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

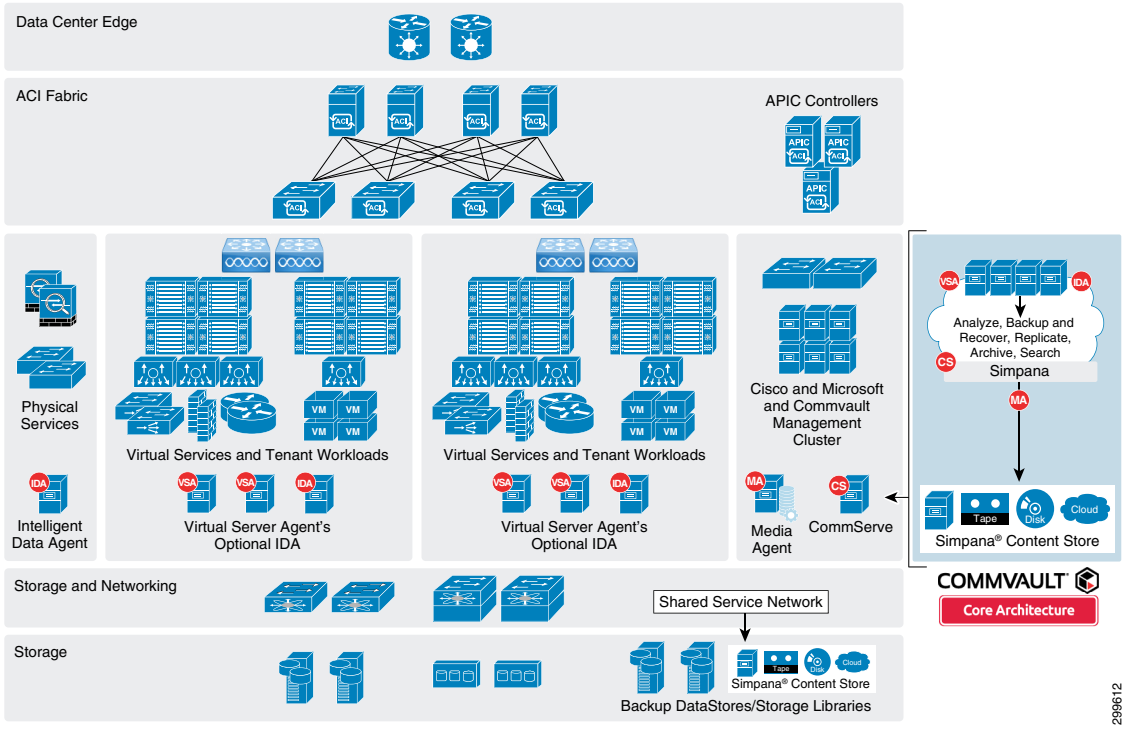

Commvault Simpana Solution Components

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

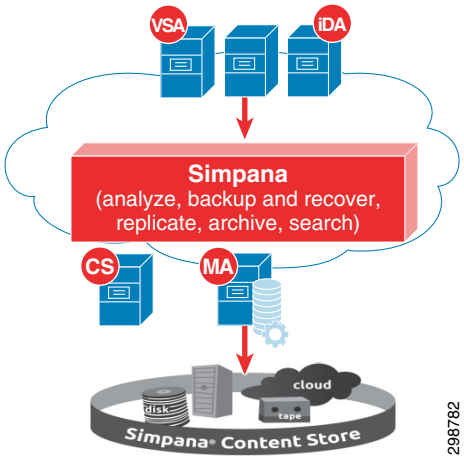

A logical representation of the Cisco BaaS solution/architecture resources for a single Cloud Service Provider is shown in Figure 3-3.

Figure 3-3 Cisco BaaS Solution Logical Topology

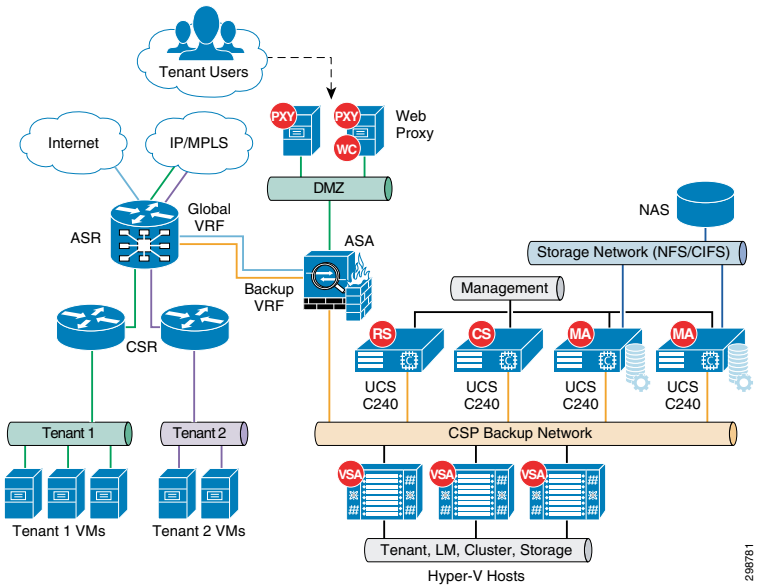

A typical CCA-MCP CSP environment will have tenants running applications—for example, Oracle, Exchange, SQL—in physical or virtual environments. Those apps are typically accessed by clients via an IP network, and leveraging production storage (either SAN, NAS, or even DAS).

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

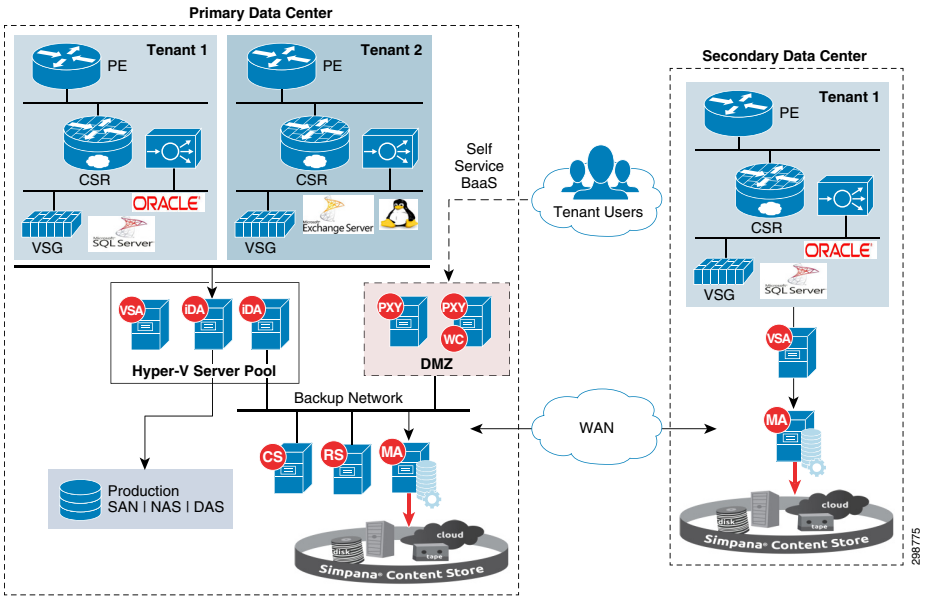

In a multitenant environment, each customer is mapped as a separate CCA-MCP tenant where the necessary network security is provided and traffic segregation is maintained.

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

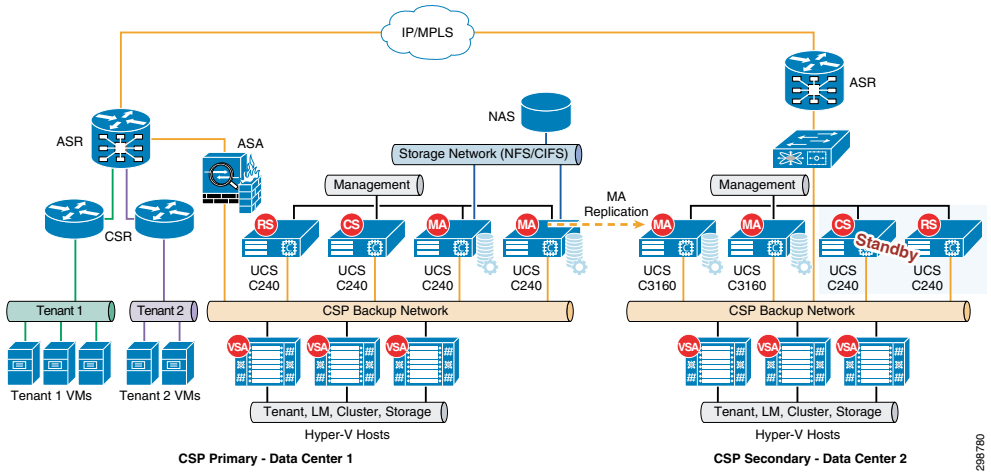

The data replicated to the secondary CSP data center can be used to support scenarios such as primary data center failure, spinning up the VMs along with production data to support additional use cases such as Test/Dev and analytics, etc.

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

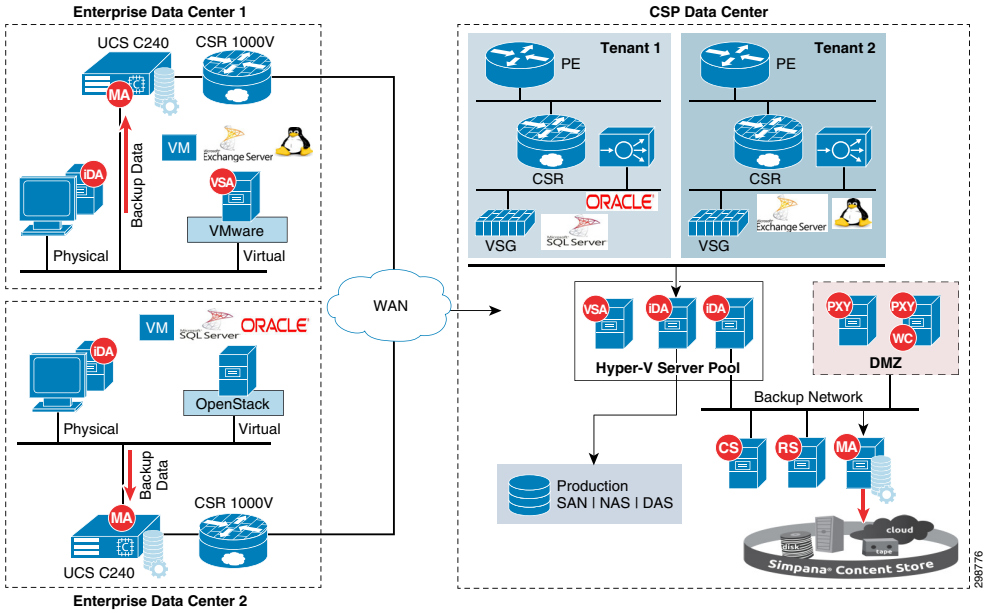

Remote BaaS without Local Retention

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

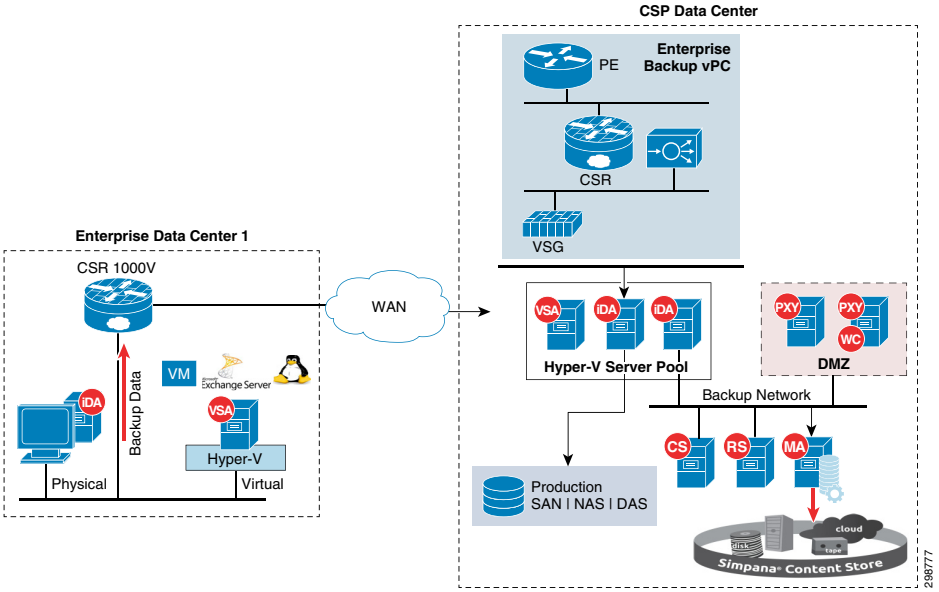

WAN Connectivity

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

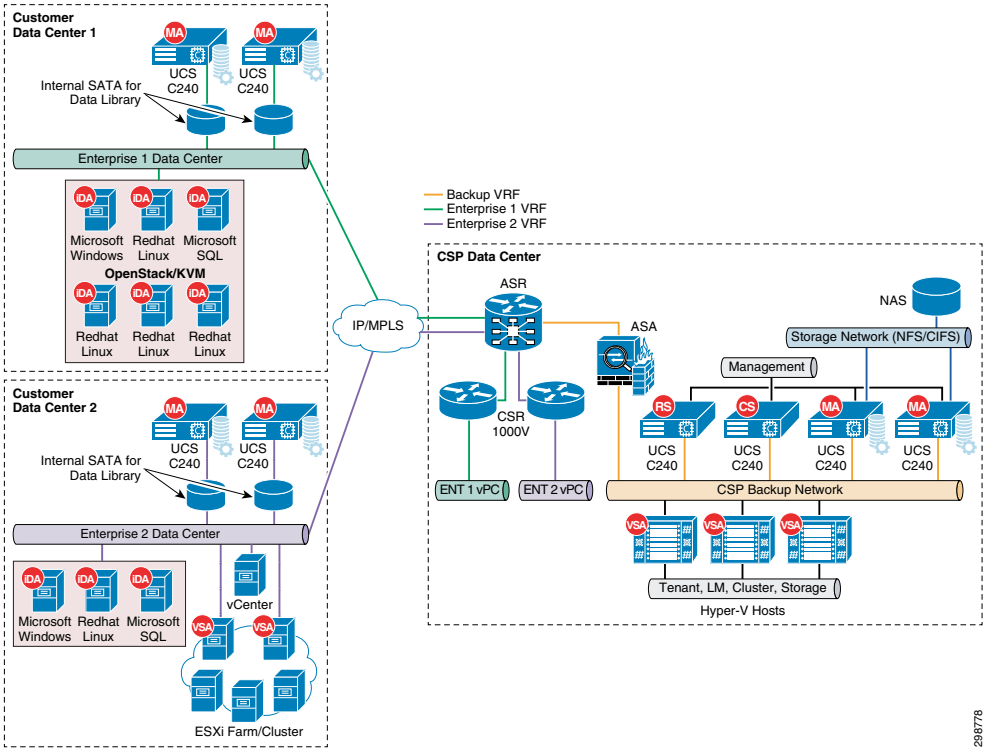

The BaaS Architecture includes a dedicated Backup VRF on the ASR 9000 WAN router and have all the Commvault components deployed behind an ASA firewall connected to this VRF. In the CCA-MCP Solution, there is a ASA firewall cluster, and used in multi-context mode – one context can be used for BaaS or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

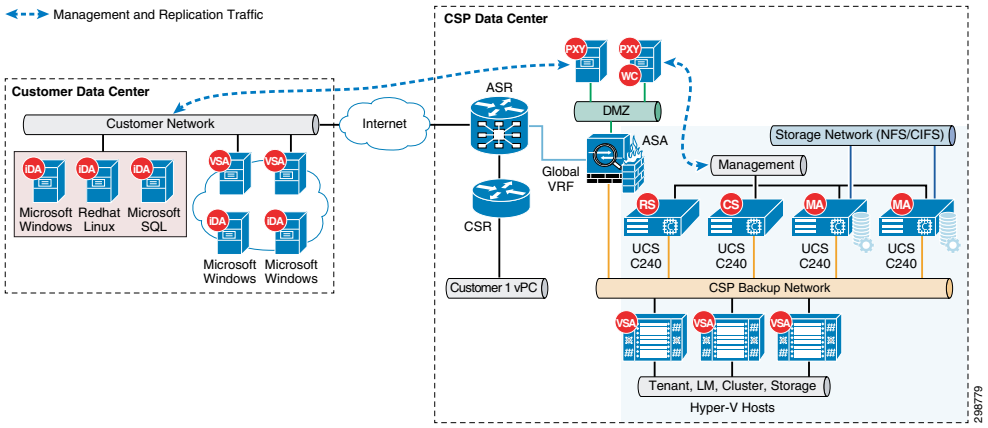

Cisco Storage Server as Converged MA

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

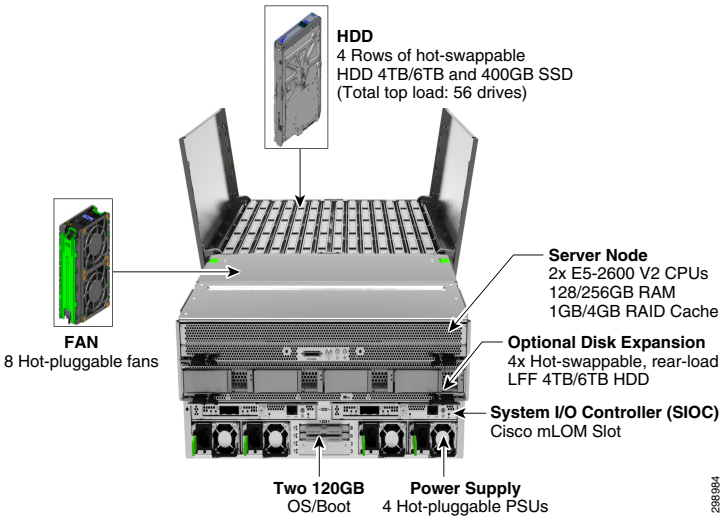

Commvault Architecture/Design Considerations

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

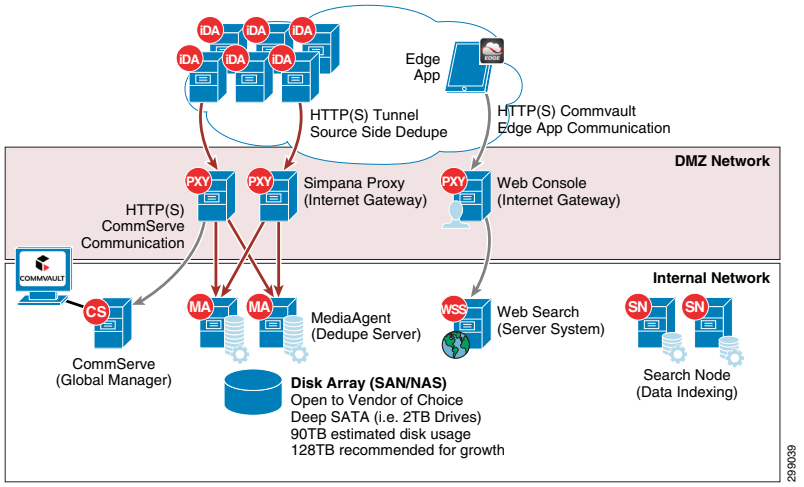

First, from a security perspective Simpana® utilizes certificate based authentication between Simpana® components and client computers. This protects against a variety of networking attacks such as spoofing. Secondly, Simpana® provides the ability to have a dedicated interfaces or shared networking interfaces 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

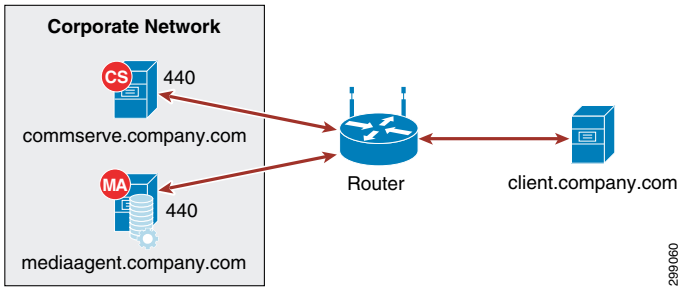

- Port-forwarding gateways - There are cases where direct connectivity setups do not work. Consider the case of the CommServe and MediaAgent being located inside a company's internal network, with the entire network being exposed to the outside world through a single IP address. Typically, this IP address 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

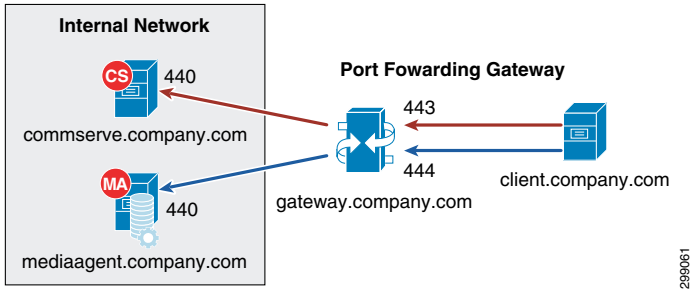

- The perimeter network (also known as a DMZ) using a Simpana proxy - Simpana proxy is a special proxy configuration where a dedicated iDataAgent is placed in a perimeter network and the firewalls are configured to allow connections (from inside and outside networks) into the perimeter network. The proxy, 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

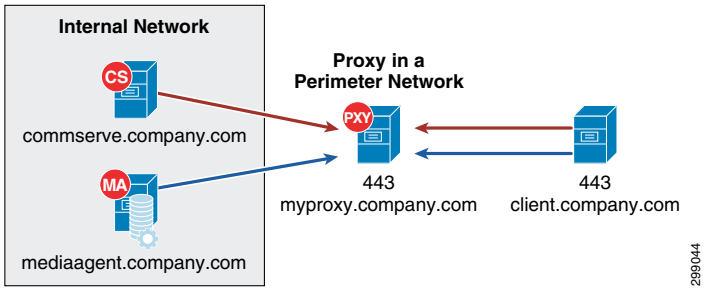

- HTTP proxies (including WiFi connections)—Consider the scenario where you are in a public location like a coffee shop, airport, hotel, or other such remote locations where Internet access is using public WiFi through a HTTP proxy. If you are a roaming user who travels frequently, you might operate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

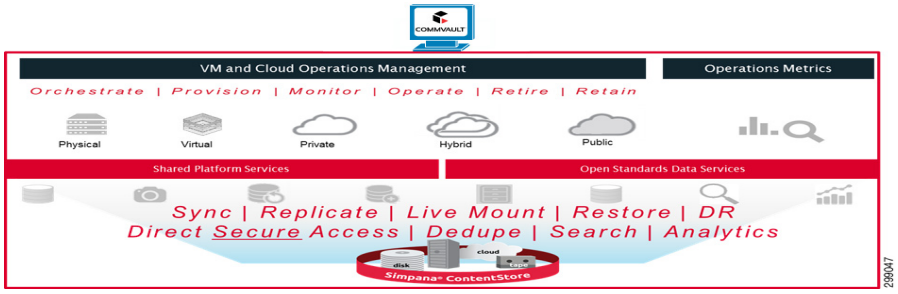

Commvault Simpana's policy based methodology allows a Service Provider to manage multiple different types of data from multiple different platforms for multiple different customers efficiently and securely. The Commvault Simpana Storage Policies act as a channel for backup and restore operations. Its 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

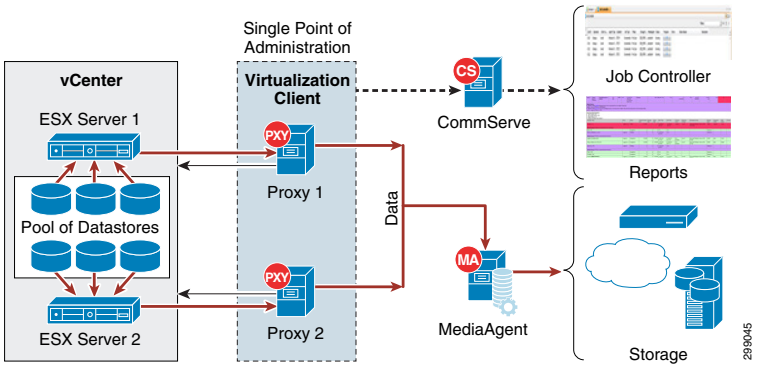

Virtual Server iDataAgent for Hyper-V

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

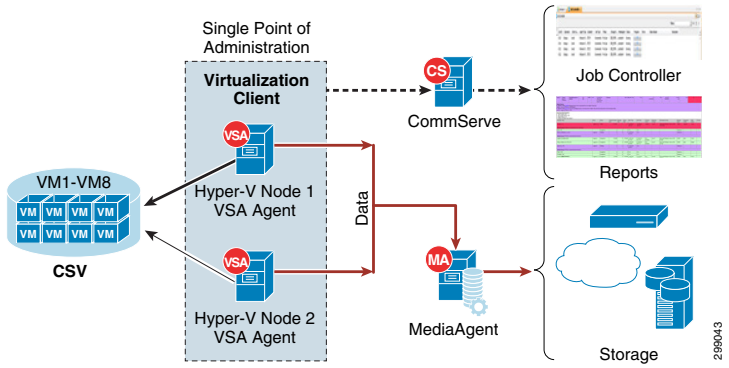

Application iDataAgent

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.