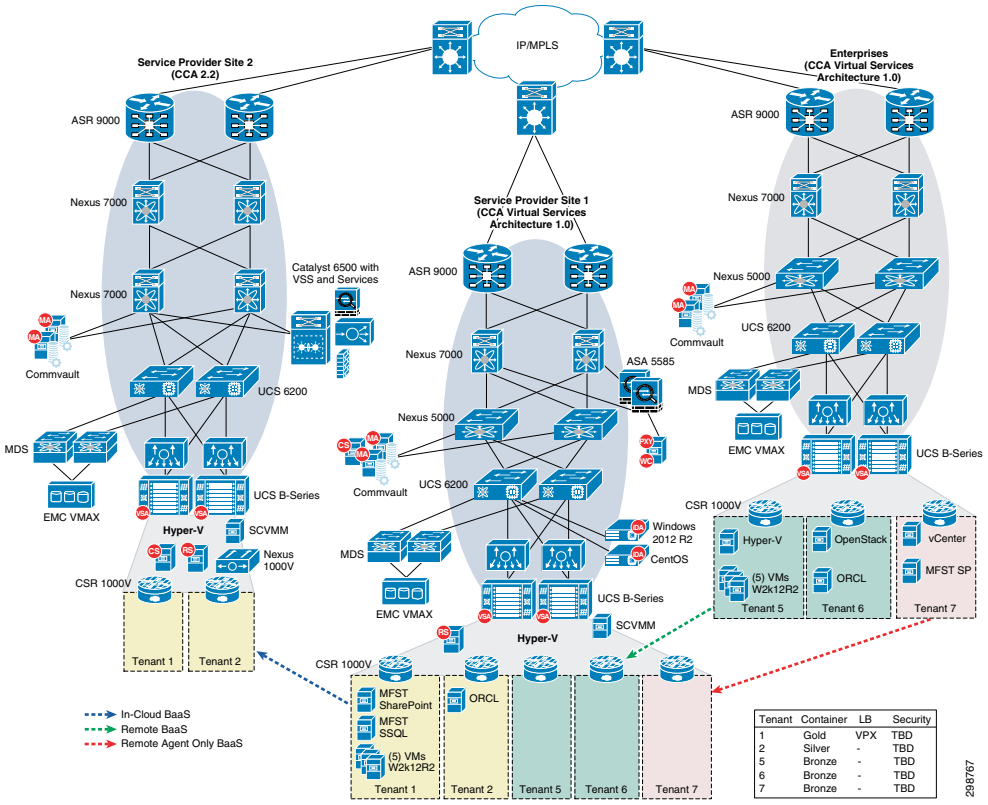

Solution Components

Table 4-1shows the components used during the validation testing. Refer to component information for the CCA-MCP Architecture 1.0.2 infrastructure.

Table 4-1 Solution Components

Function	Component	Version	Notes
Networking	• Cisco ASR 9000 • Cisco Nexus 7009, 7004 (7018 and 7010 not in SUT, but also valid options) • Sup2E, F2E, Sup2 • F2-Series 1 and 10 Gbps Ethernet cards • Cisco Nexus 5548 • Cisco Nexus Fabric Extender 2248TPE	Refer to CCA-MCP Architecture	CCA VSA 1.0 infrastructure used in Cloud Service Provider site SP1.
Services	ASA 5585-X	9.2(2)4	Firewall in Cloud Provider Site SP1.
	Citrix NetScaler VPX Load Balancer	10.5 Build 54.9.nc	Load balancer in tenant workload network.

Table 4-1 Solution Components (continued)

Function	Component	Version	Notes
Compute	• UCS B200 M3 • UCS 5108 chassis • UCS 2208XP I/O module • UCS VIC M82-8P • 96GB DDR3 1600MHz	UCSM 2.2(1b)	Used in all sites for hypervisor hosts, including Hyper-V, OpenStack, and VMware ESXi.
	• UCS C240 M3S • Intel Xeon E5-2630 (or 2620) dual socket, 12 cores • RAM: DDR3 1600MHz • UCS VIC 1225 • LSI 9271-8i Mega-RAID SAS HBA	• BIOS: 2.0.3.0 • FW: 2.0(3i) • VIC: 4.0(1e)	Bare metal server for Commvault servers: • CommServe in SP1 (CS-1) • MediaAgents in SP1 (MA-1, MA-2) • Reporting Server in SP1 (RS-1) • MediaAgents in Enterprise Tenant 6 (MA-1, MA-2)
	• UCS C3160 • Intel Xeon E5-2620 dual socket, 12 cores • 128GB DDR3 1600MHz • UCS VIC 1227 • Cisco RAID controller for C3X60	• BIOS: 2.0.2a.0 • FW: 2.0(2c) • VIC: 4.0(1b)	Bare metal server for Commvault MediaAgents in SP2 (MA-1, MA-2).
Storage	EMC VMAX	-	Provides storage for hypervisor hosts.
Virtualization	Microsoft Hyper-V	2012 R2	Used in both Cloud Service Provider sites and Enterprise Tenant 5 site.
	RHEL OpenStack	Icehouse	Used in Enterprise Tenant 6 site.
	VMware ESXi	5.1	Used in Enterprise Tenant 7 site.
	Cisco Cloud Services Router 1000V	• IOS-XE 03.13.01.S • IOS 15.4(3)S1	Virtual router in tenant workload network.
Data Management	CommmVault Simpana	V10 SP10	Single Unified Data Management Platform

SP1 Site Overview

The SP1 site was built to serve as the primary service provider data center in this solution architecture. It was integral in all of the use cases and replication scenarios.

IaaS Architecture

The data center infrastructure for the SP1 site is based on CCA-MCP Architecture 1.0 The CCA-MCP is a fully tested reference design that can be leveraged by enterprises and service providers to deploy an infrastructure that is efficient, secure, resilient, agile, simple, and scalable.

Cisco UCS

The SP1 site used both UCS B-Series and C-Series compute hardware. The B-Series were used to deploy the Microsoft Hyper-V environment, including the infrastructure and production Hyper-V hosts. The C-Series were used to deploy the Commvault CommServer, MediaAgents, and Reporting Server.

B-Series

The UCS B200 M3 servers were managed by the Cisco UCS Manager (UCSM) running release 2.2(1b). Two Service Profile templates were created on the UCSM, one for the infrastructure (management) hosts and one for the production tenant traffic) hosts. Both templates required certain BIOS settings to be explicitly configured as required for Hyper-V installed on an Intel-based server. The Execute Disabled Bit and the Virtualization technology (VT) settings were enabled in a BIOS policy that was referenced in the Service Profile templates. In addition to the BIOS settings, the Adapter Policy was set to “Windows”, and Fabric Failover was enabled.

Networking

For the infrastructure Hyper-V host, three of the four configured vNICs were used by the host to enable connectivity to the host management, cluster, and live migration VLANs. None of the virtual machines on the Hyper-V host required connectivity to these VLANs. On the UCSM, these VLANs were configured to be the native VLANs (untagged) on these vNICS and IP addresses were configured on all of these interfaces via the Windows OS on the Hyper-V host. The infrastructure management VLAN101 was used by multiple virtual machines and was therefore connected to a logical switch through the Infrastructure vNIC on the B200 M3. This allowed infrastructure management servers, such as Microsoft System Center Virtual Machine Manager (SCVMM) and Active Directory, to connect into the management VLAN.

Figure 4-2 shows the Service Provider 1 with the Hyper-V host infrastructure.

Figure 4-2 SP Site 1 Using Native Hyper-V Networking for Infrastructure Hyper-V Hosts

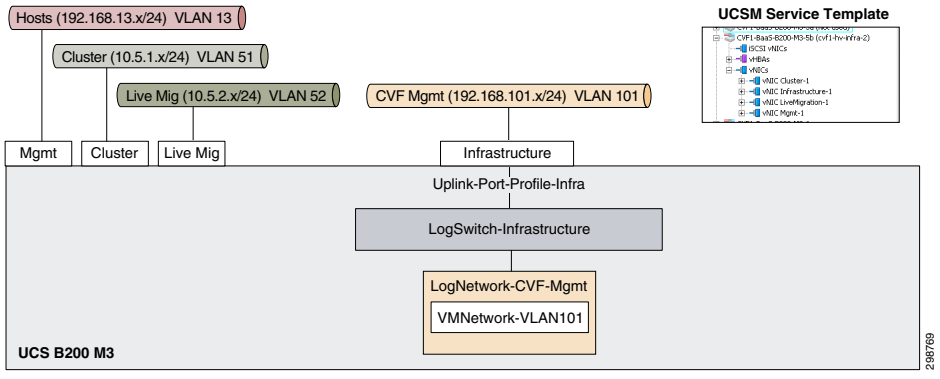

For the production Hyper-V hosts, native Hyper-V networking was originally configured and then later migrated to Cisco Nexus 1000V for Hyper-V. For the native Hyper-V networking configuration, six vNICS were configured in the UCSM Service Profile template.

The management vNIC on the production hosts carried two VLANs which were used by virtual machines on the hosts. A logical switch was configured to allow connectivity for the virtual machines to VLAN13 and VLAN101. The Hyper- V hosts required an interface in the host management VLAN13, so a virtual Ethernet interface was created in the Hyper-V logical switch via the Hyper-V Manager on each of the Hyper-V hosts.

Two of the six configured vNICs were used by the host to enable connectivity to the cluster and live migration VLANs. None of the virtual machines on the Hyper-V hosts required connectivity to these VLANs. On the UCSM, these VLANs were configured to be the native VLANs (untagged) on these vNICS and IP addresses were configured on these interfaces via the Windows OS on the Hyper-V hosts.

For the tenant traffic on the production Hyper-V hosts, a logical switch was created with two uplinks teamed together at the Microsoft OS level and configured for fabric failover at the Fabric Interconnect level. All twenty tenant VLANs were configured to be carried on these vNICs on the UCSM with tagging enabled. For each tenant, a logical network was configured and associated with only the VLANs that are used by the tenant.

The sixth vNIC on the Hyper-V host was used for Commvault replication and DMZ traffic. A Hyper-V logical switch was configured to allow virtual machines the needed access to the replication or DMZ VLAN. Only the DMZ Web Proxy server had an interface in the DMZ, with all access to and from it, restricted by the ASA firewall. Alternatively, the VLANs were included in the tenant uplinks.

Figure 4-3 shows Service Provider 1 using the Hyper-V host for production.

Figure 4-3 SP Site 1 Using Native Hyper-V Networking for Production Hyper-V Hosts

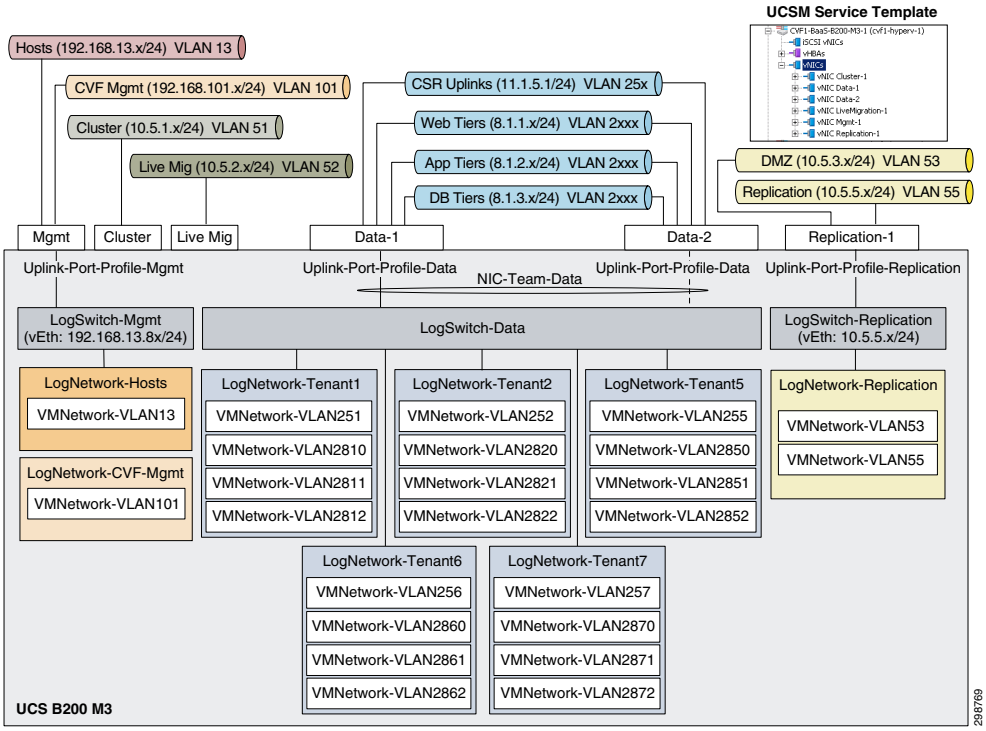

CCA-MCP architecture utilizes native Hyper-V networking configuration for the production (tenant) traffic rather than Nexus 1000v, however, this lab configuration was implemented with Cisco Nexus 1000V. The new logical networks, switches, and VLANs from the Nexus 1000V were associated with the new vNICs and, once operational, the interfaces on the virtual machines were changed in the SCVMM. The management VLANs remained on the native Hyper-V networking. The configuration details of the Nexus 1000V will be discuss in the Cisco Nexus 1000V section below.

Figure 4-4 shows the Nexus 1000V logical network hierarchy.

Figure 4-4 Nexus 1000V Logical Network Hierarchy

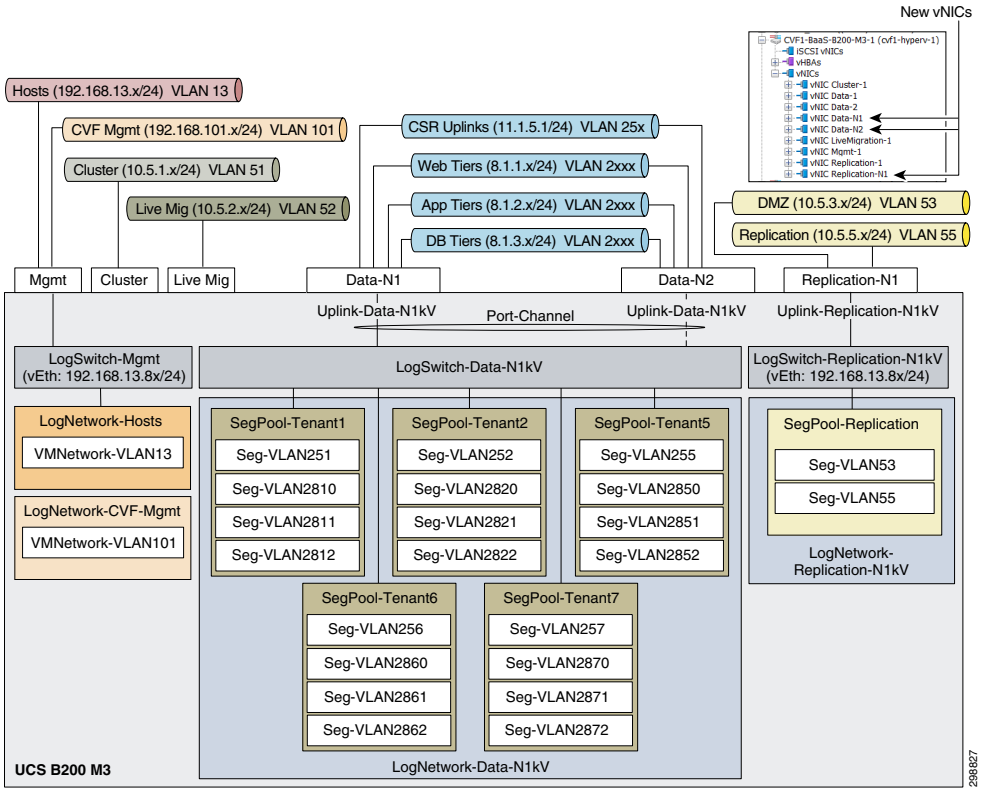

Storage

Two vHBA interfaces were configured in the Service Profile template for the production Hyper-V hosts to allow FibreChannel access to remote storage. Three LUNs were made available to the Hyper-V hosts, 1GB for the cluster quorum drive, 200GB for the first shared cluster volume, and a 2TB for the second shared cluster volume. There were multiple paths between the vHBAs and the target LUNs, so MPIO with the appropriate drivers for the storage must be configured in Windows 2012 R2. If MPIO is not enabled, a target LUN will appear multiple times in Disk Management in the Windows OS.

C-Series

The UCS C240 M3 servers were tested running version 2.0(3d), and later upgraded to 2.0(3i). The C240 M3 is managed by the UCSM or managed locally in standalone mode. Only the latter method supports multiple RAID configurations, which was required on the Commvault MediaAgent servers. Therefore, all of the C240 M3 were configured for standalone mode.

RAID Configuration

Figure 4-5 shows the physical drives that were used on the C240 M3 for the MediaAgent. There were five physical drives available, three 380 GB SSD and two 285 GB HDD.

Figure 4-5 Cisco UCS C240 M3 Physical Drives

The HDDs were configured for RAID1 and used for the Windows 2012 R2 operating system and MediaAgent software. The SSDs were configured for RAID5 and used for backup data. Figure 4-6 shows the M3 RAID configurations.

Figure 4-6 Cisco UCS C240 M3 RAID Configurations

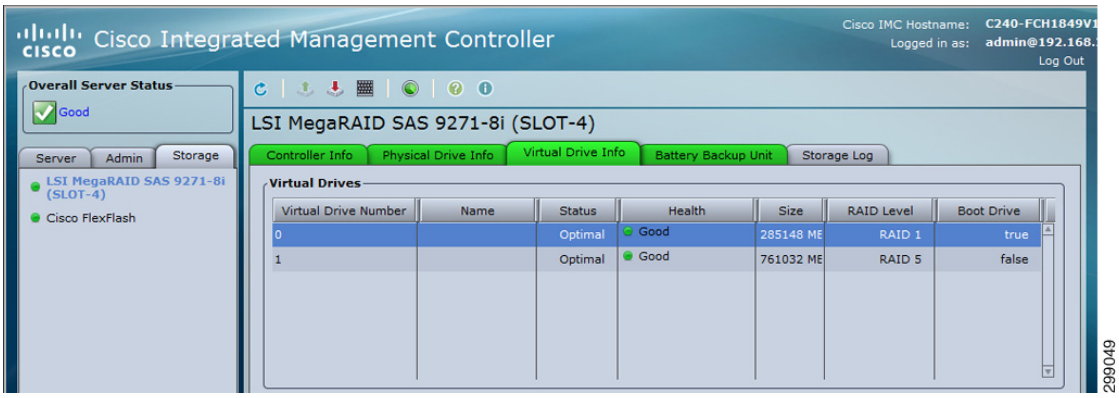

Networking

As for the network connections, the Commvault components were connected to the infrastructure management via the LOM (LAN on Motherboard) and connected to the replication network via the UCS VIC1225. A virtual Ethernet (vNIC) interface was configured and connected to the Nexus 5548 access layer switch. For the SP1 site, the vNIC interface and the Nexus switchport were configured as access switchports in VLAN55. Alternatively, the vNIC and Nexus switchport could be configured as trunk interfaces with a native VLAN configured for VLAN55. The latter approach was used in the SP2 site.

Note During validation testing, an issue was discovered that impacted the C240 servers from receiving broadcast packets. The issue was isolated to the VIC 1225 network driver in release 2.0(3d) (CSCur44975) and was resolved in the VIC driver in release 2.0(3i).

Microsoft Hyper-V

The Hyper-V cluster in the SP1 site was deployed on three B200 M3 blades to host the tenant and Commvault virtual machines. An additional blade was deployed to run Hyper-V in a non-clustered environment to host the infrastructure virtual machines. The following components were used:

- **Windows Server 2012 R2**—Datacenter version used for Hyper-V hosts, both Datacenter and Standard versions used for virtual machines.
- **Active Directory Server 2012**—Used for authentication and DNS for the cluster.
- **SQL Server 2012 R2**—Used for the relational database.
- **System Center Virtual Machine Manager 2012 R2**—Used to manage the infrastructure and virtual machines of the Hyper-V cluster.
- **Failover Cluster Manager**—Used to proof, create, and monitor the Hyper-V cluster.
- **Hyper-V Server 2012 R2**—Native hypervisor that enables virtualization of the x86-64 architecture. Installed as a role inside of Windows 2012 R2.

Note Configuration and management of the Hyper-V environment can be challenging and the recommendation is to manage from the top down using SCVMM first, Failover Cluster Manager second, and Hyper-V Manager last.

Creating the Hyper-V Cluster

Before the Hyper-V cluster was created, the following prerequisites were satisfied.

- A Windows Active Directory and DNS need to be available to the hosts.
- The Cisco UCS blades need to have Windows 2012 R2 Datacenter installed with Cisco VIC network and storage drivers installed. Without the VIC drivers, the OS will not be able to recognize the network and storage interfaces.
- If multiple data paths exist between the server and remote storage, the Multipath I/O feature is installed from Windows Server Manager and the MPIO driver for the storage vender is manually installed, if not included in Windows already.

Refer to Cisco UCS B-Series Blade Servers Windows Installation Guide for installing Windows on the Cisco UCS B-Series servers.

Once all the software and drivers were installed, basic configuration of the Windows OS was performed, including time adjustments, Windows updates, adding to local domain, firewalls, and so on. Multiple NICs were configured on the host; for example, the configuration interfaces in management were configured as were the replication VLANs (this is a minimal requirement).

From the Windows Server Manager interface, the Hyper-V and Failover Clustering features were installed on all servers that were hosts in the cluster. With the new features installed, Windows Update was used to load the available updates.

Storage configuration is dependent on the deployment environment, but there should be shared storage available to all hosts via SCSI, FibreChannel, etc. A quorum (or witness) disk can be used for the quorum configuration, which is a small LUN formatted for NTFS or ReFS. Large NTFS LUNs can be configured as Cluster Shared Volumes (CSVs) to store virtual machines configuration files and virtual hard drives. During the verification testing of this solution, several LUNs were created on a remote storage system and made available to the Hyper-V hosts. This included a small volume (1 GB) that was used by the cluster as a quorum drive, and two large LUNs (200GB, 2TB) that were used to store virtual machine files and virtual hard drives.

Refer to Failover Clustering Hardware Requirements and Storage Options for hardware requirement and storage options for failover cluster.

On one of the servers, the Failover Cluster Manager was opened and the Create Cluster Wizard was used to verify the servers. The wizard ran through multiple suites of tests to make sure that all networking, storage, and other requirements were met before the cluster was actually created. If any issues were discovered by the wizard, the issues were fixed, and the wizard was rerun. There were a large number of online resources available to help resolve any of the many issues that may arise during the validation.

Note There was one issue that was observed for which the fix was difficult to find. A shared disk used for the cluster (originally formatted for NTFS) can show up as a RAW format on some hosts. There was a persistent reservation on the disk and if formatting is attempted, it may display an error "the requested resource is in use". Refer to Microsoft Technet for details on command **Clear-ClusterDiskReservation -Disk <#>** to resolve this issue.

When the validation passed, the cluster was created. A cluster name and virtual IP address were configured and registered in the local Active Directory and DNS server. By default, all available storage was added to the cluster with the smallest shared volume used as the quorum disk.

Configuration changes and monitoring of the failover cluster was done from the Failover Cluster Manager on any of the hosts in the cluster.

Figure 4-7 shows the Failover Cluster Manager window.

Figure 4-7 Failover Cluster Manager

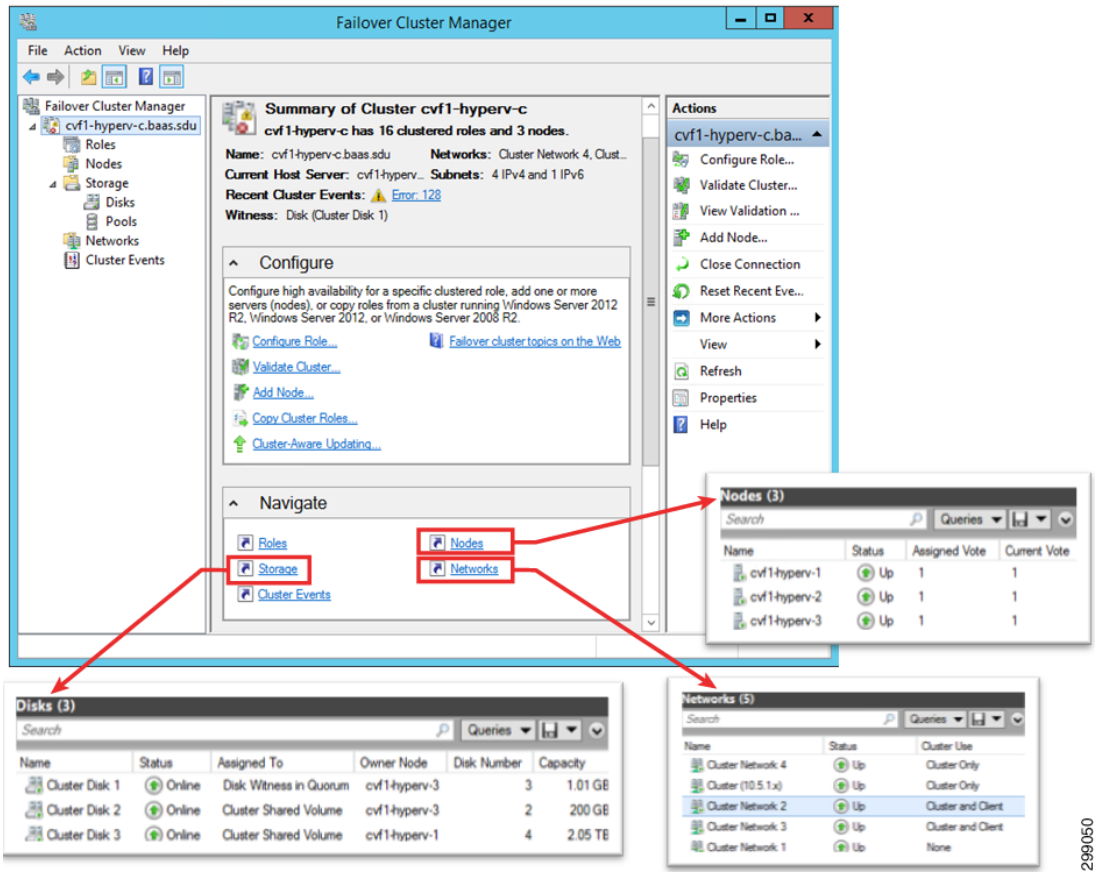

System Center Virtual Machine Manager

SCVMM is the central management interface for the Hyper-V virtualized data center. It can be used to configure and manage the Hyper-V hosts, networking resources, and storage resources to deploy virtual machines and services in a private cloud. SCVMM consists of the following components:

- **VMM Management Server**—Server that runs the VMM service, which processes commands to control the VMM database, library server, and hosts.
- **VMM Database**—Microsoft SQL database used to store configuration information (for example, virtual machines and service templates).
- **VMM Console**—The user interface into the VMM server.
- **VMM Library**—The library server hosts shared folders that are used as a repository for virtual hard disks, ISOs, templates, profiles, etc.).
- **VMM Command Shell**—Windows PowerShell with VMM cmdlets.
- **VMM Self-Service Portal**—Optional website to allow restricted access to cloud resources.

SCVMM can be deployed on a physical server or as a virtual machine. A SCVMM virtual machine deployed in a Hyper-V cluster would provide high availability and it could be deployed inside a cluster that it will manage. Once deployed, new Hyper-V hosts and clusters can easily be added to the SCVMM with all host, storage, network, and VM information for the host and/or cluster available via the SCVMM.

Native Hyper-V Networking

In SCVMM, logical networks were created to associate multiple VLANs into groups. For example, in the verification testing of the solution, a logical network was created for the management VLANs, each group of tenant VLANs, and the Commvault replication VLAN. Figure 4-8 shows the logical network for Tenant 1, which associates the CSR 1000V uplink (VLAN 251) with the three tiers for web, application, and database (VLAN 2810-2812).

Figure 4-8 SCVMM Logical Network Sample

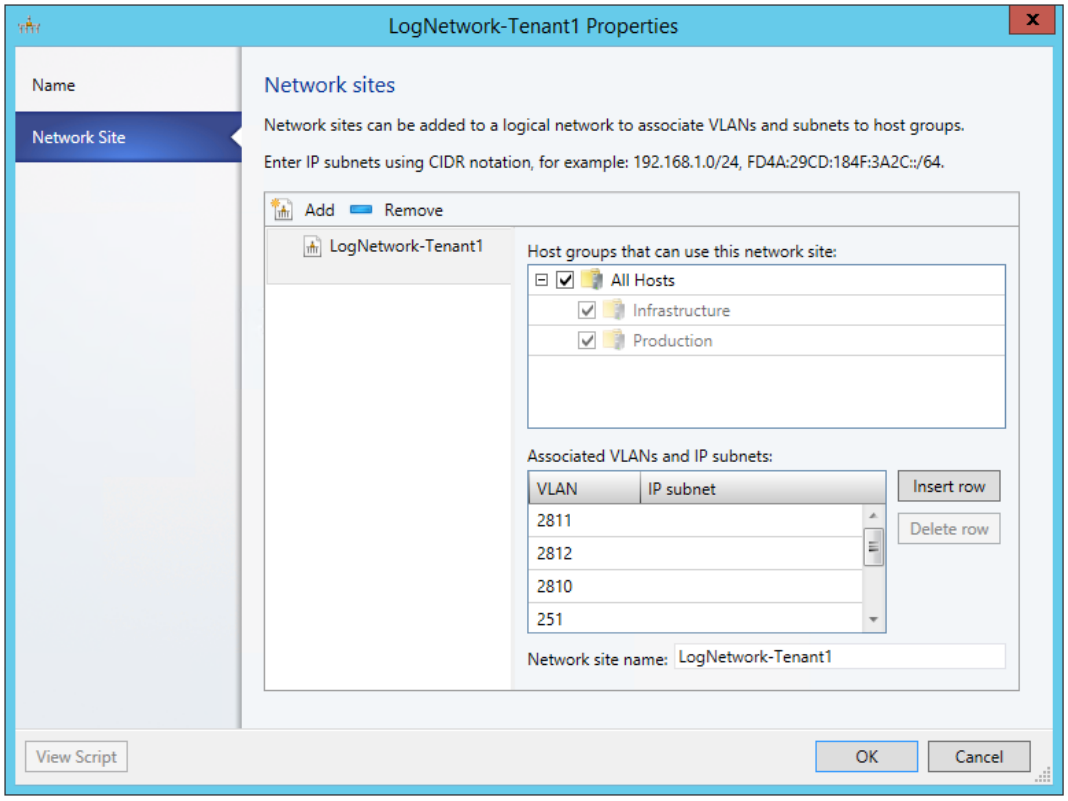

Once the logical networks were configured, uplink port profiles were created and associated with the logical networks. The uplink port profile was configured to allow select VLANs to be received and forwarded across an uplink. For example, in the verification testing of the solution, uplink port profiles were created for the management uplink, the data uplink for tenant traffic, and the replication uplink for the Commvault traffic. Figure 4-9 shows the port profile to the tenant data traffic, which includes all tenant logical networks.

Figure 4-9 SCVMM Uplink Port Profile Sample

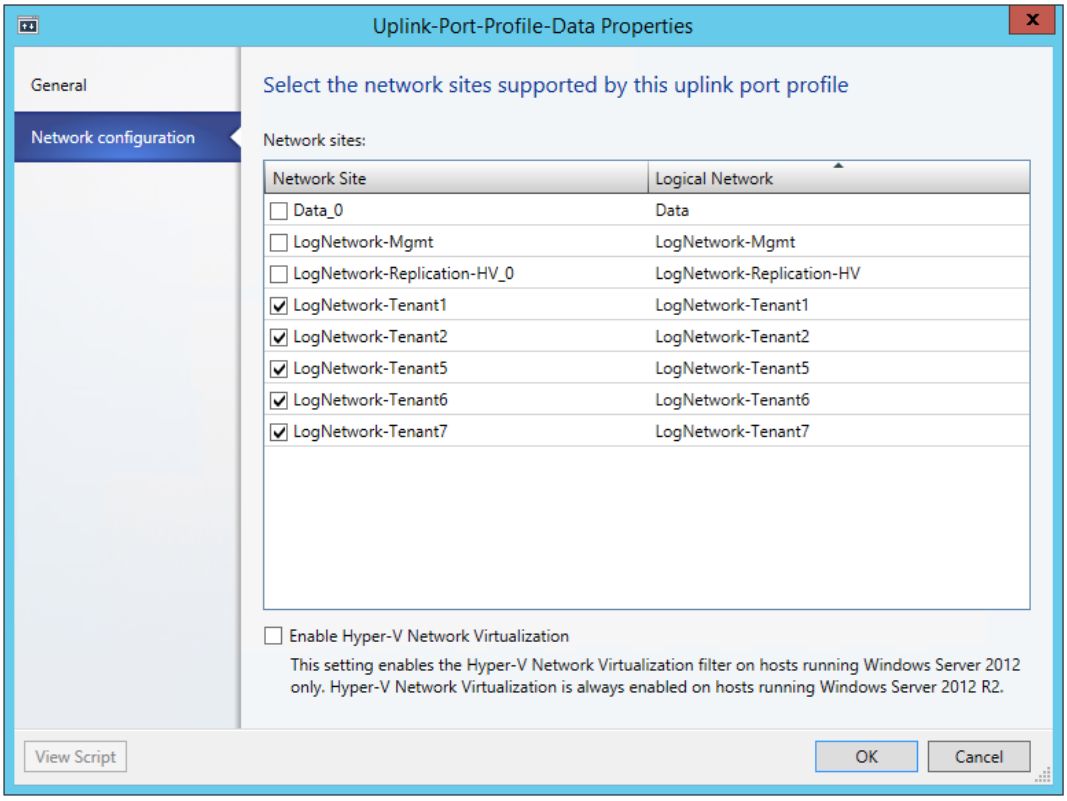

Next, logical switches were created on the Hyper-V hosts to provide an access layer for the virtual machines. The logical switch must be associated with the correct uplink port profile, which in turn associates the correct logical networks. For example, in the verification testing of the solution, three logical switches were created, one for management, one for all tenant data traffic, and one for Commvault replication traffic. In Figure 4-10, the logical switch for the tenant data traffic was associated with the uplink port profile for tenant traffic.

Figure 4-10 SCVMM Associating Logical Switch with Uplink Sample

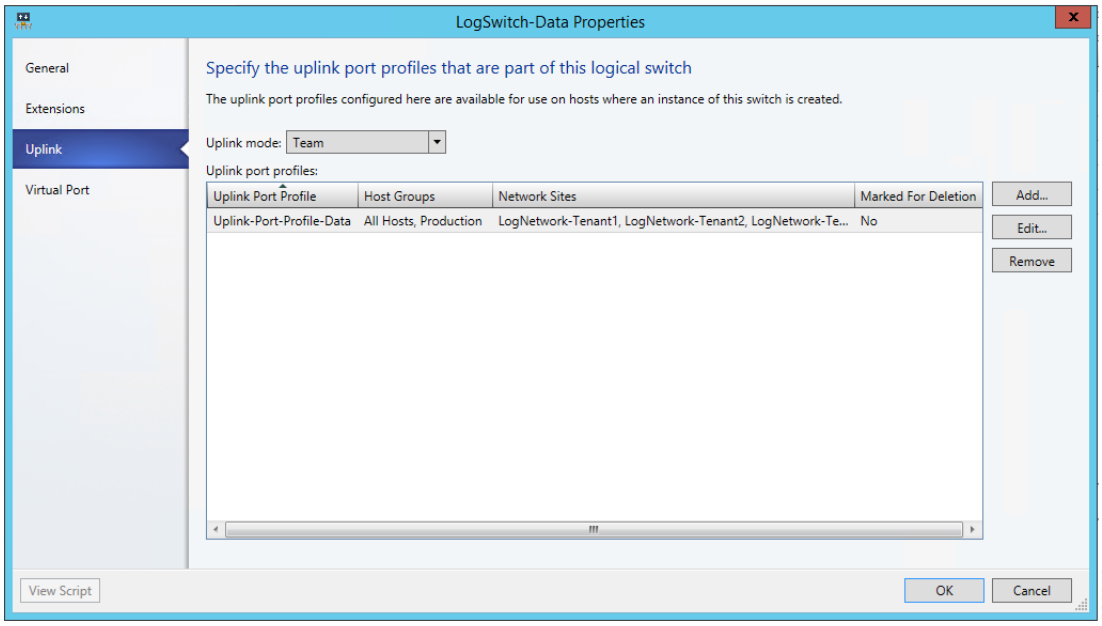

Next, the virtual machine networks were created with references to the logical network and VLAN. The VM Network was associated with the virtual machines to provide the desired access to the data center network (Figure 4-11).

Figure 4-11 SCVMM Virtual Machine Network Sample

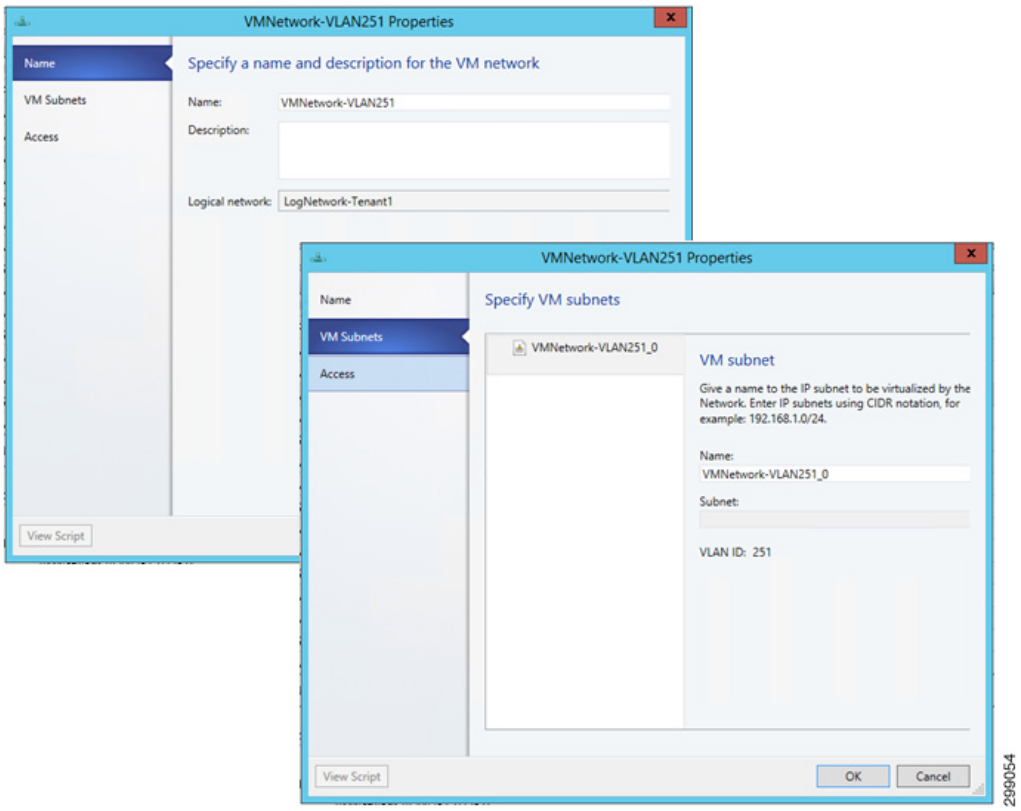

Now that the virtual networking is configured, the logical switches can be associated with the Hyper-V hosts. Figure 4-12 shows an example of the SCVMM host virtual switch.

Figure 4-12 SCVMM Host Virtual Switch Sample

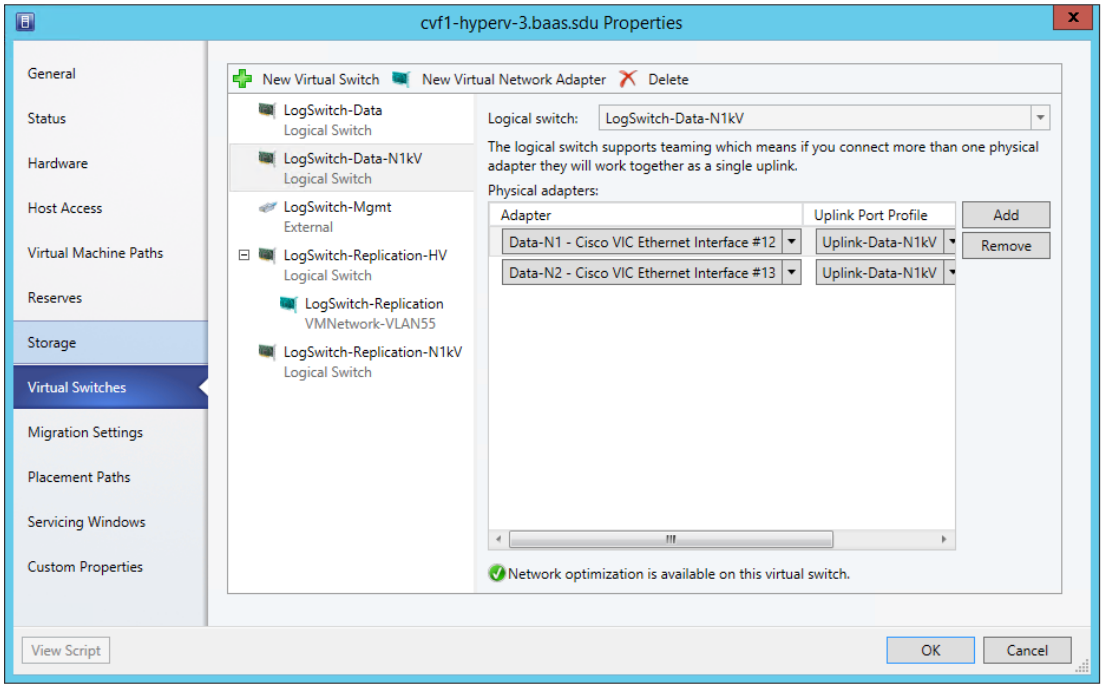

To provide network access to a VM, the VM properties were configured from the SCVMM. Under Hardware Configuration, a new network adapter created or an existing one was selected and the appropriate VM network was specified. The VLAN ID and Virtual switch also need to be selected from the dropdown menus. Figure 4-13 shows the SP1 Tenant1 CSR 1000V properties.

Figure 4-13 SCVMM Virtual Machine Network Properties

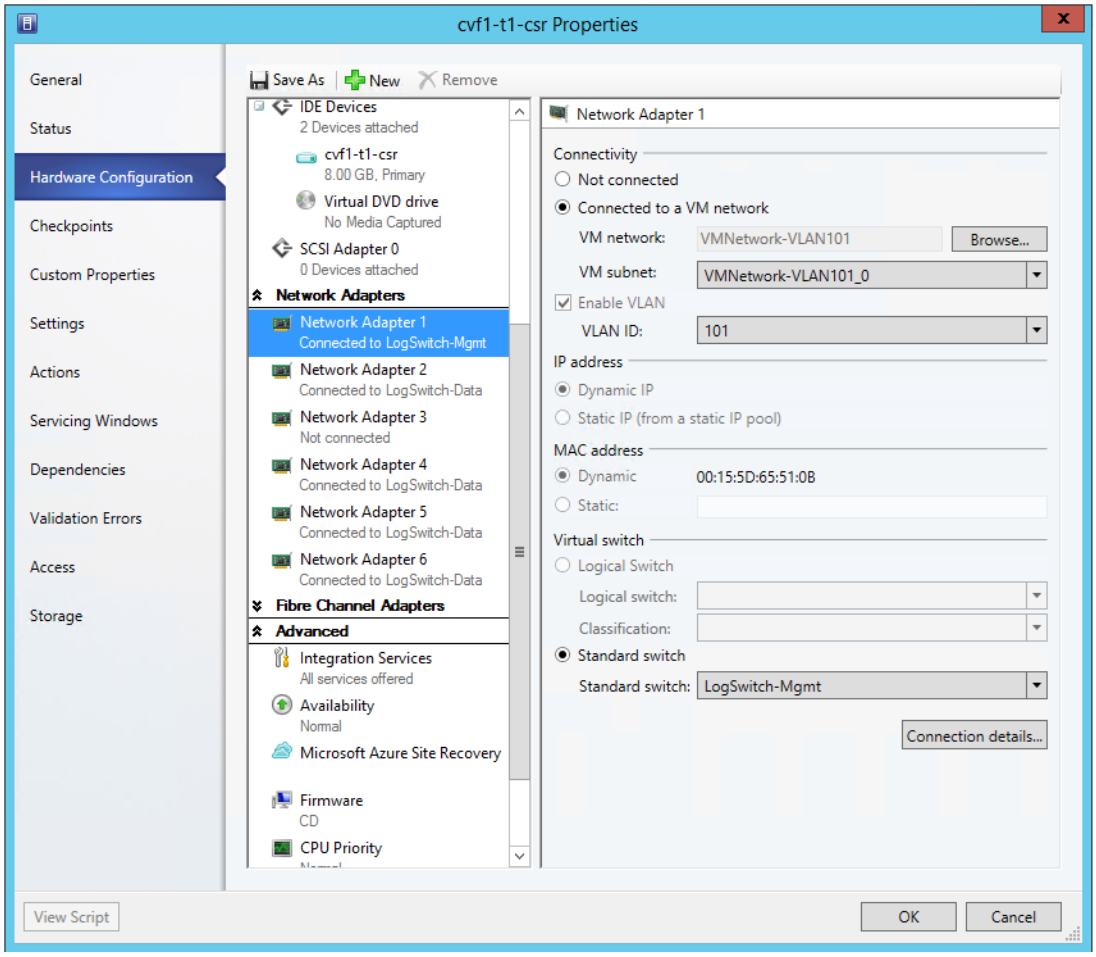

Refer to Cisco Unified Computing System with Microsoft Hyper-V Recommended Practices and Microsoft Technet Library - System Center Virtual Machine Manager 2012 for details on deploying Microsoft Hyper-V on the Cisco Unified Computing System.

Resource Reservation for Commvault VSA

For a Hyper-V deployment the Virtual Server Agent is installed directly on the host blade. This does not need to be installed on all blades in a cluster, but at a minimum there should be two for redundancy. The VSA monitors traffic from all VMs hosted in the cluster and performs backup and restore operations in conjunction with the Media Agent and CommServ Manager servers. Since the agent is loaded directly on the host, it can compete for resources with the VMs hosted on that blade. There are default resource reservations for a Hyper-V host, but they are minimal and not sufficient to ensure adequate operation.

The B200 M3 blades have 2 CPUs, a total of 16 cores and 32 threads. These resources are displayed as 32 logical processors under the Hyper-V Host>Properties>General tab.

Figure 4-14 Hyper-V Host System Information

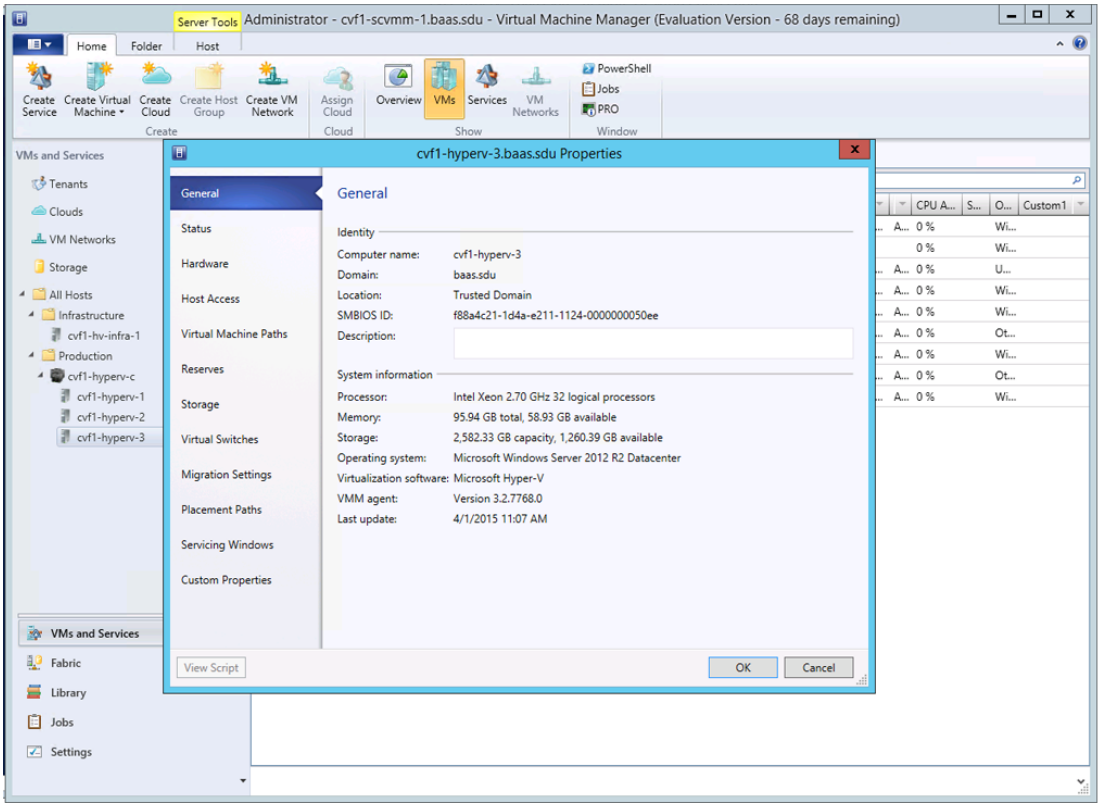

Under the Reserves window you need to check the box to ‘Override host reserve settings’. In this solution, the following resources were reserved for hosts cvf1-hyperv-2 and cvf1-hyperv-3 in Cloud Service Provider SP1 (Figure 4-15).

Figure 4-15 Hyper-V VSA Resource Reservations

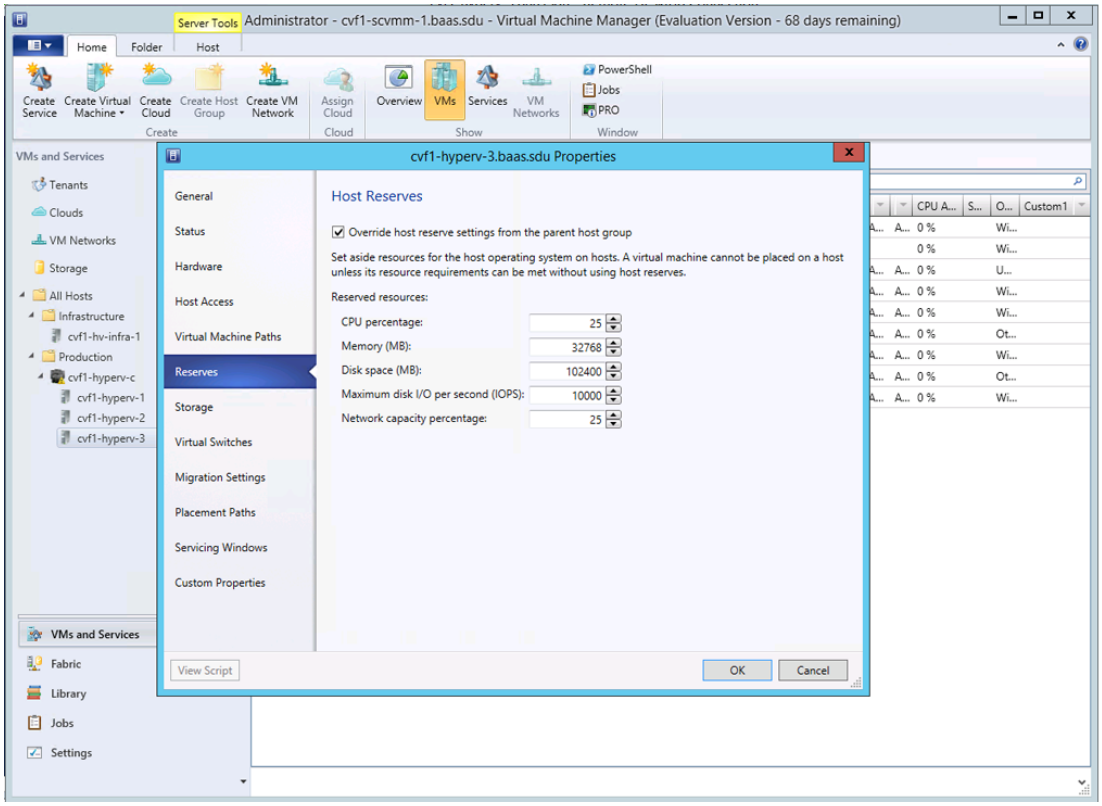

Refer to Commvault documentation for more information regarding VSA requirements.

Cisco ASA 5585 Firewall

In this system solution, the ASA 5585 is used as a perimeter firewall to provide access to the Commvault Web Proxy Server located in a DMZ and restricted access to the internal backup network in Cloud Services Provider SP1. The ASA is configured in a routed, multi-context mode with active/backup failover configured.

The DMZ Commvault web proxy server is used to allow customers Self Service Backup and Recovery Services without direct access to the inside backup network. The ASA is configured to allow access from outside sources to the web proxy server through a restricted set of ports. From the DMZ only Web Proxy originated traffic is allowed to access the internal Commvault components on the inside network. The DMZ subnet is configured with private IP addressing, so static NAT is configured to provide access to outside sources.

Communication is required between the Commvault components from different CSP locations connecting to the inside backup network behind the ASA firewall. Traffic originating from these locations over MPLS, or other methods, are allowed bi-directional communication through the ASA to the internal Commvault components.

The replication and Commvault management traffic routes to the inside backup network in CSP1 through the ASA from a backup VRF created on the ASR 9000. The ASR is configured using Import/Export route descriptors (RD) to learn these routes, along with static routes to and from the ASA. Tenant VRF specific routes are also shared for instances of when a specific server (SQL, etc.) is selected

for backup and needs to communicate directly to the Commvault components on the inside network. An alternate solution could have the server communicate with the DMZ Proxy server instead of directly to the internal backup network.

Figure 4-16 shows these traffic streams.

- Internet User is not allowed access to the Commvault components, so it uses the DMZ Web Proxy to manage Backup and Restore operations.
- The DMZ Web Proxy server acting on behalf of the Internet users is allowed to communicate with the internal Commvault components, providing a layer of separation and security.
- Enterprise traffic both management and data replication needs bi-directional communication with the internal Commvault components. This traffic routes over a L3VPN/MPLS cloud, through the Backup VRF on the WAN router to reach the ASA Firewall, which has a local interface to the internal backup network.
- For administrators in Tenant 1 or for application servers running the iDA agent, these communicate with the internal backup network through the Backup VRF, which includes the backup network and private tenant routes.

  If including the private tenant routes is not desirable in the Backup VRF, communication can be achieved through the same path, but through the DMZ Web Proxy. This flow is not depicted in the illustration.

Figure 4-16 Traffic Flows Through ASA Firewall

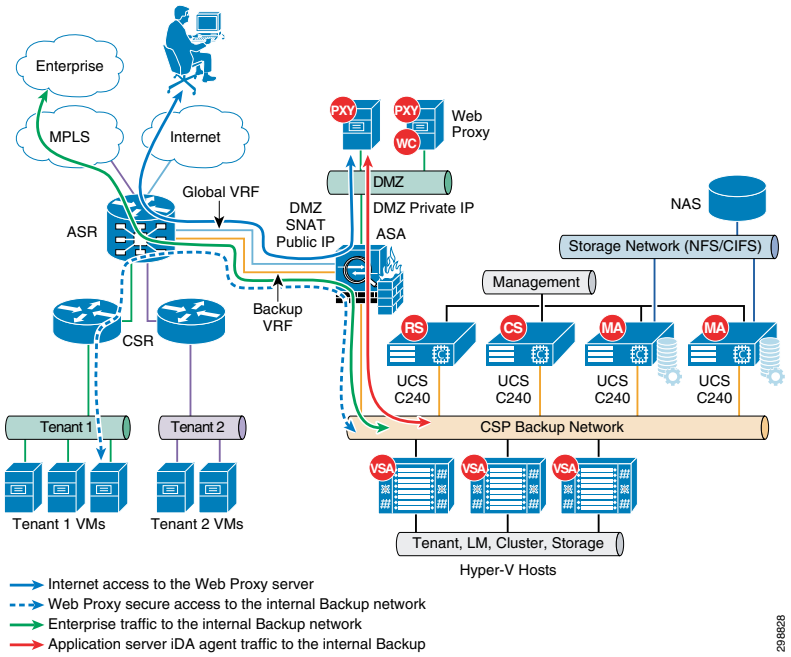

Cisco Cloud Services Router (CSR) 1000V

Beginning with Cisco IOS XE Release 3.12S, the Cisco CSR 1000V supports installation on the Microsoft Hyper-V hypervisor using Windows Server 2012 R2. The Cisco CSR 1000V installation on Microsoft Hyper-V requires the manual creation of a VM in Hyper-V and then installing the CSR software using a vDVD mapped to the CSR ISO image file. The amount of resources required for the

CSR depends on the licensed feature set and throughput requirements. For the validation testing, five virtual interfaces were configured, one for management, one for the uplink in the access layer, and three for the tenant workloads (web tier, app tier, and database tier).

A separate CSR 1000V instance was deployed for each of the four tenants in the SP1 site that use Hyper-V or VMware ESXi hypervisors. The CSR 1000V did not support the OpenStack environment. The CSR 1000V was configured as the gateway for each of the VLANs used in the tenant workloads. This included the web, application, and database VLANs for each tenant. Since Layer 3 connectivity was used between the tenant CSRs at each site, unique IP subnets were configured to enable proper routing between sites. The CSRs used BGP routing to learn IP subnets connected to the remote site's CSR.

Remote BaaS enterprise tenants (Tenants 5 and 7) used an IPsec tunnel between the CSR 1000V in the SP1 site and another tenant-specific CSR 1000V at the Enterprise site. The tunnel provided secure communications between the workload VLANs at each site. For the In Cloud BaaS tenants (Tenants 1-2), an IPsec tunnel was not configured since the network path between the sites was considered secure.

Caveats

When the Cisco CSR 1000V is installed on a Microsoft Hyper-V cluster, the interface numbers can change after a Hyper-V host failover event to a new host server or live migration. In both cases, the condition is not seen until after a reboot. The following steps can be taken to mitigate this issue:

1. Prior to executing a live migration enter the **clear platform software vnic-if nvtable** command.
2. The command can also be successful if executed after the failover, but only before the config is saved or the VM restarted.
3. Configuring static MAC addresses for the network interfaces.

In the event that the interfaces have been renumbered and the IP addressing is removed, the following steps can be used to recover.

1. Execute **clear platform software vnic-if nvtable** command.
2. Copy saved config to startup config.
3. Reboot the CSR.

Refer to the following links for details of the Cisco CSR 1000V Series Cloud Services Router:

- Cisco CSR 1000V Series Cloud Services Router Release Notes
- CSR 1000v Series Cloud Services Router Software Configuration Guide
- Installing the Cisco CSR 1000V in Microsoft Hyper-V Environments

Cisco Nexus 1000V

Cisco Nexus 1000V Series switches provide a comprehensive and extensible architectural platform for virtual machine and cloud networking. These switches are designed to accelerate server virtualization and multitenant cloud deployments in a secure and operationally transparent manner.

In this solution, the Nexus 1000V is deployed as a distributed virtual switch running version 5.2(1)SM3(1.1) and integrates with Microsoft SCVMM 2012 R2 to allow deployment of virtual machines and to manage virtual networking.

This section is a high-level overview of the steps executed to install and integrate the Nexus 1000V into the test topology. The overview is based on the steps used in this validation project and are not meant to be a complete set of steps for installation.

Refer to Cisco Nexus 1000V Install and Upgrade Guide for installation details.

To reduce complexity in CCA-MCP architecture, Nexus 1000V was replaced with the native Hyper-V virtual switch later in the design. However, this BaaS Commvault lab testing stayed with the N1kv component.

Installation Files and Template

When installing the SCVMM components, the referenced filename for the MSI file is incorrect for this release. For this step, execute the Nexus1000V-NetworkServiceProvider-5.2.1.SM3.1.1.0.msi file from the <zip package>\VMM directory on the SCVMM server.

The VEM software is a MSI file that needs to be copied (**do not execute**) into the following location on the SCVMM server: C:\ProgramData\Switch Extension Drivers. By default, the directory C:\ProgramData is hidden.

VSM Installation

In the Configured Hardware window you can either leave the Availability to Normal or set it to High. In this deployment active and standby VSMs were configured using Nexus 1000V high availability on a standalone Hyper-V host. Preferably two hosts should be used when configured in this manner to prevent a single point of failure.

If the VSMs are deployed in a Hyper-V cluster the availability mode should be set to High, which allows Hyper-V to spawn a replacement VM on another host if the original one should fail. It is important that both VSM VMs do not reside on the same Hyper-V host, so that service can be restored much faster than through the Hyper-V HA feature. To ensure this separation use an anti-affinity class. Some important items are:

- The VSM template installs three network adapters. In the Select Networks window choose the management network that is shared between the hosts and VSM for all three interfaces.
- In the Add Properties window, select Always turn on the virtual machine, then configure Delay start up 30 seconds, Turn off virtual machine, and Exclude virtual machine from optimization actions.
- After the VSM boots and the initial configuration applied you should unmap the VSM .iso image from the Virtual DVD drive.

After completing these steps, you are prompted to log into the VSM. Access the VSM via SSH using the IP address configured in the VSM installation section. The following objects need to be created on the VSM:

- Logical Network
- Network Segment Pool
- IP Pool Template
- Network Segment
- Virtual Ethernet Port Profile
- Ethernet Port Profile
- Network Uplink

Logical Network—A single logical network could be configured, but in this solution two were created; one for the Commvault replication traffic and the other for tenant data traffic. This is imported into the SCVMM under Fabric>Logical Networks.

```
nsm logical network LogNetwork-Replication-N1kV
nsm logical network LogNetwork-Data-N1kV
```

Network Segment Pool – Is a container for all of the Network Segments that are created for each VLAN. The Network Segments link to this pool and it links them to the appropriate Logical Network.

```
nsm network segment pool NetSegPool-Tenant1
   member-of logical network LogNetwork-Data-N1kV
```

IP Pool Template—Create an IP address pool for each VLAN. These are the IP addresses that could be dynamically assigned to your VMs, if you are using the SCVMM for automation. Though we did not dynamically assign these IP addresses, values must be configured; otherwise the Nexus 1000V will not integrate properly with the SCVMM. If you are using overlapping IP address space for private tenant networks, the same IP Pool Template can be imported into multiple Network Segments.

```
nsm ip pool template IP-Pool-Web-Tier
  address family ipv4
  network 8.1.1.0/24
  ip address 8.1.1.200 8.1.1.201
  default-router 8.1.1.1
```

Network Segment—A Network Segment needs to be created for each VLAN and made a member of the appropriate Network Segment Pool. The **ipsubnet** command must be configured first and match the network configured in the imported IP Pool Template.

```
nsm network segment type vethernet NetSeg-VLAN2810
  switchport access vlan 2810
  member-of network segment pool NetSegPool-Tenant1
  ipsubnet 8.1.1.0/24
  ip pool import template IP-Pool-Web-Tier
  switchport mode access
  publish network segment
nsm network segment type vethernet NetSeg-VLAN2811
  switchport access vlan 2811
  member-of network segment pool NetSegPool-Tenant1
  ipsubnet 8.1.2.0/24
  ip pool import template IP-Pool-App-Tier
  switchport mode access
  publish network segment
```

Virtual Ethernet Port Profile—To apply policies to the VMs, create a port-profile for them. This port-profile is not used for access, but as a collection of policies. Apply service policies for QoS or to bind a service, such as a VSG. Examples are displayed below, but were not used in verification testing.

```
port-profile type vethernet T1-APP
   no shutdown
   state enabled
   publish port-profile
port-profile type vethernet T20-2011
  service-policy input tenant-vm-qos
  org root/PT20
  vservice node Extended-PT20 profile PT20-Profile
  no shutdown
  state enabled
  publish port-profile
port-profile type vethernet All-TENANTS
   no shutdown
   state enabled
   publish port-profile
```

Ethernet Port Profile—An Ethernet port profile is created to apply any needed policies to the Network Uplink. This profile needs to be imported into the NSM Network Uplink configuration shown in the next step. If none is created a default policy is automatically imported into the Network Uplink. Below is an example of an uplink port profile with egress QoS marking applied. In verification testing, the default port profile was used.

```
port-profile type ethernet N1KV-Uplink-Policy
  service-policy output hyperv-egress-remark
```

```
channel-group auto mode on mac-pinning
no shutdown
state enabled
```

Network Uplink—This is used as the uplink from the host when the virtual switch is created. Each uplink is configured to allow certain segment pools. The allow network segment pool command tells SVCMM that these segments are allowed out this uplink. This is needed for the host to see the segments. If an Ethernet Port Profile is created, then it needs to be imported here; otherwise it uses the default profile.

```
nsm network uplink Uplink-Data-N1kV
  import port-profile uplink_network_default_policy
  allow network segment pool NetSegPool-Tenant1
  allow network segment pool NetSegPool-Tenant2
  allow network segment pool NetSegPool-Tenant5
  allow network segment pool NetSegPool-Tenant6
  allow network segment pool NetSegPool-Tenant7
  publish network uplink
```

After the Nexus 1000V configuration is complete a communication channel needs to be configured to integrate the switch into the SCVMM. Follow the steps in the Connecting SCVMM to VSM section of the Cisco Nexus 1000V Install and Upgrade Guide.

The following section discusses steps necessary to apply the Nexus 1000V configuration to the Hyper-V hosts and VMs.

- Create Logical Switch (Figure 4-17)
- Create VM Networks
- Add Virtual Switch to Hosts
- Assign VM Networks to VM network interfaces

On the SCVMM, under Fabric>Networking>Logical Switches create a Logical Switch.

Figure 4-17 Logical Switch Creation

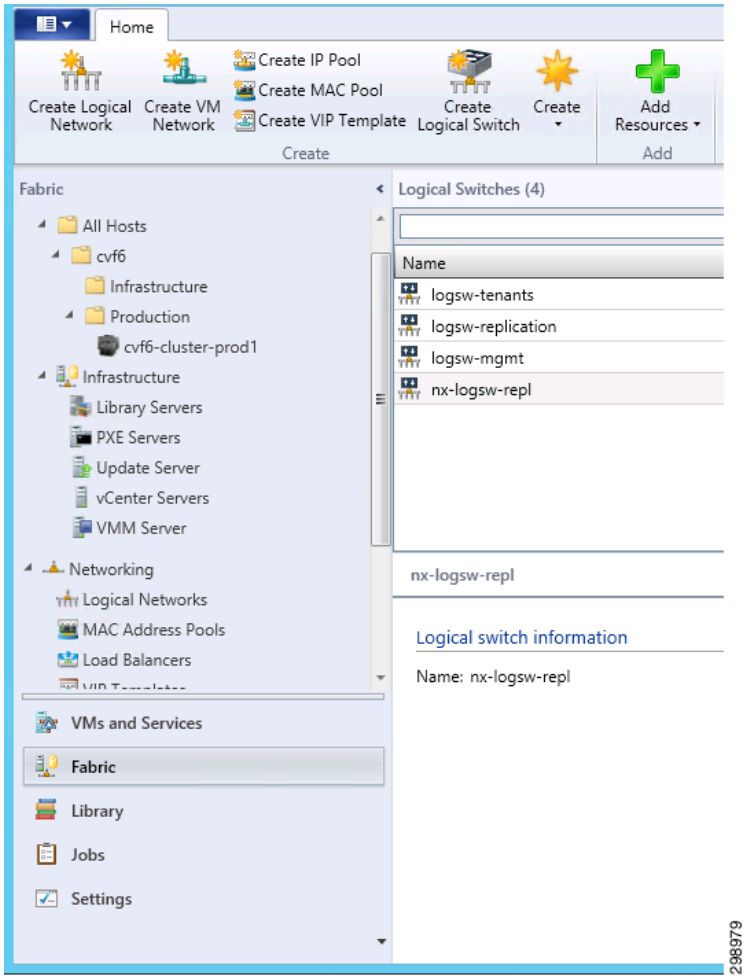

In the Extensions window leave the Microsoft Windows Filtering Platform checked and select the Nexus 1000V Forwarding extension (Figure 4-18).

Figure 4-18 Logical Switch Extensions

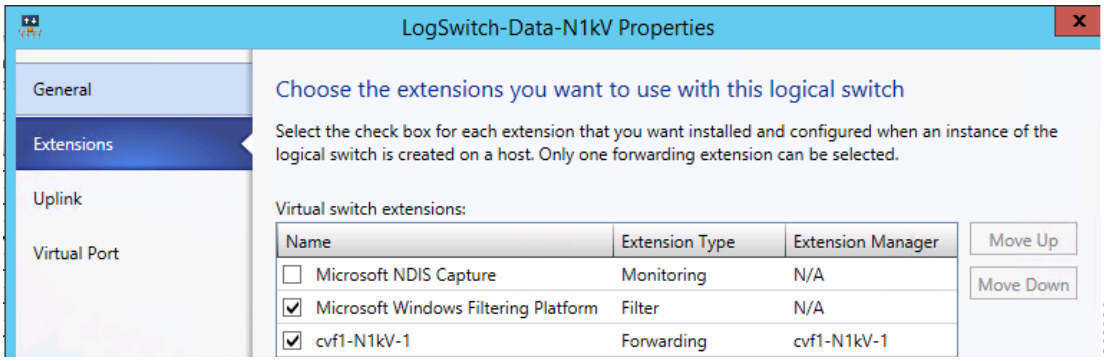

In the Uplink window, under the Uplink port profiles section, add the NSM network uplink that was configured on the Nexus 1000V. All of the network segment pools that were allowed will show up under the Network Sites column.

The vNIC interfaces from the UCS host will be linked to the Logical Switch when a new Virtual Switch is configured at the Hyper-V host level.

Note The Uplink mode must be set to Team even if only one interface is configured.

Figure 4-19 Logical Switch Extensions

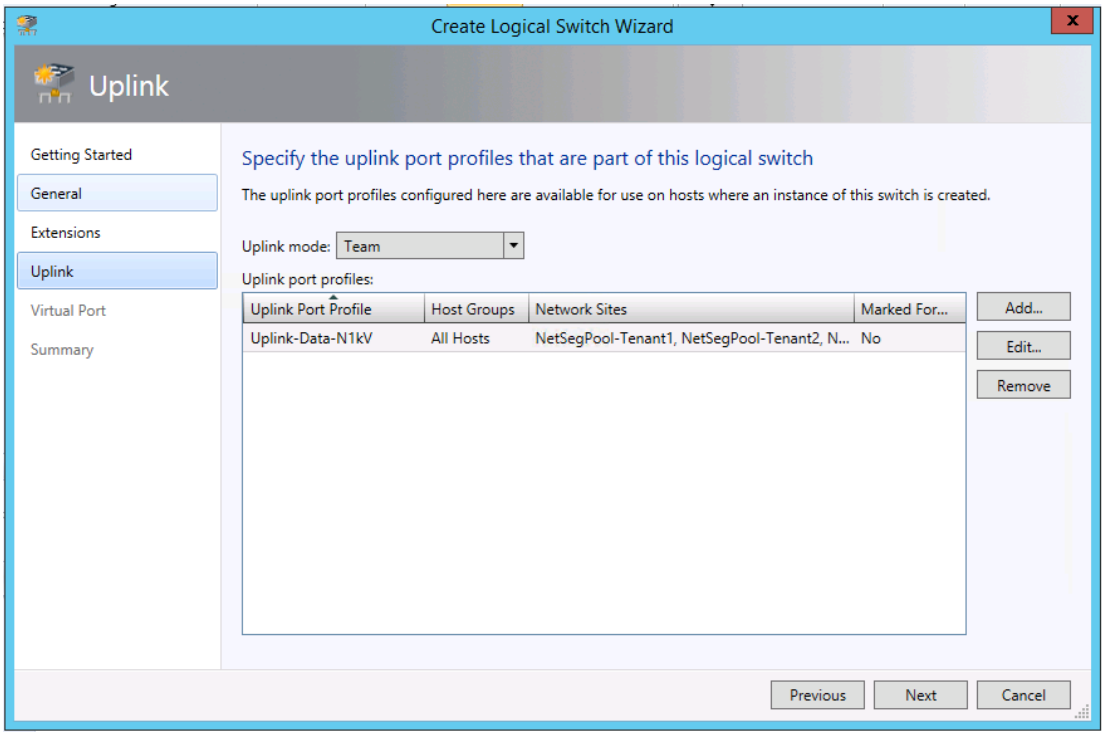

In the Virtual Port window perform the following:

Note Steps 1- 4 are represented by arrows and the remaining steps use highlighted boxes.

1. Click on Add to include a virtual port.
2. Click on Browse to open the Port Profile Classification window.
3. Click on Create Port Classification to create one.
4. Type in a name and click Ok.
5. In the Add Virtual Port window browse and select the port classification that was created.
6. Select the Nexus 1000V that was created.
7. Select a port profile (vEthernet port profile) if one was created. If not, a defaultAccess profile will exist.
8. Hit Ok and Finish to complete the Logical Switch configuration.

Figure 4-20 Logical Switch Creation

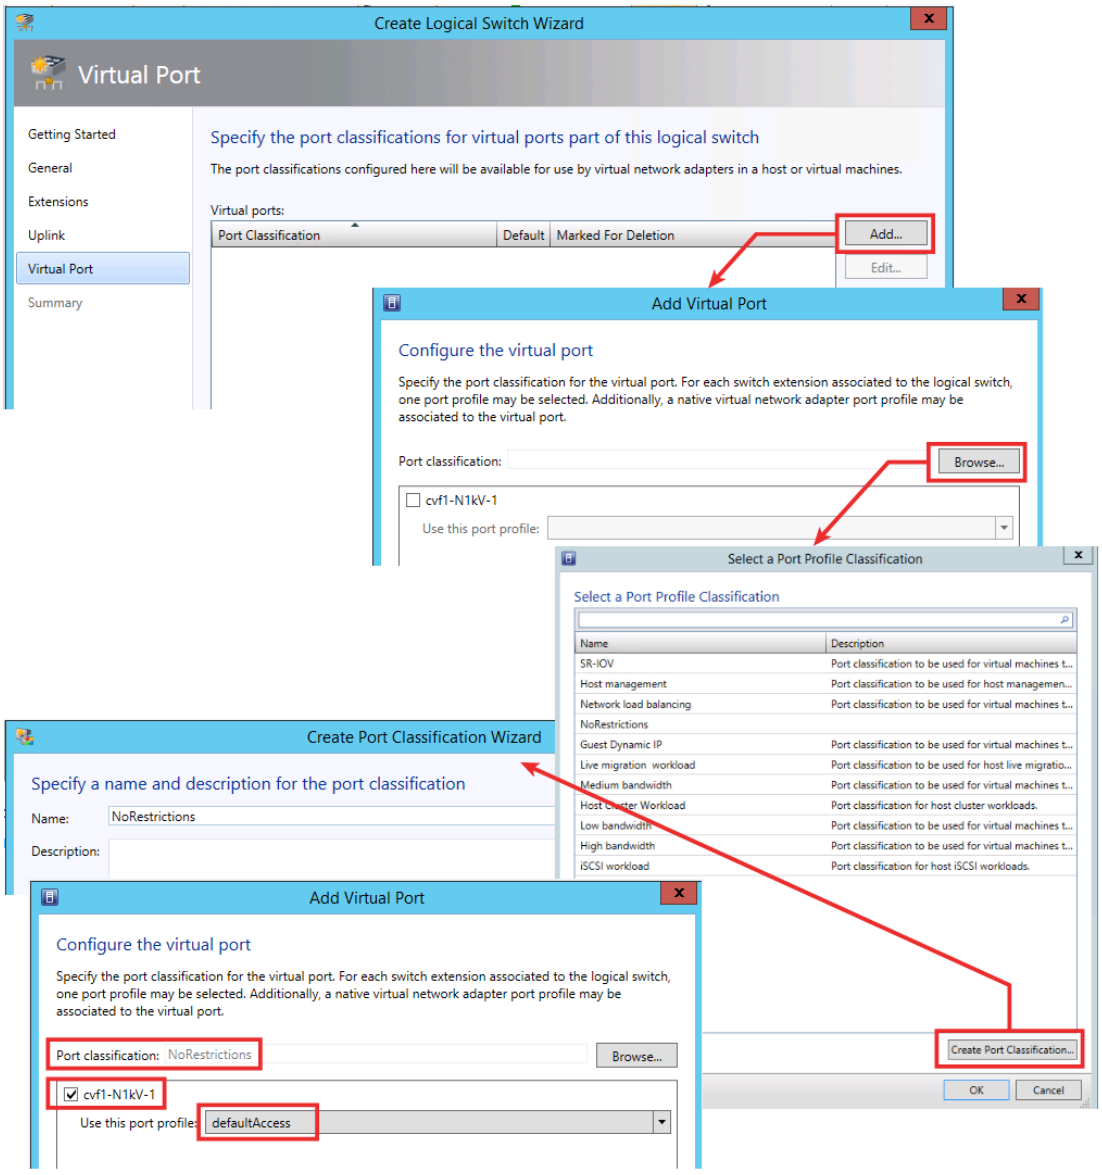

VM Networks now need to be configured. These are created under VMs and Services>VM Networks>Create VM Network.

1. Name the network and select the Logical Network that was configured through the Nexus 1000V CLI.
2. Select the Nexus 1000V instance that was created.
3. For Isolation Options click on Specify an Externally Supplied VM Network and use the pulldown arrow to select a Network Segment that was created through the Nexus 1000V CLI. Only the segments that are assigned to the Logical Network in Step 1 are available and as they are chosen, they are removed from the selection list.
4. Hit Next and then Finish.

Figure 4-21 VM Network Creation

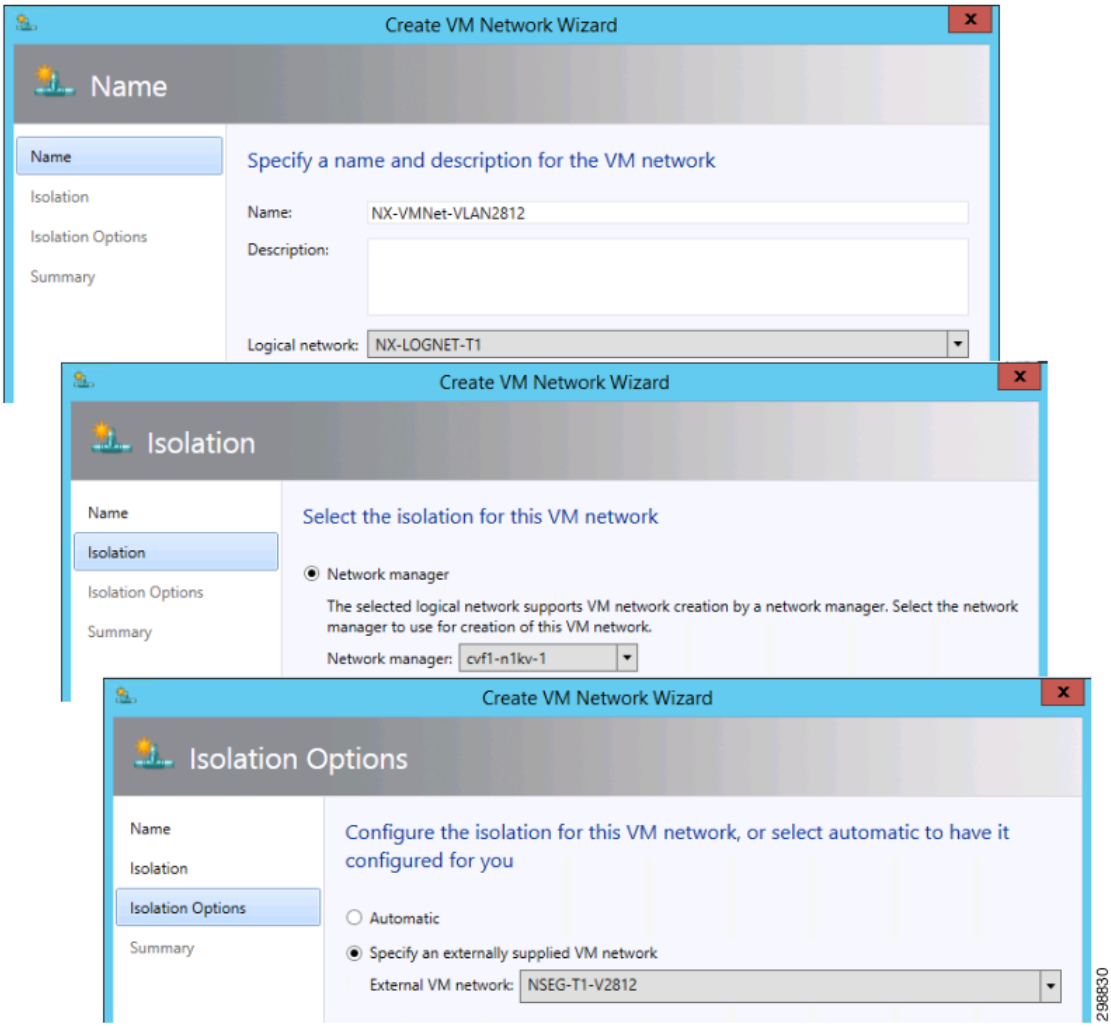

A new virtual switch needs to be added to each host in the cluster. This connects the Nexus 1000V to the host blade by installing the VEM software. From the VMs and Services>All Hosts, right click on the Host>Properties>Virtual Switches and then proceed with the following steps.

1. Click on New Virtual Switch>New Logical Switch.
2. From the pulldown arrow select the Logical switch that was created in prior steps.
3. Click on Add to include all of the vNICs that are configured from the UCS for this host blade and Virtual switch to use. For the Data Uplink there are two interfaces.
4. Select the Uplink Port Profile from the pulldown. Only the NSM Network Uplinks that were created through the Nexus 1000V CLI will be available.
5. Click Ok to save the Virtual Switch configuration.

Figure 4-22 Host Virtual Switch Config

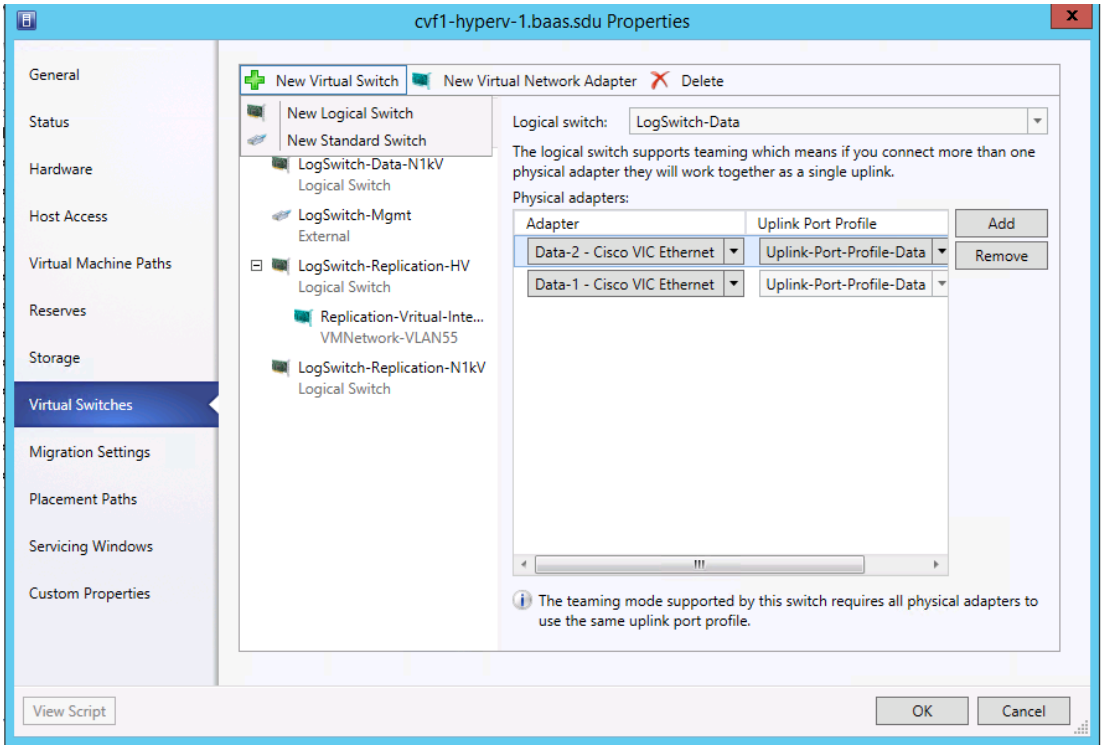

Once the configuration job for the Virtual Switch is complete, the VEM software has been automatically downloaded to the host and shows up as a module on the Nexus 1000V

```
cvf1-N1kV-1# show module
Mod  Ports  Module-Type                      Model              Status
---  -----  -------------------------------  -----------------  ------------
1    0      Virtual Supervisor Module        Nexus1000V         active *
3    288    Virtual Ethernet Module          NA                 ok
4    288    Virtual Ethernet Module          NA                 ok
5    288    Virtual Ethernet Module          NA                 ok
Mod  Sw                  Hw
---  ------------------  ------------------------------------------------
1    5.2(1)SM3(1.1)      0.0
3    5.2(1)SM3(1.1)      Windows Server 2012 R2 - Datacenter (6.3.9600, 6.40)
4    5.2(1)SM3(1.1)      Windows Server 2012 R2 - Datacenter (6.3.9600, 6.40)
5    5.2(1)SM3(1.1)      Windows Server 2012 R2 - Datacenter (6.3.9600, 6.40)
```

Create the Virtual Switch on all of the hosts in the cluster before assigning VM Networks to a VM's interfaces. To use network resources from the Nexus 1000V, click on Hardware Configuration under the VM's Properties page. Then perform the following steps.

1. Click on Connected to VM network and browse for the correct selection. This will fill in the selection for VM subnet using the Network Segment that is configured on the Nexus 1000V.
2. If the Enable VLAN is active, select it and verify that the VLAN is correct.
3. The Logical switch will be filled in. Make sure that the information is accurate.
4. The Classification field will be empty. You must select an entry for proper operation. Choose the Classification that was created in earlier steps.
5. Click Ok to save the change.

Figure 4-23 VM Interface

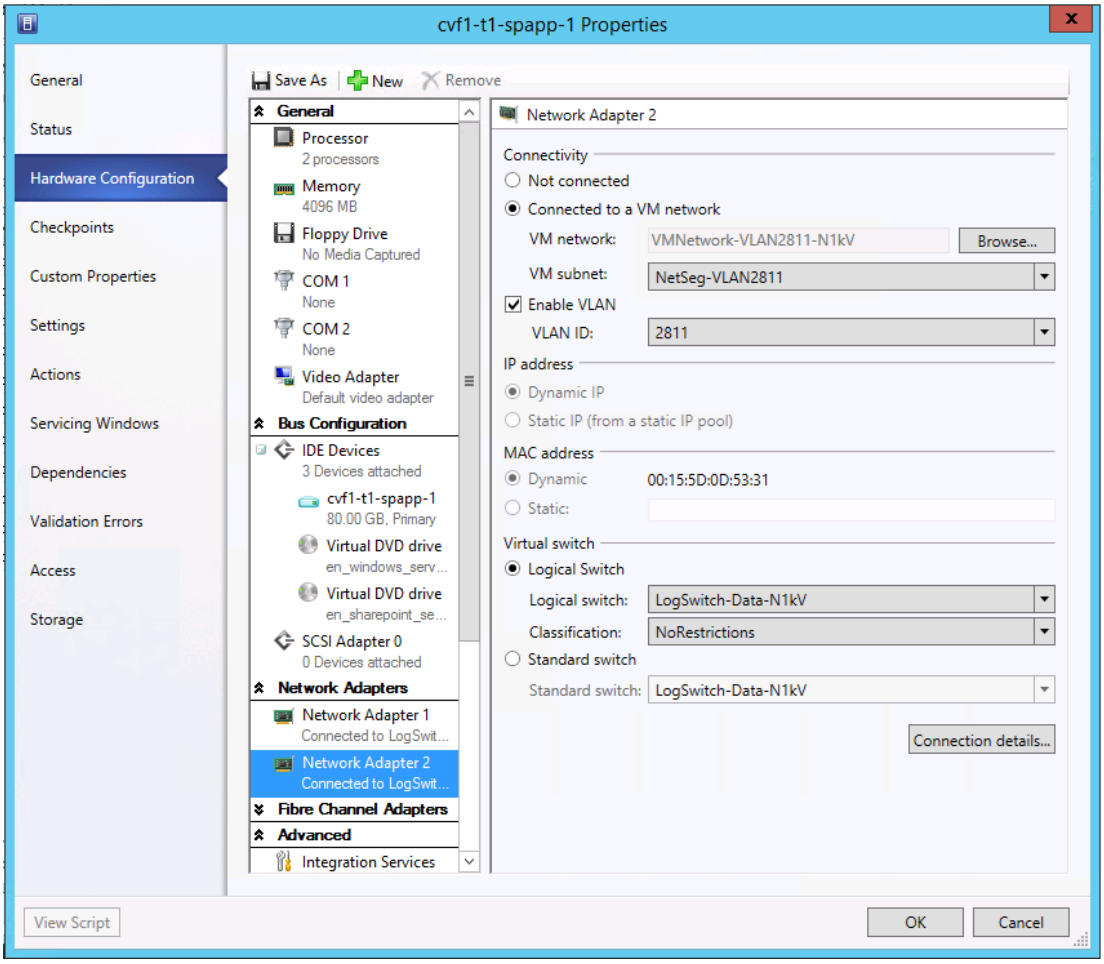

In some instances, the Hyper-V host may need to have an interface in a VLAN for management, monitoring, or replication traffic purposes. For example, in the validation testing, the Hyper-V hosts had an interface in the replication VLAN, which was used by the CommServe Manager to install a VSA on the hosts.

To configure an interface, select New Virtual Network Adapter under the Virtual Switches tab of the host properties window. The name you give the adapter will show up as a network connection on the host blade.

Figure 4-24 Host Virtual Interface Configuration

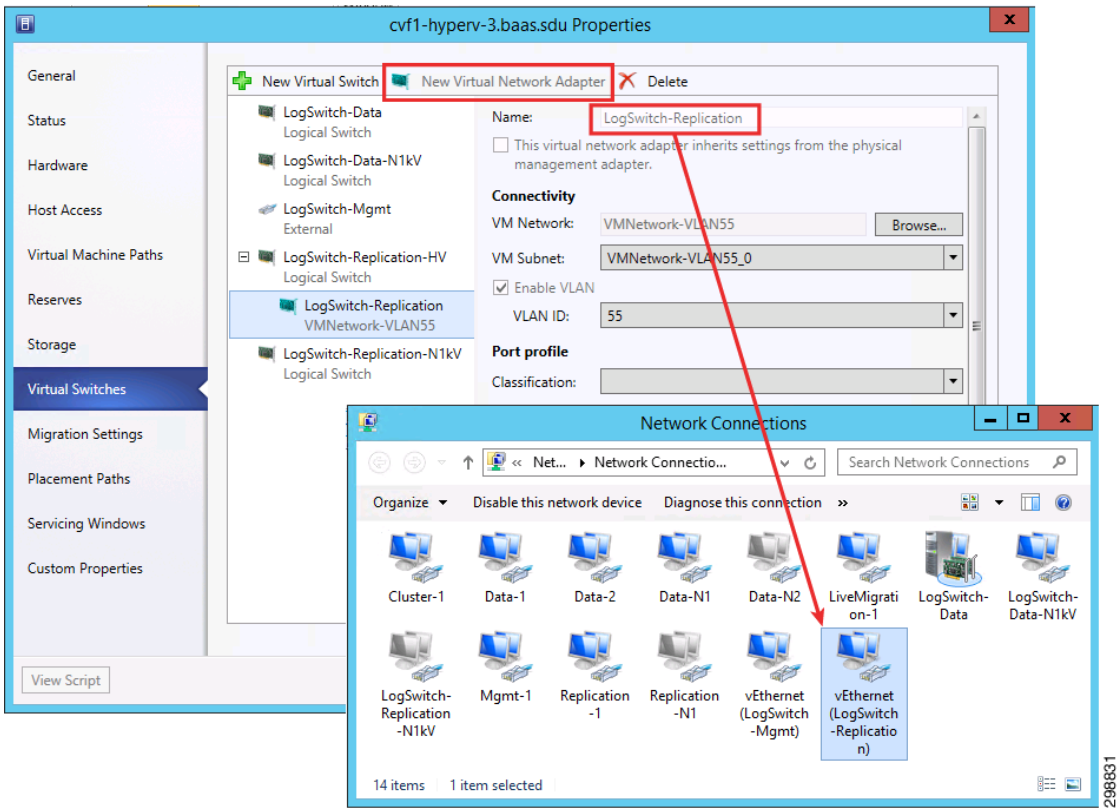

When adding a New Virtual Network adapter to the Hyper-V logical switch, configure and commit the switch before adding the New Virtual Network adapter. If it is configured at the same time, it could result in an error.

Refer to the following documents for details on deploying the Cisco Nexus 1000V in a Microsoft Hyper-V environment:

- Cisco Nexus 1000V Switch for Microsoft Hyper-V Deployment Guide
- Nexus 1000v Quickstart with Hyper-V Server Configuration Example
- Cisco Nexus 1000V for Microsoft Hyper-V Troubleshooting Guide, Release 5.2(1)SM3(1.1)

SAN Storage

Due to lab equipment availability, EMC VMAX Fibre Channel storage was presented to the Hyper-V cluster hosts in the SP1 site as various LUNs (datastores). The actual size and naming of these datastores was irrelevant to this project as these FC storage specifics were out of scope. CCA-MCP uses NetApp SAN storage in its architecture design.

Commvault Components

This section details the hardware used as Commvault infrastructure.

Compute

As discussed in the Architecture section above, each Commvault component is designed to be sized for the type of environment being supported. Here we will describe how each building block was sized and configured for the validation environments.

```
CommServe – cvf1-cs-1
    UCS C240 M3S; Dual Socket 12 Core Intel Xeon E5-2630
    256GB DDR3 1600MHz
    BIOS: 1.5.4f.0; FW: 1.5(4); VIC: 2.2(1b)
    LSI 927-8i Mega-RAID SAS HBA
        (2) 500GB HDD (RAID1)
            Windows Server 2008 Standard R2 SP1
            CVLT Simpana CommServe Software
    CVLT MSSQL Database
    UCS VIC 1225
        192.168.1.87 – 1GbE Mgmt Network
        10.5.5.87 – 10GbE Backup Network
MediaAgent(s) – cvf1-ma-1; cvf1-ma-2
    UCS C240 M3S; Dual Socket 12 Core Intel Xeon E5-2630L
    96GB DDR3 1333MHz
    BIOS: 2.0.3.0; FW: 2.0(3d); VIC: 4.0(1e)
    LSI 927-8i Mega-RAID SAS HBA
        (2) 250GB HDD (RAID1)
            Windows Server 2012 Standard R2
            CVLT MediaAgent Software
    CVLT IndexCache
(3) 400GB SSD (RAID5)
    CVLT Short-Term DDB
    CVLT Long-Term DDB
    UCS VIC 1225
        192.168.1.85,86 – 1GbE Mgmt Network
        10.5.5.85,86 – 10GbE Backup Network
Reporting Server – cvf1-rs-1
    UCS C240 M3S; Dual Socket 16 Core Intel Xeon E5-2650
    192GB DDR3 1600MHz
    BIOS: ??; FW: ??; VIC: ??
    LSI 927-8i Mega-RAID SAS HBA
        (2) 1TB HDD (RAID1)
            Windows Server 2012 Standard R2
            CVLT Simpana Enterprise Metrics Reporting Software
    CVLT Reporting MSSQL Database
    UCS VIC 1225
        192.168.1.88 – 1GbE Mgmt Network
        10.5.5.88 – 10GbE Backup Network
Web Proxy – cvf1-dmz-1
    Hyper-V presented Virtual Server
    8 vCPU; 32GB vRAM
    vDisk
        80GB HDD
            Windows Server 2012 Standard R2
            CVLT Simpana Web Proxy Software
    CVLT Simpana Web Console Software
    vAdapters
        192.168.1.95 – 1GbE Mgmt Network
            ASA NAT - cvf1-dmz-proxy
        10.5.3.95 – 10GbE Backup Network
```

Storage

Network Attached Storage (NAS) will be used for the Disk Library storage required to retain the data for the tenants in the SP1. NAS is being used to allow for the sharing of the Disk Library between the two MediaAgents, providing failover and load balancing capabilities.

```
Disk Library Disk: NFS server (CentOS)
192.168.13.23 (Mgmt)
10.5.5.149 (Replication)
UCS-C240-M3S
Running CentOS 6
(5) 1TB drives in RAID 5 (configured at h/w level)
NFS mount at /exports/files
```

Networking

Generally each server within the CommCell will have two different network adapters, one strictly for management access to the server and one much faster network adapter used for the backup data traffic. The validation environment was setup following that standard with 1GbE adapters with 192.168.x.x addresses for the management network and 10GbE adapters with 10.5.x.x addresses for the backup network. Backup or replication traffic that is traversing the backup network will be utilizing TCP ports 8400 – 8402, as well as randomized ports for data transfer, which was limited to 32768 – 65535 when crossing the firewall to get to SP2 or the Enterprise site. The Backup Network will also be used for the presentation of the NFS File Shares from the NAS.

To allow for communication from external customer networks a Commvault Simpana Web Proxy will be used. The Commvault Simpana Web Proxy is a special proxy configuration where a dedicated iDataAgent is placed in a perimeter network, and the firewalls are configured to allow connections (from inside and outside networks) into the perimeter network. The proxy, which is the agent running in the perimeter network authenticates, encrypts, and proxies accepted tunnel connections to connect the clients operating outside to clients operating inside.

The Simpana proxy acts like a Private Branch Exchange (PBX) that sets up secure conferences between dial-in client calls. With this setup, firewalls can be configured to disallow straight connections between inside and outside networks.

Figure 4-25 shows a perimeter network setup where a client from outside communicates to the CommServe and MediaAgent operating in an internal network through the Simpana proxy.

Figure 4-25 Commvault Simpana Web Proxy

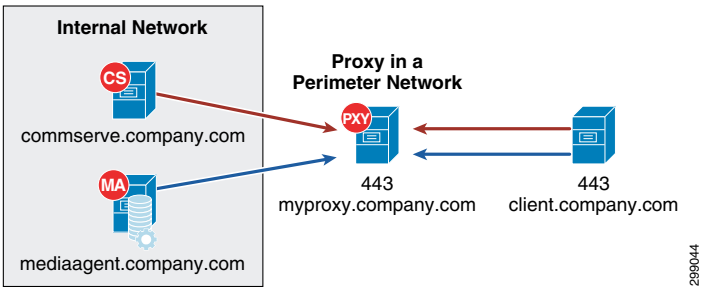

Physical Server Protection

This section describes the iDataAgents and their capabilities for the physical server types being protected at this location.

Windows File System iDataAgent

The Windows File System iDataAgent provides unified protection and recovery for file system data residing on Windows clients. In addition to complete protection of file system data for disaster recovery, the Windows File System iDataAgent also provides more granular backup and recovery options that operate seamlessly with your data protection. Further options for deduplication, job-management and reporting help to make sure that all of your file system data are easily traceable and retrievable whenever the need arises. The Windows File System iDataAgent offers the following key features:

- **Simplified Data Management**

  The Windows File System iDataAgent enables easy management of all the Windows systems in your environment, by providing a singular approach to manage the data using the same unified console and infrastructure.

- **Point-in-Time Recovery**

  In the event of a serious system failure, such as the breakdown of hardware, software, or operating systems, the Windows File System iDataAgent provides point-in-time recovery of files at any given time.

- **System State**

  The Windows File System system state contains many components and services that are critical to recovery of the Windows operating system. The system state is backed up and restored as part of Windows File System iDataAgent backup and restore.

- **Office Communications Server**

  The Office Communications Server contains data, settings, and metadata that are critical to data protection operations residing in both the File System and SQL databases. In order to fully protect the Office Communications Server, the OCS data settings and the OCS metadata must be backed up.

- **Backup and Recovery Failovers**

  In the event that a MediaAgent used during a backup or recovery operation fails, the operation will automatically resume on an alternate MediaAgent from the point of failure. This is especially useful for backups and restores of large amounts of file system data.

  In the event, that a network goes down, the backup and recovery jobs are resumed on alternate data paths. Similarly, in the event of a device failure, the jobs are automatically switched to alternate disk and tape drives.

- **Efficient Job Management and Reporting**

  You can view and verify the status of the backup and recovery operations from the Job Controller and Event Viewer windows within the CommCell Console. You can also track the status of jobs using reports, which can be saved and easily distributed. Reports can be generated for different aspects of data management. You also have the flexibility to customize the reports to display only the required data and save them to any specified location in different formats. For example, you can create a backup job summary report to view completed backup jobs at-a-glance . In addition, you can schedule these reports to be generated and sent by email without user intervention.

- **Block Level Deduplication**

  Deduplication provides a smarter way of storing data by identifying and eliminating the duplicate items in a data protection operation.

  Deduplication at the data block level compares blocks of data from multiple files against each other. For example, if two or more files contain blocks of data that are identical to each other, then block level deduplication eliminates storing the redundant data thereby reducing the size needed for storage. This way also dramatically reduces the number of backup data copies on both disk and tapes.

- **Add-On Components**
  - 1-Touch

    1-Touch recovery helps to recover a crashed system in the least amount of time. By automatically rebuilding the operating system, you can recover systems with defective components such as inaccessible volumes or crashed disks. You don't need to reinstall the individual software packages or operating systems manually.
  - Simpana OnePass

    The Simpana OnePass is an integrated File System agent that backs up and archives the qualified data. Simpana OnePass also reclaims backup storage space when files and stubs are deleted on the primary storage.

Linux File System iDataAgent

Simpana software provides a simplified end-to-end protection of file system data residing on all the Linux computers in your enterprise. In addition to complete protection of file system data for disaster recovery, it also provides a robust and comprehensive backup and recovery solution with significant speed performance and efficient use of disk and tape drives. It also assists you in full system rebuilds and eliminates recovery failures. The Linux File System iDataAgents offers the following key features:

- **Simplified Data Management**

  The Linux File System iDataAgents enables easy management of all the Linux systems in your environment, by providing a singular approach to manage the data using the same unified console and infrastructure.
- **Point-in-Time Recovery**

  In the event of a serious system failure, such as the breakdown of hardware, software, or operating systems, the Linux File System iDataAgent provides point-in-time recovery of files at any given time.
- **Backup and Recovery Failovers**

  In the event that the MediaAgent used for the backup or recovery operation fails, the backup is automatically resumed on alternate MediaAgents. In such cases, the backup or restore job does not restart from the beginning, but resumes from the point of failure. This is useful for backups and restores of large amount of file system data.

  In the event, that a network goes down, the backup and recovery jobs are resumed on alternate data paths. Similarly, in the event of a device failure, the jobs are automatically switched to alternate disk and tape drives.
- **Efficient Job Management and Reporting**

  You can view and verify the status of backup and recovery operations from the Job Controller and the Event Viewer within the CommCell Console. You can also track the status of the jobs using Reports, which can be saved and distributed. Generate reports for different aspects of data management. Customize the reports to display only the required data and save them to a specific location in different formats. For example, you can create a backup job summary report to view the completed backup jobs.

  You can schedule, generate and send the Reports via email without user intervention.
- **Block-Level Backup**

  Block-level backup is a faster method to back up data because only the extents that contain data are backed up, rather than the entire files.

Block-level backups provide better performance over file system backups and disk image-based backups if the file system has a large number of small files by reducing the scan times. Also, when compared to file system incremental backups, block-level incremental backups run faster and back up less data if the file system has very large files. By default, block-level backups are performed using native snaps, but they can be configured to function with hardware snap engines.

- **Block Level Deduplication**

  Deduplication provides a smarter way to store data by identifying and eliminating the duplicate items in a data protection operation.

  Deduplication at the data block level compares blocks of data against each other. If an object (e.g., file, database) contains blocks of data that are identical to each other, then block level deduplication does not store the redundant data, which reduces the size of the object in storage. This reduces the size of the backup data copies on both the disk and tapes.

- **Add-on Components**
  - 1-Touch

    1-Touch recovery helps to recover a crashed system in the least amount of time. By automatically rebuilding the operating system, you can recover systems with defective components such as inaccessible volumes or crashed disks. You don't need to reinstall the individual software packages or operating systems manually.

  - Simpana OnePass

    The Simpana OnePass is an integrated File System agent that backs up and archives the qualified data. Simpana OnePass also reclaims backup storage space when files and stubs are deleted on the primary storage.

  - Content Indexing and Search

    Content Indexing and Search enables users to content index their data and later search the data from a user-friendly web interface. The users can also perform restore operations or other advanced actions on the searched data.

Virtual Server Protection

This section describes the iDataAgents and their capabilities for the virtual environments being protected in this location.

Virtual Server iDataAgent for Hyper-V

The Virtual Server iDataAgent, or VSA, provides a unified protection and recovery vehicle for all virtual machine data in a Hyper-V cluster or a standalone Hyper-V server. In addition to complete protection of entire virtual machines for disaster recovery, the Virtual Server iDataAgent also provides granular backup and recovery options. The additional options such as customized automatic discovery, deduplication, and reporting ensure all your virtual machine data is easily traceable and retrievable whenever the need arises.

Figure 4-26 displays various components of Virtual Server iDataAgent and how the data flows:

Figure 4-26 Virtual Server iDataAgent Components and Data Flows

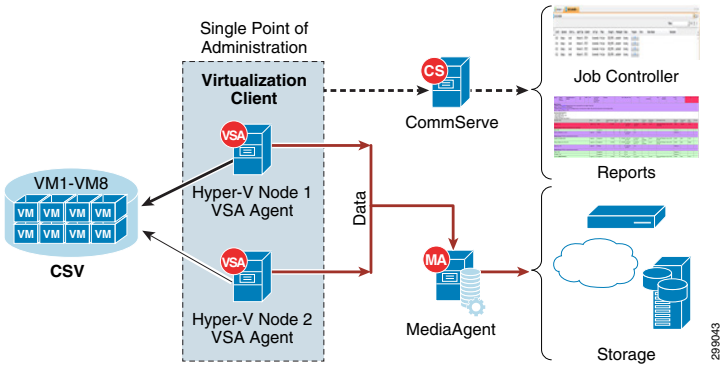

The Virtual Server iDataAgent offers the following key features:

- **Virtualization Client**

  The Virtualization client serves a single point of administration for all cluster nodes. All administration actives such backups, restores, schedules, reports etc can be performed from the Virtualization Client.

- **Protection of Migrated Virtual Machines**

  The cluster will be seamlessly backed up even if the virtual machine migrates from one host to another. The incremental backup cycles will also be maintained.

- **Automatic Protection for All Virtual Machines in a Cluster or Server**

  Once you configure the Virtual Server iDataAgent for a cluster or a server, all the virtual machines in the Hyper-V Cluster or server are automatically selected for backup. This behavior is designed to ensure all virtual machines are backed up.

- **Customized Discovery**

  If you want to backup only specific virtual machines in a cluster or a server, you can set criteria to search the virtual machines and automatically select them for backup. This is useful in environments where virtual machines are frequently added, or removed . You can easily browse and select objects such as CSV or Hosts to set criteria for automatic discovery.

- **Customized Filters**

  If you want to exclude specific virtual machines from the backup, you can set criteria to filter virtual machines.

- **Report**

  You can view a detail report about each backed up virtual machine. It contains information such as size of the data, status of integration services, guest operating system, the host on which the virtual machine is running etc. You can view all this information from the CommCell console when the backup job is running. It appears on the Virtual Machine Status tab of the View Job Details dialog box.

- **IntelliSnap Backup**

  IntelliSnap Backup enables you to create a point-in-time snapshot of a virtual machine by temporarily stilling the data, taking a snapshot, and then resuming live operations. IntelliSnap backups work in conjunction with hardware snapshot engines.

- **Block Level Deduplication**

  Deduplication provides a smarter way of storing data by identifying and eliminating the duplicate items in a data protection operation.

Deduplication at the data block level compares blocks of data against each other. If virtual machines contains blocks of data that are identical to each other, block level deduplication eliminates storing the redundant data and reduces the size of the data in storage. This dramatically reduces the virtual machine backup data copies on both the disk and tapes.

Application Protection

This section describes the iDataAgents and their capabilities for the applications and databases protected in this location.

Microsoft SQL Server iDataAgent

The Microsoft SQL Server *i*DataAgent provides a simplified end-to-end backup and recovery solution for SQL data in your enterprise. The product can be used to perform both full system rebuilds and granular recovery of the data. Some of the key features of the SQL iDataAgent are:

- **Full Range of Backup Options**

  The SQL *i*DataAgent provides the flexibility to backup the SQL database from different environments. You can perform a full or incremental backup of the entire instance, individual databases or files and file groups, and the transaction logs at any point of time as described below:

- **Database Backups**

  You can backup both the system and user-defined databases. You can comprehensively backup all the databases in an instance or schedule backups for the individual databases. You can also auto-discover new databases to comprehensively manage the backup of all databases in your environment.

- **Transaction Log Backups**

  Transaction log backups captures the transaction log whether the transaction was committed or not. The use of transaction log backups make point in time recovery possible. You can restore to any point in time within the transaction log.

- **File and File Groups Backups**

  Files or file group backups allows you to backup individual files or file groups. This functionality can be critically important, especially for large databases. Whereas a full database backup captures all files of a given database, file and file group backups allow you to back up selected portions of a database individually. As with database backups, the system provides the option of performing full, differential, and transaction log backups of file and file groups. Note that when running a transaction log backup for a File/File Group subclient, the database log is automatically backed up.

- **Advanced SQL Server Restore Capabilities**

  The SQL *i*DataAgent provides the ability to recover databases or entire SQL instance. There is no mounting, no recovery wizards, no extra steps needed – the software takes care of it all. This includes the following abilities:

  - Full or Partial Restore databases
  - Restore and replay transaction logs
  - Set Database state during restore (Recovery, Standby, No Recovery)
  - Point-in-time recovery

- **Efficient Job Management and Reporting**

  You can view and verify the status of SQL backup and recovery operations from the Job Controller and Event Viewer windows within the CommCell Console. You can also track the status of the jobs using Reports, which can be saved and easily distributed. Reports can be generated for different

aspects of data management. You also have the flexibility to customize the reports to display only the required data and save them to any specified location in different formats. For example, you can create a backup job summary report to view at-a-glance the completed backup jobs.

In addition, you can also schedule these reports to be generated and send them on email without user intervention.

- **Backup and Recovery Failovers**

  In the event that a MediaAgent used for the backup or recovery operation fails, it is automatically resumed on alternate MediaAgents. In such cases, the backup or restore job will not restart from the beginning, but will resume from the point of failure. This is especially useful for backups and restores on large SQL databases.

  In the event, that a network goes down, the backup and recovery jobs are resumed on alternate data paths. Similarly, in the event of a device failure, the jobs are automatically switched to alternate disk and tape drives.

- **Block Level Deduplication**

  Deduplication provides a smarter way of storing data by identifying and eliminating the duplicate items in a data protection operation.

  Deduplication at the data block level compares blocks of data against each other. If an object (file, database, etc.) contains blocks of data that are identical to each other, then blocvk level deduplication eliminates storing the redundant data and reduces the size of the object in storage. This way dramatically reduces the backup data copies on both the disk and tapes.

- **SnapProtect Backup**

  SnapProtect Backup enables you to create a point-in-time snapshot by temporarily quiescing the data, taking a snapshot, and then resuming live operations. SnapProtect backups work in conjunction with hardware snapshot engines.

Microsoft SharePoint iDataAgent

Microsoft SharePoint Server and Windows SharePoint Services include components that are backed up by the SharePoint Server iDataAgent, as well as data which must be backed up using the File System iDataAgent. SharePoint database files can also reside on separate SQL servers; to secure this data, you must back up these files using the appropriate SQL Server iDataAgent. The SharePoint entities that can be backed up by the system are described in detail in the following sections.

Figure 4-27 SharePoint Web/Application Tier

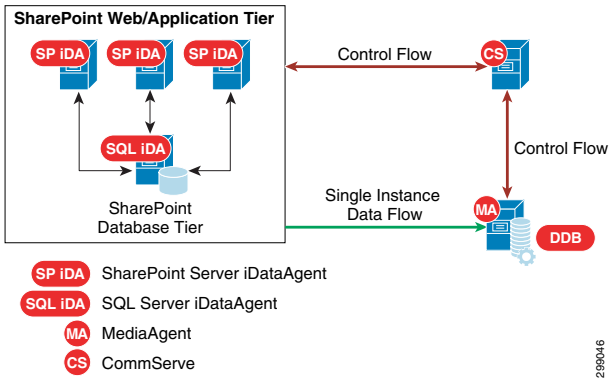

The SharePoint Server iDataAgent offers the following key features:

- **Full Fidelity Backup and Recovery**

Document level backups allows capturing of system oriented metadata such based on parameters such as modified by or version information in a canned manner so that this backed up information can be brought back as it was backed up. In place merge options determines the current state of the target entities and intelligently determines what to restore. The customization on top of base template can also be seamlessly captured during backups.

- **Simplified Backup and Recovery Strategy**

  Prior to v10, there were many options to choose from within a particular level of protection. Also there were less number of choices/strategies which optimized each level of protection. The backup and restore procedures in v10 have simplified the overall offerings as well as complexities and therefore the iDataAgent is more resilient against failures.

- **Fast Backup and Recovery for Content Database**

  Snapshot of Content databases allows for faster backups of the SharePoint data. By using the offline browse and IntelliSnap feature, the same granularity of data can be recovered quickly from a snapshot.

- **Protect all SharePoint Components**

  All entities in the File System or IIS related to SharePoint are protected along with all logical SharePoint components.

- **RBS Aware Backup for Content Database**

  Content database backups are fully RBS aware* and protect RBS blobs along with databases regardless of backup methods (streaming, VSS, or IntelliSnap).

Note *Supports only Microsoft provided File Stream RBS provider.

- **Efficient Job Management and Reporting**

  You can view and verify the status of the backup and recovery operations from the Job Controller and Event Viewer windows within the CommCell Console. You can also track the status of the jobs using reports, which can be saved and easily distributed. Reports can be generated for different aspects of data management. You also have the flexibility to customize the reports to display only the required data and save them to any specified location in different formats. For example, you can create a backup job summary report to view at-a-glance the completed backup jobs. In addition, you can also schedule these reports to be generated and send them on email without user intervention.

- **Block Level Deduplication**

  Deduplication provides a smarter way of storing data by identifying and eliminating the duplicate items in a data protection operation.

  Deduplication at the data block level compares blocks of data against each other. If an object (file, database, etc.) contains blocks of data that are identical to each other, then block level deduplication eliminates storing the redundant data and reduces the size of the object in storage. This way dramatically reduces the backup data copies on both the disk and tapes.

- **Content Indexing and Search**

  Content Indexing and Search enables users to content index their data and later search the data from a user-friendly web interface. The users can also perform restore operations or other advanced actions on the searched data.

Oracle / Oracle RAC iDataAgent

The Oracle RAC iDataAgent provides a simplified end-to-end backup and recovery solution for Oracle databases in your enterprise without using multiple subclients and storage policies. Itv also allows you to load-balance Oracle backups and restores across multiple database nodes. The product can be used to perform both full system rebuilds and granular recovery of data and logs. Some of the key features are:

- **Full Range of Backup and Recovery Options**

  The Oracle RAC iDataAgent provides the flexibility to backup the Oracle database in different environments. This is very essential since the Oracle database is always subject to constant changes.

  You can perform a full or incremental backup of the entire database or individual data files and tablespaces, and archive logs at any point in time. The following section describes the backups that can be performed in different environments.

  In configurations where there is no RMAN catalog, the Oracle control file can be used as an alternative.

- **Offline Backup**

  When the database is mounted, and not open or available for use, perform a full backup of the database without the archive logs. Use the offline backup when the data is consistent, and there are no transactions in the database.

- **Online Backup**

  When you cannot bring down the database to perform an offline backup, use the online backup method. Perform full or incremental backups when the database is online and in ARCHIVELOG mode. Use online backups to perform a point-in-time restore of the database.

  You can also backup the archive logs only when the database is online. These logs can be applied to an online backup to recover the database to the current point-in-time.

  You can also protect the non-database files and profiles using the appropriate File System iDataAgent.

- **Selective Online Full Backup**

  You can backup and store copies of valid data from a source copy of a specific storage policy to all or one active secondary copy within a storage policy to provide a better tape rotation. An online full backup job is copied to a selective copy if the full backup job cycle completes successfully. This allows you to select, store and protect your valuable data on a secondary copy for future restores.

- **Effective Management of Backups and Restores**

  This iDataAgent allows you to group any desired number of Oracle iDataAgent instances under one or more Oracle RAC database logical entities. As such, Oracle backups and restores as well as other job types and functions (including Data Aging, Scheduling, Job Management) are all consolidated and easy to manage. This allows you to maintain your data irrespective of whether you add or remove Oracle iDataAgent instances from the RAC database.

- **MultiStreaming for Accelerated Backups and Restores**

  You can configure various resources on the RAC nodes and include more number of streams for accelerated backup and restore.

- **Backup and Recovery Failovers**

  In the event that the MediaAgent used for the backup or recovery operation fails, the backup is automatically resumed on alternate MediaAgents. In such cases, the backup or restore job does not restart from the beginning, but resumes from the point of failure. This is useful for backups and restores of large amount of file system data.

In the event, that a network goes down, the backup and recovery jobs are resumed on alternate data paths. Similarly, in the event of a device failure, the jobs are automatically switched to alternate disk and tape drives.

Also, this iDataAgent automatically checks the status of each Oracle instance during a backup or restore and allocates RMAN channels only for the instances that are active. Therefore, even if a specific instance fails, the backup or restore will continue.

- **Efficient Job Management and Reporting**

  You can view and verify the status of backup and recovery operations from the Job Controller and the Event Viewer within the CommCell Console. You can also track the status of the jobs using Reports, which can be saved and distributed. Generate reports for different aspects of data management. Customize the reports to display only the required data and save them to a specific location in different formats. For example, you can create a backup job summary report to view the completed backup jobs.

  You can schedule, generate and send the Reports via email without user intervention.

- **Block Level Deduplication**

  Deduplication provides a smarter way to store data by identifying and eliminating the duplicate items in a data protection operation.

  Deduplication at the data block level compares blocks of data against each other. If an object (e.g., file, database) contains blocks of data that are identical to each other, then block level deduplication does not store the redundant data, which reduces the size of the object in storage. This reduces the size of the backup data copies on both the disk and tapes.

- **Command Line Support**

  In addition to using the CommCell Console, you can perform backup and restore operations from the command line using XML scripts or Qcommands. The options for backup and restore can be selected from the CommCell Console and saved as an XML file. Also, there are many downloadable XML file templates available in the documentation, which can be used to perform specific operations from the command line interface.

SP2 Site Overview

The Cloud Service Provider SP2 site was built to serve as the secondary CSP data center in this validation effort. It was integral in the In-Cloud BaaS use case where two tenants were deployed in SP1 with backups being done remotely to SP2.

IaaS Architecture

The data center infrastructure for the SP2 site is based on CCA-MCP, The CCA-MCP solution is a reference design that can be leveraged by enterprises and service providers to deploy an infrastructure that is efficient, secure, resilient, agile, simple, and scalable.

Cisco UCS

The Cloud Service Provider site SP2 used both UCS B-Series and C-Series compute hardware. The B-Series were used to deploy the Microsoft Hyper-V environment, including the infrastructure and production Hyper-V hosts. The C-Series were used to deploy the Commvault MediaAgents.

B-Series

The UCS B200 M2 and M3 servers were managed by the Cisco UCS Manager (UCSM) running release 2.1(1f). The Service Profile templates were configured to be the same as those described for Cloud Server Provider site SP1. Refer to SP1 Site Overview, page 4-3 for a description of how the B-Series were deployed.

C-Series

The UCS C3160 M3 servers were running release 2.0(2a). The C3160 M3 can only be managed in standalone mode and cannot be managed by the UCSM. The server had a UCS VIC 1227 configured with two vNICs for network connectivity and a RAID controller with (14) 4TB HDD.

RAID Configuration

Figure 4-28shows the RAID configurations that were used on the C3160 M3 for the MediaAgent. The first RAID was configured for RAID1 and was used for the Windows OS and Commvault MediaAgent software. The second RAID was also configured for RAID1 and used for a cache. The last RAID included ten physical drives in a RAID10 configuration and was used for the database.

Figure 4-28 Cisco UCS C3160 M3 RAID Configurations

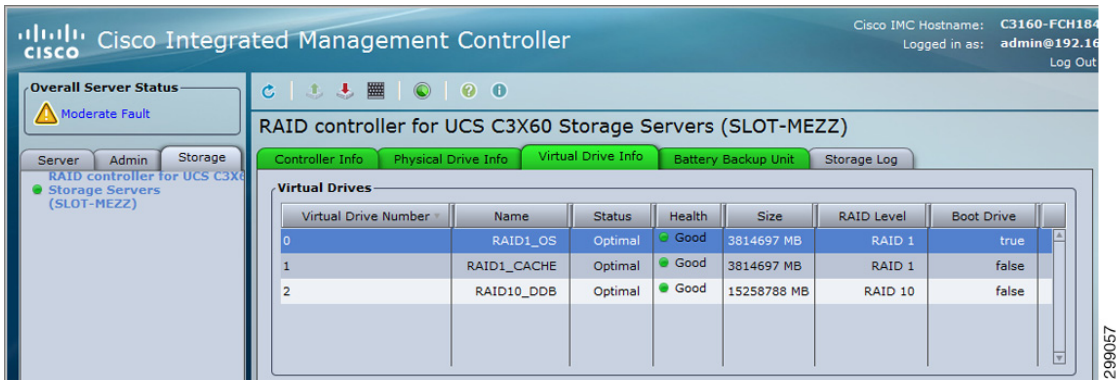

Networking

The Commvault MediaAgents were connected to the infrastructure management via the LOM (LAN on Motherboard) and connected to the replication network via the UCS VIC1227. A virtual Ethernet (vNIC) interface was configured and connected to the Nexus 5548 access layer switch. For the SP2 site, the vNIC interface and the Nexus switchport were configured as trunk ports with a native VLAN set to VLAN56. Alternatively, the vNIC and Nexus switchport could be configured as access ports configured for VLAN56. The latter approach was used in the SP1 site.

Microsoft Hyper-V

The Hyper-V implementation in SP2 was similar to the implementation in SP1. A Hyper-V cluster was created for the management virtual machines and a production cluster was created for the tenant virtual machines. Virtual networking was initially implemented using native Hyper-V as per CCA-MCP deployment and automation model, but this lab implementation used an alternative approach by using Cisco Nexus 1000V for the production hosts by repurposing the UCS vNICs to the Nexus 1000V.

VSA agents are installed on two Hyper-V hosts in the cluster (cvf6-hyperv-3 & cvf6-hyperv-4) with the same resource Reserves that were allocated in SP1.

Refer to the SP1 section for an overview of the Hyper-V implementation and references. Refer to the Cisco Nexus 1000V later in this section for a discussion on how the migration of virtual networking was executed.

Cisco Cloud Services Router (CSR) 1000V

As in SP1, a unique CSR 1000V instance was deployed for each of the two tenants in SP2. The CSR 1000V was configured as the gateway for each of the VLANs used in the tenant workloads. In the validation testing, this included the web, application, and database VLANs. Since Layer 3 connectivity was used between the tenant CSRs at each site, unique IP subnets were configured to enable proper routing between sites. The CSRs used BGP routing to learn IP subnets connected to the remote site's CSR. For the In-Cloud BaaS tenants (Tenants 1-2), an IPsec tunnel was not configured since the network path between the two CSP sites was considered secure.

Cisco Nexus 1000V

In SP2, native Microsoft Hyper-V switching was initially deployed per CCA-MCP automation and deployment model. At a later time, this was changed over to a Nexus 1000V distributed switch. The installation and integration is similar to the roll out for Cloud Services Provider site SP1. Refer to SP1 Site Overview, page 4-3 for details.

SAN Storage

Due to lab availability, EMC VMAX Fibre Channel storage was presented to the Hyper-V cluster hosts in the SP2 site as various LUNs (datastores). The actual size and naming of these datastores was irrelevant to this project as these FC storage specifics were out of scope.

CCA-MCP uses NetApp SAN storage in its architecture design.

Commvault Components

This section details the hardware used as Commvault infrastructure.

Compute

As discussed in the Architecture section, each Commvault component is designed to be sized for the type of environment being supported. Here we will describe how each building block was sized and configured for the validation environments.

```
MediaAgent(s) - cvf6-ma-1; cvf6-ma-2
    UCS C3160; Dual Socket 12 Core Intel Xeon E5-2620
    128GB DDR3 1600MHz
    BIOS: 2.0.2a.0; FW: 2.0(2c); VIC: 4.0(1b)
    Cisco RAID Controller for C3X60
        (2) 250GB HDD (RAID1)
            Windows Server 2012 Standard R2
            CVLT MediaAgent Software
 (2) 4TB SAS (RAID2)
    CVLT IndexCache
    CVLT Short-Term DDB
    CVLT Long-Term DDB
(8) 4TB HDD (RAID10)
    CVLT Disk Library Storage
    UCS VIC 1227
        192.168.60.96,97 - 1GbE Mgmt Network
```

```
        10.5.6.96,97 – 10GbE Backup Network
DR CommServe – cvf6-cs-1
    Hyper-V presented Virtual Server
    12 vCPU; 64GB vRAM
    vDisk
        100GB HDD
            Windows Server 2012 Standard R2
            CVLT Simpana CommServe Software
        300GB HDD
            CVLT MSSQL Database
    vAdapters
        192.168.60.71 – 1GbE Mgmt Network
        10.5.6.71 – 10GbE Backup Network
DR Reporting Server – cvf6-rs-1
Hyper-V presented Virtual Server
    8 vCPU; 16GB vRAM
    vDisk
        100GB HDD
            Windows Server 2012 Standard R2
            CVLT Simpana CommServe Software
        500GB HDD
            CVLT Reporting MSSQL Database
    vAdapters
        192.168.60.71 – 1GbE Mgmt Network
        10.5.6.71 – 10GbE Backup Network
```

Storage

Direct Attached Storage (DAS) is being used for the Disk Library storage with these two MediaAgents. Above and beyond the disk required for the OS, Index Cache, and DDB, there are 8 4TB Drives configured in a RAID10 array that will be used for the Disk Library space. Since the disk is not shared between the two MediaAgents, cvf6-ma-1 will be used as the target for the Long-Term data replication, while cvf6-ma-2 will be used as the target for the Short-Term data replication.

Networking

As previously mentioned, each server within the CommCell will have two different network adapters, one strictly for management access to the server and one much faster network adapter used for the backup data traffic. The validation environment was setup following that standard with 1GbE adapters with 192.168.x.x addresses for the management network and 10GbE adapters with 10.5.x.x addresses for the backup network. Backup or replication traffic that is traversing the backup network will be utilizing TCP ports 8400 – 8402, as well as randomized ports for data transfer, which was limited to 32768 – 65535 when crossing the firewall to get to SP1.

Virtual Server Protection

This section describes the iDataAgents and their capabilities for the virtual environments being protected in this location.

Refer to SP1 Site Overview, page 4-3 for Hyper-V.

Enterprise Site Overview

The Enterprise site was built to serve as the client enterprise data centers in this testing effort.

Architecture

For this release of the BaaS solution, the Enterprise data center architecture was unspecified and intended to be a generic enterprise architecture. In actuality, we leveraged an existing CCA-MCP test lab topology and used multitenancy to cover the various tenants required for validation testing.

Tenant 5 (Microsoft Hyper-V)

The Hyper-V implementation in Enterprise client Tenant 5 was similar to the implementation in SP1. A Hyper-V cluster was created for the management and tenant virtual machines.

VSA agents are installed on two Hyper-V hosts in the cluster (cvf8-hyperv-2 & cvf8-hyperv-3) with the same resource Reserves that were allocated in SP1. Refer to the SP1 section for an overview of the Hyper-V implementation and references.

A unique tenant-specific CSR 1000V instance was deployed for the Tenant 5. The CSR 1000V was configured as the gateway for each of the VLANs used in the tenant workloads. In the validation testing, this included the web, application, and database VLANs. An IPsec tunnel was configured between the CSR 1000V and another tenant-specific CSR 1000V at the SP1 site. The tunnel provided secure communications between the workload VLANs at each site.

Tenant 6 (RHEL OpenStack)

The Enterprise Tenant 6 site used OpenStack Icehouse built on RHEL 7 for compute. The tenant was built using two identical UCS B200 M3 servers in the Enterprise site, the specs for which are shown in Table 4-2.

Table 4-2 Tenant 6 Host Specifications

Hostname	CPU	Memory	Local Storage	Networking
cvf8-t6-ostck-1	2x 8 core	96 GB	2x 300GB SAS	2x vNIC
cvf8-t6-ostck-2	2x 8 core	96 GB	2x 300GB SAS	2x vNIC

Each host had two vNICs, one for data VLANs and one for the management VLAN, as shown in Table 4-3.

Table 4-3 Tenant 6 Host Networking

Hostname	vNIC	VLAN(s)
cvf8-t6-ostck-1	vNIC Data-1	1, 2121, 2122, 2123, 58
	vNIC-Data-2	8
cvf8-t6-ostck-2	vNIC Data-1	1, 2121, 2122, 2123, 58
	vNIC-Data-2	8

The management vNIC was used to provide connectivity to the Horizon management portal while the data vNIC was used to provide connectivity to the internal VLANs and the Commvault replication network.

The host cvf8-t6-ostck-1 was deployed as the OpenStack control node, running the base management services such as Horizon, Glance, Ceilometer, etc. The host cvf8-t6-ostck-2 was deployed as the OpenStack compute node, where the instances (virtual machines) would be deployed.

Figure 4-29, a screenshot of the network topology taken from the Horizon dashboard, shows how this OpenStack tenant was configured. Four networks were configured. Three (Web, App, DB) were configured to mimic a tenant's 3-tier application deployment. The fourth network was the replication network used for backup of the instances via the Commvault Simpana software.

Five instances (called cvf8-test-centos-[1,2,3,4,5]) were deployed in this OpenStack environment, each running CentOS 7. Each instance was connected to two networks, the Replication network and the Web network. A Neutron router was configured to provide Layer 3 connectivity between the three application tiers.

Figure 4-29 Tenant 6 Openstack Network Topology

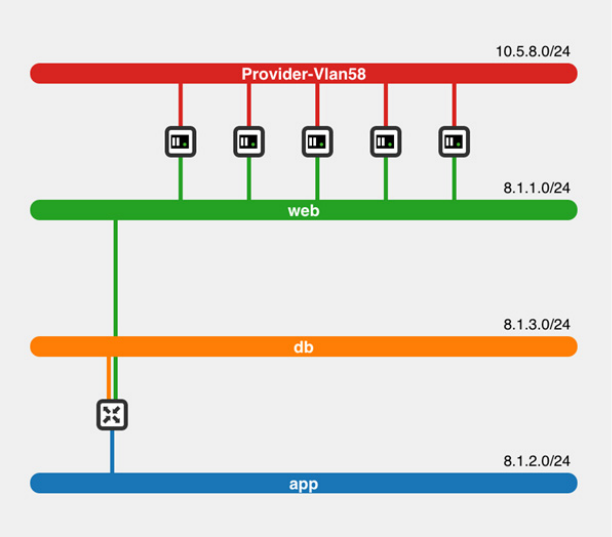

In Tenant 6, the tenant compute was isolated within the OpenStack cloud environment, with its only connectivity to the CSP clouds being on VLAN 58, the Commvault replication network. The gateway for VLAN 58 was configured using HSRP on the Enterprise aggregation switches.

Tenant 7 (VMware vSphere)

Enterprise client Tenant 7 was the only site to use vSphere. vSphere version 5.1 with vCenter, ESXi hosts, and Standard Switches were configured to deploy the required Commvault components, Cisco CSR 1000V, and client virtual machines.

VMware also requires the use of VSA agents. These are not directly loaded on the VMware hosts, but are Windows VMs that are configured with the Virtual Server Agent software. In this solution, Figure 4-30 shows resources reserved, VSA agents cvf8-t7-test-win-1 and cvf8-t7-test-win-1, in Enterprise Tenant 7.

Figure 4-30 VMware VSA Resource Allocation

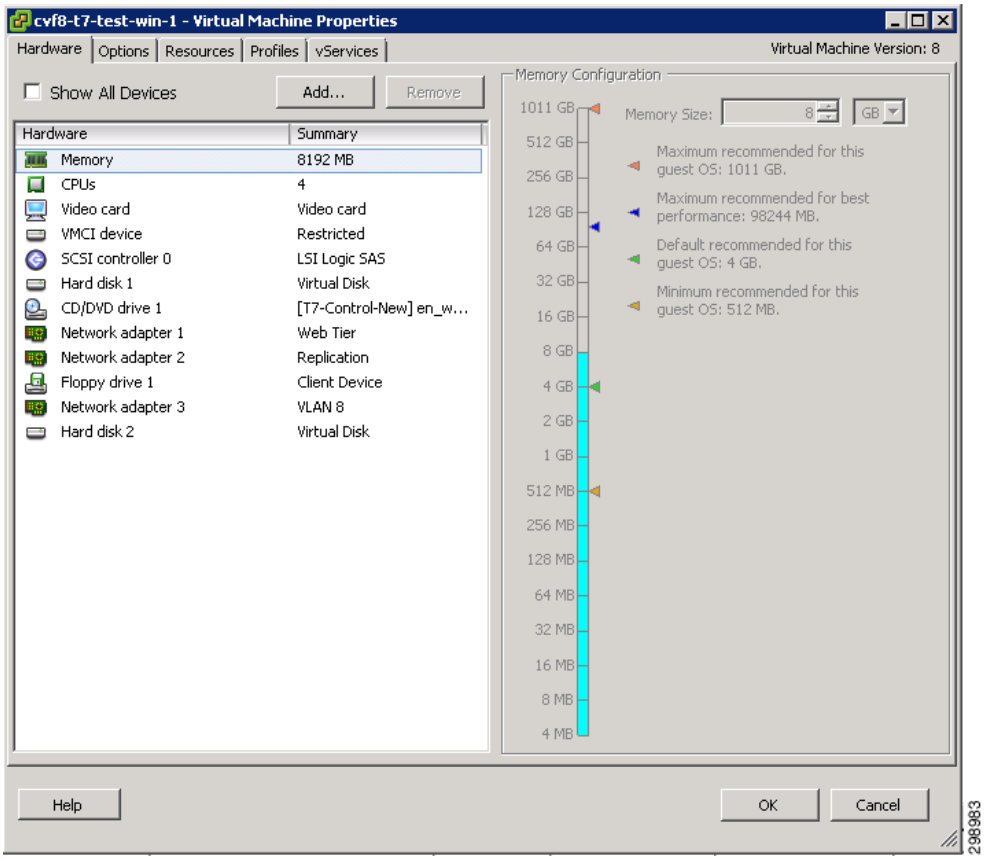

Refer to the following link for more information on Commvault Virtual Server Agent for VMware.

VMware Building Block Guide—Virtual Server Agent for VMware

A unique tenant-specific CSR 1000V instance was deployed for Tenant 7. The CSR 1000V was configured as the gateway for each of the VLANs used in the tenant workloads. In the validation testing, this included the web, application, and database VLANs. An IPsec tunnel was configured between the CSR 1000V and another tenant-specific CSR 1000V at the SP1 site. The tunnel provided secure communications between the workload VLANs at each site.

SAN Storage

Due to lab availability, EMC VMAX Fibre Channel storage was presented to the Hyper-V cluster hosts, OpenStack hosts, and the VMware hosts in the Enterprise site as various LUNs (datastores). The actual size and naming of these datastores was irrelevant to this project as these FC storage specifics were out of scope.

Commvault Components

This section details the hardware used as Commvault infrastructure.

Compute

As discussed in the Architecture section above, each Commvault component is designed to be sized for the type of environment being supported. Here we will describe how each building block was sized and configured for the validation environments.

```
MediaAgent(s) – cvf6-ma-1; cvf6-ma-2
    UCS C240 M3L; Dual Socket 12 Core Intel Xeon E5-2630L
    96GB DDR3 1600MHz
    BIOS: 2.0.3.0; FW: 2.0(3i); VIC: 4.0(1e)
    LSI 9271-8i Mega-RAID SAS HBA
        (2) 3TB SAS (RAID1)
            Windows Server 2012 Standard R2
            CVLT MediaAgent Software
 (2) 3TB SAS (RAID1)
    CVLT IndexCache
    CVLT Short-Term DDB
    CVLT Long-Term DDB
(8) 3TB HDD (RAID10)
    CVLT Disk Library Storage
    UCS VIC 1227
        192.168.8.147,148,149 – 1GbE Mgmt Network
        10.5.8.147,148,149 – 10GbE Backup Network
```

Storage

Direct Attached Storage (DAS) is being used for the Disk Library storage with these two MediaAgents. Above and beyond the disk required for the OS, Index Cache, and DDB, there are 8 4TB Drives configured in a RAID10 array that will be used for the Disk Library space. Since the disk is not shared between the two MediaAgents, cvf6-ma-1 will be used as the target for the Long-Term data replication, while cvf6-ma-2 will be used as the target for the Short-Term data replication.

Networking

As previously mentioned, each server within the CommCell will have two different network adapters, one strictly for management access to the server and one much faster network adapter used for the backup data traffic. The validation environment was setup following that standard with 1GbE adapters with 192.168.x.x addresses for the management network and 10GbE adapters with 10.5.x.x addresses for the backup network. Backup or replication traffic that is traversing the backup network will be utilizing TCPPorts 8400 – 8402, as well as randomized ports for data transfer, which was limited to 32768 – 65535 when crossing the firewall to get to SP1.

Tenant 5 however will be configured to gain access to the CommServe and MediaAgent via the Web Proxy Server as mentioned in the SP1 section. This will limit network activity to Port 8403, with the option of allocating additional ports to increase performance.

Virtual Server Protection

This section will describe the iDataAgents and their capabilities for the virtual environments being protected in this location.

RHEL OpenStack

Commvault Simpana doesn't currently have hypervisor level integration with RHEL OpenStack, but virtually all standard File System and Application Agents can be install directly in the Guest Operating System. For this case we'll be using the Linux File System iDA, see description in SP1.

Virtual Server iDataAgent for Hyper-V

Refer to SP1 Site Overview, page 4-3.

Virtual Server iDataAgent for VMware

The Virtual Server Agent for VMware provides a unified protection and recovery vehicle for all virtual machine data in your vCenter. In addition to complete protection of entire virtual machines for disaster recovery, the Virtual Server Agent provides more granular backup and recovery options. Options such as customized automatic discovery, deduplication, and reporting ensure all your virtual machine data is easily traceable and retrievable whenever the need arises.

Figure 4-31 displays components and data flows for the Virtual Server Agent.

Figure 4-31 Virtual Server Agent Components and Data Flows

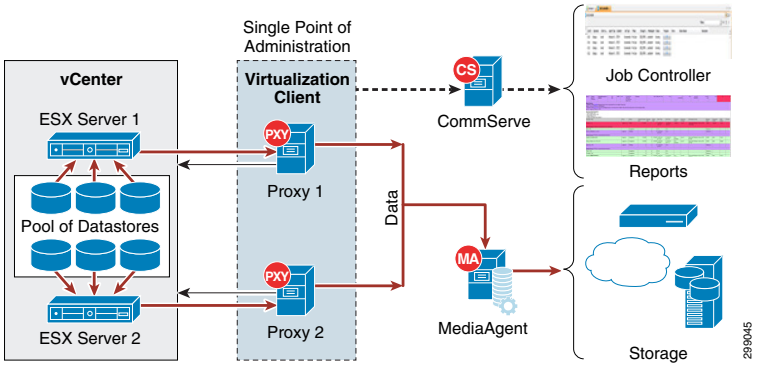

The Virtual Server Agent offers the following key features:

- Single Administration Point

  The virtualization client serves a single point of administration for all proxies. The proxy computer can be any computer where Virtual Server Agent is installed. All administration activities such as backups, restores, schedules, and reports can be performed from the virtualization client.

  The virtualization client also enables proxy teaming. The proxy teaming ensures the fault tolerant backup. It will be useful when you want to perform the backup for a large number of virtual machines in a limited backup window

- Automatic Protection for All Virtual Machines in a vCenter

  Once you configure the Virtual Server Agent for a vCenter, all the virtual machines in the vCenter are automatically selected for backup. This behavior is designed to ensure all virtual machines are backed up.

- Customized Discovery

  If you want to backup only specific virtual machines in a vCenter, you can set criteria to search the virtual machines and automatically select them for backup. This is useful in environments where virtual machines are frequently added, or removed . You can easily browse and select vCenter objects such as hosts, datastores, resource pool etc. to set criteria for automatic discovery.

- Customized Filters

  You can exclude specific virtual machines or disks on the virtual machines from the backup. You can add filters for the virtual machines or disks.

- Detailed Reporting for Each Virtual Machine Backup

You can view a detail report about each backed up virtual machine. It contains information such as name of the proxy which performed the backup of the virtual machine, size of the backup data, type of backup etc. You can view all this information from the CommCell console when the backup job is running. It appears on the Client Status tab of the View Job Details dialog box.

- vCloud Backup and Restore

  You can protect the vCloud-specific attributes of a virtual machine and restore it back to a vCloud.

- Point-in-Time Snapshots of Virtual Machines

  IntelliSnap Backup enables you to create a point-in-time snapshot of a virtual machine by temporarily quiescing the data, taking a snapshot, and resuming live operations. IntelliSnap backups work in conjunction with hardware snapshot engines.

- Block Level Deduplication

  Deduplication provides a smarter way of storing data by identifying and eliminating the duplicate items in a data protection operation.

  Deduplication at the data block level compares blocks of data against each other. If virtual machines contains blocks of data that are identical to each other, block level deduplication eliminates storing the redundant data and reduces the size of the data in storage. This dramatically reduces the virtual machine backup data copies on both the disk and tapes.

  Changed Block Tracking (CBT) is a VMware feature that can be used to optimize backups of virtual machines by reading only the allocated and modified portions of a virtual disk. CBT is automatically enabled for virtual machines running on hardware version 7 or higher.

Site Interconnect Overview

For the purposes of this validation, the interconnect between the CSP and Enterprise sites was treated as a black box, since it was not in scope for this project.

Each of the two CSP sites and the Enterprise site (where the three Enterprise tenants were configured) were connected (single-homed) over IP into a generic Internet cloud which consisted of a pair of Cisco Catalyst 6500 multilayer switches. IP-based routing (BGP and OSPF) was configured to enable routing between sites.

IPsec tunneling was configured for Tenants 5 and 7 to connect the tenant containers in the Enterprise site to their respective containers in the SP1 site. The CSR 1000V routers deployed in each tenant container provided the endpoints for these encrypted tunnels. No encryption was configured for Tenants 1 and 2 between the SP1 and SP2 sites, following the assumption that the CSP would have encryption configured between sites at another layer. Refer to SP1 Site Overview, page 4-3 for a discussion on the Cisco CSR 1000V.

Figure 4-32shows the interconnections between the Enterprise and SP1 in addition to SP1 to SP2. The routing table entries for the CSR pairs learned through BGP are shown next to each CSR. Routes for the local data center are learned via OSPF and are not shown here.

Figure 4-32 Site Interconnect Using Cisco Cloud Services Router (CSR) 1000V

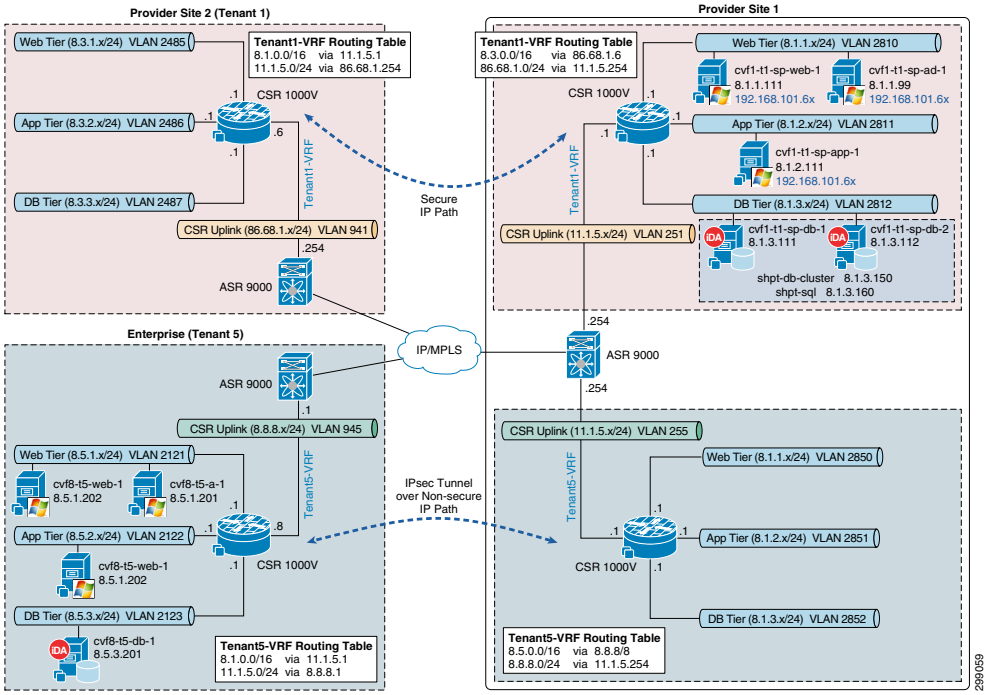

In Tenant 6, the tenant compute was isolated within the OpenStack cloud environment, with its only connectivity to the CSP clouds being on VLAN 58, the Commvault replication network. The gateway for VLAN 58 was configured using HSRP on the Enterprise aggregation switches.

Use Case 1 (In-Cloud BaaS): Implementation Details

The first use case is Cloud Backup and is focused on Enterprise customers that have their entire environment hosted at a Cloud Service Provider site. In such a case, the Enterprise production environment is hosted at one CSP site and the CSP leverages Commvault MediaAgent to perform local backups within that site as well as perform remote backups to a Commvault MediaAgent at a secondary site. Two tenants were created to proof this use case. Both of the tenants were managed within a single Hyper-V environment in the SP1 site. Figure 4-33 shows the environment used to proof this use case. Each iDA will interact with the CommServe and appropriate MediaAgent on a schedule basis to execute data protection jobs. The Primary copy of data will be stored locally on the MediaAgent that executed the backup job, but will then be DASH (deduplicated replication) copied to a MediaAgent in the secondary data center to provide for Disaster Recovery. The CVLT Database and Reporting Server are also protected and prepped for recovery at the secondary data center. Restoration of the data that is protected can be achieved at either data center.

Figure 4-33 Use Case 1 (In-Cloud BaaS) Implementation Topology

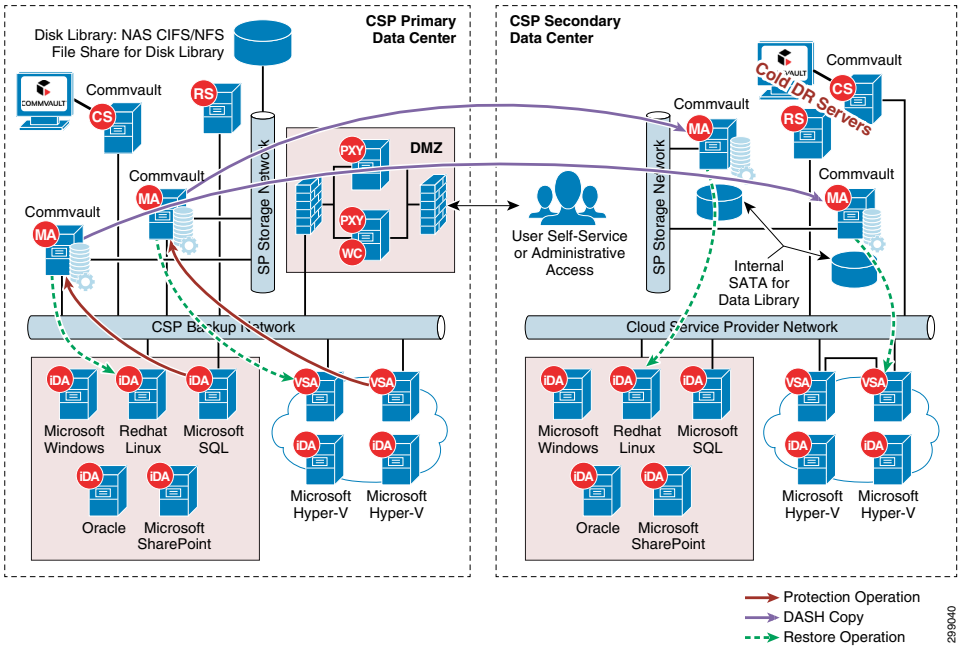

Tenant Details

Tenants 1 and 2 were configured virtually identically, with the exception of the applications represented within each.

Tenant 1 Application

In Tenant 1, the application used for testing purposes was Microsoft SharePoint with a Microsoft SQL database. A stand-alone instance of SharePoint 2013 was installed on a VM with a local database. The VM was configured with a single interface in the tenant workload web tier. Connectivity between the VM and the Commvault MediaAgent was enabled by the Cisco ASA 5585, with BGP route leaking to enable routing between the replication VRF and tenant VRFs.

In addition to the SharePoint instance, a Microsoft Windows 2012 R2 Failover Cluster was created and then SQL 2012 R2 was installed on the cluster. The cluster consisted of two Windows 2012 R2 VMs, each with their own 40GB C: drives and two additional shared volumes. The first shared volume was a 1GB quorum drive that was used for management of the cluster. The second shared volume was a 100GB data drive that was used for the database data.

In the VM configuration, these two shared drives were configured on a second SCSI controller that required a sharing mode to be enabled to allow sharing across multiple Hyper-V hosts. The sharing mode can only be configured from the Failover Cluster Manager on the Hyper-V host where the SQL VMs are running. The two VMs were configured with two interfaces each, one for the tenant workload database VLAN and a second for the replication VLAN. The virtual IP addresses for the cluster and for the SQL instance were configured on the database VLAN. The SQL database was populated with a test database and SQL scripts were used add/remove/count data in the database.

Tenant 2 Application

In Tenant 2, the application used for testing purposes was Linux Oracle 11g. A single Oracle RedHat Linux 6.6 VM was created and then Oracle 11g was installed on the VM. The VM was configured with two interfaces, with one connected to the tenant workload database VLAN and the other connected to the replication VLAN. The Commvault components in SP1 were also configured on the replication VLAN and this enabled Layer 2 connectivity between the MediaAgents and the Oracle database VM.

Commvault Configuration

This section details how the Commvault Infrastructure and Data Protection agents were configured for this use case.

Commvault Infrastructure

This section describes how each component of the Commvault Simpana is installed and configured.

CommServe Server

The CommServe is the primary server in the data protection environment, so it of course is the first server that is built out. There is also a stand-by CommServe that will be pre-staged in a recovery data center, its installation process is the same as the primary. On a Windows Server 2008 Standard server with the disk laid out for the Commvault CommCell Database, disable the firewall and install IIS and NET Framework 3.5.x, which are prerequisites to the Commvault CommServe software installation. Choose the custom installation method once starting the installation wizard. Select the CommServe, Consoles, Web Server, and Work Flow platforms. Further in the Wizard Choose the name for the CommServe.

Figure 4-34 CommServe Server Setup

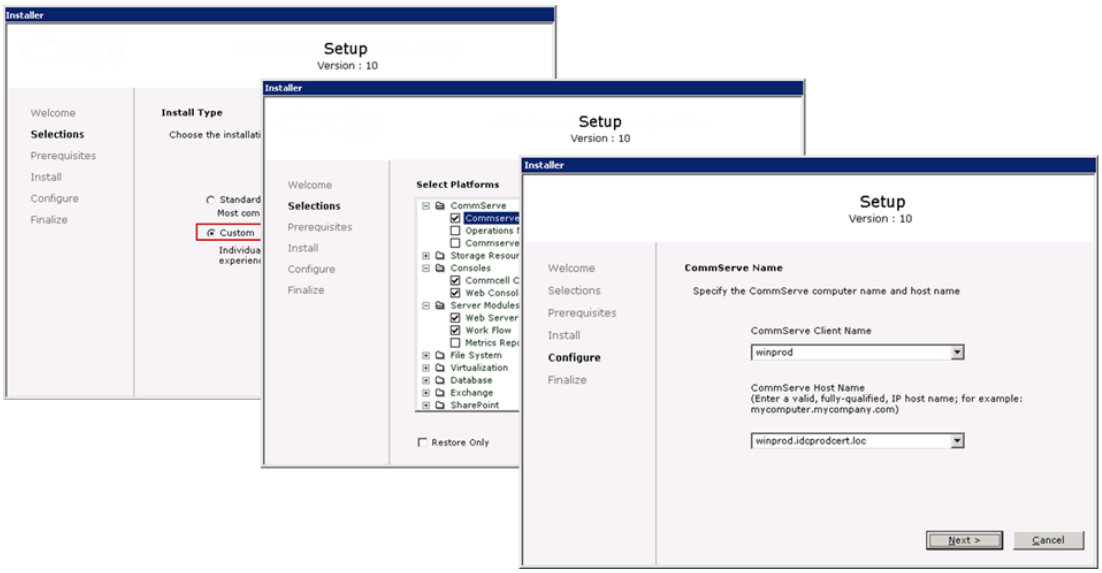

Also within the wizard, choose the location where the Commvault CommCell DB will reside, generally higher speed disk in an isolated array, setup the first Administrative User Account, and configure the location of where the Software Cache, used to install other Commvault infrastructure servers.

Figure 4-35 Setup Administrative User Account

Once the installation is complete the Commvault Administrative Console will provide access to all Commvault activities (Figure 4-36).

Figure 4-36 Commvault Administrative Console

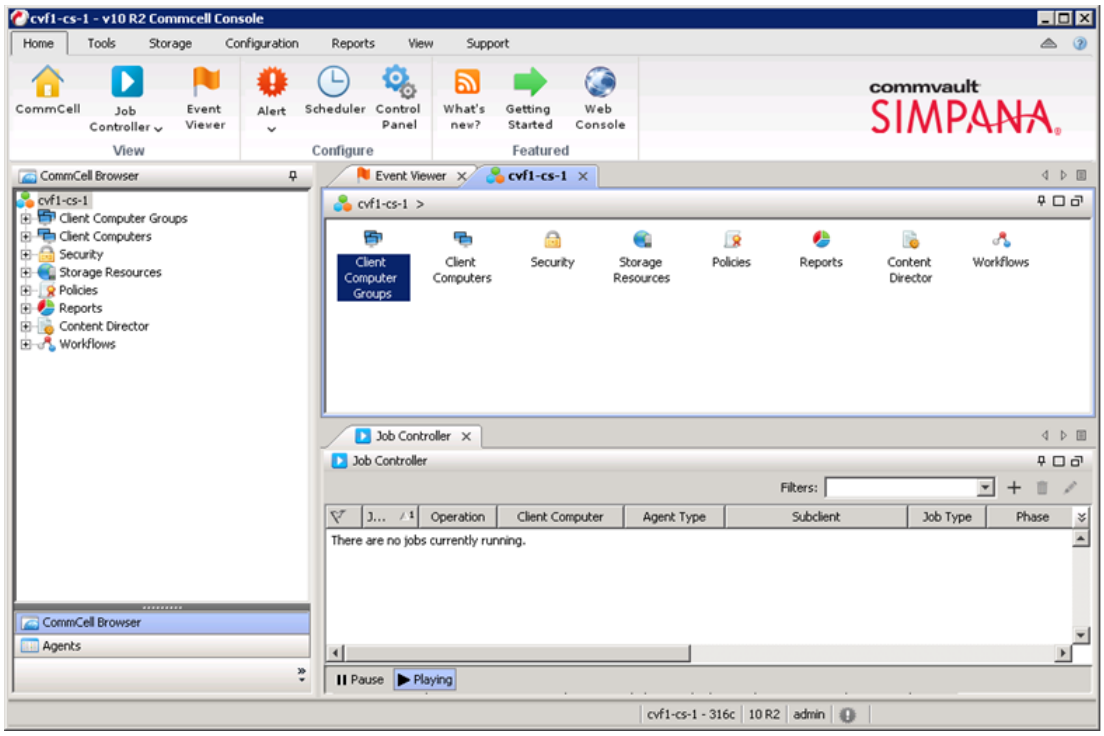

The only remaining step is to configure Commvault Disaster Recovery Backup Configuration (Figure 4-37), which is what is required in the event of a loss of the CommServe server itself. Generally this would be to a network drive or a shared volume mounted from the DR CommServe.

Figure 4-37 Commvault Disaster Recovery Backup Configuration

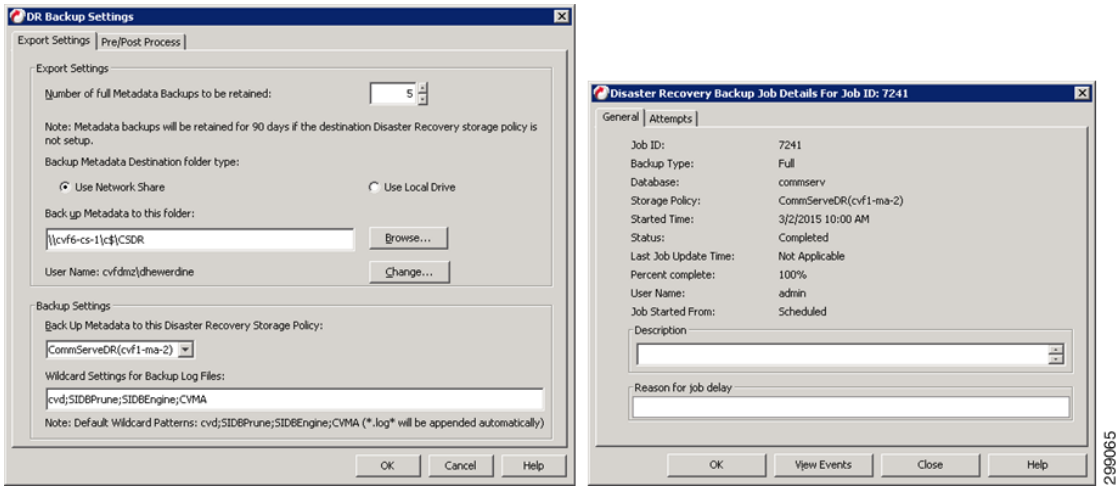

Refer to Commvault documentation for more information the CommServe.

MediaAgent Server(s)

The MediaAgent server is the data mover within the CommCell so it is the next server that needs to be built out. All connections to the actual storage, as well as the control of the deduplication is handled by the MediaAgent. The MediaAgents should have disk volumes created for the Operating System, CVLT DDB and Index Cache. They are then loaded with Windows 2012. The installation of the MediaAgent software can be pushed from the CommServe. In this case there are two MediaAgents, cvf1-ma-1 and cvf1-ma-2.

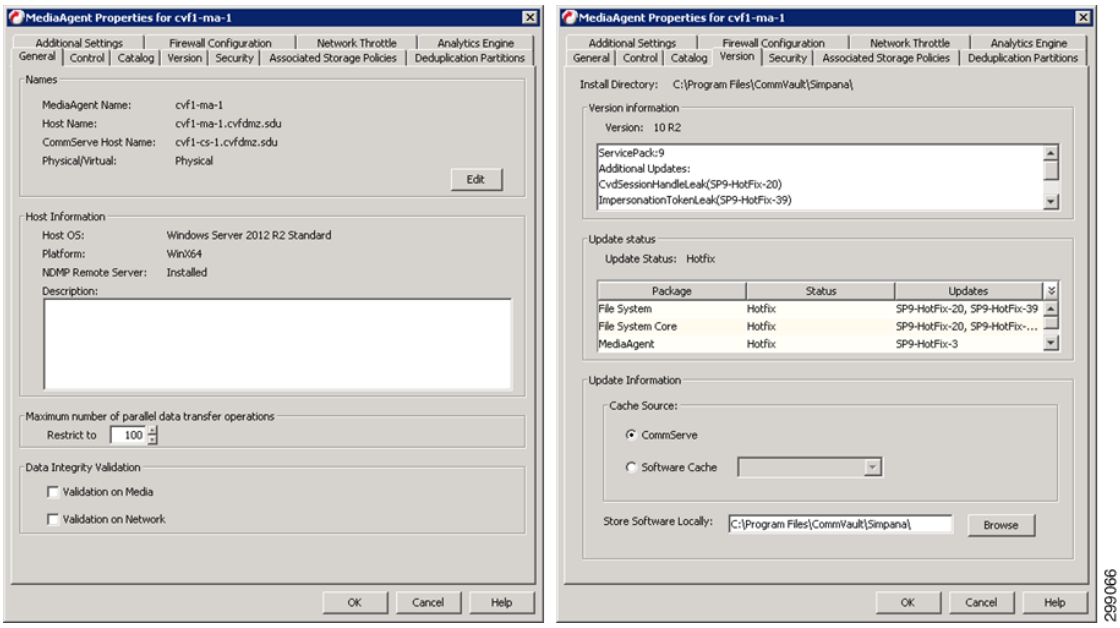

A Shared Disk Library is created from NFS File Share on both CVF1-MA-1 and CVF1-MA-2.

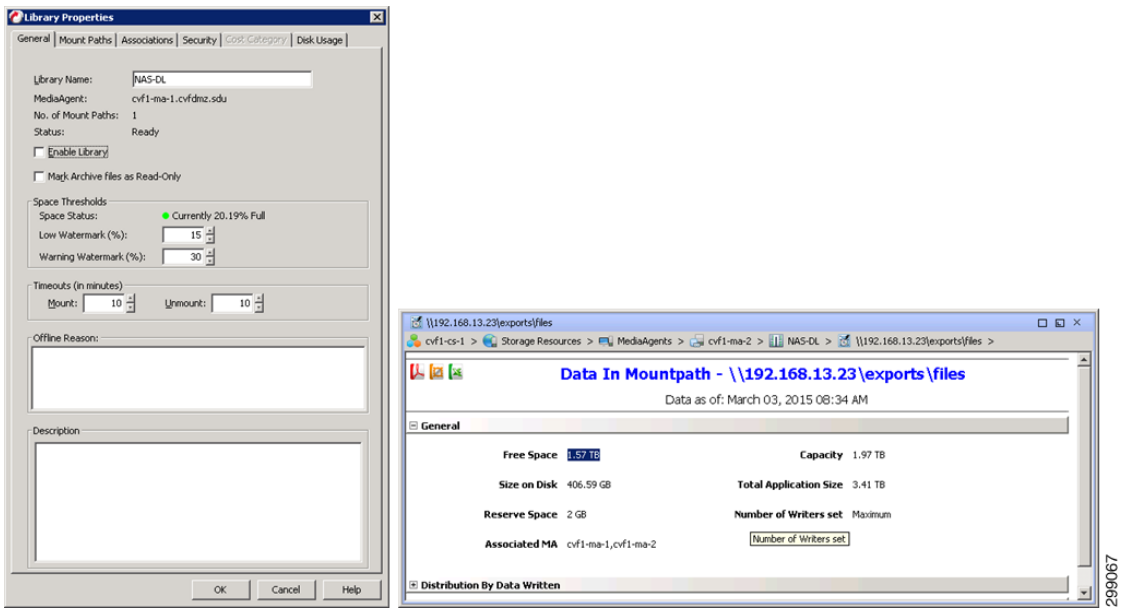

A Global Partitioned DDB is then created to frontend the NAS-DL just created.

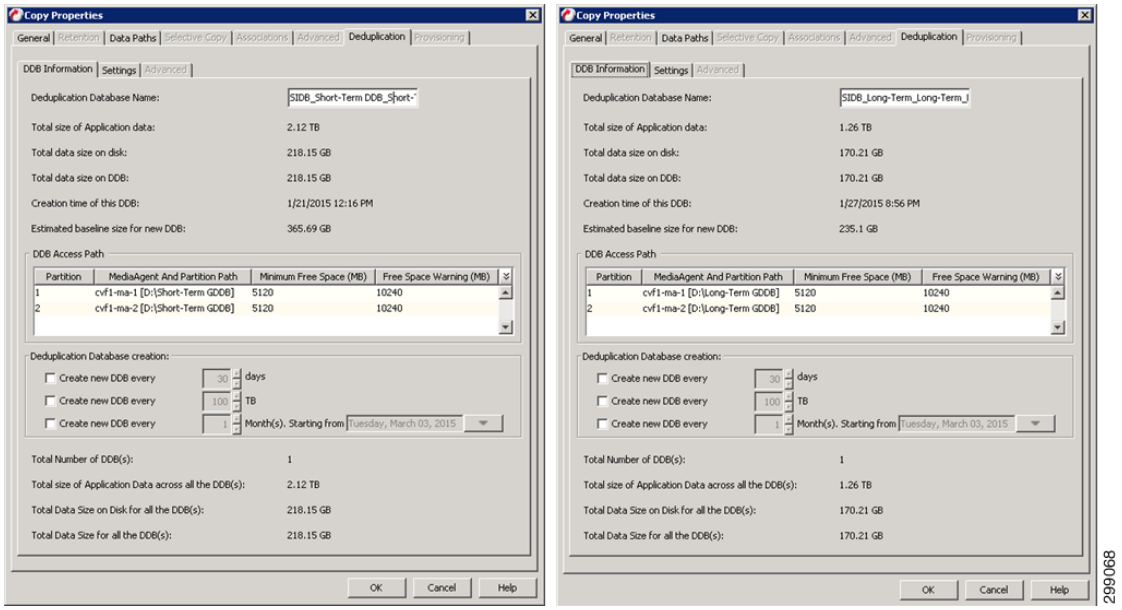

Finally the Tenant Storage Policies are configured.

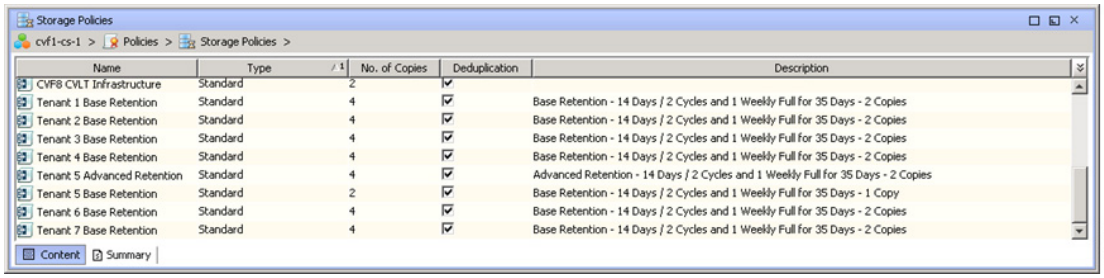

Refer to Commvault documentation for details on the MediaAgents.

Enterprise Metrics Reporting Server

The Commvault Simpana Enterprise Metrics Reporting Server collects data from all CommCells within an Enterprise or Service Provider's environment, allowing for a wide breath of reporting capabilities. The Commvault Simpana Custom Report tool allows for the creation of very customized reports, as well as access to the Commvault Software Store, where written reports are made available. This server will have a separate MSSQL DB for the reporting data, so it should have some high speed disk allocated for this. It will be a Windows Server and require NET 3.5 and IIS before installing Commvault. There is also a stand-by server at a recovery site, its installation process is the same. Then the server can be installed, patched, and configured.

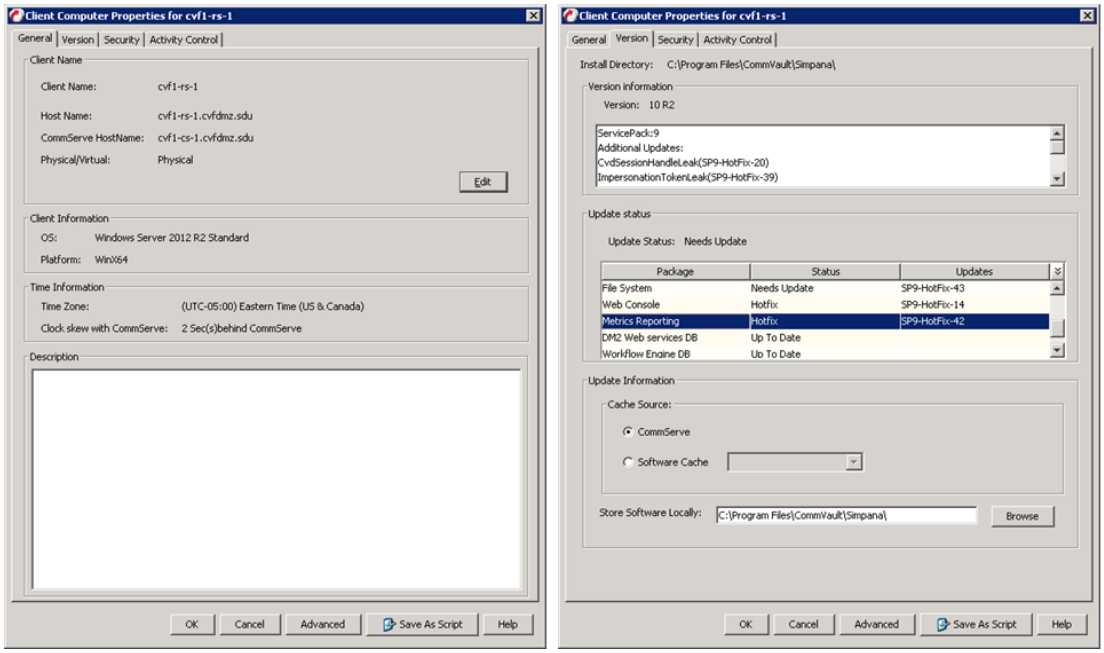

The Enterprise Metrics Reporting is configure via the administrative control panel.

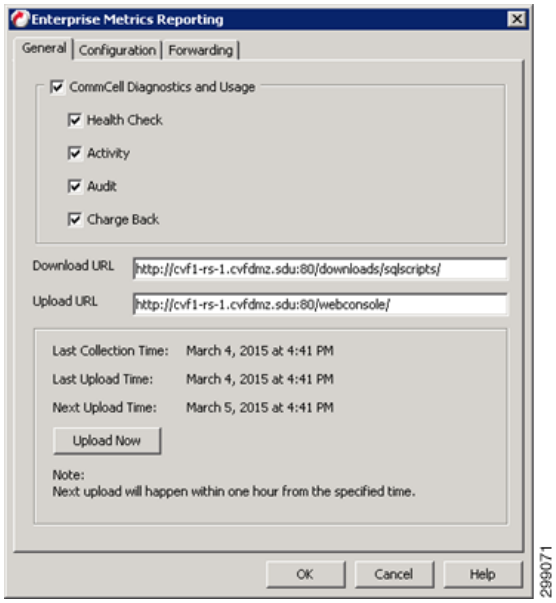

Refer to Commvault documentation for details on the Enterprise Metrics Reporting Server.

Web Proxy and Web Console Server

Proxies are an important component of service provider's network security configuration to reduce the number of ports opened and provide secure data transfer between service provider and tenant. In this case a Shared Proxy is used as it will provide a single proxy with multiple tenants pooled together, be located in the service provider's DMZ, and prevents the service provider's infrastructure from being internet facing,. This server will also function as the Web Console allowing access to the environment to administrators and users, without exposing the CommCell to the outside world.
CVF1-DMZ-PROXY-1 has the File System iDA installed and is patched to the latest version.

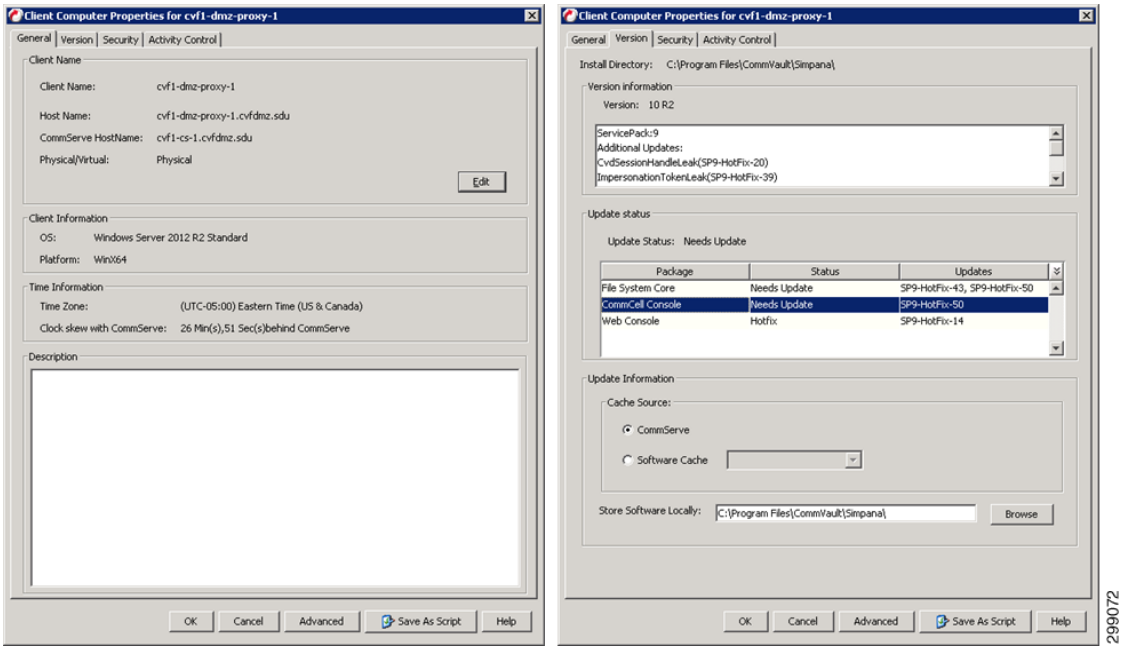

Configure the Web Proxy Server Firewall settings.

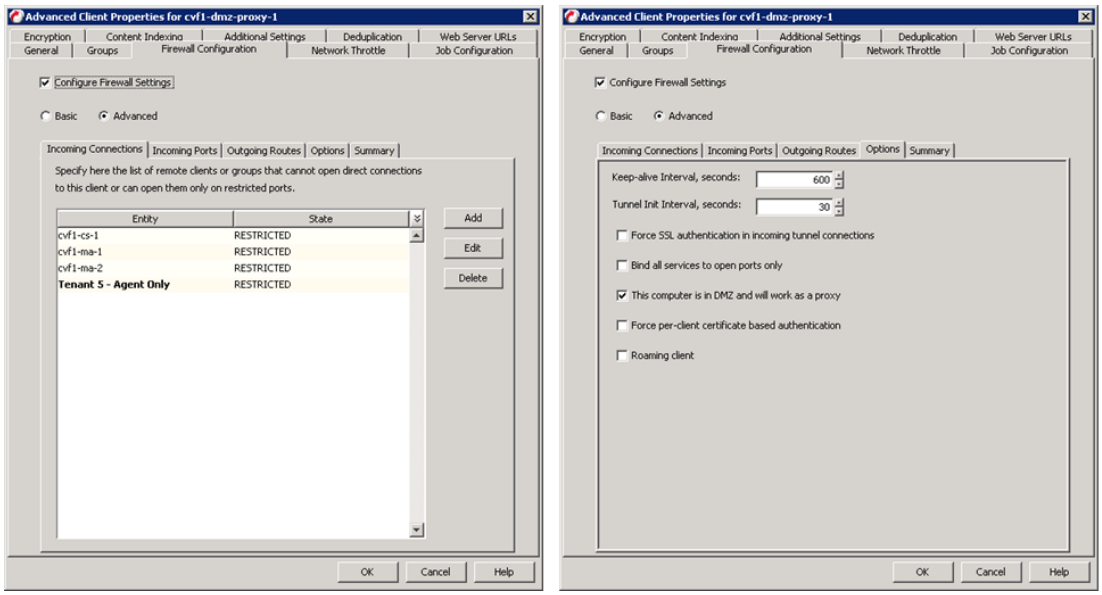

The CommServe firewall settings will then be updated

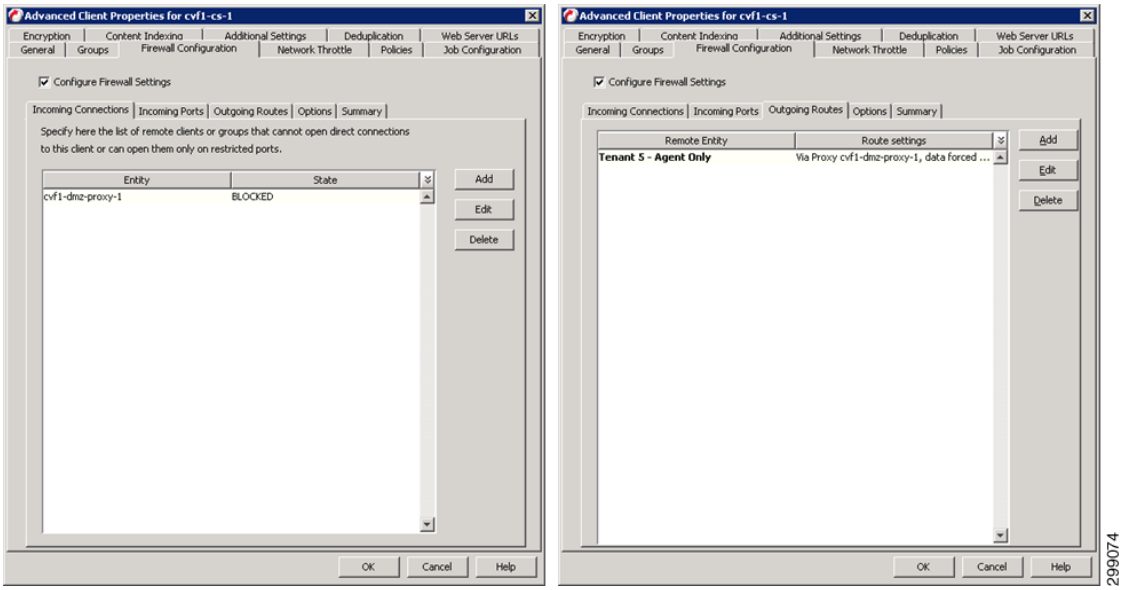

As well as the MediaAgent(s), that will be visible via the Web Proxy, Firewall settings.

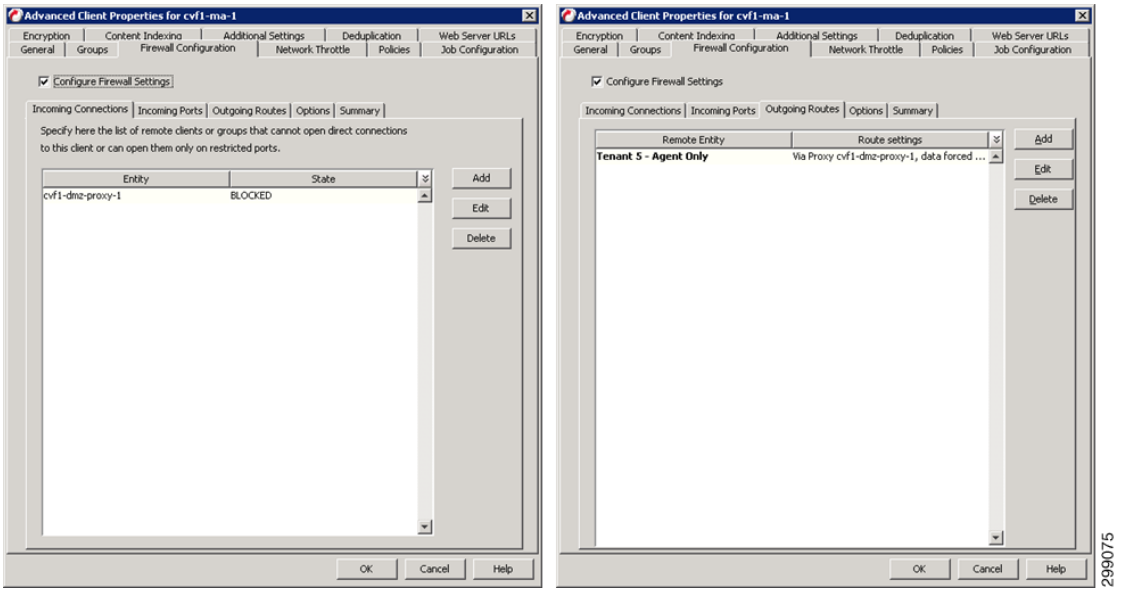

Finally, the Tenant Client Computer Group that will be accessing the CommCell via the Web Proxy, Firewall settings.

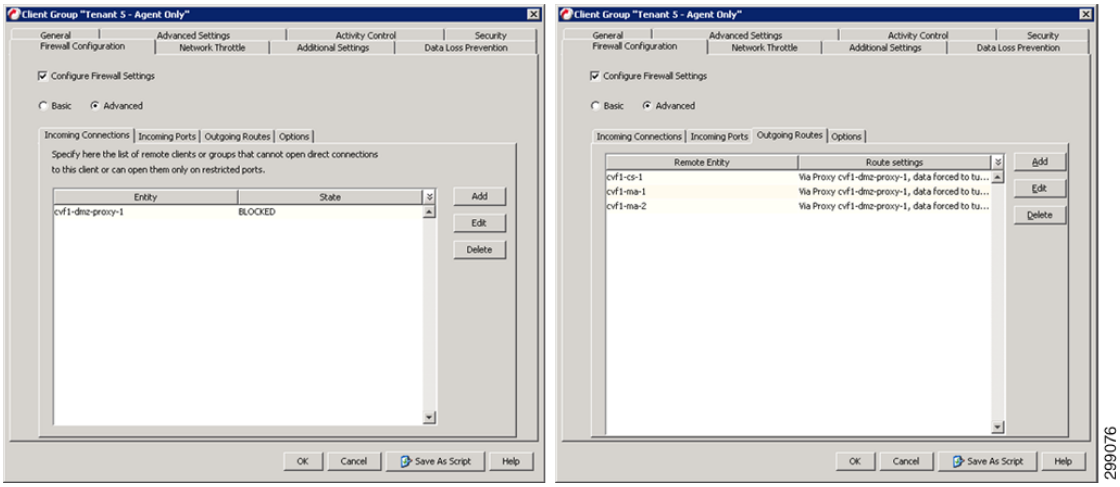

Use the eMail and Web Server configuration button in the control panel to configure the new Web Console.

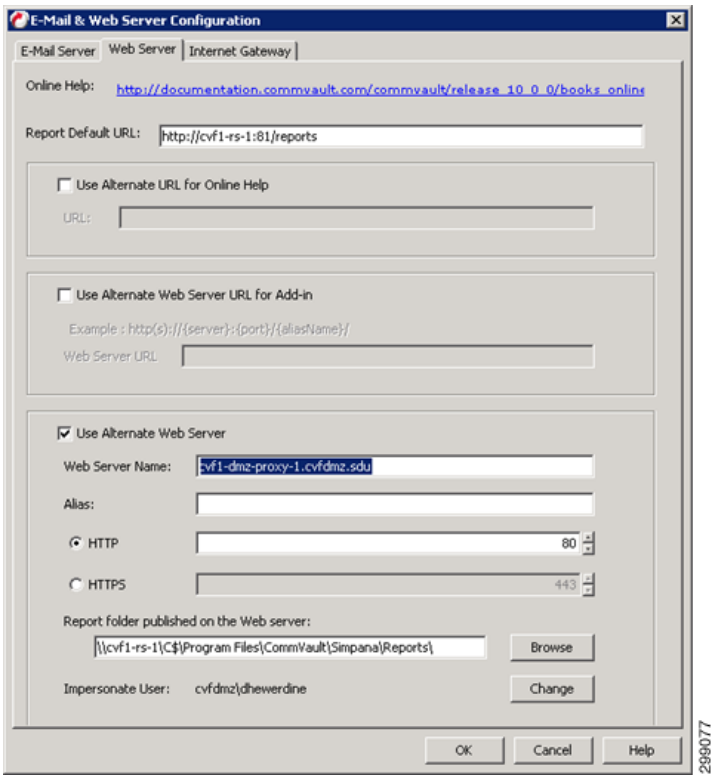

Access the Web Console via any browser.

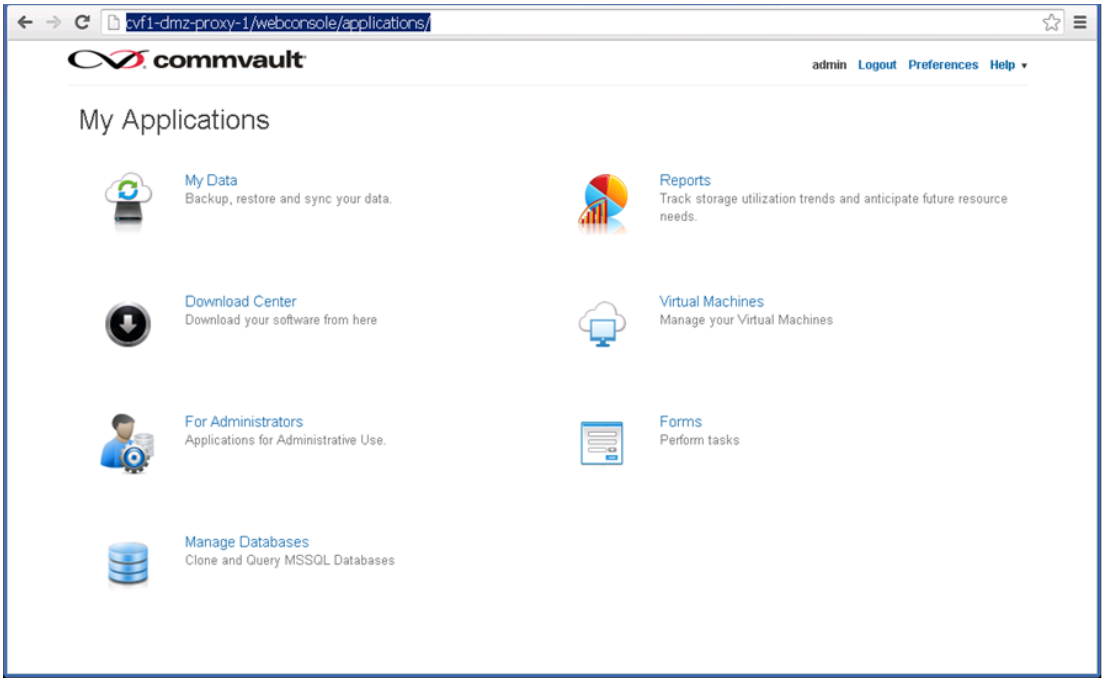

Refer to Commvault documentation for details on the Web Proxy.

Refer to Commvault documentation for details on the Web Console.

Physical Server Protection Configuration

This section details how the iDataAgents were configured to provide data protection to the physical servers in this use case.

Windows File System iDA

The Windows File System iDA will be installed on each Window server requiring data protection from the File System level. Once the Client requiring protection is defined to the CommCell, the appropriate iDA can then be installed and configured.

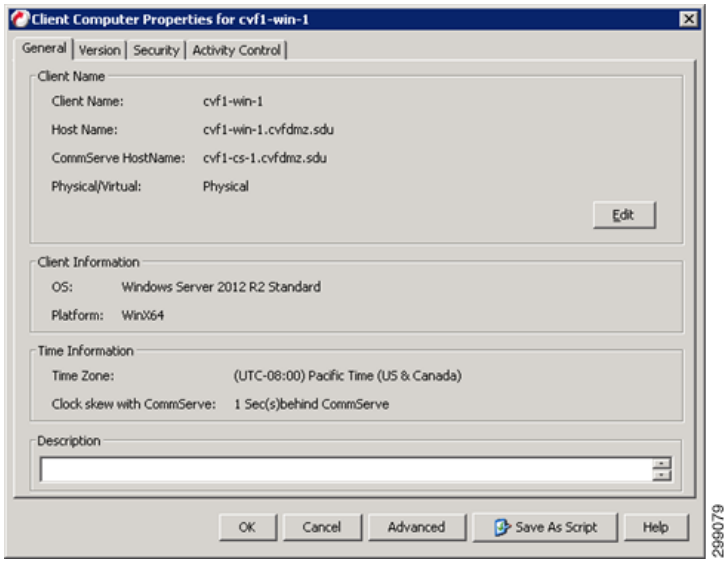

The Default File System Subclient will then be assigned to the proper Tenant Retention Storage Policy.

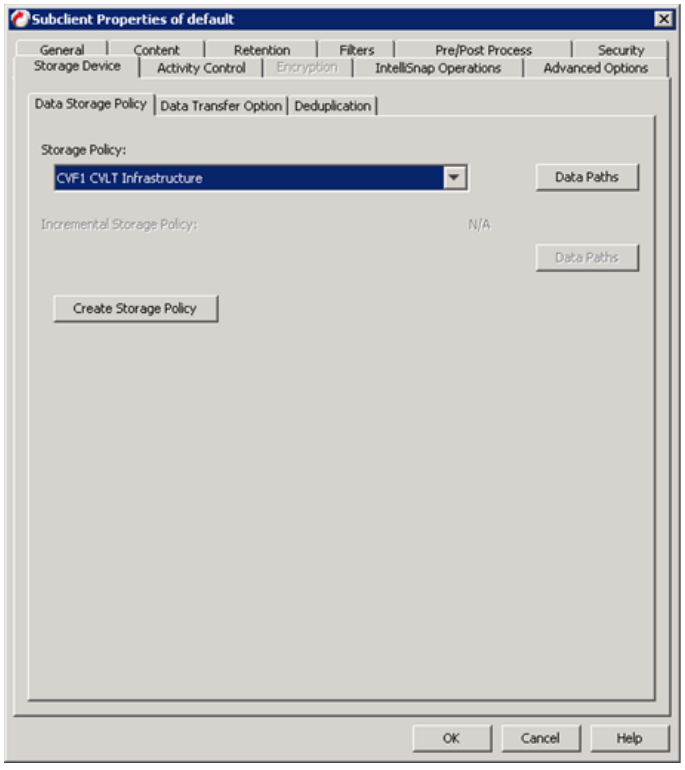

The Client is then associated to the proper Schedule Policy. The protected clients will be running an incremental backup nightly, while generating a synthetic full backup once a week.

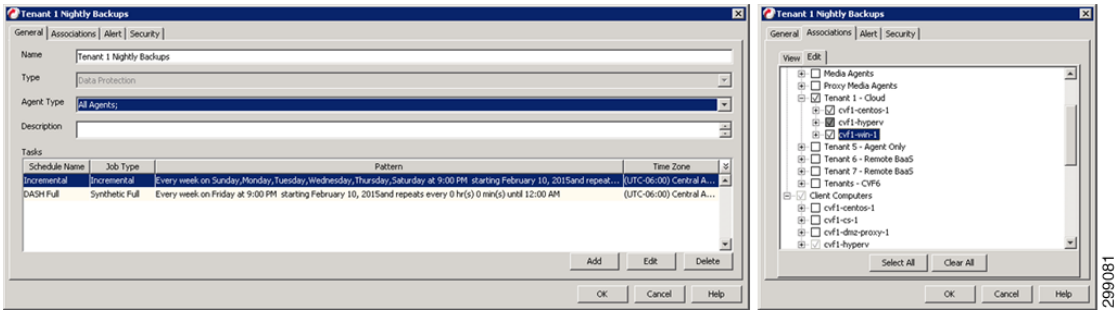

Alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy. Select the alert criteria to use.

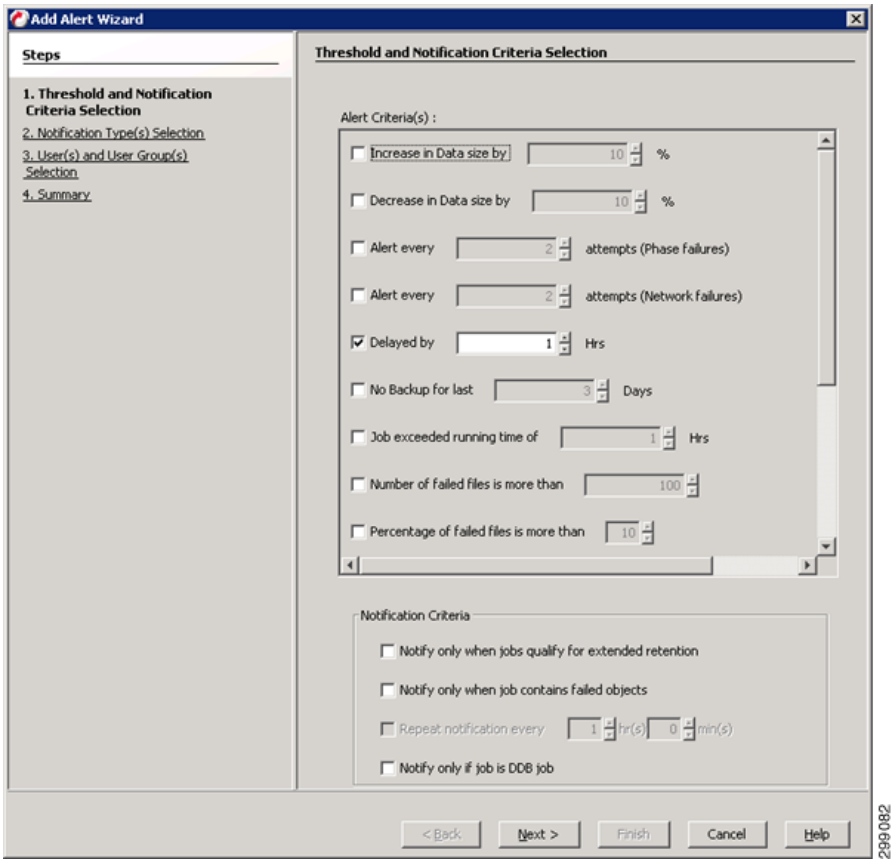

Select the notification method.

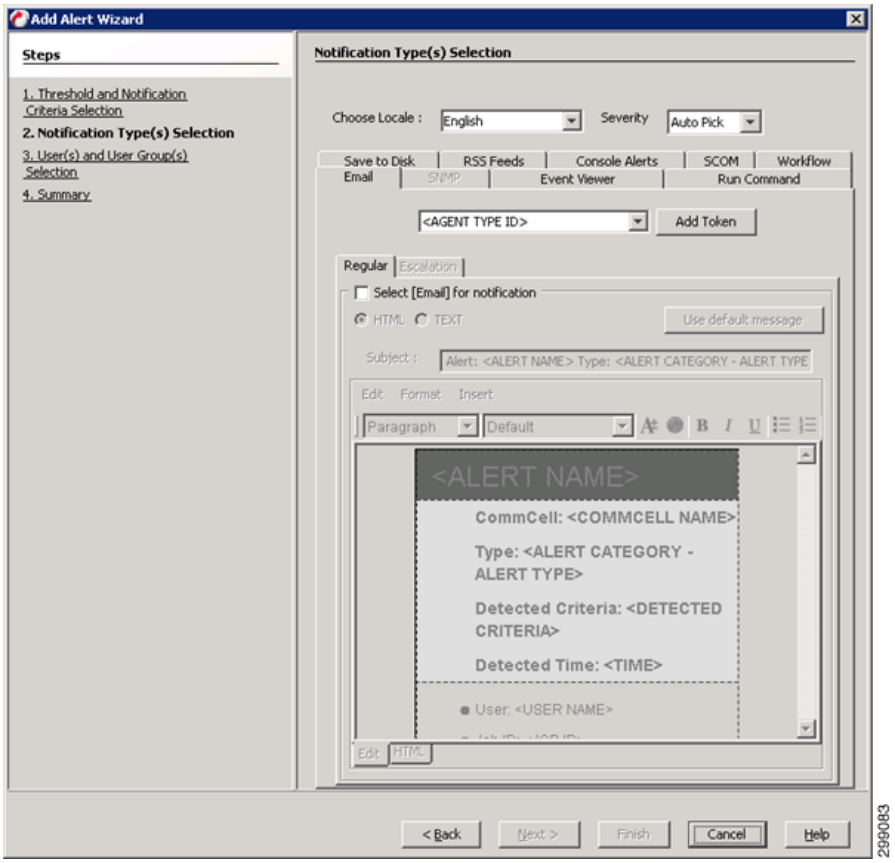

Select the Users or User Groups that will receive the alert.

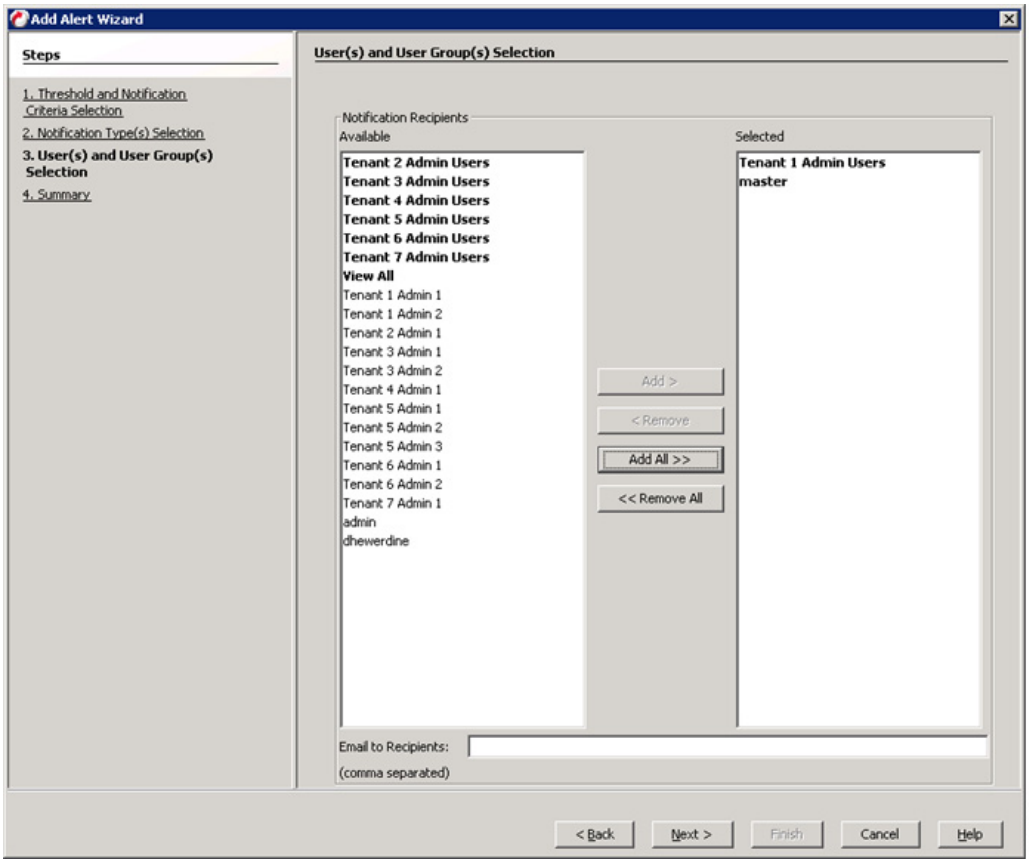

Refer to Commvault documentation for details on the Windows File System iDA.

Linux File System iDataAgent

The Linux File System iDA will be installed on each Linux server requiring data protection from the File System level. Once the Client requiring protection is defined to the CommCell and then the appropriate iDA can be installed and configured.

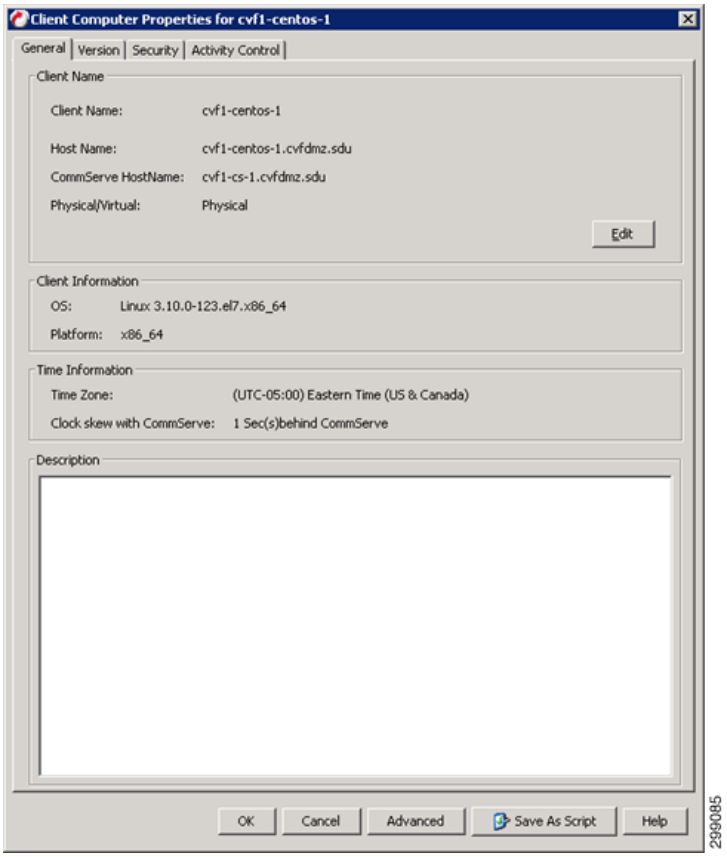

The Default File System Subclient will then be assigned to the proper Tenant Retention Storage Policy.

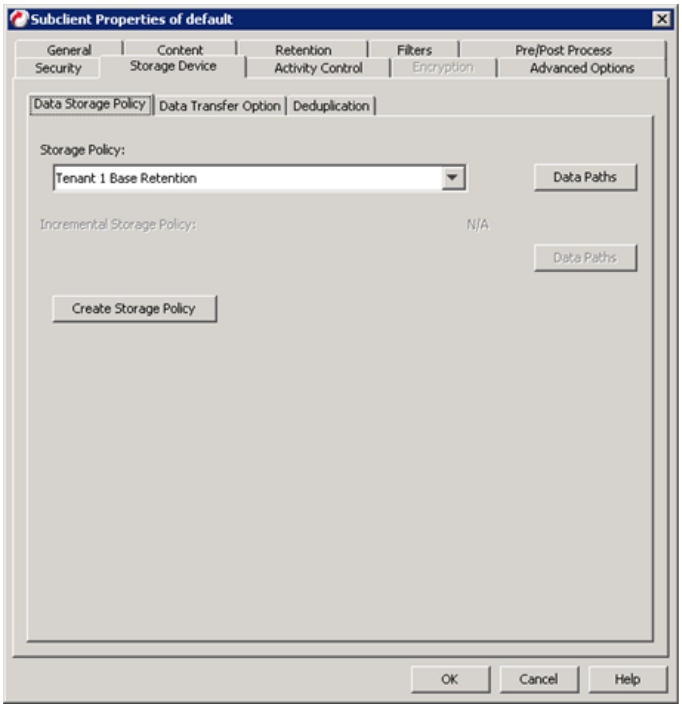

Now the Client is associated to the proper Schedule Policy. The protected clients will be running an incremental backup nightly, while generating a synthetic full backup once a week and alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy. Refer to the Alert configuration for Windows File System iDA, page 4-62, above.

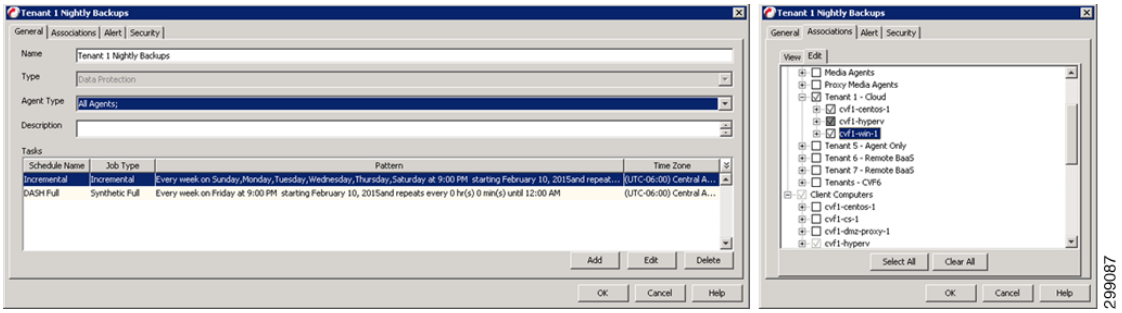

Refer to Commvault documentation for details on the Linux File System iDA.

Virtual Server Protection Configuration

This section details how the iDataAgents were configured to provide data protection to the virtual environments in this use case.

Virtual Server iDataAgent for Hyper-V

The VSA for Hyper-V is installed on one or more of the Hyper-v servers in the Hyper-V Cluster. The Hyper-V server that will be used as the VSA is defined to the CommCell, the required iDAs are installed and updated to the latest levels.

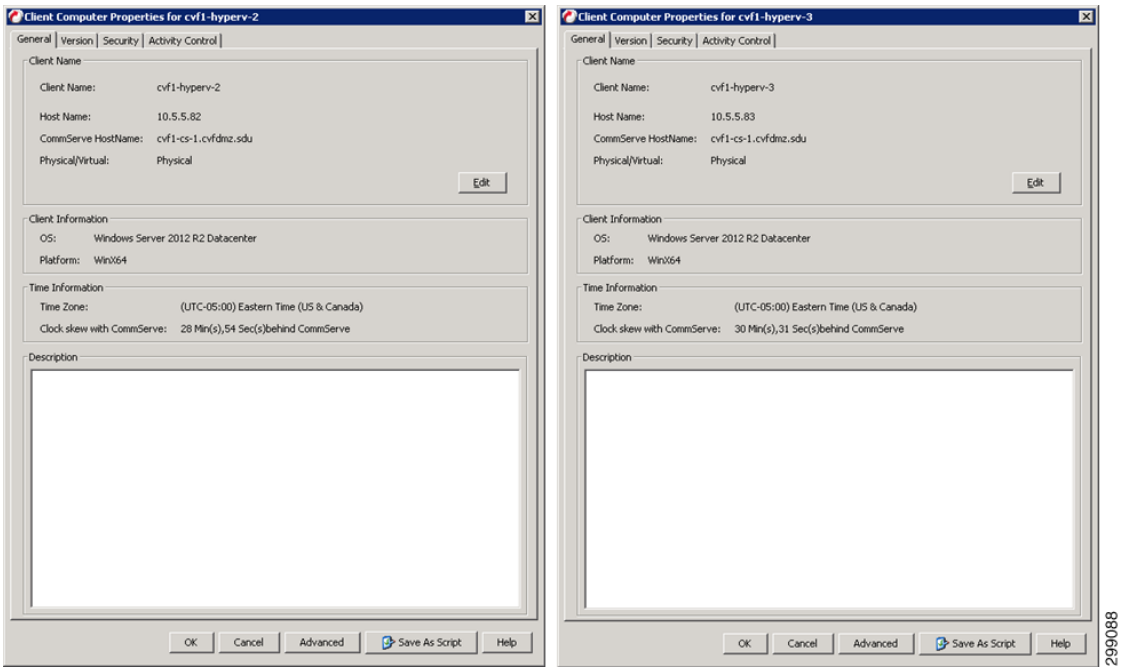

Once the VSA are installed, define a new Hyper-V Client, providing credentials and designating the VSA Servers to use.

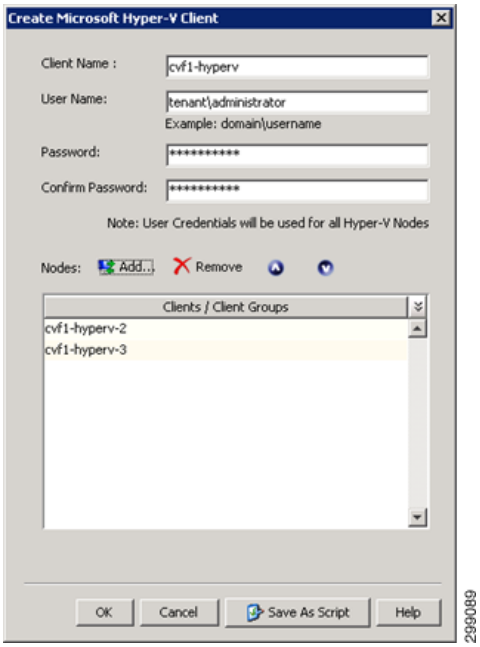

Create a new subclient for the Tenant Guests and select the VMs that will be protected for this Tenant.

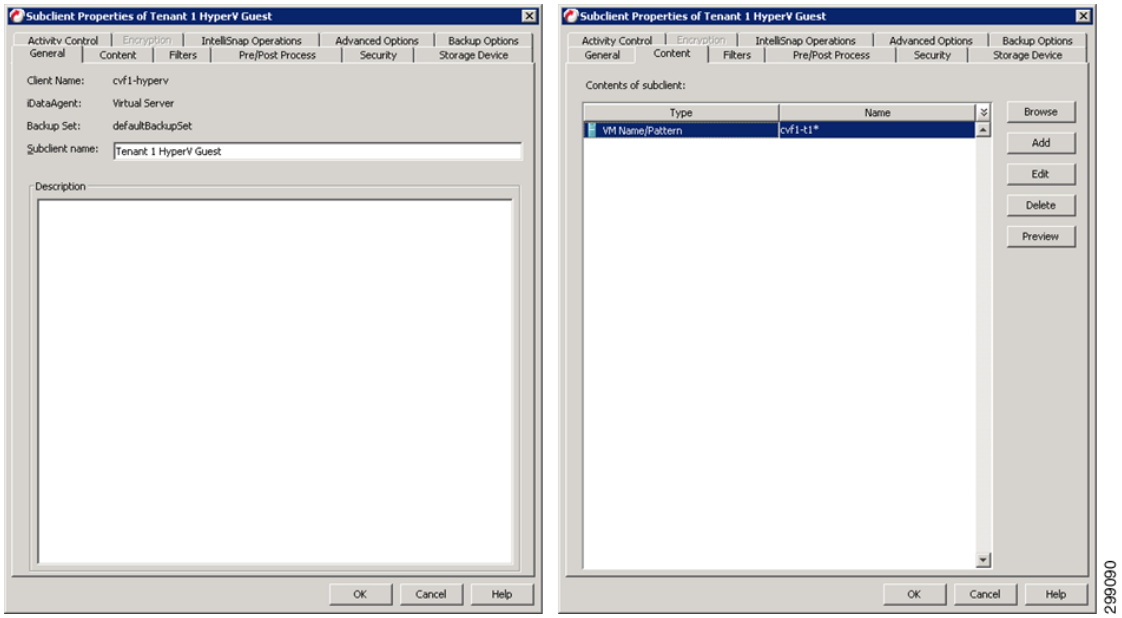

Select the appropriate Tenant Storage Policy.

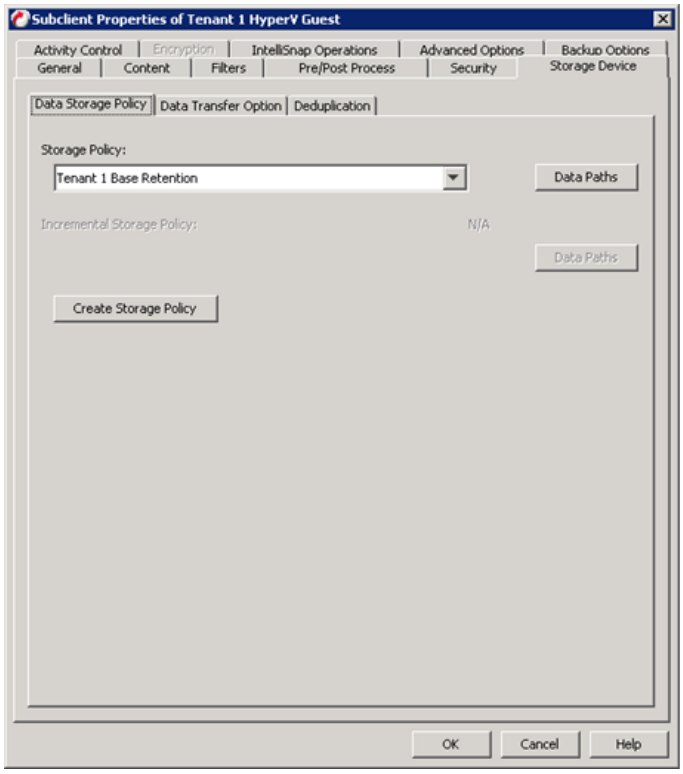

Finally the Client is associated to the proper Schedule Policy and alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy. Refer to the Alert configuration for Windows File System iDA, page 4-62, above.

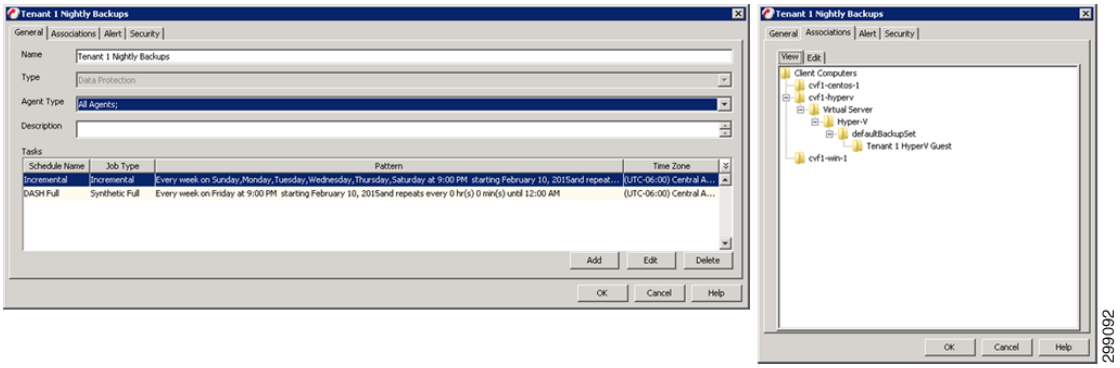

Refer to Commvault documentation for details on the Virtual Server iDA for Hyper-V.

Application Protection Configuration

This section details how the iDataAgents were configured to provide data protection for the applications and databases in this use case.

Microsoft SharePoint iDA

The Microsoft SharePoint iDA is going to be installed on any SharePoint server, whether they are virtual or physical servers, that require SharePoint Farm, Site, or Document level data protection. For this use case, SharePoint has all of its components installed on the same server. That server requiring protection is defined to the CommCell and then the appropriate iDAs can be installed, updated, and configured.

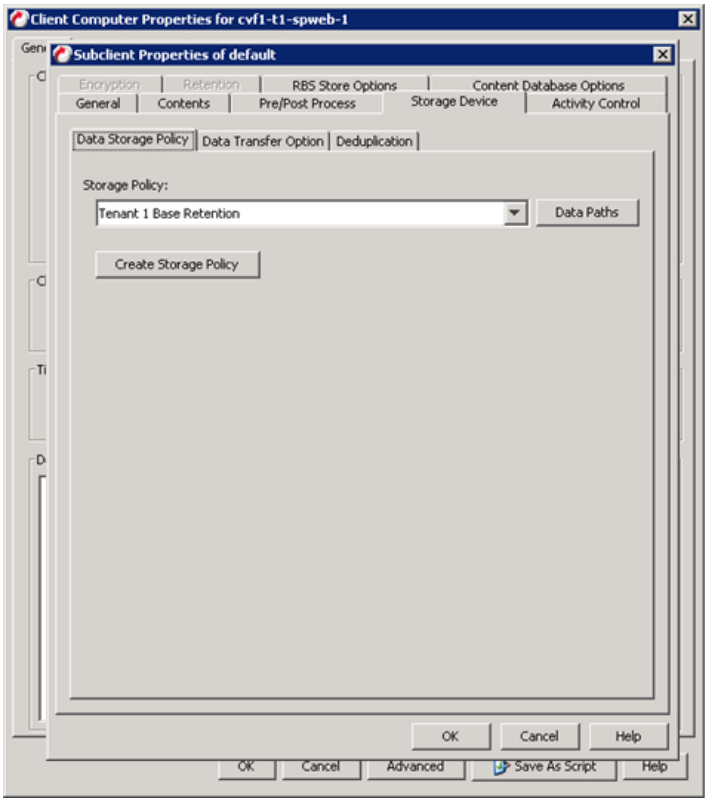

SharePoint has two levels of backups Database and Document. First configure the SharePoint Database default subclient to protect the entire Farm.

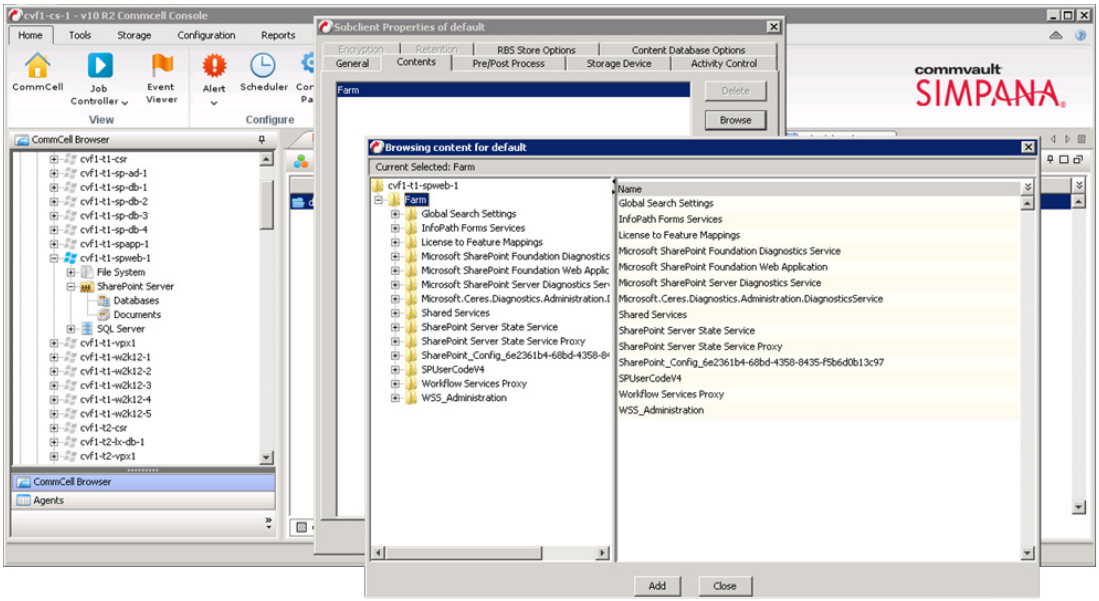

Assign it to the proper tenant Storage Policy.

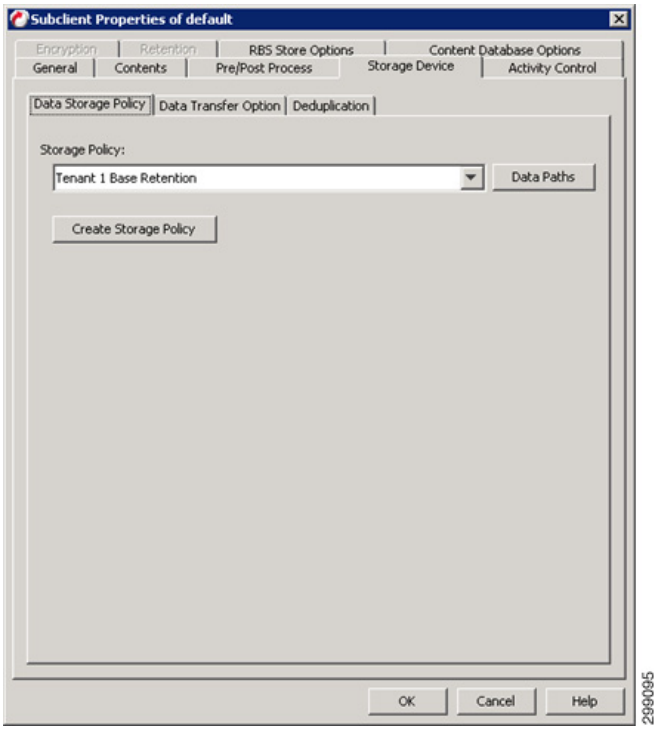

Configure the SharePoint Document default subclient to protect All Folders.

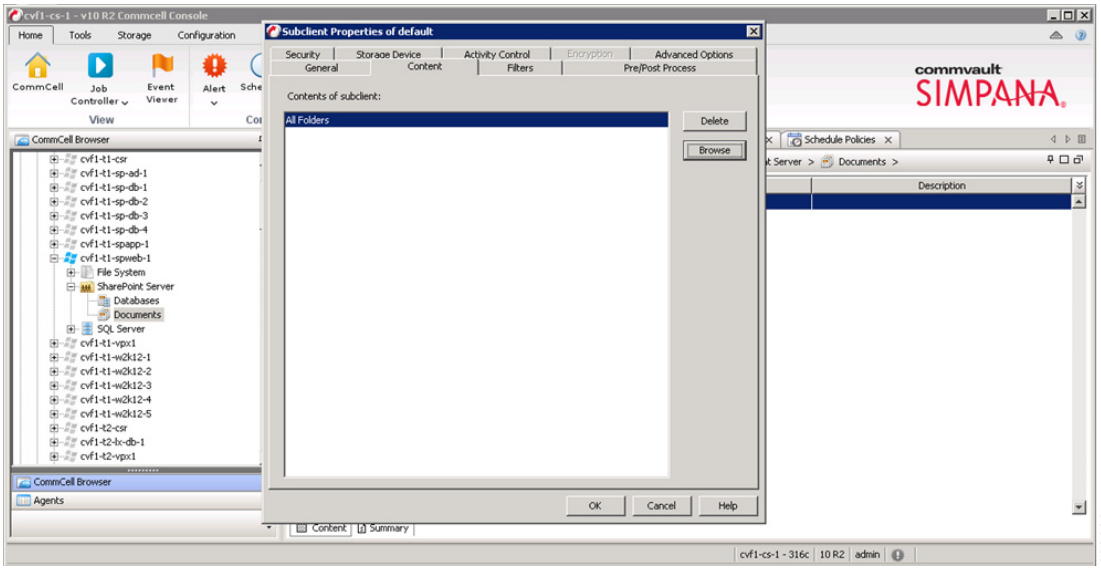

Assign it to the proper tenant Storage Policies:

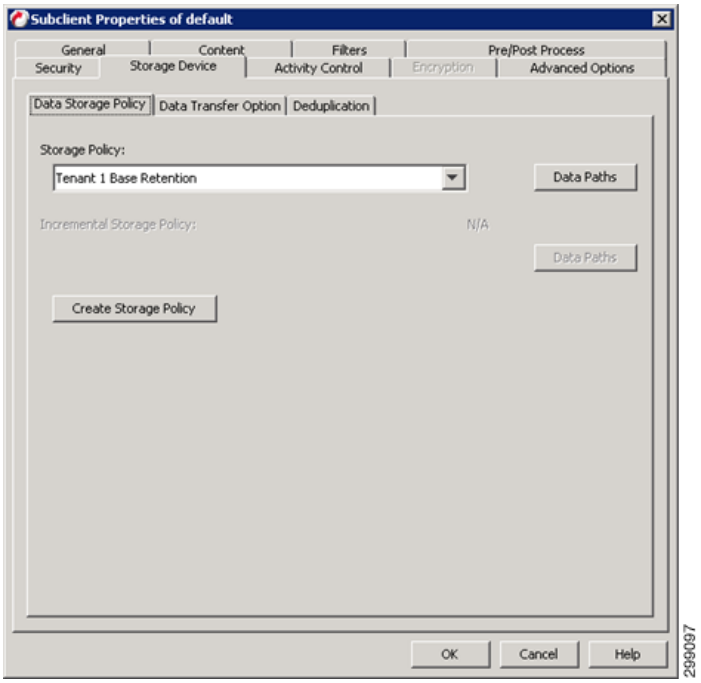

The Client is now associated to the proper Schedule Policy. SharePoint can have multiple different protection methods. In this example we have a Schedule Policy for the Database protection and another Schedule Policy for the Document protection. Alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy. Refer to the Alert configuration for Windows File System iDA, page 4-62, above.

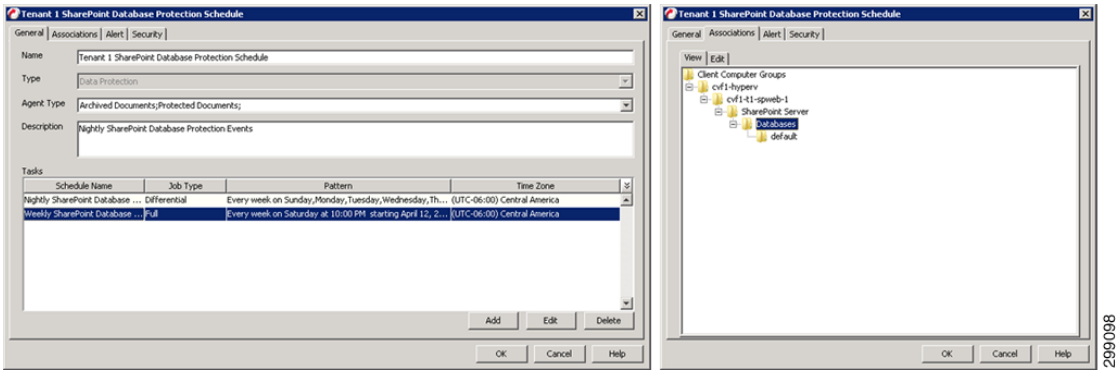

Refer to Commvault documentation for details on the Microsoft SharePoint iDA.

Microsoft SQL Server iDA

The Microsoft SQL Server iDA is going to be installed on any SQL server, whether they are virtual or physical servers, that require SQL DB, Table, and Log level data protection. For this use case, a two node cluster has been configured to run SQL. The servers requiring protection are defined to the CommCell and then the appropriate iDAs can be installed, updated and configured.

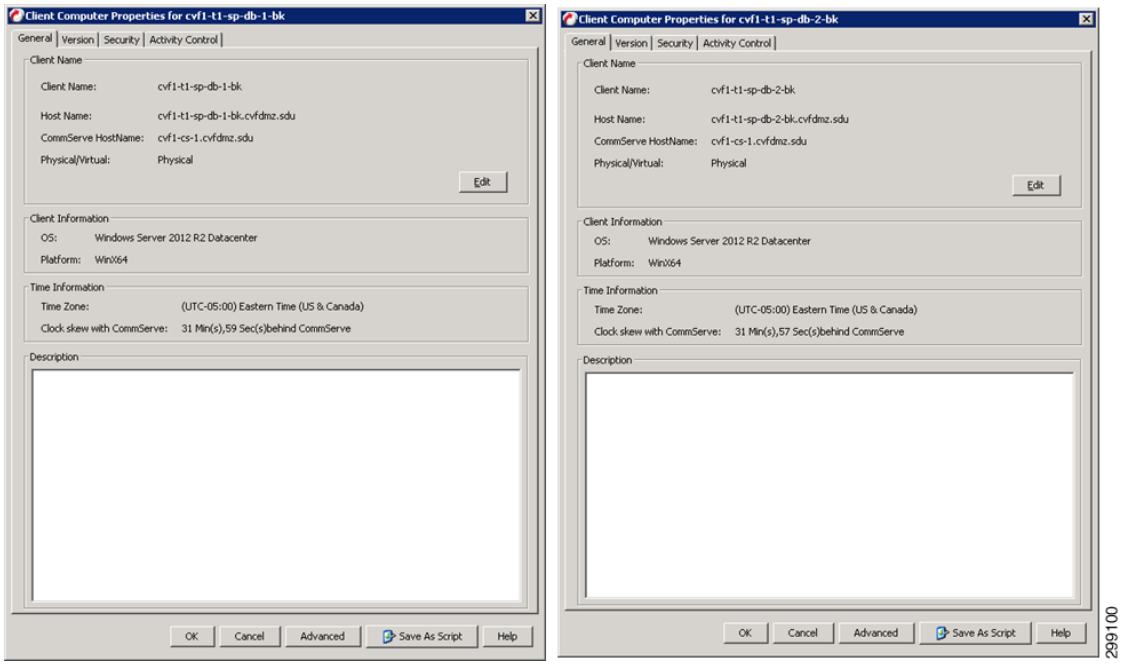

Define a new Windows Cluster Client, providing credentials and designating the VSA Servers to use.

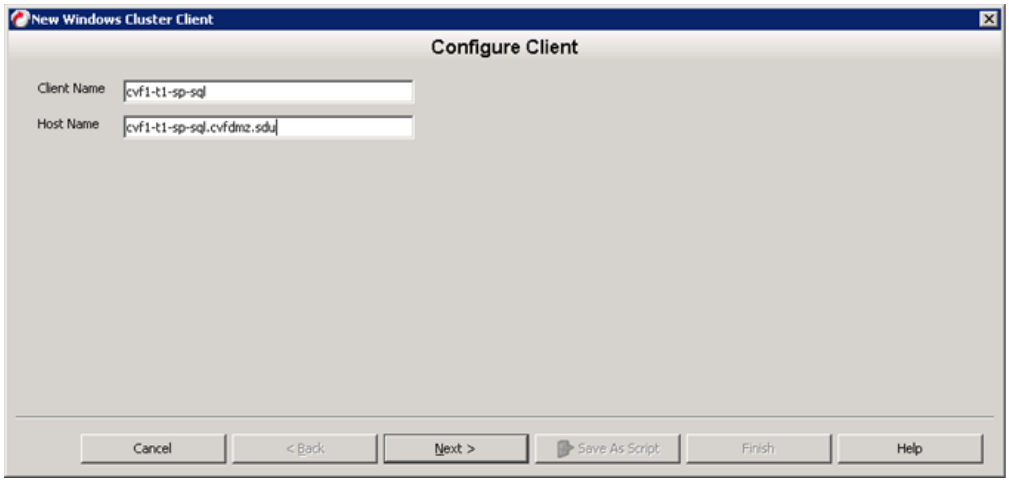

Define the servers that are part of the cluster.

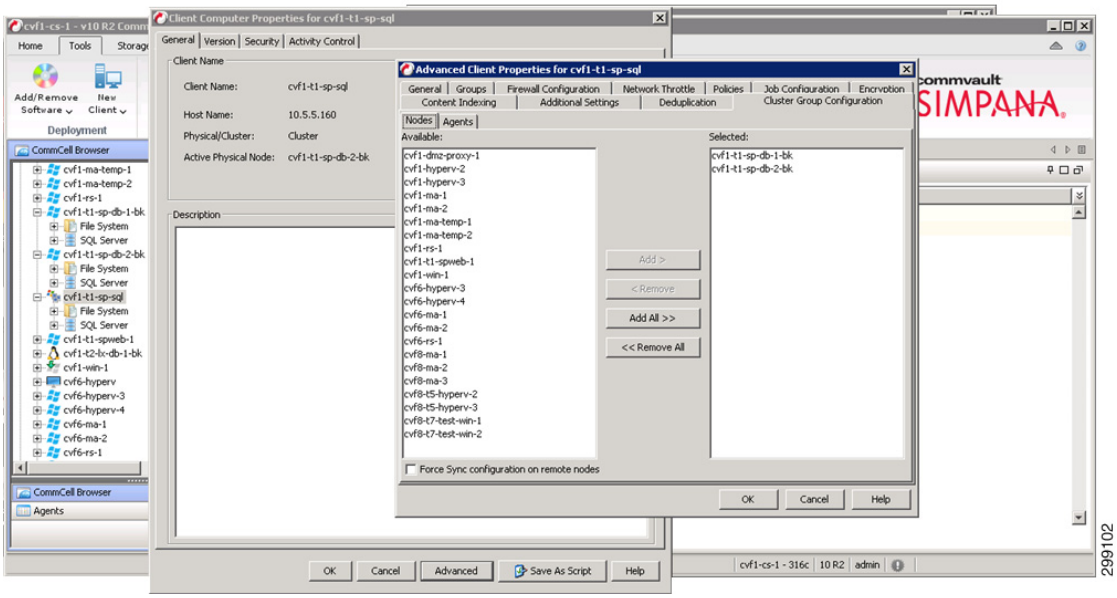

Define the iDAs that will be protecting this cluster, in this case just the SQL iDA.

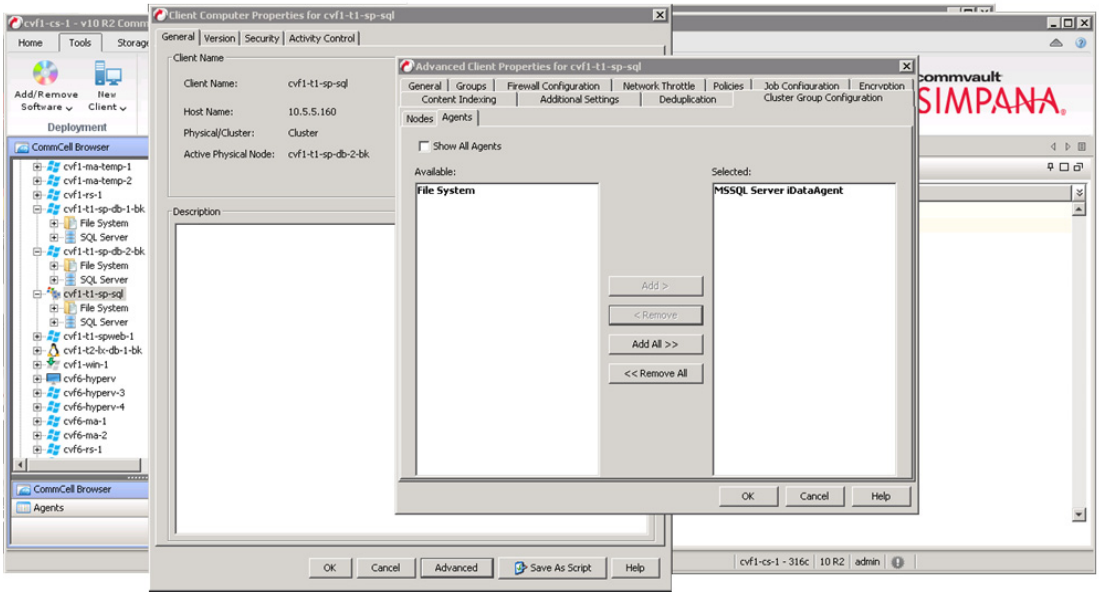

Create a new SQL Subclient.

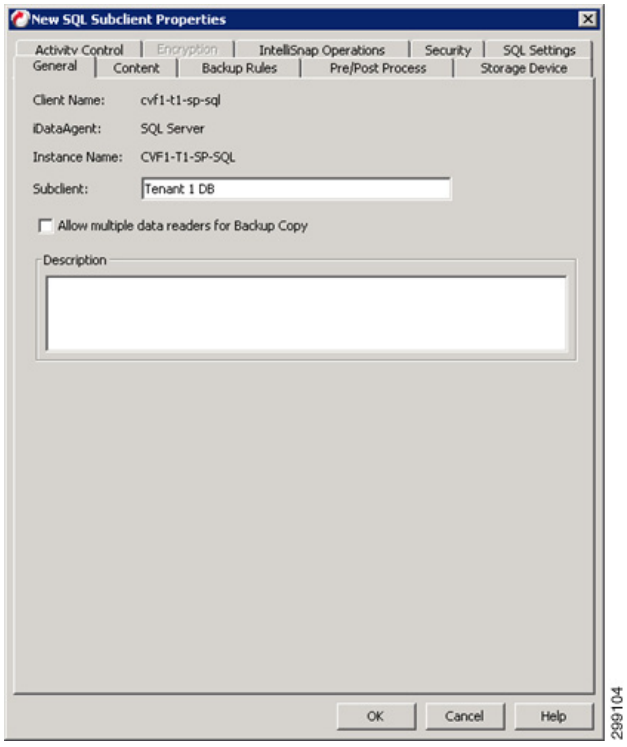

Assign the Databases to the new subclient.

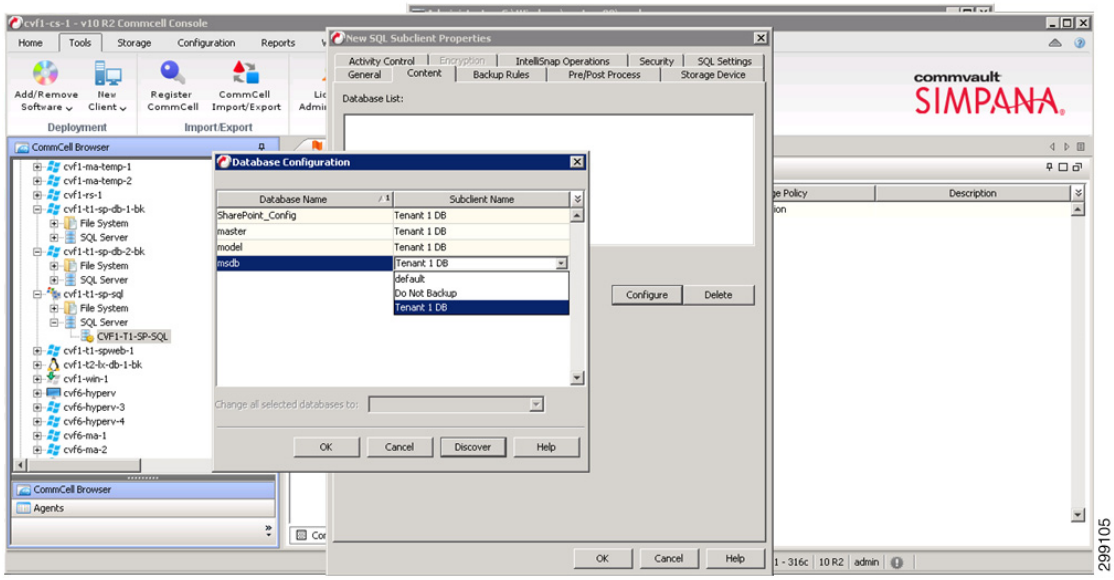

Assign this subclient to the proper Tenant Storage Policy.

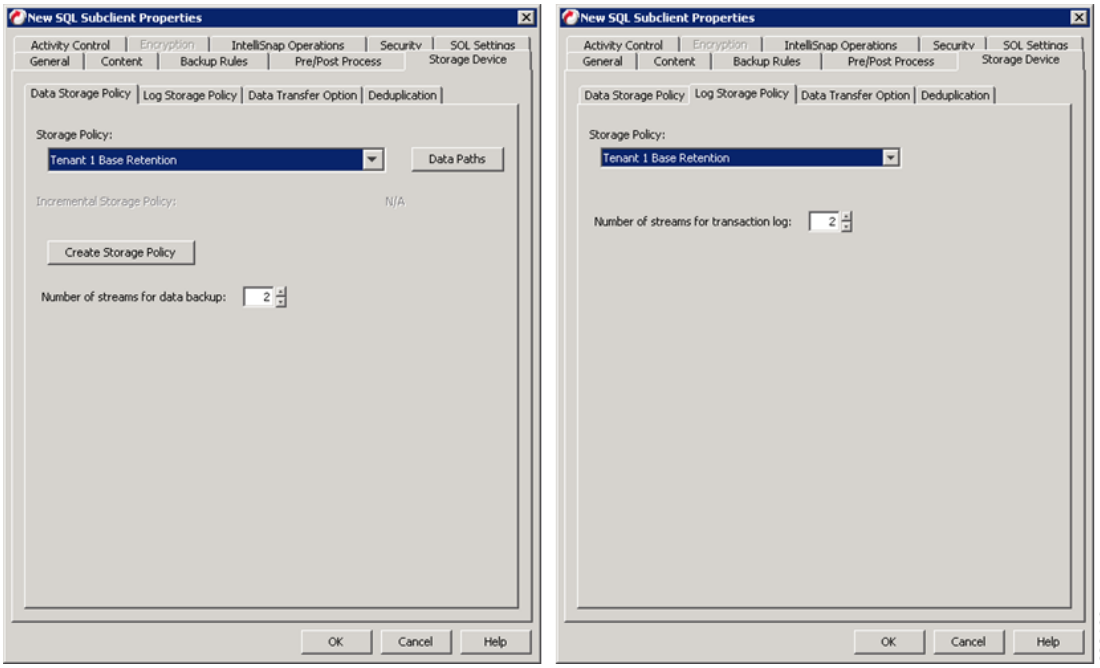

Assign this Subclient to the proper Schedule Policy and alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy. Refer to the Alert configuration for Windows File System iDA, page 4-62, above.

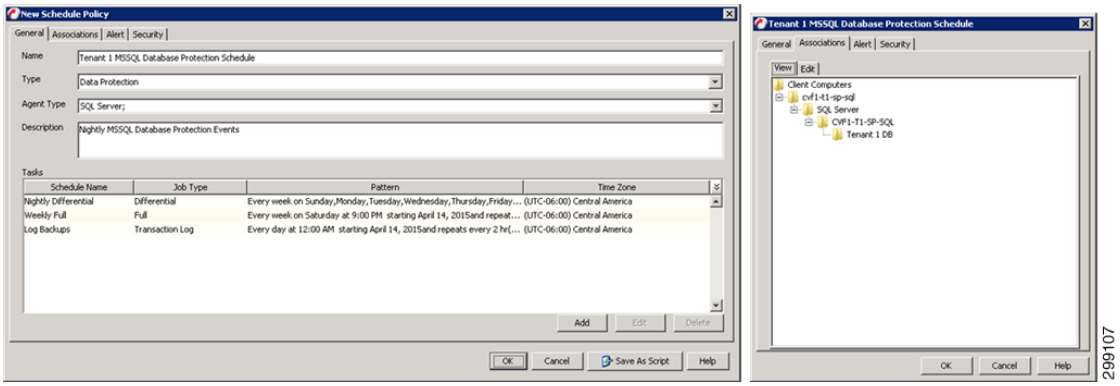

Refer to Commvault documentation for details on the Microsoft SQL Server iDA.

Oracle / Oracle RAC iDataAgent

The Oracle / Oracle RAC iDA is going to be installed on any Oracle servers, whether they are virtual or physical servers, that require database, table, and log level data protection. In this use case a single Oracle instance on a standalone server was used. Once the Client requiring protection is defined to the CommCell and then the appropriate iDA can be installed, updated and configured.

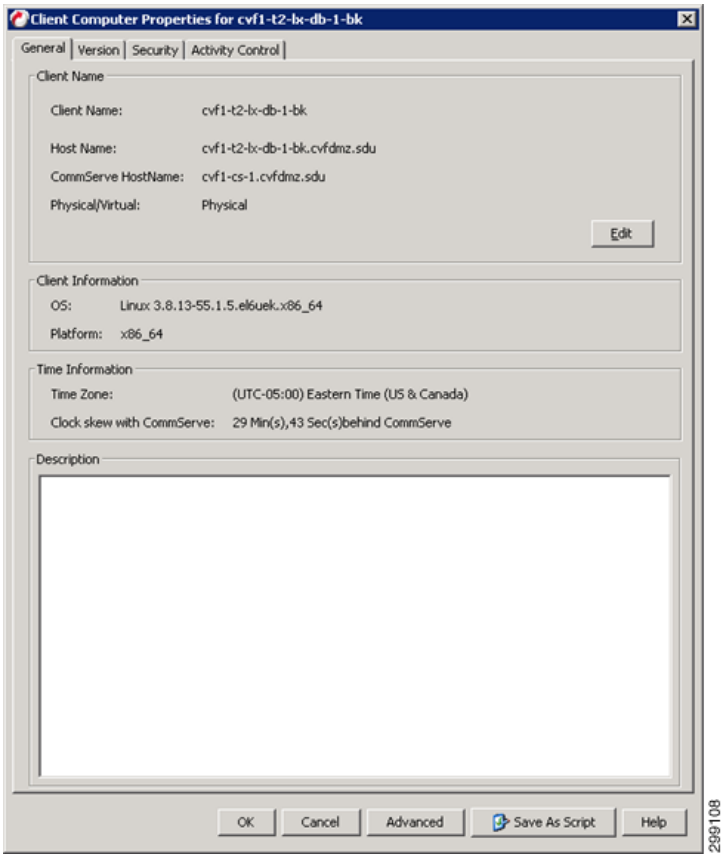

Create a new Subclient for the database being protected.

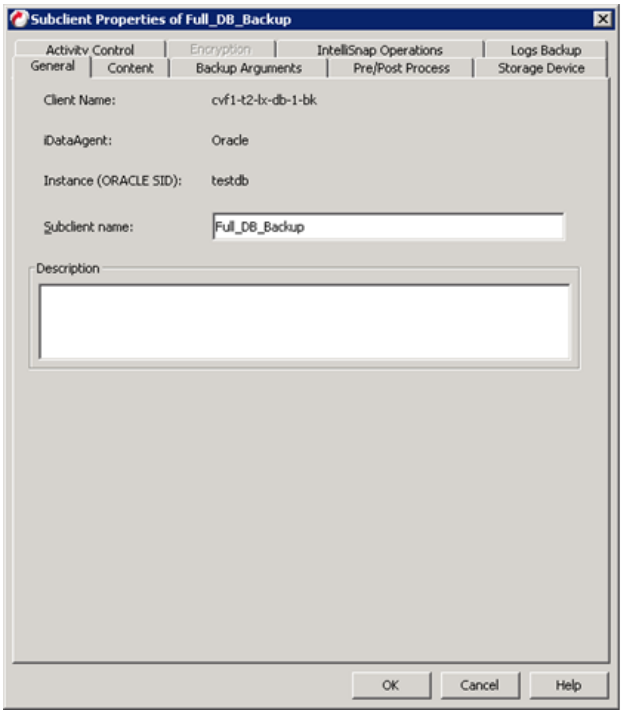

Select the Database to protect and how it will be protected.

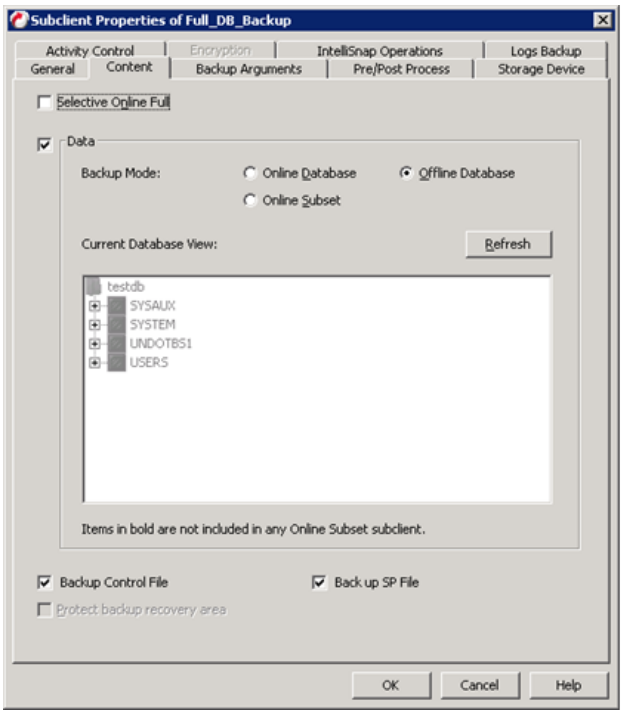

Assign the Subclient to the proper Tenant Retention Storage Policy.

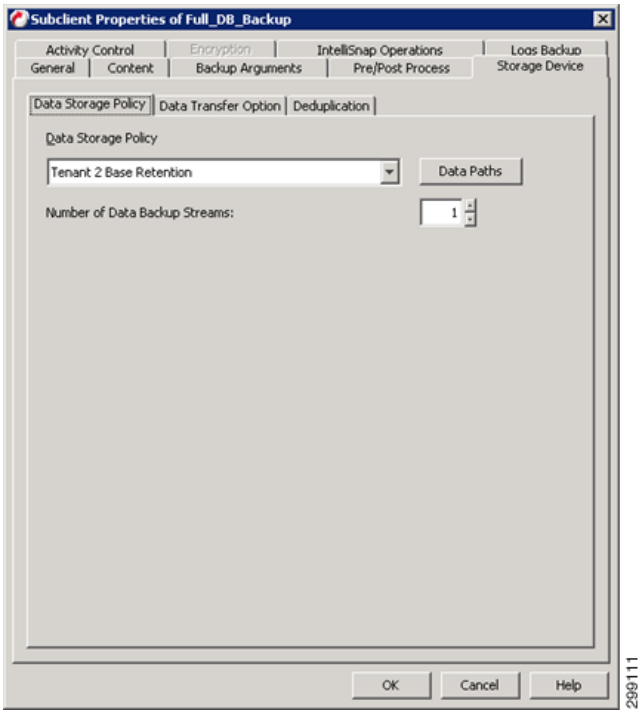

The Client is now associated to the proper Schedule Policy. The protected clients will be running an incremental backup nightly, while generating a synthetic full backup once a week. Alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy. Refer to the Alert configuration for Windows File System iDA, page 4-62, above.

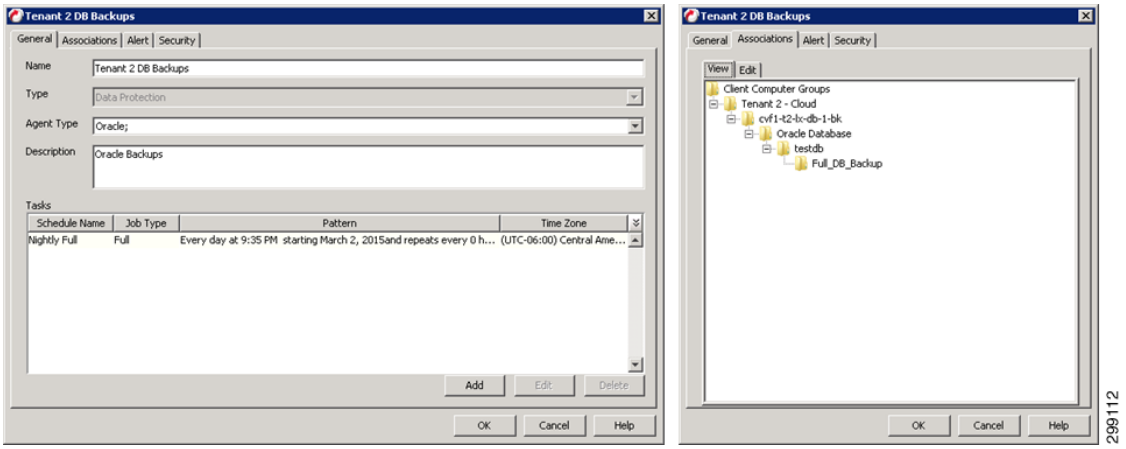

Refer to Commvault documentation for details on the Oracle RAC iDA.

Refer to Commvault documentation for details on the Oracle iDA.

Use Case 2 (Remote BaaS): Implementation Details

The second use case is Remote BaaS, where the tenant production environment is hosted at the Enterprise's own site and backed up locally within that site and remotely to the SP1 cloud. Two tenants were created to proof this use case. In the Enterprise site, Tenant 6 was configured with OpenStack and Tenant 7 was configured with VMware. Both of these tenants were backed up locally and also into Cloud Provider site SP1.

Figure 4-38 shows the environment used to proof this use case. This use case utilizes the pre-existing SP1 environment as the primary CVLT Infrastructure, so in this use case we focused on the customer locations.. Since a local copy of the data was stored at the customer's data center, integration of the MediaAgents was required. There was a MediaAgent pair in the Tenant 6 environment and a single MediaAgent in the Tenant 7 environment. The Tenant 6 environment was an RHEL OpenStack environment being protected, while Tenant 7 consisted of a VMWare environment. Each iDA interacted with the CommServe and local MediaAgent on a schedule basis to execute data protection jobs. The Primary copy of data was stored locally on the Local MediaAgent that executed the backup job, but then DASH (deduplicated replication) copied to a MediaAgent in the SP1 data center to provide for Disaster Recovery. Restoration of the data that was protected could be executed at either the customers' data center or the SP1 data center.

Figure 4-38 Use Case 2 (Remote BaaS) Implementation Diagram

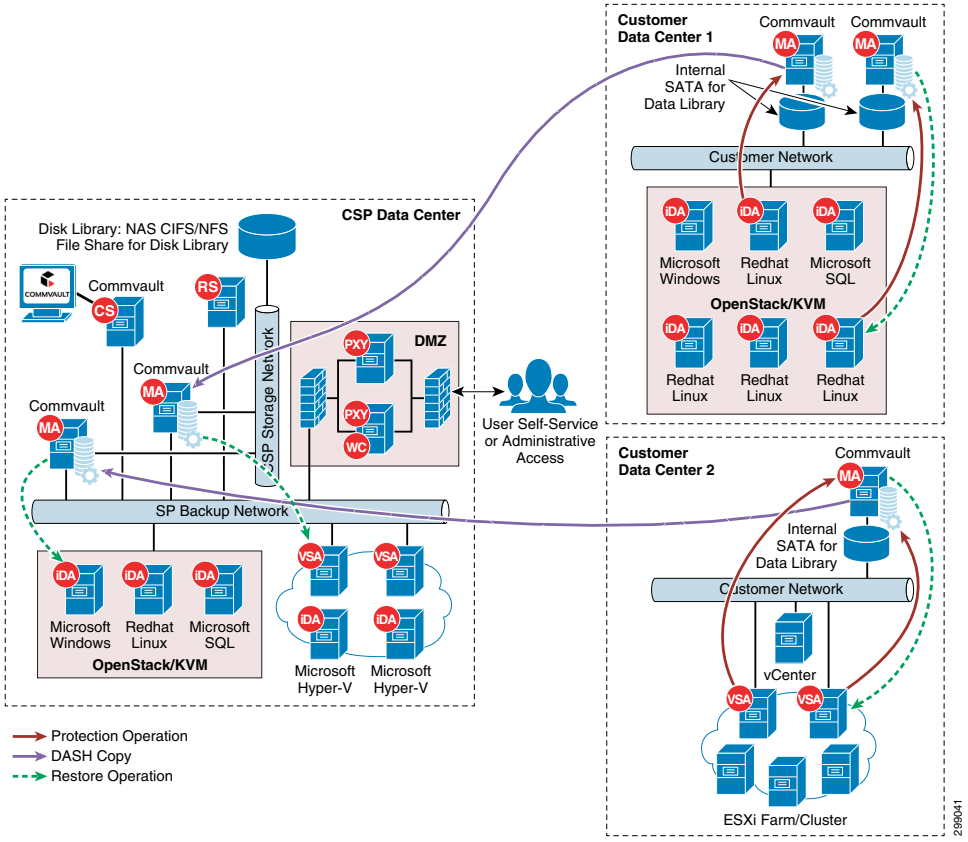

Tenant Details

Tenant 6 (RHEL OpenStack) Applications

In Tenant 6, five Linux Centos servers were used for testing purposes. There were no applications installed on these servers.

Tenant 7 (VMware) Applications

In Tenant 7, three Windows Server 2012 Standard servers were used for testing purposes. There were no applications installed on these servers.

Commvault Configuration

This section will detail how the Commvault Infrastructure and Data Protection agents were configured for this use case.

Commvault Infrastructure

This section will describe how each component of the Commvault Simpana is installed and configured.

MediaAgent Server(s)

The MediaAgent server is the data mover within the CommCell so it is the next server that needs to be built out. All connections to the actual storage, as well as the control of the deduplication is handled by the MediaAgent. The MediaAgents should be have disk volumes created for the Operating System, CVLT DDB and Index Cache. Then loaded with Windows 2012. The installation of the MediaAgent software can be pushed from the CommServe. In this case there are two MediaAgents, cvf8-ma-1, cvf8-ma-2, and cvf8-ma-3.

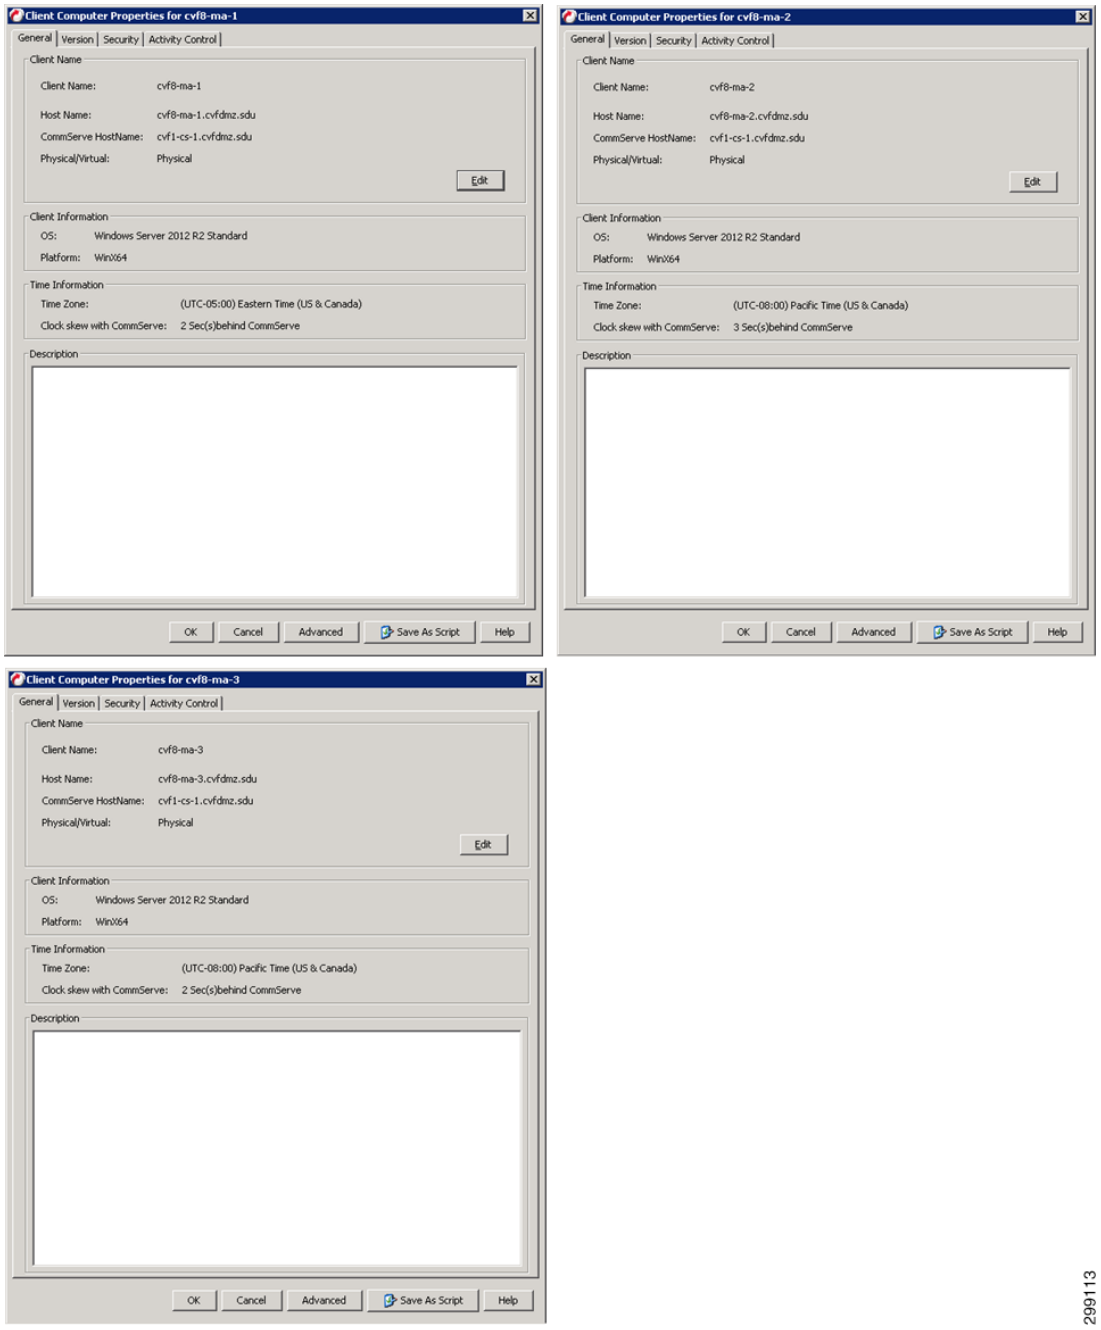

Then Disk Libraries are created from the local drives on all three MediaAgents.

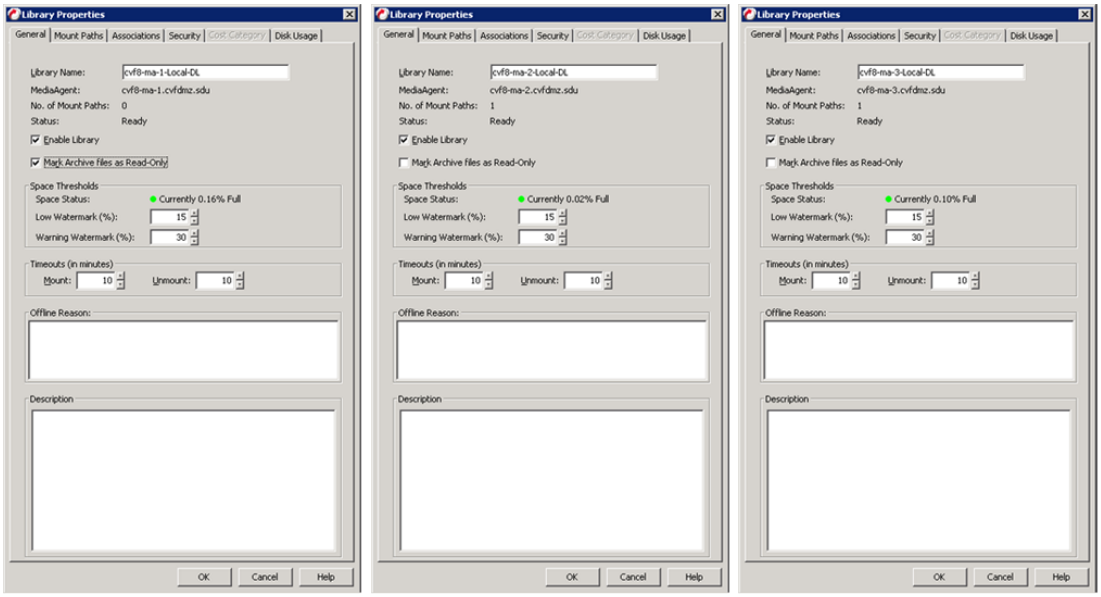

Global Partitioned DDBs are created to frontend the libraries created for CVF8-MA-1 and CVF8-MA-2 as they are working in a pair.

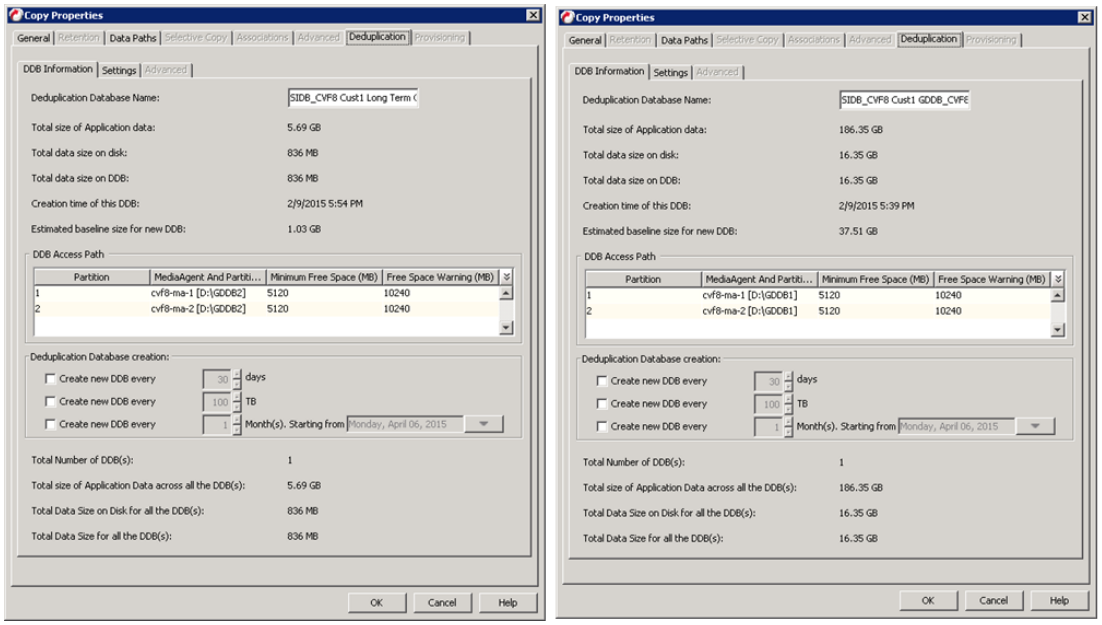

Standard DDBs are created to frontend the libraries created for CVF8-MA-3 as it is working as a standalone MediaAgent.

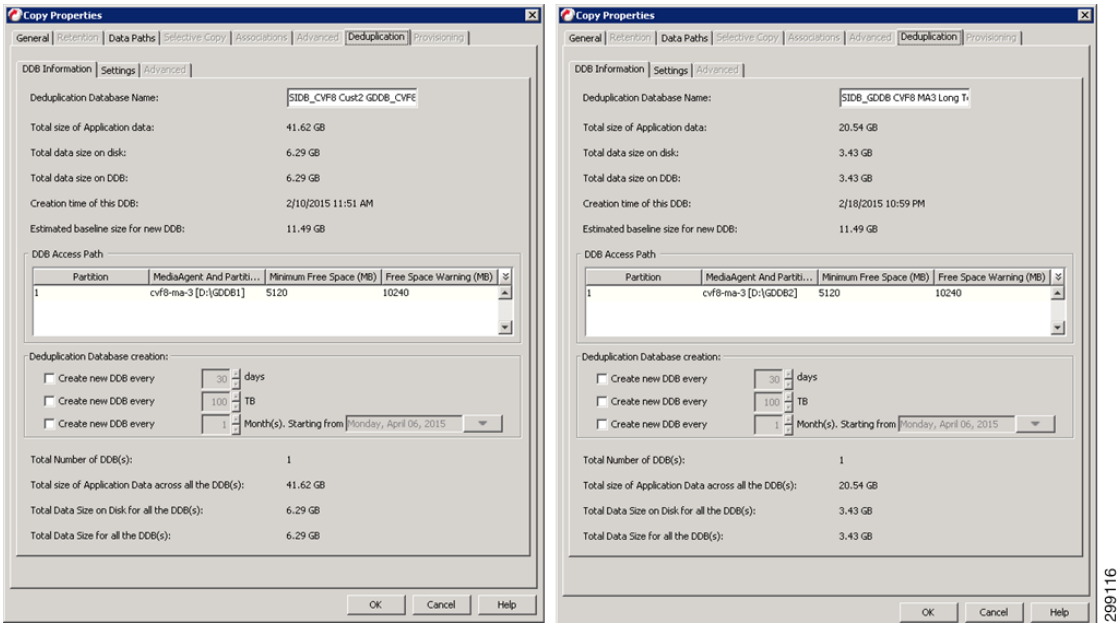

Finally the Tenant Storage Policies are configured.

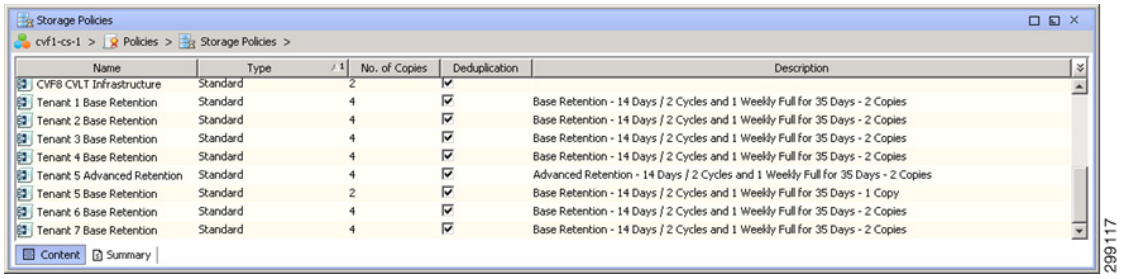

Refer to Commvault documentation for details on the MediaAgents

Virtual Server Protection

This section details how the iDataAgents were configured to provide data protection to the virtual environments in this use case.

RHEL OpenStack—Linux File System iDataAgent

The Linux File System iDA will be installed on each Linux server requiring data protection from the File System level. Once the Client requiring protection is defined to the CommCell and then the appropriate iDA can be installed, updated, and configured.

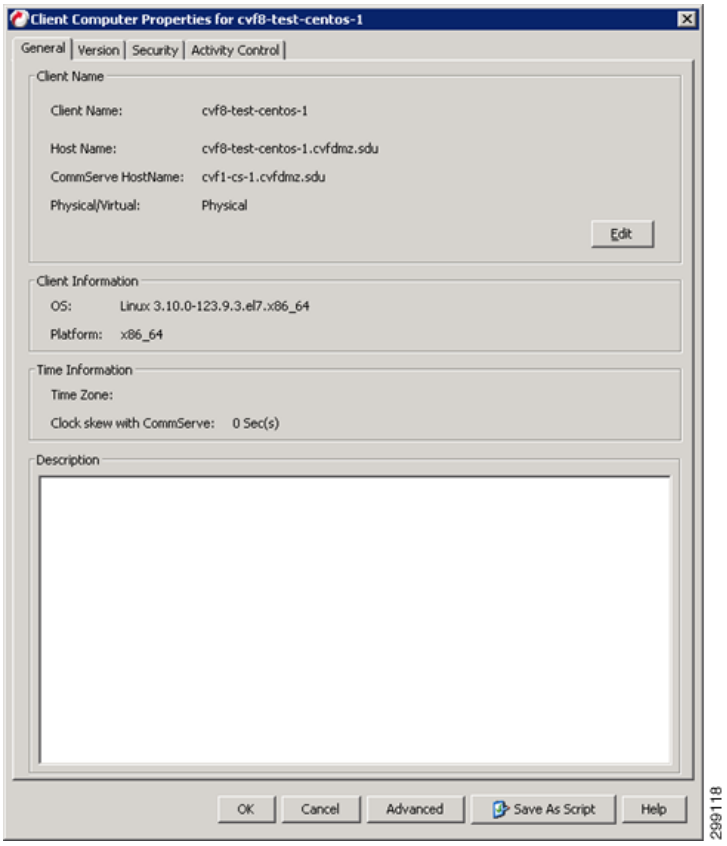

The Default File System Subclient is then assigned to the proper Tenant Retention Storage Policy.

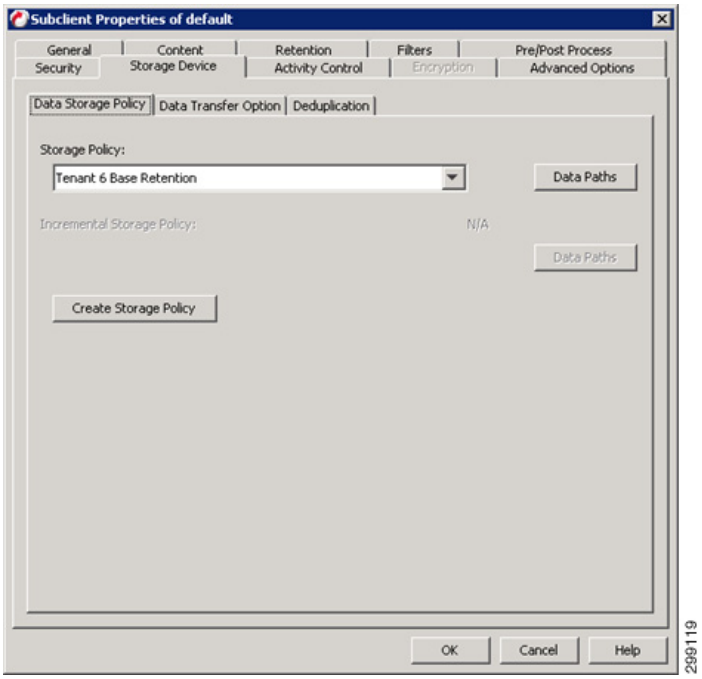

The Client is now associated to the proper Schedule Policy. The protected clients will be running an incremental backup nightly, while generating a synthetic full backup once a week and alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy.

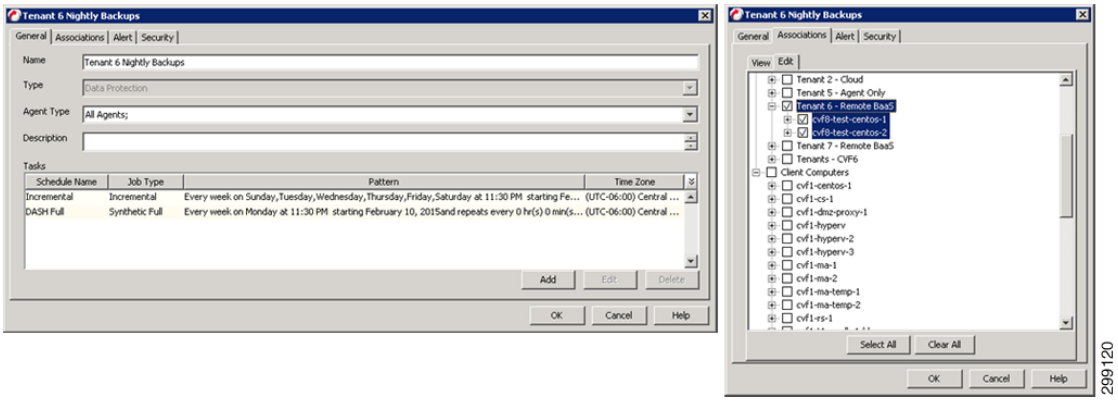

Refer to Commvault documentation for details on the Linux File System iDA.

Virtual Server iDataAgent for VMware

The VSA for VMware is installed on one or more virtual guests running Windows Server 2012. The Guests that will be used as the VSA is defined to the CommCell, the required iDAs are installed, updated and configured.

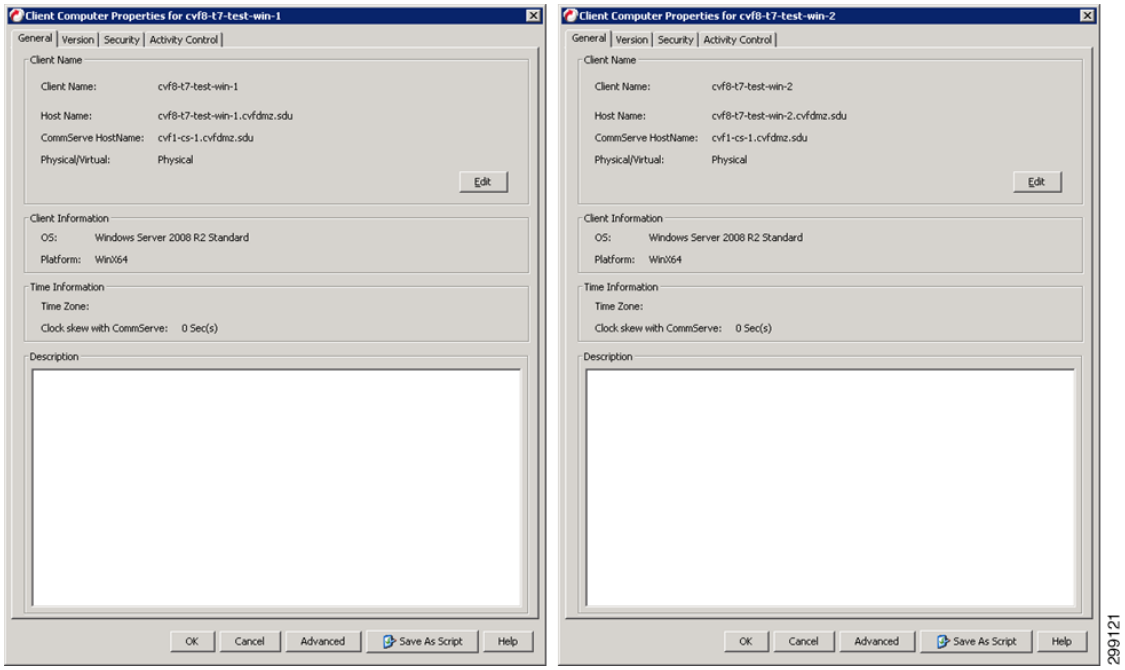

Once the VSA clients are installed, define a new VMware vCenter Client, providing credentials and designating the VSA Servers to use

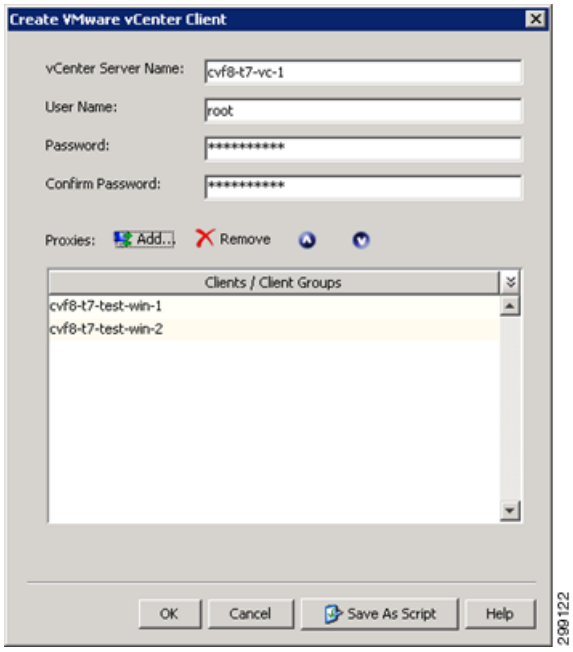

Create a subclient for the Tenant Guests and select the Guests that will be part of that Tenant.

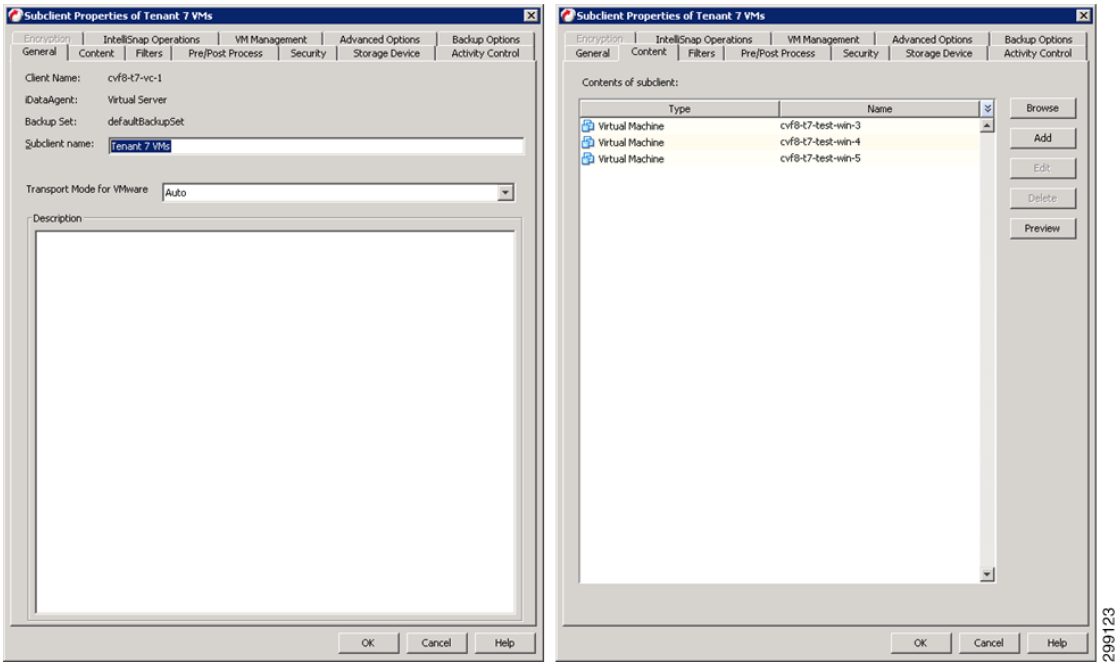

Choose the appropriate Tenant Storage Policy:

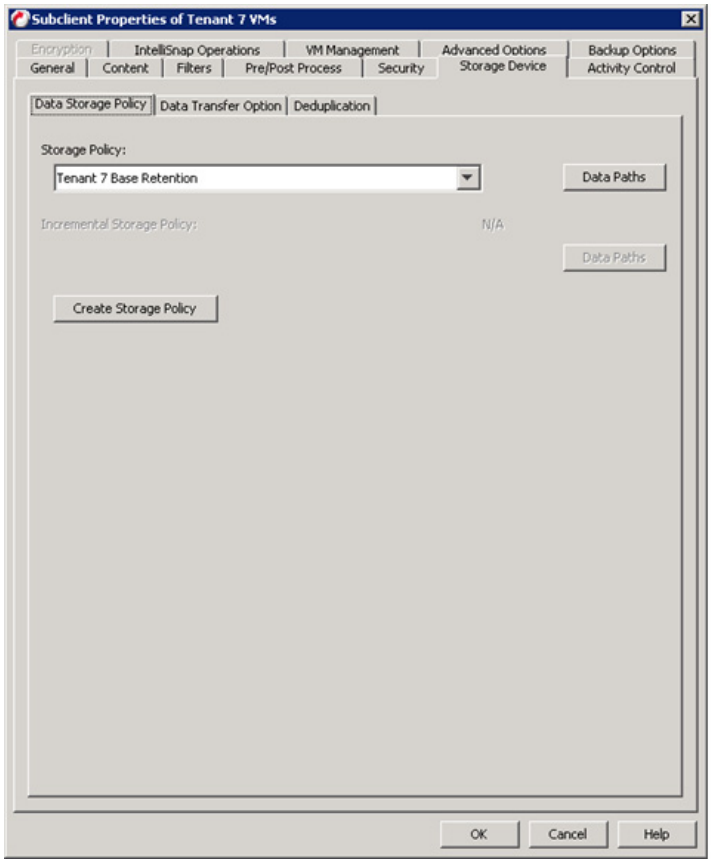

The Client is now associated to the proper Schedule Policy and alerts can be configured against a given Schedule Policy allowing the control of who gets alerts for each different policy. Refer to the Alert configuration for Windows File System iDA, page 4-62, above.

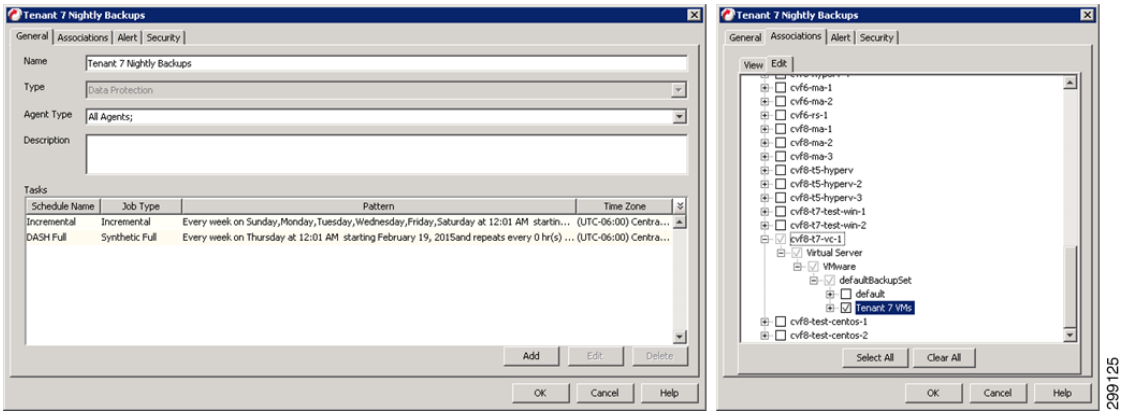

Refer to Commvault documentation for details on the Virtual Server iDA for VMware

Use Case 3 (Remote BaaS Without Local Data Retention): Implementation Details

The third use case is Remote BaaS Without Local Data Retention, where the tenant production environment is hosted at the Enterprise's own site and backed up remotely to the SP1 cloud. One tenant was created to proof this use case. In the Enterprise site, Tenant 5 was configured with Hyper-V. The virtual machines in Tenant 5 were backed up into Cloud Provider site SP1. Commvault MediaAgents were not deployed at the Enterprise site, but Commvault agents were deployed to enable the remote protection. This use case utilizes the pre-existing SP1, and optionally the SP2, environment(s) as the primary CVLT Infrastructure, so in this use case we focused on the customer location. As this use case is not going to provide a local copy of the data within the customers' data center, there is no further CVLT infrastructure integration that must be completed.

Figure 4-39 shows how the data flowed. Each iDA interacted with the CommServe and MediaAgent at SP1 on a schedule basis to execute data protection jobs. The Primary copy of data was stored on the MediaAgent in the SP1 data center that executed the backup job. Optionally, the data was later DASH (deduplicated replication) copied to a MediaAgent in the SP2 data center to provide a secondary copy outside of the SP1 data center. Restoration of the data that is protected could be executed at either the customer's data center across the network, at the SP1 or the SP2 data centers.

Figure 4-39 Use Case 3 (Remote BaaS Without Local Data Retention) Implementation Topology

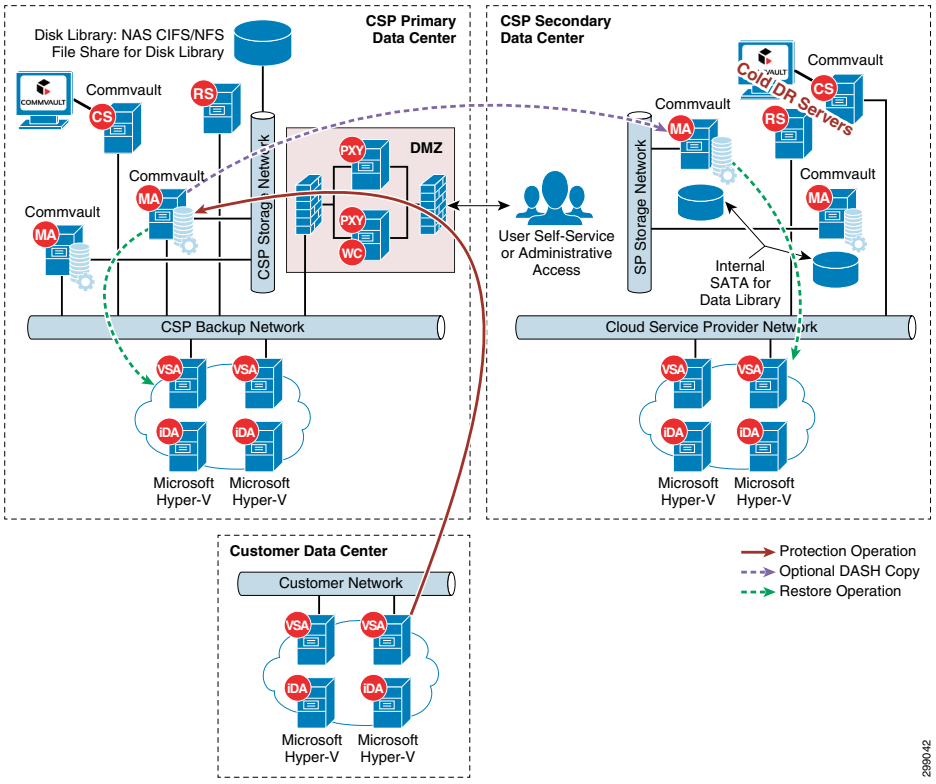

APPENDIX A

Best Practices/Caveats

This section highlights best practices and caveats that were discovered or encountered during the validation testing.

Design & Implementation Best Practices

This section discusses best practices for this solution. Review documents links in Commvault Simpana, page A-1 for complete coverage of best practices.

During the validation testing, Hyper-V networking was originally implemented in two sites. These sites were later changed to Nexus 1000V using two different methods. The first method repurposed the vNIC interfaces on the Hyper-V hosts from Hyper-V switch uplinks to Nexus 1000V uplinks. This method required removing VM and Hyper-V network configurations, implementing the Nexus 1000V configuration, then reconfiguring the VM configurations.

The second method added vNICs to the Hyper-V hosts (UCS B200 M3) Service Profile on the UCSM, which required reboots. The Nexus 1000V was then implemented in parallel to the Hyper-V networking. VMs were then moved from Hyper-V to the Nexus 1000V with little impact. Neither of these methods resulted in problem-free migrations and disruptions to the Hyper-V cluster while troubleshooting issues did occur. If the Nexus 1000V is going to be deployed in an environment, it might be desirable to deploy it at the start to avoid this migration.

As noted in this document, Nexus 1000V was replaced with the native Hyper-V virtual switch later in the CCA-MCP design to reduce complexity. However, this BaaS Commvault lab testing stayed with the N1kv component.

Commvault Simpana

This section provides sizing and best practice guides.

CommServe

- CommServe System Requirements
- CommCell Performance Tuning Guide
- CommCell Scalability Guide

MediaAgent

Building the MediaAgents in pairs, using partitioned DDBs and NAS file shares provides the best availability within the environment.

Separate SSD based arrays for the DDB and Index Cache provides the best performance and growth potential for each individual MA.

- MediaAgent System Requirements
- DeDuplication Building Block Guide
- DeDuplication Best Practices

Web Proxy

- Commvault Firewall Best Practices

Reporting Server

- Private Metrics Reporting Server system requirements

Virtual Server iDataAgent for Hyper-V

- System Requirements
- Best Practices

Virtual Server iDataAgent for VMware

- CVLT VSA Building Block Guide
- Best Practices

Windows File System iDA

- Best Practices

Linux File System iDA

- Best Practice

Microsoft SharePoint iDA

- Best Practices
- Microsoft SQL Server iDABest Practices

Microsoft SQL Server iDA

- Best Practices

Oracle / Oracle RAC iDA

- Best Practices—Oracle
- Best Practices—Oracle RAC

Caveats

This section discusses solution caveats.

Cisco CSR 1000V

Interface Renumbering After Moving to New Hyper-V Host—When the Cisco CSR 1000V is installed on a Microsoft Hyper-V cluster, the interface numbers can change after a Hyper-V host failover event to a new host server or live migration. In both cases, the condition is not seen until after a reboot. The following steps can be taken to mitigate this issue.

Prior to executing a live migration enter the **clear platform software vnic-if nvtable** command.

The command can also be successful if executed after the failover, but only before the config is saved or the VM restarted.

Configuring static MAC addresses for the network interfaces.

In the event that the interfaces have been renumbered and the IP addressing is removed, the following steps can be used to recover.

1. Execute **clear platform software vnic-if nvtable** command.
2. Copy saved config to startup config.
3. Reboot the CSR.

Cisco Nexus 1000V

Migrating VM from Hyper-V Switch to Nexus 1000V—After migrating to the Nexus 1000V, when configuring existing VM Network Adapter interfaces that were previously configured for a Hyper-V switch, the change may not complete or the VM may fail to start. In either case, you may need to remove the existing interfaces and create new ones.

Cisco UCS C240 M3

Broadcast Packets—During validation testing, an issue was discovered that impacted the C240 servers from receiving broadcast packets. The issue was isolated to the VIC 1225 network driver in release 2.0(3d) and was resolved in the VIC driver in release 2.0(3i). Refer to CSCur44975 for more details.

APPENDIX B

Technical References

This appendix lists all of reference matter used during the design and implementation of the solution.

Cisco

- Cisco Cloud Architecture for the Microsoft Cloud Platform: Infrastructure Foundation Guide, Release 1.0
- Cisco UCS B-Series Blade Servers Windows Installation Guide
- Cisco Nexus 1000V Switch for Microsoft Hyper-V Deployment Guide
- Nexus 1000v Quickstart with Hyper-V Server Configuration Example
- Installing the Cisco CSR 1000V in Microsoft Hyper-V Environments
- Cisco Nexus 1000V Install and Upgrade Guide
- Cisco CSR 1000V Series Cloud Services Router Release Notes
- CSR 1000v Series Cloud Services Router Software Configuration Guide

Commvault

- Commvault Simpana v10 Books Online

Microsoft

- Microsoft Applications on Cisco UCS
- Cisco Unified Computing System with Microsoft Hyper-V Recommended Practices
- Failover Clustering Hardware Requirements and Storage Options
- Microsoft Technet Library—System Center Virtual Machine Manager 2012

OpenStack

- How can OpenStack standardize cloud computing?
- OpenStack: Packaged by and for the CentOS community
- OpenStack

APPENDIX C

Terms and Acronyms

The following tables list industry vs Commvault descriptors, and technical terms and acronyms used in this document.

Industry and Commvault Terminology

Commvault and its Simpana® data management software incorporates many industry standard and unique features. It may be difficult for a non-Commvault user to compare well-known components such as a management server with Commvault branded terminology. Table C-1 showing industry terms with Commvault branded terminology.

Table C-1 Service Support Roles

Industry Term	Commvault Branded Term
Group and User Permission	Capabilities & User Actions
Agent	iDataAgent
Backup Server	MediaAgent
Desktop & Laptops	Edge Devices
Management Server	CommServe
Backup Environment	CommCell
Laptop, Server, and/or Virtual Machine containing 1 or more Agents	Client or Client Computer
Selection of data on a Client to be managed uniquely	Sub-client

Table C-1 Service Support Roles (continued)

Industry Term	Commvault Branded Term
Secure network routing	• Firewall Configuration • Direct Connections using port tunnels • Port-forwarding gateways • The perimeter network (also known as a DMZ) using a Simpana® proxy • HTTP proxies (including WiFi connections) • Combinations of these
Collection of settings Retention Storage logical target Number of data copies Storage lifecycle policy Storage configuration	Storage Policy

Terminology

Table C-2 defines technology terms and acronyms used in this document or by Commvault.

Table C-2 Terms and Definitions

Term	Definition
Application	A term of convenience that encompasses applications, databases, even VMs.
Archiving	Copying computer data.
Auxiliary Copy	Copy (backup) of computer data.
BaaS	Backup-as-a-Service
Backup	Archived computer data.
Backup Set	A group of subclients defining data to be backed up by the iDataAgent.
Backupset	A group of subclients which includes all of the data backed up by the Virtual Server Agent.
BGP	Border Gateway Protocol
CCA-MCP	Cisco Cloud Architecture for Microsoft Cloud Platform
Client	Software that accesses a remote service on another computer.
CommCell	A Commvault Simpana data protection environment made up of at least a CommServer, a MediaAgent, and some number of clients to be protected.
CommServe	The Commvault Simpana primary server that is responsible for client management, scheduling, data retention, job history, and media management.

Table C-2 Terms and Definitions (continued)

Term	Definition
CSP	Cloud Service Provider
CSR (1000V)	Cisco Cloud Services Router 1000V
Data Set	Collection of data
DC	Data center
Discovery	Searching Virtual Machines in the vCenter, Hyper-V cluster or server based on a specific criteria.
DR	Disaster Recovery
Drive Pool	Collection of storage devices.
DVS	Distributed Virtual Switch
FC	Fibre Channel
GE	Gigabit Ethernet
Guest Host	A virtual machine.
Guest OS	The operating system running on the virtual machine (such as Windows or Linux).
HA	High Availability
HDD	Hard Disk Drive
IaaS	Infrastructure-as-a-Service
iDataAgent	Provides unified protection and recovery for most common operating systems, databases, and applications.
Index Cache	Catalog of all data that has been protected and where it is retained within the CommCell.
ISR	Integrated Services Router
LAN	Local Area Network
Library	Logical collection of disk or tape used to store the data protected within the CommCell.
MediaAgent	The workhorse of the environment that manages deduplication database and the data transmission between clients and storage media.
Mount Path	The definition of the path used to write to the disks or tape drives within a Library.
MPLS	Multi Protocol Label Switching
MPLS	Multi Protocol Label Switching
NAS	Network Attached Storage
NAT	Network Address Translation
NFS	Network File System
Node	The Hyper-V server on which the *i*DataAgent is installed. This computer facilitates most of the data movement from the Hyper-V server to the backup media. Such computers are referred to as a **Node**.
OSPF	Open Shortest Path First
PE	Provider Edge

Table C-2 Terms and Definitions (continued)

Term	Definition
Proxy Computer	A physical computer separate from the host computer on which the Virtual Server Agent is installed. This computer facilitates most of the data movement from the host computer to the backup media. In some cases, this computer may be a virtual machine installed on the host computer. Such computers are referred to as a **Proxy**.
Restore	Recovering stored data from backup or archive.
ROBO	Remote Office/Branch Office
SAN	Storage Area Network
SLA	Service License Agreement
SLB	Server Load Balancing
SP	Service Provider
SSD	Solid State Drive
Storage Policy	Procedural guidelines for housing, securing, and archiving data.
STP	Spanning Tree Protocol
Subclient	A logical entity that uniquely defines a unit of data, or set of virtual machines, to be backed up.
SVI	Switch Virtual Interface
Tenant	Consumer of backup or any cloud services.
UCS	(Cisco) Unified Computing System
UCSM	(Cisco) UCS Manager
Vendor	The virtualization software being used (such as VMware).
Virtual Client	A logical entity that serves as a single point of administration for all proxies or servers in a Hyper-V cluster.
Virtual Server Agent Virtual Server iDataAgent	A software module that performs backup and restore of virtual machine data.
VLAN	Virtual LAN
VM	Virtual Machine
vPC	Virtual PortChannel
VPC	Virtual Private Cloud
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
WAN	Wide Area Network