usenix
the advanced computing systems association

ISBN-13: 978-1-931971-83-6
90000
9 781931 971836

24 Arrasjid, Lin, Veeramraju, Kaplan, Epping, and Haines Cloud Computing with VMware vCloud Director

24 | Short Topics in System Administration

Jane-Ellen Long, Series Editor

Cloud Computing

with VMware vCloud Director

usenix
the advanced computing systems association

John Y. Arrasjid
Ben Lin
Raman Veeramraju
Steve Kaplan
Duncan Epping
Michael Haines

Foreword by Paul Maritz, VMware CEO

Booklets in the Series

#24: **Cloud Computing with VMware vCloud Director,** by John Y. Arrasjid, Ben Lin, Raman Veeramraju, Steve Kaplan, Duncan Epping, and Michael Haines

#23: **A Sysadmin's Guide to Navigating the Business World,** by Mark Burgess and Carolyn Rowland

#22: **Job Descriptions for System Administrators, 3d Edition,** by Tina Darmohray

#21: **Foundation for Cloud Computing with VMware vSphere 4,** by John Arrasjid, Duncan Epping, and Steve Kaplan

#20: **Running the Numbers: System, Network, and Environment Monitoring,** by Daniel V. Klein and John Sellens

#19: **Configuration Management with Bcfg2,** by Narayan Desai and Cory Lueninghoener

#18: **Deploying the VMware Infrastructure,** by John Arrasjid, Karthik Balachandran, Daniel Conde, Gary Lamb, and Steve Kaplan

#17: **LCFG: A Practical Tool for System Configuration,** by Paul Anderson

#16: **A System Engineer's Guide to Host Configuration and Maintenance Using Cfengine,** by Mark Burgess and Æleen Frisch

#15: **Internet Postmaster: Duties and Responsibilities,** by Nick Christenson and Brad Knowles

#14: **System Configuration,** by Paul Anderson

#13: **The Sysadmin's Guide to Oracle,** by Ben Rockwood

#12: **Building a Logging Infrastructure,** by Abe Singer and Tina Bird

#11: **Documentation Writing for System Administrators,** by Mark C. Langston

#10: **Budgeting for SysAdmins,** by Adam Moskowitz

#9: **Backups and Recovery,** by W. Curtis Preston and Hal Skelly

#8: *Replaced by #22*

#7: **System and Network Administration for Higher Reliability,** by John Sellens

#6: **A System Administrator's Guide to Auditing,** by Geoff Halprin

#5: **Hiring System Administrators,** by Gretchen Phillips

#4: **Educating and Training System Administrators: A Survey,** by David Kuncicky and Bruce Alan Wynn

#3: **System Security: A Management Perspective,** by David Oppenheimer, David Wagner, and Michele D. Crabb, and edited by Dan Geer

#2: **A Guide to Developing Computing Policy Documents,** edited by Barbara L. Dijker

24 *Short Topics in*
System Administration

Jane-Ellen Long, Series Editor

Cloud Computing with VMware vCloud Director

John Y. Arrasjid, Ben Lin, Raman Veeramraju,

Steve Kaplan, Duncan Epping, and Michael Haines

Published by the USENIX Association
2011

ISBN 978-1-931971-83-6

To purchase additional copies, see http://www.sage.org/pubs/short_topics.html.

The USENIX Association
2560 Ninth Street, Suite 215
Berkeley, CA USA 94710

http://www.usenix.org/

Contents

Figures and Tables

Figures

Tables

Acknowledgments

The authors work for VMware and INX. The opinions expressed here are the authors' personal opinions. Content published was not read or approved in advance by VMware and does not necessarily reflect the views and opinions of VMware. This is the authors' book, not a VMware book.

Thank you to Duncan Epping and Steve Kaplan for participating again on our third book in the USENIX Association Short Topics series. Thank you to Ben Lin, Raman Veeramraju, and Michael Haines for joining our team in creating this work and lending additional field experience.

A special thank you to Daniel Conde (co-author of *Deploying the VMware Infrastructure*) for his exhaustive review and various contributions, including the vision of VMware's view of cloud computing. Thank you to Mahesh Rajani and David Baldwin for providing review and additional content. Thank you to Matthew Wood for providing editorial support, Bret Bechis for his help in legal approval, Izhar ul Haq for his cover design, and Sam Charrington for permitting use of "The Blind Men and the Cloud."

Thank you to my management team (Michael "Dino" Cicciarelli; Matthew Stepanski; Enis Konuk; Carl Eschenbach; and Stephen Beck) for their support on this project. Thank you to Paul Maritz for his leadership and for providing the Foreword to this book.

Thank you to other contributors and reviewers of this work. These include Howie Xu, Catherine Fan, Aidan Dalgleish, Serge Maskalik, Patrick Chang, Bill Call, Amrainder Singh, Ananda Kammampati, Sanjay Patnaik, Deji Akomolafe, Phil Callahan, Chris Colotti, Daniel Crider, Dushyanth Nataraj, Hemma Prafallchandra (HyTrust), and others.

Thank you to Jane-Ellen Long for editorial support and to Jane-Ellen Long, Casey Henderson, and Jennifer Peterson for their production support and great ideas! Thank you to the USENIX Association for helping this work see the light of day and for supporting my participation as a member of the USENIX Board of Directors. Thank you to Dan Klein for allowing me to teach (since 2005) and for being a great friend. I encourage readers to experience USENIX and LISA personally to see why I believe they can influence your present and future in this field.

Thank you to the families and friends of the authors.

I dedicate this book to Amy, Catherine, Sofi, Lila, Mom, and Dad.

John Y. Arrasjid, VCDX
VMware Inc., Principal Architect, Cloud Services,
and
USENIX Association, Board of Directors

Foreword

Before you read this book, it's useful to have some context on VMware's view of cloud computing and to place this in a historical context that helps us understand why the cloud is becoming the new computing paradigm and why it's important to simplify IT so that businesses can focus on core functions, not on the plumbing.

Cloud = How, Not Where

Cloud computing is about *how* computing is getting done, not *where* it is getting done. The great promise of the cloud is that it will enable things to get done faster and more cheaply—by removing and hiding complexity. The popular press often associates cloud computing with workloads running off-premises at an external, public computing provider. VMware has a broader view of cloud computing that does not tie computing to a location and advocates a much more flexible form of computing that spans locations and enables greater application development agility and portability by presenting a common platform with consistent management. But to appreciate this view of cloud computing, we should first ask how things got so complex in the first place.

How IT Got Complex, and How We Intend to Simplify It

Computer systems have continuously evolved to balance the capabilities of the technology era against the requirements of the users, whether they are programmers or end users, resulting in trade-offs that made sense at the time.

Batch processing mainframes above all optimized the use of the scarce and expensive computing resources. People were willing to be inconvenienced as long as every scarce CPU cycle was used effectively. Timesharing systems were designed to give the illusion that each user had their own computer, along with access to shared file storage, but the compromise was that the response time deteriorated as more demands were placed on the shared computers. Since computers were still relatively expensive, end users were willing to live with those constraints.

As computer components declined in price, thanks to the semiconductor revolution, distributed computing and personal computers arose. Each user or group of users benefited from control over their own machine, although they lost convenient sharing of data and also encountered relatively low CPU utilization rates, which wasted the full potential of machines. Client-server systems arose to address these limitations by providing highly interactive user interfaces on personal computers, along with centrally managed servers with shared data and processing.

But the cost was the need to manage ever increasing complexity. Now there were many more independently movable pieces. Keeping track of the interdependencies led to dramatic increases in operational cost. Attempts to fix this have resulted in even more layers and more complexity, turning into a truly Sisyphean task. We are now paying for the sins of our past. Inertia and the desire to retain compatibility have kept the IT industry on this path, since the costs of switching to a new paradigm were considered too high to offset the costs of complexity.

This very inertia often does not allow us to truly recognize the opportunity for change. But, standing back, it is now becoming apparent that a new model of computing offers hope.

The VMware View of Cloud Computing

VMware's view of cloud computing is twofold. First, we stitch together compute resources so as to appear as one large computer behind which the complexity is hidden. By coordinating, managing, and scheduling resources such as CPUs, network, storage, and firewalls in a consistent way across internal and external premises, we create a flexible cloud infrastructure platform. This platform includes security, automation and management, interoperability and openness, self-service, pooling, and dynamic resource allocation. In the view of cloud computing we are advocating, applications can run within an external provider, in internal IT premises, or in combination as a hybrid system—it matters how they are run, not where they are run.

In the past, people have used terms such as *utility*, *grid*, or *on-demand* computing to describe these approaches to computing. These are not new terms: comparisons of computing to a public utility such as a telephone system date back to the early 1960s in academia and were certainly popularized by industry in the 1990s. However, these systems did not achieve the level of popularity originally hoped for.

Some historical perspective will also help us understand why this form of computing is needed today and why cloud has become the new paradigm. Super-servers built on clusters of commodity hardware components were foreseen many years ago. What was not apparent at that time was that new software is required to exploit its capabilities. It was assumed that conventional systems software and applications could run on top of the superservers, but now it is apparent to us that conventional software does not fully exploit the servers' capabilities. Many off-the-shelf applications did not dynamically scale, and it was difficult to reconfigure them to meet new demands. Although new specialized software architectures can be developed to take advantage of these systems, they will not be compatible with existing applications. Is there a way we can address these two potentially conflicting needs?

The emergence of extremely high-performance, low-cost standard hardware, virtualization, and modern programming platforms now allow us to do this. We can apply the hardware dividend that the continuing semiconductor revolution gives us to the purpose of eliminating complexity while retaining compatibility.

VMware's vision for cloud computing addresses this. The virtualization technology at its foundation enables it to accommodate existing applications and extends them with additional system services and a programming model to form the basis of a new model of computing.

The vSphere system has the flexibility to make existing workloads run well–in many cases, better than physical systems. The first book in this series (*Deploying the VMware Infrastructure* [2008]) explains how vSphere enabled server consolidation and solutions such as novel forms of disaster recovery.

The virtualization system is a key foundation for the cloud computing system. The magic lies in its ability to encapsulate applications, along with associated middleware and operating systems, in a black box. Once we have applications encapsulated, we can then jack the boxes up, figuratively speaking, slide different system services underneath them, and even slide the black boxes around in real time to take full advantage of the underlying capabilities in a transparent way. This is what enables us to run existing applications in a more efficient, flexible way—cutting the tentacles of complexity that bind applications to the underlying infrastructure.

VMware vSphere has evolved to provide additional capabilities to form this new infrastructure layer for cloud computing. This was described in the second book in this series (*Foundation for Cloud Computing with VMware vSphere 4* [2010]).

The next step is to take system resources and aggregate them to an even larger scale, to place resources such as networking, firewalls, and storage under its control, and to add appropriate management to support a more advanced form of cloud computing. By securely delivering these resources as virtual datacenters, organizations can efficiently deliver these resources to users. This book discusses this next step in the evolution of VMware technology.

In Closing

This book describes the foundation for this form of cloud computing. It is by no means the end of the story. Although we have not yet completely achieved our vision, we have created a solid basis for cloud computing and are working on more innovations. Simplification of IT is a large and formidable problem to solve, but that is our goal. I hope you enjoy reading this book and learning how VMware is taking on this challenge.

Paul Maritz
Chief Executive Officer, VMware Inc.

The Blind Men and the Cloud

Sam Charrington

It was six men of Info Tech
To learning much inclined,
Who went to see the Cloud
(Though all of them were blind),
That each by observation
Might satisfy his mind.

The First approached the Cloud,
So sure that he was boasting,
"I know exactly what this is…
This Cloud is simply Hosting."

The Second grasped within the Cloud,
Saying, "No it's obvious to me,
This Cloud is grid computing…
Servers working together in harmony!"

The Third, in need of an answer,
Cried, "Ho! I know its source of power—
It's a utility computing solution
Which charges by the hour."

The Fourth reached out to touch it,
It was there, but it was not.
"Virtualization," said he.
"That's precisely what we've got!"

The Fifth, so sure the rest were wrong
Declared "It's sass [*sic*] you fools,
Applications with no installation
It's breaking all the rules!"

The Sixth (whose name was Benioff),
Felt the future he did know,
He made haste in boldly stating,
"This *IS* Web 3.0."

And so these men of Info Tech
Disputed loud and long,
Each in his own opinion
Exceeding stiff and strong,
Though each was partly in the right,
And all were partly wrong!

Based on "The Blind Men and the Elephant," by John Godfrey Saxe, Appistry, Inc., 2008 (original post: http://www.appistry.com/blogs/sam/the-blind-men-and-cloud)

1. Introduction to Cloud Computing

What is cloud computing? Is it an "as-a-service" enabler? Is it the next generation of virtualization? What technologies are used for the cloud? Is the cloud only about technologies or is it a new form of IT and business enablement? Does the cloud provide more alignment between business and IT? How can an organization effectively implement a cloud computing model?

This Short Topics book provides use cases, design considerations, and technology guidance to answer these questions. It is a companion to volume 21 of the USENIX Short Topics in System Administration series, *Foundation for Cloud Computing with VMware vSphere 4.* Since the first book was published, new VMware vSphere™ features have become available that impact previous recommendations about supporting the cloud layer. These updates are covered in Chapter 4, "Foundation for Cloud Computing."

The intended audience is those interested in learning about VMware cloud computing products and solutions. Content on third-party technologies is also included where appropriate. The information provided will help current VMware users deploy and utilize cloud computing platforms.

By working closely with customers and partners, we've gained keen insight into how VMware technology maps to the cloud. VMware vCloud™ Director is VMware's first offering in the cloud management space. Additional VMware technologies are rapidly evolving to support the cloud vision. As companies move to cloud computing, some of the same challenges encountered during the adoption of virtualization will resurface. Advances in technology bring additional complexity and difficulty, as architectures, processes, and skills must evolve accordingly. This book offers solutions to some of those challenges.

Cloud Computing

Cloud computing concepts may seem elusive at first. The term itself has different connotations for different people—depending on the perspective, some or all attributes of cloud computing could satisfy an organization's business requirements.

Focusing on any single aspect of the cloud does not reveal the true nature of cloud computing. To quote Sam Charrington, "When we try to define the cloud based on some subset of the technologies used to implement it, we risk missing the forest for the trees." Asking ten people to define "cloud computing" may yield ten different answers, because each perspective is different and the term is continually evolving. However, various research efforts have led to convergence of taxonomy, leading to the broader acceptance of key attributes.

Many motivating factors have led to the emergence of cloud computing. Businesses require services that include both infrastructure and application workload requests, while meeting defined service levels for capacity, resource tiering, and availability. IT delivery often necessitates costs and efficiencies that create a perception of IT as a hindrance, not a strategic partner. Issues include underutilized resources, over-provisioning or under-provisioning of resources, lengthy deployment times, and lack of cost visibility. Virtualization is the first step towards addressing some of these challenges by enabling improved utilization through server consolidation, workload mobility through hardware independence, and efficient management of hardware resources.

Cloud computing builds on virtualization to create a service-oriented computing model. This is done through the addition of resource abstractions and controls to create dynamic pools of resources that can be consumed through the network. Benefits include economies of scale, elastic resources, self-service provisioning, and cost transparency. Consumption of cloud resources is enforced through resource metering and pricing models that shape user behavior. Consumers benefit through leveraging allocation models such as pay-as-you-go to gain greater cost efficiency, lower barrier to entry, and immediate access to infrastructure resources.

The technologies covered in this book include VMware vSphere, VMware vCloud Director, the VMware vShield™ product family, and VMware vCenter Chargeback™.

Topics Covered in This Book

Chapters in this book cover the following topics:

- ❖ Chapter 2, "What Is Cloud Computing?" defines cloud computing and explains the private, public, community, and hybrid cloud computing models. The service models supported by VMware vCloud are also covered, including Infrastructure as a Service, Platform as a Service, Software as a Service, and IT as a Service.
- ❖ Chapter 3, "The Benefits of Cloud Computing," dives into the economic and other advantages of cloud computing, explaining the differences between private cloud and public cloud computing models. Guidelines for building a compelling business case for cloud computing include a list of the common variables used to generate an effective ROI (return on investment) analysis.
- ❖ Chapter 4, "Foundation for Cloud Computing," presents an update on how VMware vSphere provides the foundation for cloud computing and Service-Oriented Architecture.
- ❖ Chapter 5, "VMware vCloud and VMware vCloud Director," introduces vCloud Director as one of the key components for vCloud, which also includes VMware vSphere, VMware vShield, and VMware vCenter Chargeback. VMware vCloud Director makes broad deployment of compute clouds possible by enabling self-service access to compute infrastructure through the abstraction and orchestration of virtualized resources.

- Chapter 6, “VMware vCloud Director Virtual Datacenters,” introduces the virtual datacenter (VDC), along with various characteristics of allocation models, and describes how each of the allocation models affects the vSphere layer.
- Chapter 7, “VMware vCloud Networking,” examines vCloud network pools, network layers, organization networks, vApp networks, vShield components, and vCloud network use cases.
- Chapter 8, “VMware vCloud Storage,” discusses the increased importance that cloud computing places on storage and covers the unique challenges for heightened availability, security, compliance and regulation. It presents a modular tiered storage approach for designing an optimal cloud storage layer. A guiding set of storage design principles assists with configuration of each storage pool tier in the absence of specific customer application requirements.
- Chapter 9, “VMware vCloud Director Logging and Monitoring,” covers the methods used by VMware vCloud Director for logging and monitoring application deployment.
- Chapter 10, “VMware vCloud API,” introduces the vCloud API to administrators and users, including definition, features and benefits, and design considerations.
- Chapter 11, “vCenter Chargeback,” provides information about vCenter Chargeback architecture, integration with vCloud Director, cost models, billing policies, and design considerations.
- Chapter 12, “Applications in the Cloud,” discusses applications in the cloud and the underlying OVF (Open Virtualization Format) standard to allow portability between platforms. It covers OVF, vApps, licensing considerations, the VMware vFabric Cloud Application Platform, VMware ThinApp™, and migrations to and from the cloud.
- Chapter 13, “Scalability,” covers the requirements for scalability and performance of the vCloud Director environment.
- Chapter 14, “vCloud Security,” discusses the vCloud Director security model, including securing applications, the perimeter, user access, and the datacenter. VMware vSphere and vCloud security functions are discussed in detail, along with vShield.
- Chapter 15, “Business Resiliency,” covers business resiliency in both the cloud and the virtualization layers, including redundancy, the management cluster, resource groups, and vApp backup and recovery.
- The Appendix discusses some of the vCloud Director ecosystem contributions to integrated computing stacks and orchestration tools.
- The Glossary provides brief definitions of terms and acronyms used by VMware that are applicable to vCloud.
- The References include publications and documentation, as well as online communities and recommended books.

2. What Is Cloud Computing?

Cloud computing is a style of computing that enables on-demand network access to a shared pool of scalable and elastic infrastructure resources. The term *cloud computing* originates from the standard network diagram where a cloud is used to represent the abstraction of a complex networked system such as the Internet. The concept of delivering computing resources through the network has evolved as a result of the success of cloud-based applications, widespread availability of broadband Internet access, and mainstream adoption of server virtualization technology. Figure 1 depicts a consumer using assets over a network without any knowledge of its location or how it is resourced.

User
Multi-tenant Cloud

Figure 1. Cloud Computing

Characteristics of Cloud Computing

The following definitions are provided by the National Institute of Standards and Technology (NIST).[1] Depending on the business requirements, some cloud computing characteristics may be more pertinent than others. For example, enterprises may opt to use the metered service to provide showback, rather than chargeback to internal consumers.

- *On-demand self-service*—Consumers can unilaterally provision their own computing capabilities, such as server time and network storage, automatically as needed without requiring human interaction with each service's provider.
- *Broad network access*—Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms such as mobile phones, laptops, and PDAs.
- *Resource pooling*—The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources, but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processors, memory, network bandwidth, and virtual machines.

[1] The full document is located on the NIST Cloud Computing site at http://csrc.nist.gov/groups/SNS/cloud-computing/. The definitions used are based on Version 15 of the document.

- *Rapid elasticity*—Capabilities can be rapidly and elastically provisioned, in some cases automatically, from quick scale-out and rapid release to quick scale-in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- *Measured service*—Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and the consumer of the utilized service. *Measured service* includes pay-for-use controls that apply for chargeback or showback requirements.

To adopt cloud computing, organizations should look to cloud providers for SLAs that meet their requirements—for example, security to protect private data, integration with existing infrastructure, near-term and long-term cost savings, availability of resources, and minimized risk through lock-in avoidance.

Efficiency through Utilization and Automation

Compute (CPU and memory), storage, and network resources are used within a cloud. Pooling of these resources supports higher resource utilization while achieving greater efficiency and scalability. This is adaptable for public or private clouds by providing policy-driven provisioning, deployment, and management.

Policy-driven automation of provisioning, deployment, and management leads to an agile infrastructure. This automation may require customized scripting in some cases, depending on the underlying technology and the policies that need to be implemented.

Agility with Control

Cloud computing includes self-service models for policy-based provisioning and deployment, allowing centralized IT organizations to provide a measure of control to the businesses and users they support. The infrastructure is application-aware, with built-in availability, scalability, security, and performance guarantees.

Freedom of Choice

Freedom of choice is supported by a cloud based on open standards that support mobility between different clouds. This ensures compatibility and portability of existing applications.

Types of Clouds

Deployment models for cloud computing are based on the location of where the cloud infrastructure is hosted. The most common cloud types are:

- Private
- Public

- Community
- Hybrid

Private

A private cloud is dedicated to one organization and may be on-premise or off-premise. Off-premise locations may include other corporate locations or a cloud service provider. It is a cloud computing architecture that is hosted and controlled by the enterprise behind a firewall. A private cloud may also be called an *internal cloud*, an *enterprise private cloud*, or a *corporate cloud*. It is designed to meet the internal needs of tenants (customers) within a corporation.

Resource metering is typically implemented in private cloud deployments to help determine budgets, cross-charges, and capacity planning. Whether used as chargeback or showback (showing what charges would be, but the consumer is not actually charged), the financial aspects of running a cloud must be considered for capacity planning and budgeting. If no chargeback/showback model is present, shaping user behavior to fit a cloud computing model becomes challenging. Figure 2 shows users interacting within an on-premise cloud. All external access is controlled by the firewall.

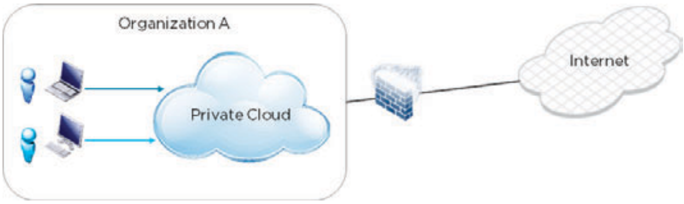

Figure 2. Private Cloud Deployment

Public

A public cloud is a cloud typically owned and managed by a cloud service provider. It is a cloud computing architecture that supports multiple tenants and may be operated by a third-party service provider. A public cloud may also be called an *external cloud* and offers access to a wide variety of Internet users (see Figure 3, next page). Most public clouds are located at the service provider's premises.

A public cloud should not be confused with an off-premise private cloud. A public cloud provider is a service provider that hosts public clouds, but may also host off-premise private clouds.

Public cloud services are provisioned dynamically using a self-service Web portal for access and control. Charges are usually tied to consumer resource usage. Services offered may range from the ability to provision new operating systems from ISO images, to packaged virtual appliances and multi-tiered applications.

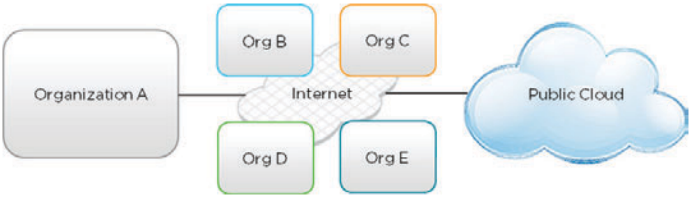

Figure 3. Public Cloud Deployment

Community

A community cloud may be shared by different organizations with a direct relationship to each other and some overlapping requirements. Community clouds include organizations with similar interests. Typically, a community cloud is also a hybrid cloud. There may be sharing of specific resources and applications, but the ownership and management are split among the members. A community cloud is often found in government organizations where different areas of the government (county, state/province, and city) need to work together.

Hybrid

A hybrid cloud incorporates a combination of clouds and may include both on-premise and off-premise resources. For example, access to a cloud handling a required Software as a Service (SaaS) application may include VPN or other special access controls. There might be multiple clouds deployed for a company in different geographic regions. Applications may exist on-premise, off-premise, or a combination of both. Enterprises with an existing private cloud may choose to provide and manage public cloud resources in a secure and scalable way. Connectivity between different clouds enabling data and application portability characterizes a hybrid cloud solution.

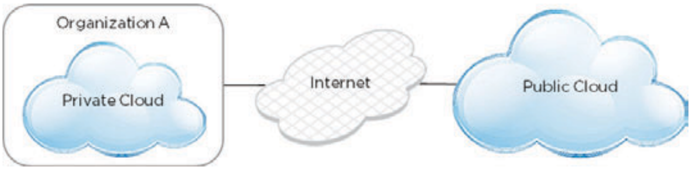

Figure 4. Hybrid Cloud Deployment

Cloud Layers and Service Models

IT-as-a-Service

IT as a Service (ITaaS) is the culmination of the various cloud computing models. It is the transformation of IT to a business-centric approach focused on outcomes such as operational efficiency, competitiveness, and rapid response. IT shifts from being a cost center to being a driver of strategic value as both production and consumption of services are optimized in a manner consistent with business demand. ITaaS provides users with seamless, secure access from any device anywhere, not only to logical pools of internal resources, but also to SaaS and public cloud providers.

Service Models (IaaS, PaaS, and SaaS)

As a whole, IT as a Service (ITaaS) is the overarching definition of cloud computing. Within ITaaS, vCloud supports Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) (see Figure 5, which shows the VMware strategy for these services).

- ❖ Infrastructure as a Service (IaaS) is a model where service providers offer infrastructure resources through the Internet or other data mediums such as MPLS to organizations on demand. Service providers own and maintain all hardware, while consumers pay on a per-usage basis, similar to utility computing. The self-provisioning aspects, combined with dynamic scaling and elasticity, provide users with the illusion of unlimited resources if they are willing to pay for them. Services and policies are governed through SLAs between the provider and consumer. Examples include Amazon EC2, BlueLock, Colt, Rackspace, and Terremark.
- ❖ Platform as a Service (PaaS) accelerates the development of applications by providing all of the facilities necessary to develop, test, deploy, and host application services. The entire development and deployment platform is hosted in the cloud and accessed through the Web browser. Examples include VMforce, Force.com, Savvis, Google AppEngine, and Microsoft Azure.
- ❖ Software as a Service (SaaS) refers to applications that are deployed over the Internet as a service. Paid services are typically subscription-based, with costs applied on a per-user or per-usage basis. Examples include Zimbra, Mozy, Dropbox, Boxine, Salesforce.com, Gmail, and Google Docs.

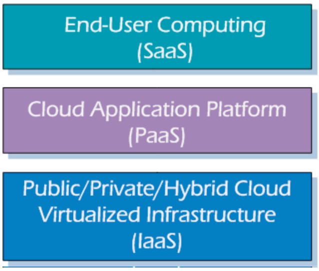
End-User Computing (SaaS)
Cloud Application Platform (PaaS)
Public/Private/Hybrid Cloud Virtualized Infrastructure (IaaS)

Figure 5. Cloud Computing Layers

Use Cases for Service Models and Workloads

Use cases for cloud computing fit the *X*-as-a-Service models. The following use cases show the linkage between these models and their use by service providers and enterprise customers.

- ❖ IaaS is the underpinning resource model for private clouds as well as for public cloud providers, offering a complete infrastructure hosting service.
- ❖ PaaS is commonly utilized by public cloud providers offering services to software developers.
- ❖ SaaS is an increasingly popular model for ISVs to offer their products to organizations on a hosted basis.

The use cases for workloads fall into the following categories (workloads may fit into multiple categories):

- *Transient workloads*—Transient workloads consist of vApps that exist for only short periods of time. Examples include demonstration environments, technical training, help desk, pre-production environments, and test/dev labs.
- *Elastic workloads*—Elastic workloads consist of vApps that fluctuate dramatically in their resource requirements. Examples include seasonal workloads such as e-commerce and tax-related applications. This category also incorporates scientific workloads that may include high-performance compute applications, gene mapping, and mathematical modeling.
- *Infrastructure workloads*—Infrastructure workloads include vApps that are considered essential for operations of a datacenter. These are deployed long-term and take advantage of vCloud capabilities. Examples include file and print servers, domain controllers, directory servers, NTP servers, and other infrastructure workloads.
- *Other workload types*—This category applies to other types of workloads which do not necessarily fit any of the previously discussed categories. Typically these include non-mission-critical IT applications that are often neglected and could benefit dramatically from self-service. Examples may include legacy servers placed in a vApp network.

Note: A *vApp* is a container for compute workload in the cloud and is the standard unit of deployment for vCloud Director. A vApp can contain one or more VMs and typically represents an encapsulated software solution.

Use Case Areas

Each of the workload types applies to additional use case areas. Examples include production environments, Web applications, marketing/brochure sites, multi-tiered Web applications, e-commerce Web sites, corporate portals and intranet sites, messaging and collaboration applications, Microsoft SharePoint, content and document management, internal wikis and blogs, pre-production environments, dev/test/stage, application evaluation, and capacity augmentation for existing environments.

Summary

Cloud computing use cases consist of transient, elastic, infrastructure, and other workload types. These workloads fit into multiple service models, including IaaS, PaaS, and SaaS. Requirements determine the use cases, service models, and workload categories for a cloud deployment.

3. The Benefits of Cloud Computing

The USENIX Short Topics book *Foundation for Cloud Computing with VMware vSphere 4* describes many benefits of vSphere, from enabling operating system consolidation onto fewer hardware platforms to defining characteristics inherent in virtual machines such as compatibility, isolation, encapsulation, and hardware independence. Cloud computing leverages the cost reductions, efficiencies, and benefits of virtualization and extends them as a transformative platform for all of an organization's computing, whether internal on a private cloud, external on a public cloud, or using a combination of both. It enables the IT organization to focus on the applications relevant to its business rather than purchasing, managing, and upgrading ever-changing infrastructure.

Cloud computing is transformative technology which enables IT as a Service. Both production and consumption of services are optimized in a manner consistent with business demand. Time-to-market, business process improvement, greater cost efficiencies, agility, scalability, security, and privacy transform IT from a cost center to a driver of strategic value.

Cloud Computing as a Transformative Platform

Acquisition of IT infrastructure for a physical datacenter tends to be project-driven. New applications and other departmental endeavors provide budget for servers, SANs, and software, filling up datacenters with a hodgepodge of technology islands. Disparate administrative processes and lack of interoperability make ongoing management difficult. Larger datacenter staffs tend to divide into functional silos of OS administration, servers, storage, network, and development, straining effective collaboration. As a consequence, traditional datacenters are beset with both high inefficiencies and high costs.

Unfortunately, virtualization is frequently deployed as yet another technology island running in parallel with the traditional physical infrastructure. Because the huge costs and inefficiencies of existing physical infrastructure continue to consume the largest share of financial and staffing resources, IT administrators inescapably view their world through a physical filter, relegating virtual machines to the status of tertiary infrastructure. Physical servers still need upgrading, rack space, switch ports, UPS slices, cabling, power, and cooling. And tasks associated with testing, adding hardware, remote access, performance monitoring, troubleshooting, patching, and capacity planning require far more time than in a virtualized datacenter.

Though virtual machines clearly reduce some costs and staffing requirements, a hybrid physical/virtual environment can lead to an overall increase in staffing demands and complexity. A mixed environment has many more objects requiring management. Even simple bottlenecks in the virtual environment commonly force IT to re-budget multiple times for additional licensing, memory, ESX hosts, or storage funds. This reactionary approach to virtualization ensures that any expansion of the environment will be slow and painful—assuming, of course, that it doesn't stall altogether.

Cloud computing forces a different approach, in which virtualization becomes the standard while physical servers become the exceptions. Rather than purchasing infrastructure in isolated chunks for specific departmental projects or applications, it is uniformly and automatically provisioned as required. Departmental purchasing budgets are replaced with a cloud costing model that combines utility-based costing with activity-based costing, enabling reduced expenses while increasing the business value of IT.

Strategic Value of IT

As important as reducing costs may be, the primary value of cloud computing for many organizations is the competitive advantage often enabled through enhancements in time-to-market, scalability, and agility. For example, a department that wants to implement a new application which requires multiple servers along with storage, networking, load balancing, and specific security attributes no longer needs to go through an arduous design, procurement, and implementation process—potentially taking several months or longer. Instead, the department can quickly self-provision the required infrastructure, which is deployed virtually and charged to the department based on a cloud costing model that may incorporate both resources actually consumed and specifically required engineering resources. While current cloud computing technology may not provide the level of automation and policy controls needed to achieve optimal application performance, future product refinements will quickly close this gap.

Cloud computing capabilities that enhance the strategic value of IT include cloud costing, rapid scalability, automation, self-service catalog, enhanced management, flexibility, availability, datacenter "greening," security, regulatory compliance, and reduced risk.

Cloud Costing

A cloud, whether private or public, provides IT as a Service, enabling business units and users to purchase only the resource (infrastructure or application) elements required for their specific tasks. VMware vCenter Chargeback provides the essential mechanism to positively influence behavior by giving users the transparency they need to make efficient consumption decisions. Chargeback enables both utility- and activity-based costing.

Rapid Scalability

Resources are shared among multiple tenants: departments that utilize a common private cloud or companies that access a public cloud. Rather than having to go through a lengthy and painful process of designing, specifying, budgeting, bidding, procuring, and configuring infrastructure for a specific project, the desired infrastructure can be quickly or instantly provisioned as necessary. The requirement for over-building to ensure adequate capacity is eliminated.

Automation

VMware vCloud Director enables automation of organization-specific tasks and processes that require significant time to create manually. The published APIs and SDKs enable further customization. VMware vCenter Orchestrator and vCloud Request Manager provide the ability to script automatic responses and actions within the cloud.

Self-Service Catalog

The self-service catalog provides business units and users with the ability to quickly provision the virtual infrastructure they need to support their applications and business requirements, enabling both higher productivity and faster time-to-market. Standardization of service catalog items reduces deployment times by reducing installation and configuration time.

Enhanced Management

The disparate technologies and equipment of a traditional datacenter inevitably results in differing management tools and processes. Using virtualization as a unifying platform and incorporating vCloud Director as a management framework enables much more efficient operation of the datacenter with significantly less effort. The IT staff is able to focus on technologies and tasks that impact the business instead of on the infrastructure.

Improved Agility

Self-service provisioning of applications and user accounts increases user productivity and quickens time-to-market, while effective use of vCenter Chargeback enables maintaining the optimal balance between user capabilities and organizational expenditures.

Public clouds offer an additional benefit in agility because organizations do not need to be concerned with finding or expanding high-quality datacenter space. However, some models of public cloud may have limits that affect their ability to adapt to an organization's specific requirements in areas ranging from regulatory compliance to supported operating systems.

Technologies such as VMware vCloud Connector enable mobility between clouds as business requirements change.

Flexibility

Cloud computing incorporates the agility benefits of virtualization and makes it available in a myriad of different ways to fit an organization's budget and culture.

Availability

The High Availability (VMware HA) and Dynamic Resource Scheduling (VMware DRS) capabilities of vSphere are extended to the cloud infrastructure. This enables superior SLAs and potentially reduced costs.

Datacenter “Greening”

Green initiatives can be realized with cloud computing. Datacenters are huge consumers of electricity, using around 1.5% of all the power produced in the United States. The multi-tenancy aspect of cloud computing, whether departmental or organizational, extends the already significant virtualization energy savings across a wider spectrum of users. Capabilities such as VMware Distributed Power Management (DPM) further automate energy efficiency. A public cloud–based disaster recovery service offers a particularly attractive solution from an energy saving perspective because disaster recovery facilities and equipment tend to remain idle nearly all the time. This type of service can potentially be safely oversubscribed to many geographically disparate organizations, thereby minimizing the facilities, equipment, energy, and staffing required.

Security

VMware security and monitoring tools enable much more granular security controls than are available in the physical world. They provide very granular security specific to a virtual machine with dedicated physical firewalls and other devices. A public cloud, though offering security capabilities beyond what is typically available in a physical datacenter, generally does not have the same degree of control as a private cloud. But certain specialty public cloud providers may be able to offer industry or application-specific security that would be difficult for an organization to replicate internally.

Use of VMware vShield and its components provides greater flexibility in securing cloud workloads, providing vApp isolation capabilities while adhering to compliance guidelines.

Reduced Risk

Risk of failure to comply with regulations such as SOX, HIPAA, PCI, and others is reduced because of the ability to significantly increase both control and monitoring of the compute environment within a private cloud.

Server failure was shown by a 2009 Webtorials IT survey to be the number one cause of outage, followed by human error. Clouds built on vSphere can leverage VMware HA (including VM and Application Monitoring) to significantly reduce the risk of an outage. Incorporating cloud-based disaster recovery methodologies provides superior recoverability in the event of a datacenter catastrophe.

A public cloud, by shifting capital expenses (CAPEX) to operational expenses (OPEX), reduces the risk of under- or over-building the virtual infrastructure or of acquiring equipment that becomes obsolete earlier than projected. Certain operational and downtime risks might also be mitigated, depending on the SLAs offered and on the reliability of the provider.

Expense Reduction

Reductions in CAPEX and OPEX enabled by virtualization are discussed in *Foundation for Cloud Computing with VMware vSphere 4*. Cloud computing both enables and expands these cost reductions, and it adds further efficiencies specific to either private or public cloud models.

Capital Expense Reduction

Server virtualization lowers infrastructure costs, but a lack of strategic planning can result in infrastructure choices that are not optimal for an enterprise virtualization implementation. The hosting platform, networking bandwidth, and storage architecture are all areas where organizations sometimes make choices to support a small virtual environment rather than a datacenter platform.

One of the advantages of a private cloud is that it forces organizations to adopt a strategic outlook when planning a virtualized datacenter on an enterprise scale. Not only are the equipment purchases reduced significantly from what would be required in a physical datacenter, but they are optimized to enable the resiliency, reliability, and scalability demanded by an effective private cloud offering.

Instead of purchasing lightly utilized resources with individual departmental budgets, a private cloud pools all resources and shares them as a service among all departments. Peak loads are accommodated without the necessity of purchasing large amounts of additional equipment for each department.

Unlike a private cloud, utilizing a public cloud eliminates CAPEX entirely and instead shifts the budget to OPEX, enabling organizations to pay for resources only as needed. This may pertain to the infrastructure related to a specific application accessed through the cloud, or to part or all of the entire datacenter or disaster recovery facility. For example, an organization might have a private cloud as its primary computing platform but seamlessly integrate into a public cloud in order to utilize additional capacity for infrequent periods of peak demand instead of purchasing the infrastructure internally.

Operational Expense Reduction

Virtual machines require significantly less management than their physical counterparts, but these staff savings can be negated by an overall increase in complexity when organizations attempt to manage both physical and virtual infrastructures, primarily utilizing the equipment, tools, and processes of the physical environment. This results, in part, because there are more objects to manage, including virtual machines, virtualization hosts, vSwitches, and vNICs. The resources are typically limited by the continuing need to contend with physical infrastructure.

Adopting a private cloud means making a commitment to virtualization technology as your datacenter standard. The virtualization platform is the standard and the remaining physical servers are exceptions. The private cloud incorporates vCloud Director to efficiently automate and manage virtual infrastructure–specific issues such as efficient provisioning of virtual infrastructure and preventing VM sprawl, and it facilitates more effective coordination among functional IT silos, as well as federation to public clouds.

A public cloud provider benefits from the same OPEX efficiencies and, as with the CAPEX savings, should be able to leverage them among many customers and potentially pass them on. This model can provide more readily available staffing resources than individual organizations can obtain by hiring or training their own datacenter staff. Public providers may have specialists in areas such as power, security, management, and so on.

Additionally, a specialized provider such as a SaaS provider may be able to leverage application-specific resources and knowledge more cost-effectively among its large client base.

Developing a Compelling Business Case for Cloud Computing

Foundation for Cloud Computing with VMware vSphere 4 provides the details behind the compelling return on investment (ROI) from implementing a VMware virtual infrastructure. The primary areas of savings result from virtualizing production datacenters, backup and disaster recovery, and desktop systems.

A private cloud typically provides an even more compelling ROI because, unlike a standard virtualization deployment, a cloud strategy demands an encompassing approach to the entire datacenter. Though an ROI approach is very practical for evaluating a private cloud versus a physical or hybrid physical and virtual datacenter, a public cloud is typically measured with financial metrics more similar to those for evaluating an outsourcing arrangement.

Measurable ROI is just one facet of building a compelling business case for cloud computing. The ability to better meet overall corporate objectives such as time-to-market, sudden growth requirements, high availability, disaster recovery, employee empowerment, and green initiatives may be as important as, or more important than, financial metrics. The effect on those objectives should be evaluated and presented to senior executives in conjunction with the economics.

ROI for a Private Cloud

An in-depth ROI analysis helps secure the funds required for a private cloud while also assisting the virtualization champions to successfully navigate the organizational resistance and politics that accompany significant change. It is generally worthwhile to engage in an ROI analysis even when the economic advantages are obvious because an analysis can reveal both costs and opportunities that might otherwise be overlooked. It also provides a baseline against which the IT organization can measure the financial aspects of its private cloud success, helping to build credibility.

An ROI analysis showing discounted cash flows on a yearly basis can convince senior management to invest in the hardware, software, and services necessary for a successful private cloud. Financial people are familiar with this format, and it allows them to easily compare the expected return from a private cloud initiative with other opportunities.

Figure 6 shows an example of an ROI analysis that resulted from implementing a private cloud. The left column shows the yearly cash flows, representing the delta between the costs of maintaining a physical datacenter and a private cloud. The right column discounts the cash flows back to the current year using the organization's internal rate for cost of capital. In this example, the organization projects an investment of $2,447,390 to implement a private cloud versus cash flow savings of $2,467,373 the first year. The investment, therefore, is paid back using the first year's cash flow savings, resulting in a payback period of 11.9 months.

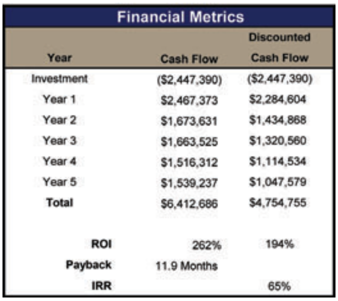

Financial Metrics

Year	Cash Flow	Discounted Cash Flow
Investment	($2,447,390)	($2,447,390)
Year 1	$2,467,373	$2,284,604
Year 2	$1,673,631	$1,434,868
Year 3	$1,663,525	$1,320,560
Year 4	$1,516,312	$1,114,534
Year 5	$1,539,237	$1,047,579
Total	$6,412,686	$4,754,755
ROI	262%	194%
Payback	11.9 Months	
IRR		65%

Figure 6. Yearly Discounted Cash Flow Analysis for a Private Cloud

The organization in this example has an internal cost of capital of 8%, meaning that the cash flows each year are discounted back to their present value using the 8% discount rate. For example, the $2,467,373 in the first year cash flow has a present value of $2,284,604. For the entire five-year period, the total discounted cash flows come to $4,754,755, resulting in a discounted return on investment of 194%.

The internal rate of return (IRR) is the discount rate that makes the net present value of all cash flows from the private cloud investment equal to zero. A high IRR indicates an attractive yield from the private cloud investments. In this case, the organization would have to earn an annual return of 65% for five years on its investment of $2,447,390 to deliver the same expected financial results as the private cloud.

Every organization has different cost structures and variables to consider when calculating estimated cash flows, but there are datacenter variables common to most organizations. For purposes of an ROI analysis, they can be categorized simply by cash flow and investment. It is often a good idea to break down the variables according to CAPEX (the costs of building the private cloud) and OPEX (the costs of operating the private cloud) to make it easier for finance people to quickly evaluate their relevance.

Cash Flow Variables

- *Server costs*—Server costs are calculated both for expected new server purchases each year and for expected upgrades to existing servers. Server costs should reflect not only the cost of the hardware but also the relevant associated costs, including tax, shipping, cabling, power whips, host bus adapters, Ethernet adapters, prorated share of UPS and generator cost, and set-up cost to rack, stack, and configure the servers. The same calculations are used for the reduced number of ESX hosts to determine the cash flow savings from a private cloud.
- *Server maintenance costs*—Annual maintenance costs per server/ESX host.
- *Network costs*—The prorated cost per server for core and distributed Ethernet switch ports. A private cloud tends to significantly reduce these costs due to server consolidation.
- *Storage costs*—The anticipated increased storage required by year to support either physical infrastructure storage or the storage associated with newly added virtual machines. Both scenarios should include the cost for SAN switches and any required software, including backup.

- *Power and cooling costs*—The estimated (or actual, if utilizing a collocation facility) cost to power and cool each server and prorated associated equipment.
- Datacenter buildout costs—Some organizations will incur costs for enabling their datacenters to accommodate an increased number of servers or power and cooling requirements. These may necessitate major construction or simply the purchase of a new air conditioning unit, UPS, or PDU.
- *Disaster recovery costs*—Both physical datacenters and private clouds will have costs associated with disaster recovery that frequently mirror the production datacenter but may also include specialty replication products such as VMware vCenter Site Recovery Manager, additional facility costs, bandwidth, and so on.
- *Microsoft Windows Server*—The cost spent on Standard, Enterprise, and Datacenter Editions of Windows Server licensing and the expected share of new servers that require licenses. Virtual instances of Windows Server Standard can be licensed, as in the physical world, on a per-instance basis. Alternatively, each license of Windows Server Enterprise on a virtualization host allows four instances of Windows Server. But Windows Server Datacenter Edition, when licensed by the underlying physical CPUs of the virtualization host, allows unlimited instances of any type of Windows Server guests.
- *Microsoft SQL Server*—The cost spent on Standard, Enterprise, and Datacenter Editions of SQL Server licensing and the expected share of new servers that will require licenses. As with Windows Server Enterprise, the Enterprise license allows for instances of SQL Server (either Standard or Enterprise), and the Datacenter Edition allows unlimited instances. With SQL Server Enterprise or Datacenter Editions, running instances can be migrated as needed across servers within a server farm as long as the number of CPUs on the target host does not exceed the number of CPU licenses.
- *IT staffing*—The average fully loaded hourly IT staff cost per server. It is generally accepted that an administrator can manage between three and six times as many virtual machines as physical, but this does not take into consideration the potential inefficiencies of running a partially virtualized datacenter. Both the breadth and automated tasks inherent in a private cloud can result in an administrator managing 10 or more times the number of virtual machines than physical servers.

Investment Variables

- ESX hosts
- Shared storage or infrastructure stacks
- Possible enhancements to the network infrastructure
- Licensing for VMware vSphere, vCenter Server, vCloud Director, vShield, vCenter Chargeback
- Licensing for other software products specifically acquired for a private cloud
- Professional services for private cloud design and implementation
- IT staff training

Financial Metrics for Evaluating a Public Cloud

A prudent approach is to prepare an ROI analysis for a private cloud and then compare the monthly net outgoing cash flows with the anticipated cash flows from utilizing a public cloud. This provides a sense of the financial implications of each alternative.

Of course, electing to utilize a public cloud to provide IT as a Service is vastly more involved than simply looking at the monthly cost. Security, regulatory compliance, agility, and corporate culture are all important considerations, among many other variables. The quality and capabilities of the specific public cloud provider also must be considered, ranging from proven ability to meet specific SLAs to long-term financial viability.

One of the advantages of a vCloud Director approach to cloud computing is the ability to seamlessly integrate both private and public cloud environments to optimize organizational objectives.

Summary

Cloud computing expands the benefits of virtualization and offers new levels of agility, flexibility, and efficiency. Presenting a discounted cash flow–based ROI analysis to senior executives along with an evaluation of how cloud computing enables corporate objectives can help build a compelling business case.

4. Foundation for Cloud Computing

VMware vSphere is a key enabler for cloud computing and Service-Oriented Architecture (SOA). Details about the VMware vSphere technology can be found in *Foundation for Cloud Computing with VMware vSphere 4*.

Virtualization is recognized as the technical foundation for building a cloud offering and provides an evolutionary approach to transitioning to cloud computing. The first step toward the private cloud for many customers will be to virtualize and consolidate existing physical workloads.

To support the cloud model, an additional layer of abstraction is introduced which sits above the physical and virtual layers, as shown below. The cloud layer is a pure abstraction layer that obscures virtualization features and hides physical hardware. Resource manipulation, rather than being handled by vSphere, is performed by vCloud Director at the cloud layer.

Cloud Layer
Virtual Layer
Physical Layer

Physical Layer

The following summarizes the major physical aspects of a typical cloud deployment, including design considerations and some mechanisms to mitigate risks inherent in a multi-tenant environment.

Compute

The physical compute layer consists of the server hardware and virtualization layer (vSphere), which provide virtual compute resources to the cloud layer. Compute resources are required for running vCloud workloads, as well as for supporting the management and infrastructure virtual machines.

Network

The physical network stack integrates with vSphere to provide connectivity through port groups and separation of traffic through VLAN trunking. VMware vSphere can take advantage of 10 GbE to simplify virtual switch configurations and increase performance of workloads, vMotion, and Storage vMotion.

Storage

The physical storage layer includes the storage arrays, storage switches, and host bus adapters. The disk layout of the storage array is important to support the tiered storage model beneficial to most cloud architectures. As storage is the root of most performance issues in virtualized environments, the storage environment should be designed to allow balanced, predictable performance.

Virtual Layer

The vSphere platform is the foundation for the VMware vCloud solution. Virtualization decouples the management of devices from the underlying physical resources, enabling many of the properties that make up cloud computing.

Compute

VMware ESXi hosts are the fundamental building blocks for vCloud. ESXi hosts are aggregated into clusters of highly available pools of compute resources.

VMware vMotion and VMware DRS balance cloud workloads dynamically across a pool of compute resources. Storage vMotion extends vMotion to allow for live migration of virtual disks. In the future, long-distance vMotion may allow for live workload migration between clouds, enabling the hybrid cloud vision.

Network

Network I/O Control (NIOC) is a vSphere 4 feature used to prioritize traffic at the vNetwork Distributed Switch. This is especially useful with 10GbE connections, providing more flexibility for shaping traffic types.

The VMware vShield product suite provides fine-grained network security for virtualized environments. These include vNIC-level application firewalls with vShield App, perimeter network security with vShield Edge, antivirus offloading with vShield Endpoint, and centralized security management through vShield Manager.

Storage

The storage layer can take advantage of NIOC to prioritize IP storage. Additionally, Storage I/O Control (SIOC) ensures that individual virtual machines do not unfairly monopolize storage resources by using a share-based resource allocation mechanism when there is contention. Contention is specified in units of milliseconds of latency for I/O.

Cloud Layer

The cloud layer is a pure virtual layer in that the resources are presented without a visible relationship to the virtual or physical layers. This abstraction aligns with the concepts of cloud computing, where the user can access resources without any knowledge of the underlying virtual and physical resource orientation.

Setting the Virtual Layer to Support the Cloud Layer

Traditional vSphere design practices apply to support the cloud layer, with a few exceptions. Some design patterns for vSphere involve creating multiple resource pools based on workload requirements, priorities, and SLAs. With the vCloud technology it is a best practice to map the provider VDC (virtual datacenter) to the cluster resource pool. Although green-field deployments for vCloud work well, there is nothing that prevents overlaying cloud infrastructure on top of existing vSphere environments. Workloads can be transitioned to the cloud and, as resources become available at the vSphere layer, those resources can be allocated to the cloud layer.

5. VMware vCloud and VMware vCloud Director

Virtualization provides many of the qualities needed to deliver cloud services. Inherent benefits of virtualization, such as improved hardware efficiency and application mobility, have led to its broad adoption throughout companies worldwide. The cloud computing market is expected to grow rapidly as technology, processes, and comprehension advance to the next stage. VMware is uniquely positioned to facilitate entry into cloud computing by leveraging core platforms to create a comprehensive cloud stack.

VMware vCloud

VMware vCloud is a common set of cloud computing services for enterprises and service providers. The core components are the ecosystem of VMware vClouds delivered by industry leading service providers and a broad set of applications relying on the foundation of VMware technologies such as vSphere and the vCloud API.

VMware's vCloud initiative was created to enable customers to work closely with VMware cloud partners, who provide reliable, enterprise-ready cloud services without vendor lock-in and proprietary tools, formats, and infrastructure. VMware vCloud has an open architecture that offers choice and flexibility for running applications in public and private clouds.

vCloud is important to enterprise customers as it enables them to leverage service providers, software vendors, and advanced VMware technology to build private clouds or flex capacity off-premise, as needed. For small and medium businesses, vCloud delivers peace of mind in knowing that the services they get from hosting/service providers for disaster recovery, test and development, or just simple infrastructure on demand will be reliable and flexible.

VMware vCloud Director

Cloud computing requires a high degree of coordination, automation, and management. Provisioning control is transferred to the end user to allow for access and deployment of compute resources. Secure multi-tenancy must be enforced, due to the shared infrastructure model. Resource metering is needed to hold users accountable for the resources consumed. In the Pay-As-You-Go model, users can consume all the resources they want, as long as they are willing to pay for them. The ideal cloud is built to an exacting set of standards to ensure compatibility and federation with future clouds.

VMware vCloud Director is a platform that makes broad deployment of compute clouds possible by enabling self-service access to compute infrastructure through the abstraction of virtualized resources.

The key components for vCloud (shown in Figure 7) are:

- VMware vSphere
- VMware vCloud Director
- vShield Manager with vShield Edge
- VMware vCenter Chargeback

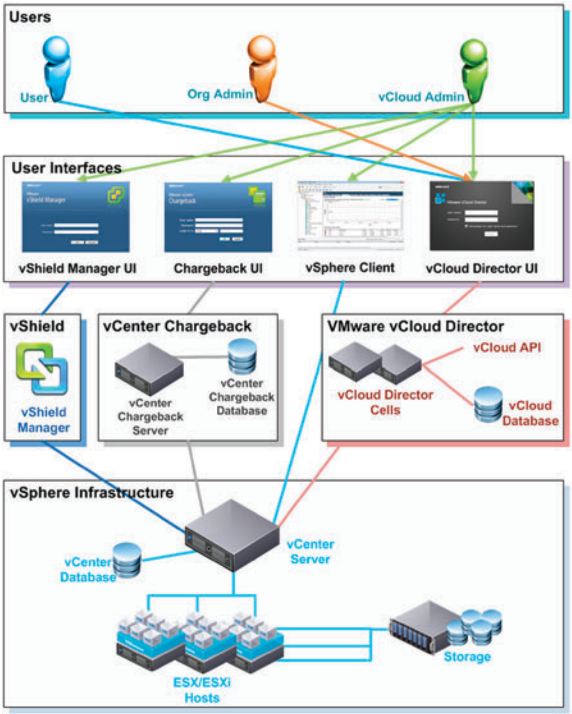

Figure 7. VMware vCloud Building Blocks

vCloud Building Blocks

VMware vCloud Director

VMware vCloud Director provides the interface, automation, and management feature set to allow enterprise and service providers to supply vSphere resources as a Web-based service. Having a consistent interface for public and private clouds eases the transition to hybrid cloud, as users only need to familiarize themselves with one system. For further simplification, all physical implementation details are masked from users, who see only a block of available infrastructure resources. Complexity is concealed from them, and they are constrained to a set of tasks dictated by their access privileges.

Think of vCloud Director as a centralized landing point for end users to access infrastructure resources through a Web browser. VMware vCloud Director provides the self-service portal that accepts user requests and translates them into tasks in the vSphere environment. The portal includes a catalog of standardized vApp templates and media files that users can leverage to provision applications. To achieve the scalability required for cloud environments, virtual resources from multiple vCenter Servers can be aggregated and presented through vCloud Director.

vShield Manager

vShield Manager (VSM) is the centralized network management interface for all vShield products, providing a single framework for comprehensive security and compliance. The technology for vShield came from the VMware acquisition of BlueLane, a security vendor focused on solutions for physical and virtual data centers. vShield Manager integrates with vCloud Director and vCenter Server to get deeper introspection into cloud networks and workloads. This authoritative knowledge is what allows for innovative approaches towards securing networks.

VSM is bundled as a virtual appliance and deploys in its own security virtual machine. Functionality is exposed through Web-based access, a vCenter plug-in, and REST APIs. VSM development is happening in concert with the next-generation cloud and virtualization platforms, as VMware is fully committed to security as a core pillar of the entire cloud stack. This includes actively working with partners on future integration plans to build on vShield as the foundation for cloud security.

vShield Edge

vShield Edge is a virtual firewall router that provides the network edge security needed to support multi-tenancy. vShield Edge security virtual appliances deploy automatically when routed or when isolated networks are created from vCloud Director. With vApp networks, vShield Edge devices are dynamically provisioned and decommissioned based on the power state of the vApp. vShield Edge devices connect the isolated, private networks of cloud tenants to outside networks through common-edge services such as firewall, NAT, and DHCP. The full version of vShield Edge adds additional features such as VPN and load-balancing capabilities.

vShield Edge is deployed as a virtual appliance, which is fast becoming the trend for network security products. There are several benefits realized:

- As a virtual appliance, vShield Edge takes advantage of the latest capabilities and improvements in the underlying x86 hardware. Special-purpose physical appliances typically must be refreshed every 5–7 years or risk becoming outdated.
- vShield Edge consolidates network services. It is able to perform firewall, NAT, DHCP, VPN, and load-balancing services. Performing these functions in software allows for quicker updates and improvements to the platform due to the rapid way in which edge devices can be provisioned and configured.
- vShield Edge takes a distributed approach to scaling. The scale-out approach allows each tenant to be protected by a set of edge devices without compromis-

ing performance. This is more fault-tolerant, since the loss of any single edge device impacts only a single tenant.

- While special-purpose ASICs and FPGAs may provide advantages at Layers 2–3, virtual machines provide similar performance for Layers 4–7 due to the type of operations and calculations performed.

vCloud API

The vCloud API is an open, RESTful interface to the vCloud. It bridges the gap between internal and external clouds by enabling consistent mobility, provisioning, and management of cloud workloads. VMware vCloud Director implements the vCloud API and generates the corresponding vSphere API calls required. The vCloud API is the key enabler for future hybrid cloud infrastructures.

vCenter Chargeback

vCenter Chargeback is used to measure, analyze, and report on costs associated with virtualized environments. The components of vCenter Chargeback include the server, database, data collectors, Web interface, and REST APIs. vCenter Chargeback is a standalone product with built-in integration with vCloud Director.

VMware vSphere

As the core building block, vSphere provides the abstraction of the physical infrastructure layer for the vCloud architecture. VMware vCloud Director layers on top of multiple vSphere instances to enable large-scale deployment, logical resource pooling, and secure multi-tenancy. All cloud workloads still execute on the underlying vSphere environment.

Resource Abstractions

VMware vCloud Director adds an additional layer of resource abstraction to enable multi-tenancy and provide interoperability between clouds that are built to the vCloud API standard.

- Physical compute, storage, and network resources are passed to the vSphere layer where resource pools, virtual switches, and datastores are created.
- Resource pools and datastores are then passed up to vCloud Director and attached to provider VDCs.
- Pure virtual compute and storage resources are exposed to users through virtual datacenter constructs.
- Users consume pure virtual resources from virtual datacenters through various allocation models.

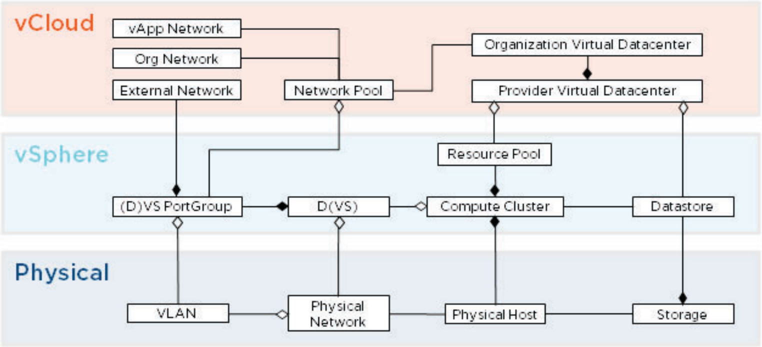

Figure 8. Relationship between Physical, Virtual, and vCloud Layers

In Figure 8, UML (Universal Markup Language) is used to model the relationships between the physical, virtual, and cloud abstractions. Aggregation is a variant of the "has a" relationship, meaning that one object contains another object. This association relationship does not have a strong lifecycle dependency on the container and is represented by a clear diamond shape (◇) at the end of the containing class. A provider VDC contains a resource pool and datastores but those datastores can be associated with other provider VDCs. If a provider VDC is destroyed, the resource pool and datastores persist. Composition is represented by a solid diamond shape (◆) at the end of the containing class and has a strong lifecycle dependency. The organization VDC cannot be used without a provider VDC. A provider VDC cannot be removed until all dependent organization VDCs are removed.

VMware vCloud Director Cell

VMware vCloud Director utilizes stateless cells that provide the functionality for higher-order resource management and structure. A vCloud Director *cell* is a single server running the vCloud Director software. There is some caching that happens at the vCloud Director cell level, such as SSL session data, but all refreshes and updates are done to information stored in the database. Cells are the main entry point of users' requests through the UI layer. All cells interact with vCenter Server as well as directly with ESX/ESXi hosts via VIM calls. One cell always has an open connection with vCenter Server and listens for updates (see Figure 9, next page).

A cell contains a set of services including the vCloud Director Web interface, the REST-based vCloud API, and the Remote Console Proxy component of vCloud Director.

To provide increased availability, vCloud Director can scale horizontally, clustering multiple cells together. Multi-cell configurations rely on an external load balancer to direct traffic among the clustered cells. Any load balancer can be used as long as it is configured for sticky sessions to reduce session database calls.

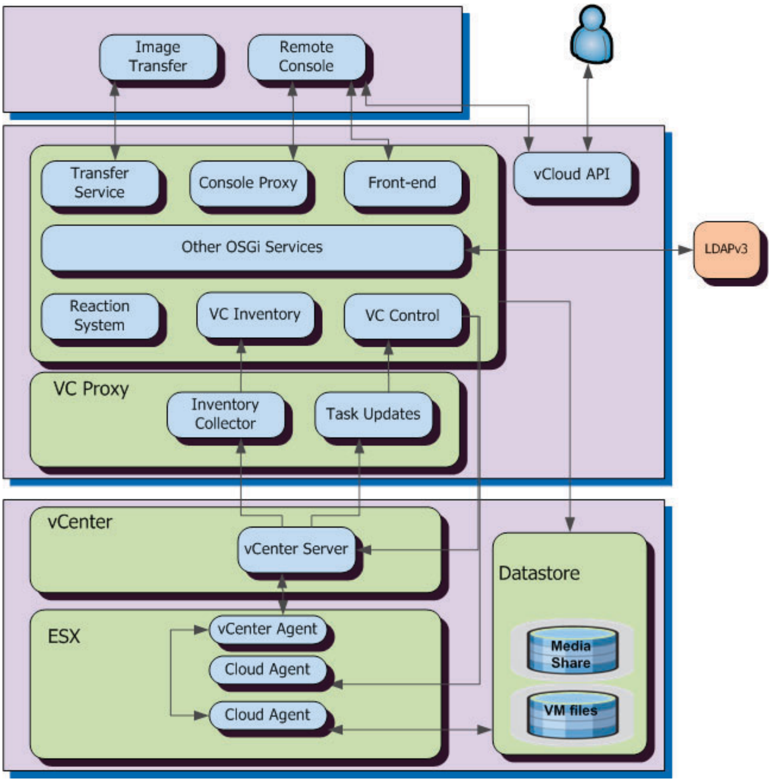

Figure 9. VMware vCloud Director Cell

The Remote Console Proxy component of vCloud Director allows a cloud user to access the console of a vApp hosted by vCloud Director over the Internet using a Web browser plug-in. The remote console access involves three components:

- VMware Remote Console browser plug-in
- VMRC client
- Remote Console Proxy

VMware vCloud Director Database

VMware vCloud Director cells point to a shared database where all configuration and state data is stored. Losing the database removes the ability to manage the cloud infrastructure, although all running workloads are unaffected. As such, the database is critical to the operation of vCloud Director. In a production environment, the database should be clustered or a hot standby should be provided.

The upper bound for the number of cells in a cluster is dependent on the number of connections to the vCloud Director database.

Management and Monitoring

Existing management and monitoring tools for vSphere can continue to be used with cloud infrastructures. Not all monitoring solutions are vCloud-aware yet as this requires leveraging the vCloud API to gain visibility into vCloud constructs. Recent VMware acquisitions such as Integrien Alive™ provide unique analytics and algorithms to analyze performance data in real-time. Future cloud environments will employ end-to-end proactive monitoring to identify sources of performance degradation before services are impacted.

Multi-Tenancy and Security

vCloud Director mediates access to the cloud environment. The organization is the unit of multi-tenancy and is a logical security boundary. Users in an organization are restricted to the vApps, catalogs, networks, and virtual datacenters that belong to that organization. VMware vCloud Director enforces network isolation for each of the networks provisioned for the cloud. vShield Edge provides network security support for multi-tenancy.

Orchestration and Workflow

VMware Orchestrator and vCloud Request Manager allow for customized process workflows with the vCloud Director platform. VMware Orchestrator is used with vSphere and can help cloud service providers enhance their cloud solutions by adding automated workflows integrated with vCloud Director. VMware vCloud Request Manager supports asset inventory and workflow, as well as a packaged approach for approval, including support for email, forms, and connectors to vCloud Director.

Summary

VMware vCloud Director is the layer of software that takes virtual resources and exposes them as cloud artifacts to end users. The compute cloud stack leverages the integration of the vCloud API, vCloud Director, vShield, and vSphere platforms to offer functionality that can be used in a variety of cloud solutions. For cloud consumers, the location of workloads is immaterial because deployment occurs without bindings to specific virtual or physical compute resources. Although the complexity of the system is hidden from end users, implementers must be aware of the challenges involved in integrating cloud components into existing environments. The same technology that simplifies the provisioning process can also disrupt existing operations such as backup and recovery. As future versions of vCloud Director are developed, commonly requested features will be added and the technology will become easier to deploy.

6. VMware vCloud Director Virtual Datacenters

VMware vCloud Director is considered to be a pure virtualization layer. It abstracts the virtual resources from vSphere in the form of virtual datacenters that contain vSphere clusters and/or resource pools. Provider virtual datacenters supply the resource abstraction of vSphere resources for vCloud use, while organization virtual datacenters partition provider VDCs and allocate resources to an organization. Different models are available for allocating and consuming virtual datacenter resources. Each type of model affects the resource abstractions in the vSphere layer.

Provider Virtual Datacenter

VMware vCloud Director abstracts resources managed by vCenter Server. The following types of resources can be used by a tenant:

- Compute
- Network
- Storage

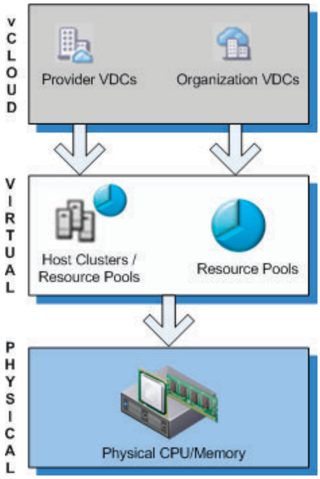

Figure 10. Cloud Resources

These resources are offered through a self-service portal that is part of vCloud Director. A cloud administrator can use the vCloud Director portal to carve up these resources as required and assign them to a tenant, referred to in vCloud Director as an *organization.*

To carve up resources, a container is created for a virtual datacenter (VDC). There are two different types of VDC:

- Provider VDC
- Organization VDC

The provider VDC is the foundation for compute and storage resources. Each provider VDC combines a single resource pool and a set of datastores. Provider VDCs are created and managed by the vCloud system administrator. Each vCloud deployment can have multiple provider VDCs and each provider VDC can deliver resources to multiple organization VDCs.

Creating a provider VDC requires selecting a resource pool. This can be a vSphere cluster, which is essentially the root resource cluster. Figure 11 shows the screen used to select a resource pool.

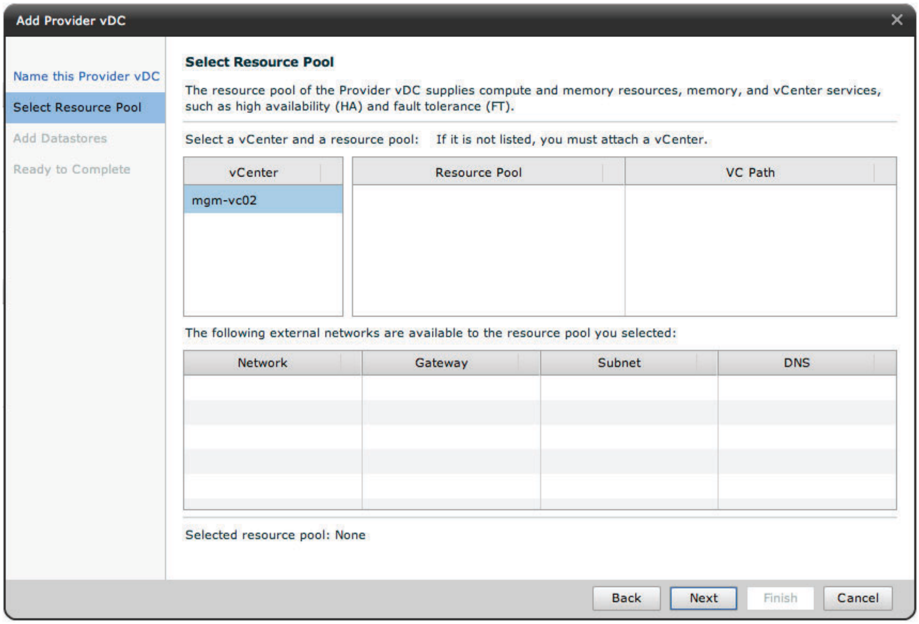

Select Resource Pool

The resource pool of the Provider vDC supplies compute and memory resources, memory, and vCenter services, such as high availability (HA) and fault tolerance (FT).

Select a vCenter and a resource pool: If it is not listed, you must attach a vCenter.

vCenter	Resource Pool	VC Path
mgm-vc02		

The following external networks are available to the resource pool you selected:

Network	Gateway	Subnet	DNS

Selected resource pool: None

Figure 11. Creation of Provider VDC

After selecting the appropriate cluster or resource pool, one or multiple datastores need to be associated with the provider VDC. Storage is often the main differentiator when it comes to service level agreements. However, the level of availability can also be part of the SLA and, subsequently, of your provider VDC when multiple clusters with a different VMware HA N+X architecture are available.

As an example, you could have a provider VDC labeled “Gold” with 15K FC disks and N+2 storage redundancy for HA, while your “Silver” provider VDC offers N+1 redundancy and runs on SATA disks. VMware vCloud Director enables this through abstraction.

Organization Virtual Datacenter

Organization VDCs are used by vCloud Director to partition provider VDCs and allocate resources to an organization. VMware vCloud Director uses vSphere resource pools as the basic construct to partition these resources.

An organization VDC is associated with a single organization and has a 1:1 relationship with a provider VDC. An organization can have multiple organization VDCs associated with it. Figure 12 depicts the scenario where a single organization has multiple organization VDCs that belong to two different provider VDCs. The two provider VDCs each have a specific SLA where, for example, Gold offers N+2 redundancy and Silver offers N+1.

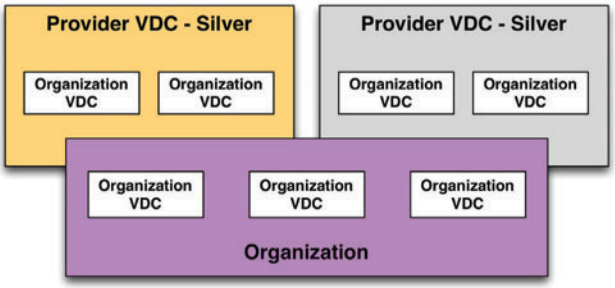

Figure 12. Provider VDC and Organization VDC (Org VDC) Relationship

Like the provider VDC, the organization VDC is a container for resources, but the way resources are allocated can be specified. A network pool can be added to an organization VDC with limits on the number of networks that can be created. You can also specify the maximum amount of storage the organization VDC can consume.

When creating an organization VDC, it is first necessary to select the VDC to which it will belong. From a vSphere perspective, both provider VDCs and organization VDCs are resource pools and have a parent-child relationship. The organization VDC inherits availability characteristics from the provider VDC to which it belongs. Figure 13 (see next page) shows the screen used to select the provider VDC.

When creating an organization VDC, choosing an appropriate allocation model is important. The allocation model not only determines how provider VDC compute resources are committed to the organization virtual datacenters, but also how the provider bills the customer for these resources.

Allocation Models

VMware vCloud Director abstracts and allocates resources. It offers multiple resource allocation methods, referred to as *allocation models*, that are defined at an organization VDC level. VMware vCloud Director has three different types of allocation models. These are listed in Table 1 along with the original description provided by vCloud Director.

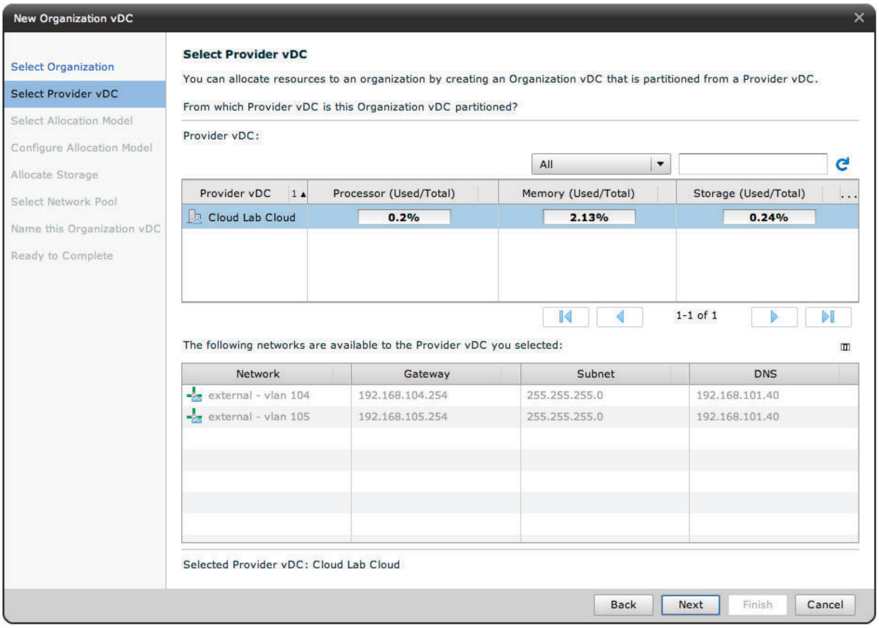

Select Provider vDC

You can allocate resources to an organization by creating an Organization vDC that is partitioned from a Provider vDC.

From which Provider vDC is this Organization vDC partitioned?

Provider vDC:

Provider vDC	Processor (Used/Total)	Memory (Used/Total)	Storage (Used/Total)
Cloud Lab Cloud	0.2%	2.13%	0.24%

The following networks are available to the Provider vDC you selected:

Network	Gateway	Subnet	DNS
external - vlan 104	192.168.104.254	255.255.255.0	192.168.101.40
external - vlan 105	192.168.105.254	255.255.255.0	192.168.101.40

Selected Provider vDC: Cloud Lab Cloud

Figure 13. Creation of an Organization VDC

Allocation Model	**Specification**
Allocation Pool	Only a percentage of the allocated resources are committed to the organization VDC. A specified percentage allows resource over-commitment.
Pay-As-You-Go	Allocated resources are only committed when users create vApps in the organization VDC. Maximum amounts of CPU and memory resources can be committed to the organization VDC.
Reservation Pool	All of the allocated resources are committed to the organization VDC.

Table 1. vSphere Infrastructure HA Cluster Configuration

Each allocation model has its own characteristics and can be placed in one of two categories: VM or resource pool.

Reserving or limiting resources on either a per-VM level or on a resource pool level distinguishes each of the allocation models. VMware vCloud Director guarantees resources to a VM or to an entire pool by setting a reservation. VMware vCloud Director also caps resource over-allocation by setting limits. Both reservations and limits are set on a VM or on a resource pool level depending on the selected allocation model.

Figure 14 shows the screen used to select the allocation model.

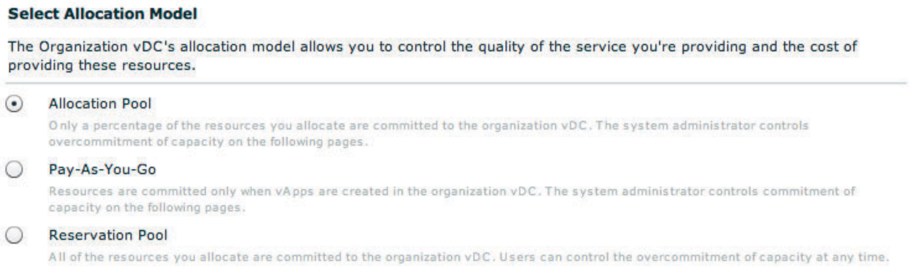

Select Allocation Model

The Organization vDC's allocation model allows you to control the quality of the service you're providing and the cost of providing these resources.

● Allocation Pool

Only a percentage of the resources you allocate are committed to the organization vDC. The system administrator controls overcommitment of capacity on the following pages.

○ Pay-As-You-Go

Resources are committed only when vApps are created in the organization vDC. The system administrator controls commitment of capacity on the following pages.

○ Reservation Pool

All of the resources you allocate are committed to the organization vDC. Users can control the overcommitment of capacity at any time.

Figure 14. Organization VDC Allocation Model Screen

Allocation Pool

The default allocation model is the Allocation Pool. This model specifies an amount of resources for the organization VDC and the amount that is guaranteed. The amount guaranteed is set as a reservation on the corresponding vSphere resource pool.

When the percentage of guaranteed resources is set to 100%, a reservation of 100% of the allocated resources is set on the resource pool. When an organization VDC is created with 10 GHz of CPU resources and a guarantee of 75%, it translates into a resource pool with a limit of 10 GHz and a reservation of 75% of that 10 GHz—in this case, 7500 MHz.

Characteristics

Each allocation model has very specific characteristics that can be placed in either a VM or resource pool category.

- The Allocation Pool is a pool of resources, of which a percentage can be guaranteed.
  - A reservation is set to guarantee resources on a resource pool level.
  - By default, the resource pool reservations are: CPU 0%, memory 100%.
  - Customer is allocated a fixed amount of guaranteed resources but has the ability to burst.
- On a per-VM level no reservation is set for CPU resources.
- On a per-VM level a reservation is set for memory resources. This reservation is based on the percentage of guaranteed resources.

The following example explains these characteristics in more detail.

Allocation Pool Example

In this example a tenant has requested 10 GHz of CPU resources with a guarantee of 25% and 10 GB of memory with a guarantee of 100%. This is shown in Figure 15 (see next page).

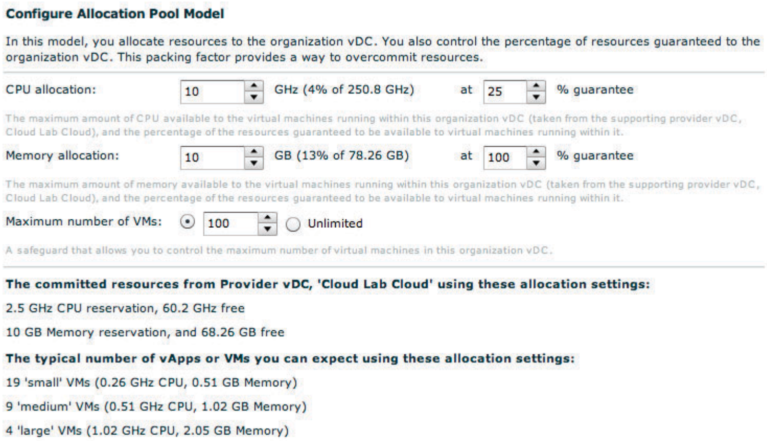

Configure Allocation Pool Model

In this model, you allocate resources to the organization vDC. You also control the percentage of resources guaranteed to the organization vDC. This packing factor provides a way to overcommit resources.

CPU allocation: 10 GHz (4% of 250.8 GHz) at 25 % guarantee

The maximum amount of CPU available to the virtual machines running within this organization vDC (taken from the supporting provider vDC, Cloud Lab Cloud), and the percentage of the resources guaranteed to be available to virtual machines running within it.

Memory allocation: 10 GB (13% of 78.26 GB) at 100 % guarantee

The maximum amount of memory available to the virtual machines running within this organization vDC (taken from the supporting provider vDC, Cloud Lab Cloud), and the percentage of the resources guaranteed to be available to virtual machines running within it.

Maximum number of VMs: ● 100 ○ Unlimited

A safeguard that allows you to control the maximum number of virtual machines in this organization vDC.

The committed resources from Provider vDC, 'Cloud Lab Cloud' using these allocation settings:

2.5 GHz CPU reservation, 60.2 GHz free

10 GB Memory reservation, and 68.26 GB free

The typical number of vApps or VMs you can expect using these allocation settings:

19 'small' VMs (0.26 GHz CPU, 0.51 GB Memory)

9 'medium' VMs (0.51 GHz CPU, 1.02 GB Memory)

4 'large' VMs (1.02 GHz CPU, 2.05 GB Memory)

Figure 15. Organization VDC Allocation Pool

The default for CPU is a 0% guarantee, but in this example it has been manually set to 25% guarantee for CPU allocation.

VMware vSphere resource pools are the constructs that are used to carve resources. The tenant's request results in a resource pool with a reservation of 2500 MHz and a limit of 10,000 MHz. As viewed in vCloud Director, 2500 MHz are guaranteed, and the tenant has the ability to burst up to 10.000 MHz. The resource pool that is created and its properties are shown in Figure 16.

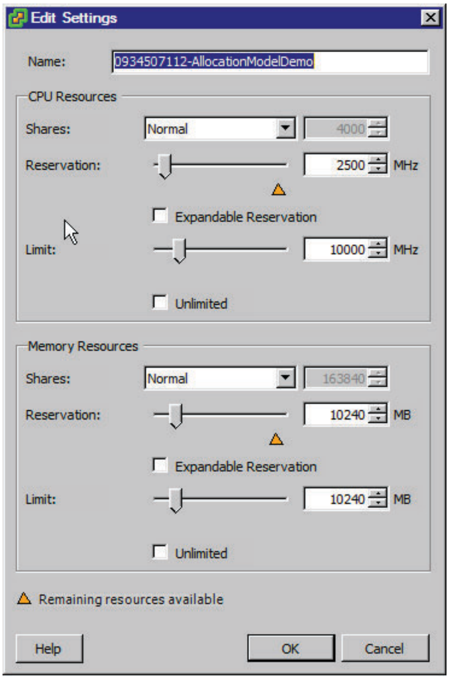

Name: 0934507112-AllocationModelDemo

CPU Resources

Shares: Normal 4000

Reservation: 2500 MHz

□ Expandable Reservation

Limit: 10000 MHz

□ Unlimited

Memory Resources

Shares: Normal 163840

Reservation: 10240 MB

□ Expandable Reservation

Limit: 10240 MB

□ Unlimited

Remaining resources available

Figure 16. Allocation Pool Resource Pool

A virtual machine contained as part of a newly created vApp is placed in the resource pool that corresponds with the chosen organization VDC. For the Allocation Pool model, the VM level has the CPU Reservation set to zero. This means that no GHz are reserved for this VM. This is shown in Figure 17.

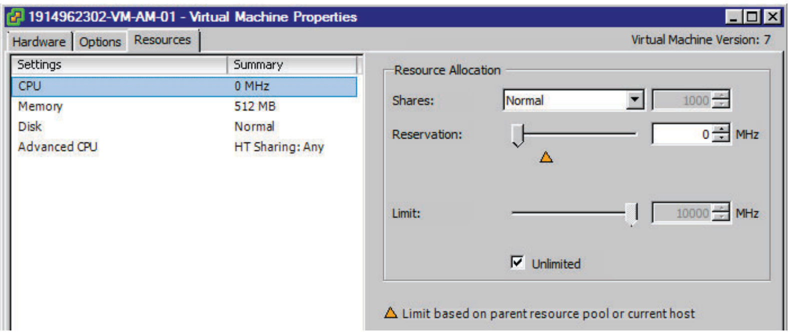

1914962302-VM-AM-01 - Virtual Machine Properties

Hardware | Options | Resources

Virtual Machine Version: 7

Settings	Summary
CPU	0 MHz
Memory	512 MB
Disk	Normal
Advanced CPU	HT Sharing: Any

Resource Allocation

Shares: Normal 1000

Reservation: 0 MHz

Limit: 10000 MHz

☑ Unlimited

⚠ Limit based on parent resource pool or current host

Figure 17. Allocation Pool VM-Level CPU Reservation and Limits

For memory, this behavior is slightly different because both a reservation and a limit are configured. The limit always equals the provisioned memory and the reservation equals the percentage of guaranteed memory as defined as part of the allocation model.

Figure 18 shows that the guaranteed amount of memory resources has been set to 100% of all available memory.

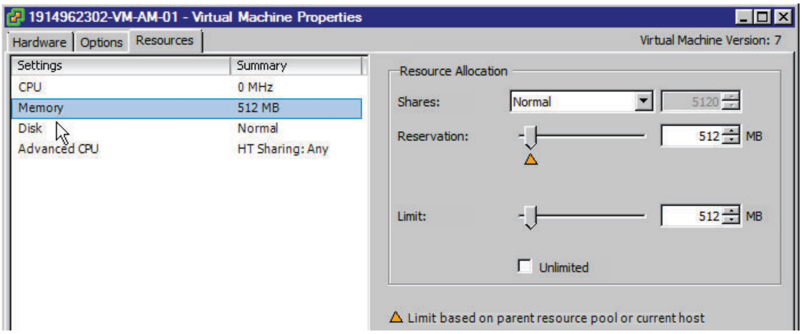

1914962302-VM-AM-01 - Virtual Machine Properties

Hardware | Options | Resources

Virtual Machine Version: 7

Settings	Summary
CPU	0 MHz
Memory	512 MB
Disk	Normal
Advanced CPU	HT Sharing: Any

Resource Allocation

Shares: Normal 5120

Reservation: 512 MB

Limit: 512 MB

□ Unlimited

⚠ Limit based on parent resource pool or current host

Figure 18. Allocation Pool VM-Level Memory Reservation and Limits

However, if we change the percentage of guaranteed resources for memory on the Allocation Pool organization VDC to 50%, the resource pool changes accordingly, as is shown in Figure 19 (see next page). Note that the reservation has been decreased from 10,240 MB to 5120 MB.

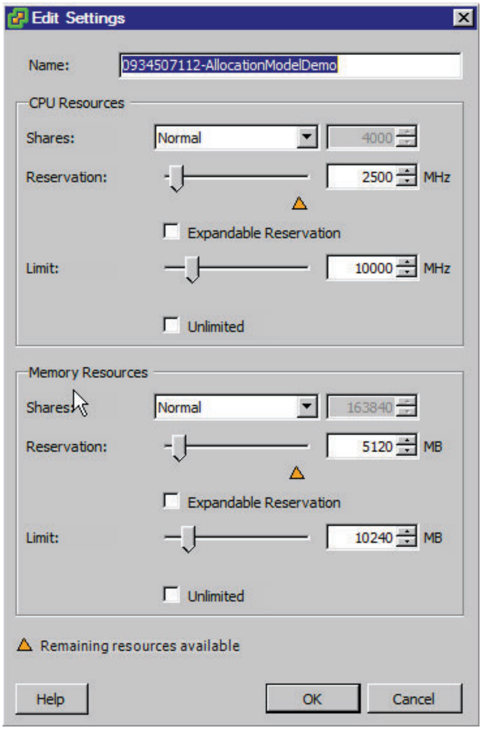

Edit Settings

Name: 0934507112-AllocationModelDemo

CPU Resources

Shares: Normal 4000

Reservation: 2500 MHz

☐ Expandable Reservation

Limit: 10000 MHz

☐ Unlimited

Memory Resources

Shares: Normal 163840

Reservation: 5120 MB

☐ Expandable Reservation

Limit: 10240 MB

☐ Unlimited

⚠ Remaining resources available

Figure 19. Changed Memory Guarantee to 50% on Resource Pool

This in turn results in a decrease of the per-VM level memory reservation to 50% of the provisioned memory. In Figure 20, the VM has been provisioned with 512 MB, of which 256 MB is reserved (guaranteed):

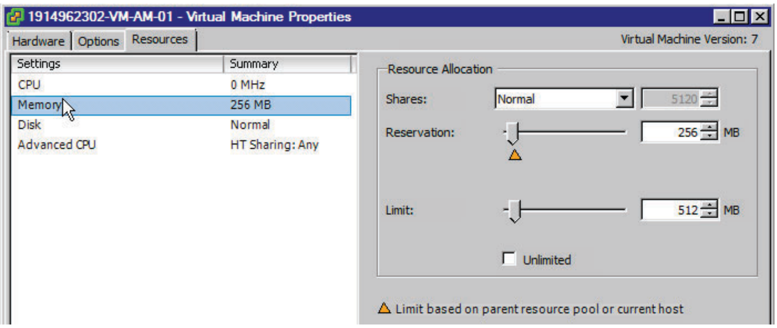

1914962302-VM-AM-01 - Virtual Machine Properties

Virtual Machine Version: 7

Settings	Summary
CPU	0 MHz
Memory	256 MB
Disk	Normal
Advanced CPU	HT Sharing: Any

Resource Allocation

Shares: Normal 5120

Reservation: 256 MB

Limit: 512 MB

☐ Unlimited

⚠ Limit based on parent resource pool or current host

Figure 20. Changed Memory Guarantee to 50% on Virtual Machine

Allocation Pool Summary

As shown in the example, the percentage of guaranteed resources impacts the implementation of the resource pool, the associated limits, and reservations. The Allocation Pool model allows the cloud provider and cloud consumer (tenant) to contract for a certain quantity of resources in the organization VDC, with the resource guarantee levels being a part of that contract.

Pay-As-You-Go

Pay-As-You-Go is the traditional model used by many enterprise environments. This allocation model allows you to specify an amount of guaranteed resource per VM, unlike the other allocation models. When the percentage of guaranteed resources is set to 100%, a reservation is set to 100% of what has been allocated to that particular VM. This model also differs from the other models by allowing limitation of the vCPU speed.

Characteristics

- Percentage of resources guaranteed on a per-VM level.
  - A reservation is set on a VM level.
  - By default, the VM reservation on CPU is 0% and memory 100%.
  - By default, the vCPU speed is set to 0.26 GHz, which means the vCPU is limited to 0.26 GHz.
- The resource pool that corresponds with the vCloud Director organization VDC is an accumulation of all reservations set on a per-VM level.

Pay-As-You-Go Example

In this example, a tenant has requested an organization VDC with a Pay-As-You-Go allocation model. The tenant has requested the vCPU speed to be set to 0.26 GHz, of which 25% needs to be guaranteed. Memory resources have been left at default values.

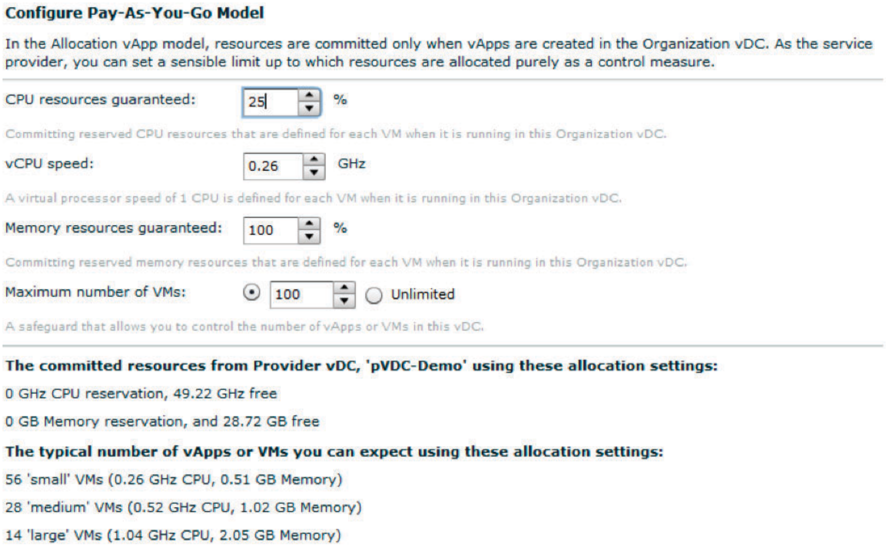

Configure Pay-As-You-Go Model

In the Allocation vApp model, resources are committed only when vApps are created in the Organization vDC. As the service provider, you can set a sensible limit up to which resources are allocated purely as a control measure.

CPU resources guaranteed: 25 %

Committing reserved CPU resources that are defined for each VM when it is running in this Organization vDC.

vCPU speed: 0.26 GHz

A virtual processor speed of 1 CPU is defined for each VM when it is running in this Organization vDC.

Memory resources guaranteed: 100 %

Committing reserved memory resources that are defined for each VM when it is running in this Organization vDC.

Maximum number of VMs: ● 100 ○ Unlimited

A safeguard that allows you to control the number of vApps or VMs in this vDC.

The committed resources from Provider vDC, 'pVDC-Demo' using these allocation settings:

0 GHz CPU reservation, 49.22 GHz free

0 GB Memory reservation, and 28.72 GB free

The typical number of vApps or VMs you can expect using these allocation settings:

56 'small' VMs (0.26 GHz CPU, 0.51 GB Memory)

28 'medium' VMs (0.52 GHz CPU, 1.02 GB Memory)

14 'large' VMs (1.04 GHz CPU, 2.05 GB Memory)

Figure 21. Organization VDC Pay-As-You-Go Model

Note that the default for CPU is 0%, but it has been manually set to 25%. Also note that the vCPU speed is set by default to 0.26 GHz, which is reflected as the vCPU limit for the VM. It is recommended that this default value be increased. When the organization VDC is created it results in the resource pool shown in Figure 22 (see next page).

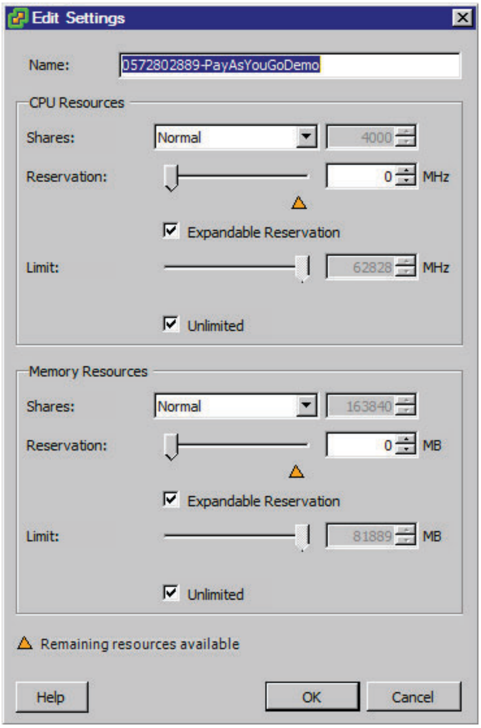

Name: 0572802889-PayAsYouGoDemo

CPU Resources

Shares: Normal 4000

Reservation: 0 MHz

☑ Expandable Reservation

Limit: 62828 MHz

☑ Unlimited

Memory Resources

Shares: Normal 163840

Reservation: 0 MB

☑ Expandable Reservation

Limit: 81889 MB

☑ Unlimited

⚠ Remaining resources available

Figure 22. Pay-As-You-Go Resource Pool before Deploying a vApp

When vApps are created, their associated reservations are accumulated into the resource pool. When a vApp is created with two VMs, each having a single vCPU and 512 MB of memory, reservations on memory and CPU are altered accordingly (see Figure 23).

On a resource pool level, a reservation of 130 MHz is set, which is 25% of 2 x 0.26 GHz. A guarantee of 100% was set on memory, which translates to 1221 MB in total.

Note: A resource pool includes the memory overhead of virtualization. See the *VMware Resource Management Guide* for more details.

The primary difference between the Pay-As-You-Go allocation model and the other allocation models is the use of limits and reservation on a per-VM level. Figures 24 and 25 show the reservation and limit that have been set as a result of the selected values.

As shown in Figure 25, a reservation of 65 MHz and a limit of 260 MHz on CPU have been defined. For memory, a 512 MB reservation and limit have been configured. If guaranteed memory resources had been configured with a value of 50%, the reservation of memory resources for this VM would have been set to 256 MB.

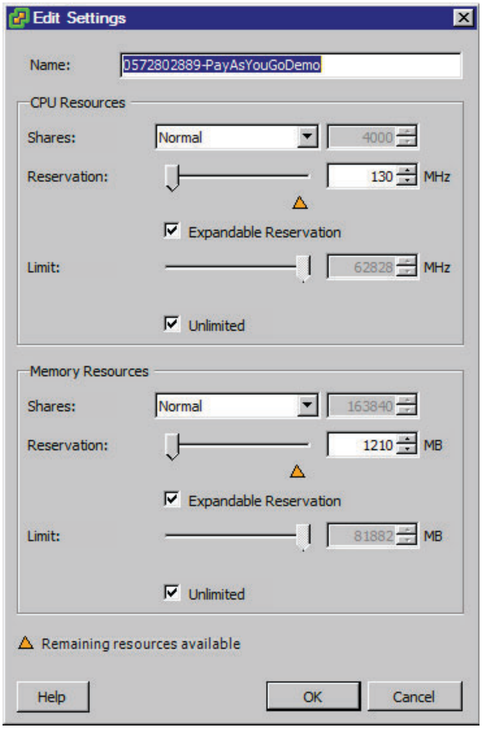

Edit Settings

Name: 0572802889-PayAsYouGoDemo

CPU Resources

Shares: Normal 4000

Reservation: 130 MHz

☑ Expandable Reservation

Limit: 62828 MHz

☑ Unlimited

Memory Resources

Shares: Normal 163840

Reservation: 1210 MB

☑ Expandable Reservation

Limit: 81882 MB

☑ Unlimited

⚠ Remaining resources available

Figure 23. Pay-As-You-Go Resource Pool after Deploying a vApp

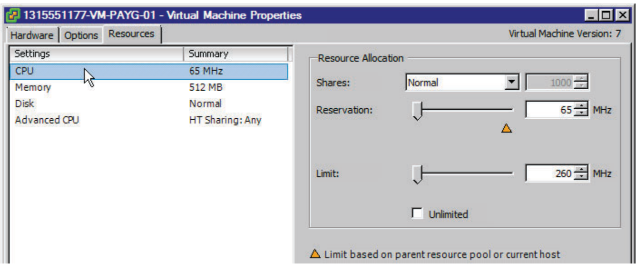

1315551177-VM-PAYG-01 - Virtual Machine Properties

Virtual Machine Version: 7

Settings	Summary
CPU	65 MHz
Memory	512 MB
Disk	Normal
Advanced CPU	HT Sharing: Any

Resource Allocation

Shares: Normal 1000

Reservation: 65 MHz

Limit: 260 MHz

☐ Unlimited

⚠ Limit based on parent resource pool or current host

Figure 24. Pay-As-You-Go VM-Level CPU Reservation and Limits

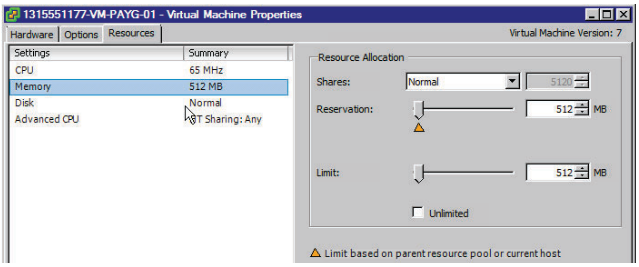

1315551177-VM-PAYG-01 - Virtual Machine Properties

Virtual Machine Version: 7

Settings	Summary
CPU	65 MHz
Memory	512 MB
Disk	Normal
Advanced CPU	HT Sharing: Any

Resource Allocation

Shares: Normal 5120

Reservation: 512 MB

Limit: 512 MB

☐ Unlimited

⚠ Limit based on parent resource pool or current host

Figure 25. Pay-As-You-Go VM-Level Memory Reservation and Limits

Pay-As-You-Go Summary

The tenant has guaranteed resources per VM and typically is charged per instantiated and powered-on virtual machine. The resource pool created as part of the organization VDC only accumulates reserved resources and does not limit the VMs. Limits are placed on a per-VM level.

Reservation Pool

The Reservation Pool allocation model is the most static model. In this model all resources are guaranteed. It can be compared to an Allocation Pool with all guarantees set to 100%.

Characteristics

- Fully guaranteed pool of resources.
  - A reservation is set to guarantee resources on a resource pool level.
  - Customer pays a fixed amount for guaranteed resources.
- No reservations or limits are set on a per-VM level for CPU.
- It provides the ability to set custom limits, reservations, and shares on a per-VM level for CPU and memory.

These characteristics are explained in more detail in the following example.

Reservation Pool Example

In this example, the tenant has requested 10 GHz of CPU resources and 10 GB of memory resources. This is shown in Figure 26:

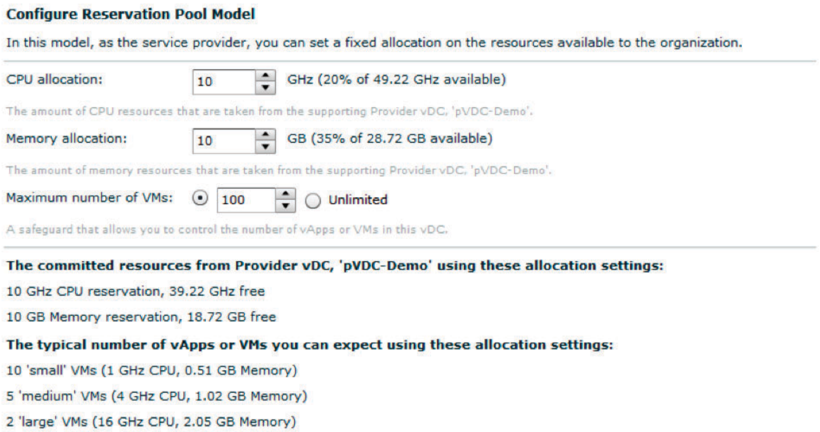

Configure Reservation Pool Model

In this model, as the service provider, you can set a fixed allocation on the resources available to the organization.

CPU allocation: 10 GHz (20% of 49.22 GHz available)

The amount of CPU resources that are taken from the supporting Provider vDC, 'pVDC-Demo'.

Memory allocation: 10 GB (35% of 28.72 GB available)

The amount of memory resources that are taken from the supporting Provider vDC, 'pVDC-Demo'.

Maximum number of VMs: ● 100 ○ Unlimited

A safeguard that allows you to control the number of vApps or VMs in this vDC.

The committed resources from Provider vDC, 'pVDC-Demo' using these allocation settings:

10 GHz CPU reservation, 39.22 GHz free

10 GB Memory reservation, 18.72 GB free

The typical number of vApps or VMs you can expect using these allocation settings:

10 'small' VMs (1 GHz CPU, 0.51 GB Memory)

5 'medium' VMs (4 GHz CPU, 1.02 GB Memory)

2 'large' VMs (16 GHz CPU, 2.05 GB Memory)

Figure 26. Organization VDC Reservation Pool

All resources are fully guaranteed. A resource pool with a reservation equal to the limit is created within vSphere, as shown in Figure 27.

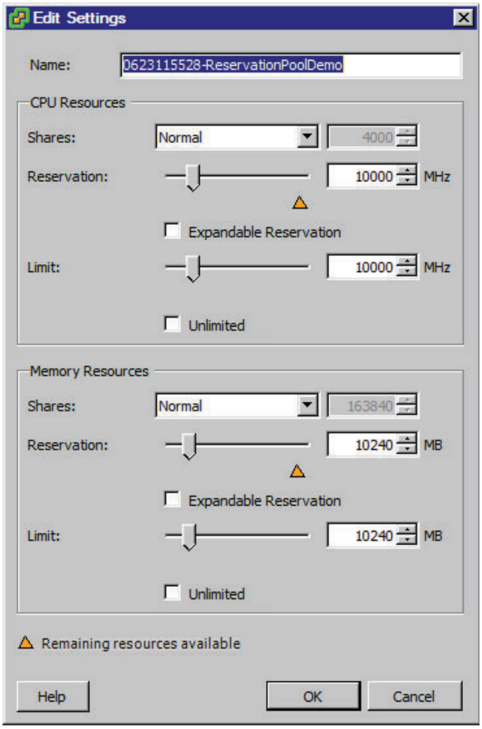

Name: 0623115528-ReservationPoolDemo

CPU Resources

Shares: Normal 4000

Reservation: 10000 MHz

☐ Expandable Reservation

Limit: 10000 MHz

☐ Unlimited

Memory Resources

Shares: Normal 163840

Reservation: 10240 MB

☐ Expandable Reservation

Limit: 10240 MB

☐ Unlimited

Remaining resources available

Figure 27. Reservation Pool Resource Pool

All resources are guaranteed on a resource pool level, with no reservations or limits set on the virtual machine. This is shown in Figure 28 and Figure 29.

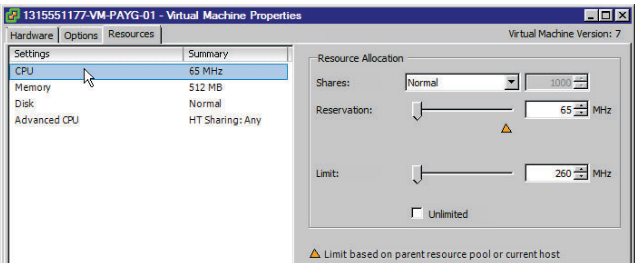

Virtual Machine Version: 7

Settings	Summary
CPU	65 MHz
Memory	512 MB
Disk	Normal
Advanced CPU	HT Sharing: Any

Resource Allocation

Shares: Normal 1000

Reservation: 65 MHz

Limit: 260 MHz

☐ Unlimited

Limit based on parent resource pool or current host

Figure 28. Reservation Pool VM-Level CPU Reservation and Limits

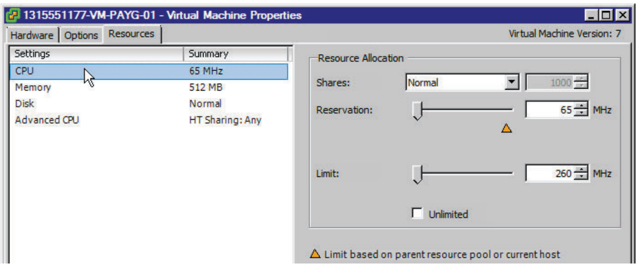

Virtual Machine Version: 7

Settings	Summary
CPU	65 MHz
Memory	512 MB
Disk	Normal
Advanced CPU	HT Sharing: Any

Resource Allocation

Shares: Normal 1000

Reservation: 65 MHz

Limit: 260 MHz

☐ Unlimited

Limit based on parent resource pool or current host

Figure 29. Reservation Pool VM-Level Memory Reservation and Limits

Reservation Pool Summary

The Reservation Pool model is very straightforward. A limit equal to the reservation is set on a resource pool level, which gives the tenant a guaranteed pool of resources. There are no limits or reservations set on a per-VM level, and this gives the customer flexibility in carving up resources.

7. VMware vCloud Networking

To provide secure multi-tenancy and resource elasticity, vCloud Director introduces new networking constructs that automate the creation of secure L2 networks for cloud tenants. The network functionality in vCloud Director provides robustness and flexibility, but can add complexity to each deployment. It is critical to understand the impact of network design decisions. In this chapter we introduce various components and layers of vCloud networking and discuss its abstraction from vSphere components. Sample use cases are also provided.

Overview

VMware vSphere networking consists of resources such as virtual switches, port groups, and vNICs. These virtual networking resources are abstracted from hardware resources such as physical switches, VLANs, and Network Interface Cards. There are two types of virtual switches, the standard virtual switch (vNetwork Standard Switch) and distributed switches (vNetwork Distributed Switches and Cisco Nexus 1000V).

The vCloud layer abstracts network resources from the virtual switches and port groups of the vSphere layer. Figure 30 illustrates how the network resources are abstracted.

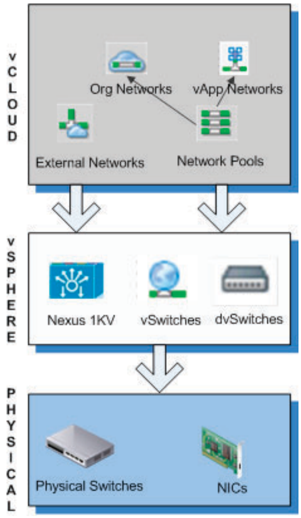

Figure 30. vCloud Network Layers

Network Pools

Network pools are a collection of isolated Layer 2 networks that can be used to create organization and vApp networks on demand. Network pools are used as a blueprint for creating new networks at the organization and vApp levels. They are available to both providers and consumers, but can only be created by the providers.

There are three types of network pools available from which the organization and vApp networks are created by vCloud Director:

- vSphere port group–backed
- VLAN-backed
- vCloud Director Network Isolation–backed (vCD-NI)

All three types of network pool can be used within the same instance of vCloud Director. However, the requirements and use cases are different for each type of network pool.

vSphere Port Group–Backed Network Pool

For vSphere port group-backed network pools, the provider is responsible for creating pre-provisioned port groups on vNetwork Standard Switch (VSS), vNetwork Distributed Switch (vDS), or Cisco Nexus 1000V virtual switches in vSphere. This can be done manually or through orchestration. Port groups are added to the network pool, which allows the creation of organization networks or vApp networks. It is recommended that the vSphere port group–backed network pools provide Layer 2 network isolation using IEEE 802.1Q VLANs with standard frame format.

vSphere port group-backed network pools support all the standard and distributed virtual switches (vNetwork Standard, vNetwork Distributed, and Cisco Nexus 1000V). The vSphere port group-backed network pool should be used in scenarios where there is a requirement for using Cisco Nexus 1000V virtual switches or when vSphere Enterprise Plus licensing is not available.

The following are some of the pros and cons of port group–backed network pools as compared to other types of network pools:

Pros:

- Versus VLAN-backed and vCD-NI-backed network pools:
  - Does not require Enterprise Plus licensing.
  - Support Cisco Nexus 1000V virtual switches.
  - Does not require MTU size changes to physical infrastructure.
- Allows use of existing features such as QoS (quality of service), ACLs (access control lists), and security. QoS and ACLs apply for Cisco Nexus 1000V.

Cons:

- All port groups must be created manually or through orchestration before they can be mapped to the network pool.
- Versus VLAN-backed and vCD-NI–backed network pools:
  - Scripting or host profiles must be used to make sure that the port groups are created consistently on all hosts, especially when using vNetwork Standard

Switches. Inconsistency between hosts could lead to vApps not being spread across hosts.

- Highest cost and complexity are associated with Cisco Nexus 1000V virtual switches.
- The port group isolation relies on VLAN Layer 2 isolation.

Figure 31 shows how a vSphere port group-backed network pool is mapped between various kinds of vSphere virtual switches and the vCloud Director networks.

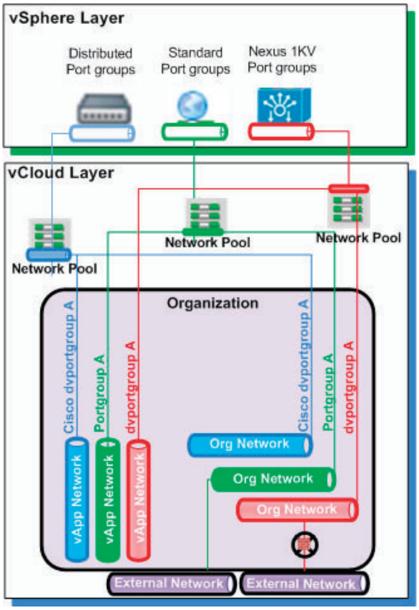

Figure 31. Port Group–Backed Network Pool

VLAN-Backed Network Pool

In VLAN-backed network pools, the provider is responsible for creating a range of VLANs on the physical network and trunking them to all ESXi hosts. The provider then maps the vNetwork Distributed Switch and the appropriate VLAN range to the VLAN-backed network pool.

Each time a vApp or an organization network is created, a dvPort group is created on the vNetwork Distributed Switch and assigned an available VLAN from the specified range. When the vApp or organization network is decommissioned, the VLAN ID is returned to the network pool so it can be reused. The VLAN-backed network pool provides network isolation using IEEE 802.1Q VLANs with standard frame format.

The creation of this network pool requires a vNetwork Distributed Switch and a range of available VLAN IDs that are trunked through the physical uplinks connected to the vNetwork Distributed Switch.

The vNetwork Distributed Switch is the only virtual switch supported for VLAN-backed network pools. VLAN-backed network pools should be used in scenarios where there are requirements for providing the most secure isolation, or to provide optional

VPN/MPLS, or any special consumer requirements, so it doesn't take up a lot of VLANs.

The following are some of the pros and cons of a VLAN-backed network pool as compared with other types of network pools:

Pros:

- Versus vCD-NI network pool—There is no need to change the Ethernet network frame MTU size, which is 1500 bytes by default.
- Versus vSphere port group-backed network pool—All of the port groups are created automatically and no manual intervention is required for consumers to create vApp and organization networks (unless the VLAN IDs are exhausted).

Cons:

- Requires VLANs to be configured and maintained on the physical switches and trunked to the NIC ports on the ESXi hosts.
- Versus vCD-NI network pool—Requires a wide range of VLANs depending on the number of vApp and organization networks needed for the environment, which may not be available.

Figure 32 shows how a VLAN-backed network pool is mapped between vSphere Distributed Virtual Switches and the vCloud Director networks.

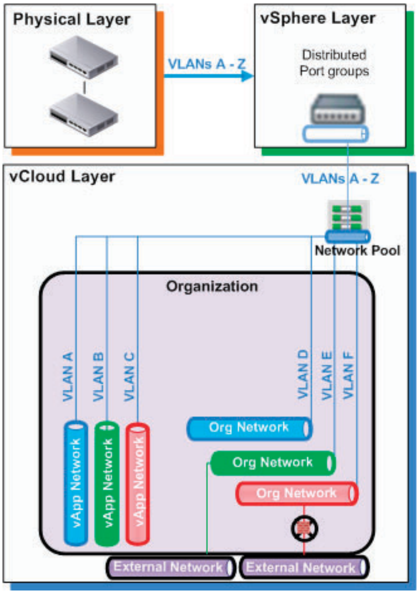

Figure 32. VLAN-Backed Network Pool

vCloud Network Isolation–Backed Network Pool

vCloud Director Network Isolation (vCD-NI) was introduced to provide an alternative technology for network isolation. It addresses issues such as VLAN sprawl by creating overlay networks using MAC-in-MAC Ethernet frame encapsulation. In vCloud Director Network, network isolation is exposed as a type of network pool. This network pool is similar to "Cross-Host Fencing" in VMware vCenter Lab Manager™.

The fence ID is used by vCloud Director to enforce isolation of each vCD-NI–backed network. vCloud Director Network Isolation adds 24 bytes for the encapsulation to each Ethernet frame, increasing the frame size to 1524 bytes. The encapsulation contains the source and destination MAC addresses of ESX hosts where VM endpoints reside, as well as the vCD-NI–backed network IDs. The ESX host strips the vCD-NI packets to expose the virtual machine source and destination MAC-addressed packet that is delivered to the destination virtual machine. All of this is transparent to the physical infrastructure as frame encapsulation purely terminates on ESXi hosts. When the guest OS and the underlying physical network infrastructure are configured with the standard MTU size of 1500 bytes, the vCD-NI-backed protocol causes frame fragmentation resulting in performance penalties. To avoid fragmentation, the MTU size should be increased by a minimum of 24 bytes on the physical network infrastructure and the vCD-NI–backed network pools. The MTU on the VM guest operating systems accessing the vCD-NI network should remain at the default of 1500 bytes.

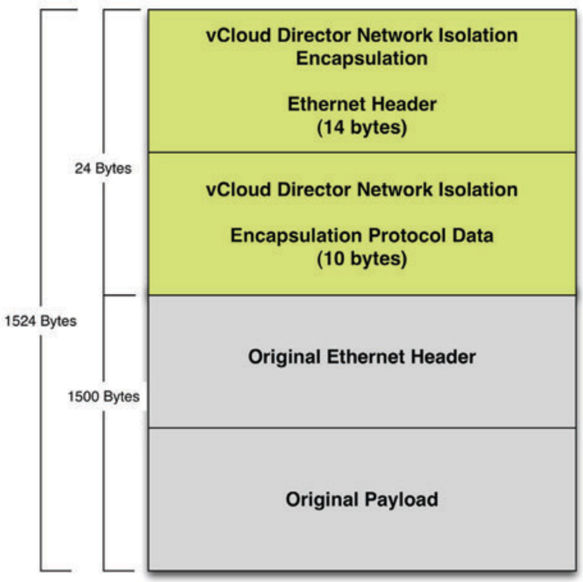
vCloud Director Network Isolation Encapsulation
Ethernet Header (14 bytes)
vCloud Director Network Isolation
Encapsulation Protocol Data (10 bytes)
Original Ethernet Header
Original Payload
24 Bytes
1524 Bytes
1500 Bytes

Figure 33. vCloud Network Isolation Frame

Each time an organization or vApp network is created, a dvPort group is created on the vNetwork Distributed Switch and assigned an available fence ID within vCD-NI. When a vApp or organization network is destroyed, the network ID is returned to the network pools so it can be available for others.

To create a vCD-NI network pool, you should specify the name of the network pool, the appropriate vDS, the total number of vCD-NI networks IDs to allocate, and the transport VLAN ID for the VLAN which carries the encapsulated traffic. After the creation of the network pool, the network pool MTU needs to be changed to 1524.

The vNetwork Distributed Switch is the only virtual switch supported for vCD-NI–backed network pools. vCloud Director NI–backed network pools are used in scenarios where there is no requirement for routed networks, when only a limited number of VLANs are available, or when the management of VLANs is problematic.

Here are some of the pros and cons of vCD-NI–backed network pools:

Pros:

- Do not require additional VLANs for creating vApp and organization networks aside from the transport VLAN.
- All port groups are created automatically. No manual intervention is required for the consumers to create vApp and organization networks unless the network pool runs out of vCD-NI–backed network IDs.
- Allow for the creation of isolated L2 networks within the same VLAN, increasing scalability.
- Do not require any special physical switch configuration.

Cons:

- Administrative overhead results from increasing the MTU size to 1524 across the entire physical network infrastructure.

Figure 34 illustrates how a vCD-NI–backed network pool is mapped between vSphere Distributed Virtual Switches and the vCloud Director networks.

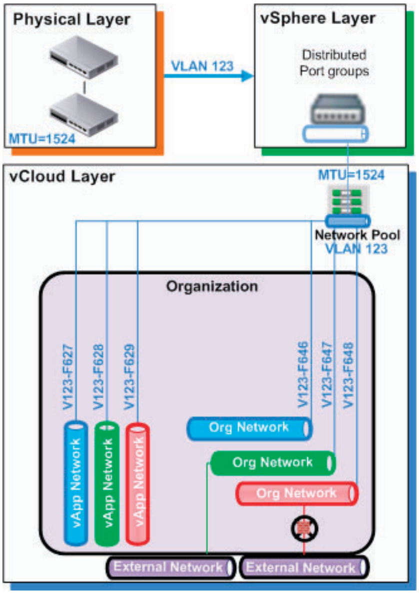

Figure 34. vCloud Director Network Isolation–Backed Network

vCloud Network Layers

There are three layers of networking available in vCloud Director:

- External networks
- Organization networks
- vApp networks

Cloud providers offer infrastructure resources to consumers such as IT organizations or internal divisions of an enterprise (e.g., a finance or marketing department). For vCloud networking, external and organization networks are created and managed by cloud providers. Consumers use those resources and have the ability to create their own vApp networks.

Each type of network is represented by a port group in vCenter Server, and each port group must characterize an isolated L2 network, either through a unique VLAN tag or using VMware vCloud Network Isolation technology.

External Networks

External networks are created by the providers and offer external connectivity to vCloud Director—they are the connections to the outside world. External networks are created by mapping a dvPort group or port group from a vNetwork Standard, vNetwork Distributed, or Cisco Nexus 1000V virtual switch at the vSphere layer. This port group must be manually created by the cloud provider in vSphere before an external network can be added in vCloud Director.

Some typical use cases for the external networks include:

- Internet access
- Provider supplied network endpoints such as:
  - IP-based storage
  - Online or offline backup services
- Backhauled networking services for consumers such as:
  - VPN access to a private cloud
  - MPLS termination

Figure 35 shows how an external network can be used as a gateway to vCloud Director to provide services.

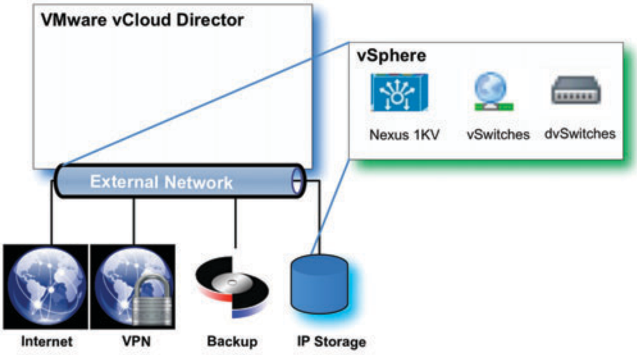

Figure 35. External Network

When providing external networks for services such as the Internet, providers allocate public IP address ranges to the consumers for inbound and outbound access. Though it is possible to create one large external network and provide it to all consumers, it is challenging to create and maintain the public IP addresses in one large IP range. It is worth considering creating one or more external networks per organization. This keeps the public IP address range separate for each consumer, allowing for easier maintenance while keeping multi-tenancy intact.

Organization Networks

Organization networks are created by the providers and are contained within the organizations, where organizations are the logical constructs of consumers. Their primary purpose is to allow multiple vApps to communicate, not only with each other, but with the external world by connecting to the external networks. In other words, organization networks bridge the vApps and the external networks.

Organization networks are provisioned from a set of preconfigured network pools that typically map a port group or dvPort group coming off a vNetwork Standard, vNetwork Distributed, or Cisco Nexus 1000V virtual switch at the vSphere layer.

Organization networks can be connected to the external networks in several ways:

- Public or direct connectivity—An organization network is bridged directly to an external network, where the deployed vApps are directly connected to the external network.
- Private or external NAT/routed connectivity—An organization network is NAT/routed to an external network where the deployed vApps are connected via a vShield Edge that provides firewall and/or NAT functionality for security.
- Private or isolated or internal connectivity—This is very similar to private or external NAT/routed connectivity, except that the organization network is not connected to the external network but is completely isolated within the organization.

Some of the typical use cases for organization networks include:

- Consumers who need access to the backhaul portion (intermediate links between the core network and the other networks) of their networking services via a trusted external network can be directly connected to an external network.
- Consumers who need access to the Internet via a non-trusted external network can be NAT/routed connected to the external network.
- Consumers who do not need any access to the public networks can use an internal organization network that is completely isolated.

Figure 36 shows how an organization network acts as a bridge between vApps and external networks.

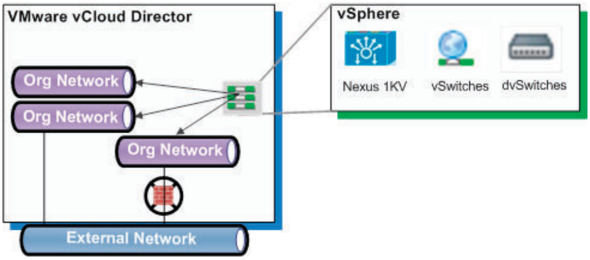

Figure 36. Organization Network

vApp Networks

vApp networks are created by consumers and are contained within vApps. A vApp is a logical entity comprising one or more virtual machines. The main purpose of vApp networks is to connect virtual machines privately within a vApp. Embedding networks inside vApps allows for consistent deployment from vApp templates, as well as the ability to fence multiple copies of vApps within an organization.

vApp networks are also provisioned from a set of preconfigured network pools, which typically map a port group or dvPort group from a vNetwork Standard, vNetwork Distributed, or Cisco Nexus 1000V virtual switch at the vSphere layer.

The vApp Networks can be connected to the organization networks in several different ways:

- Direct connectivity—A vApp network is bridged directly to an organization network, where the deployed virtual machines are directly connected to the organization network.
- Fenced connectivity—A vApp network is NAT/routed to an organization network, where the deployed virtual machines are connected to the organization network via a vShield Edge that provides firewall and/or NAT functionality for security.
- Isolated connectivity—A vApp network is completely isolated from the other vApps and the organization network. This is similar to an isolated organization network except that this is isolated only between the virtual machines in the vApp.

Typical use cases for the vApp networks include:

- Consumers who need to communicate with the virtual machines in other vApps within the same organization and with the same security requirements can be directly connected to the organization network.
- Consumers who need to communicate with the virtual machines in other vApps within the same organization, but with different security requirements, can be NAT/routed connected to the organization network. For example, production vApps and DMZ vApps within the same organization may need to communicate with each other, but through a firewall.
- Consumers who do not need to communicate with the virtual machines in other vApps can be isolated from the organization network.

Figure 37 shows how a vApp Network can be either isolated or connected to the organization network.

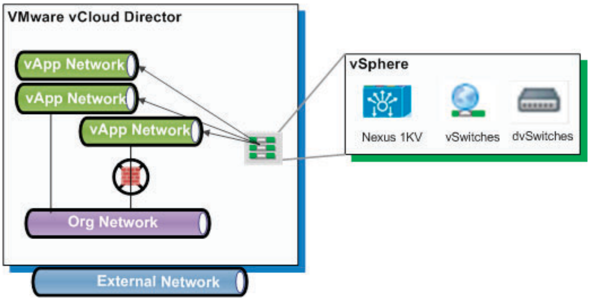

Figure 37. vApp Network

VMware vShield Components

VMware vShield is a suite of security virtual appliances built for VMware vCenter Server integration and is a critical component for protecting vCloud Director environments. The vShield components that provide security services to VMware vCloud are vShield Manager and vShield Edge. Additional vShield products that can be used peripherally are vShield App and vShield Endpoint.

After the perimeter has been secured and policies defined using vShield Edge, it is possible to create interior zoning of the VDC. This can be implemented using additional virtual networks interconnected via vShield Edge appliances or in a network-independent way using vShield App, where virtual machines can be placed on the same virtual wire and Security Groups are configured to carve up the LAN into zones.

vShield Manager

VMware vShield Manager is the centralized network management component that integrates with vCenter Server to create and manage all of the vShield service virtual machines. It is also responsible for pushing out configuration changes to the vShield Edge devices.

VMware vShield Manager is a virtual appliance deployed from an OVF package whose virtual hardware does not need any customization. This appliance should be configured with an IP address, hostname, and DNS server information to integrate it into a vCenter Server.

The vShield Manager user interface leverages the vSphere SDK to display a copy of the vSphere Client inventory panel and includes the Hosts & Clusters and Networks views. Each vShield Manager should be mapped.

vShield Edge

VMware vShield Edge devices provide network security services such as NAT, DHCP, firewall, port forwarding, and IP masquerading to an organization. vShield Edge connects isolated vApp, organization, and external networks by providing common gateway services. The vShield Edge security services can be thought of as an L3 (NAT) device in one appliance.

In the vCloud environment, the most common use case of vShield Edge is to provide multi-tenant cloud environments with perimeter security for organization VDCs.

The following are some of the standard vShield Edge services that are used in VMware vCloud environments:

- Firewall—The vShield Edge firewall provides L3/L4 stateful firewall support. The vShield Edge firewall can match on flow IP address (source and destination), Layer 4 ports, and ICMP packets. The vShield Edge firewall also provides Application Level Gateway (ALG) to support FTP firewalling.
- Network Address Translation (NAT)—NAT provides separate controls for source and destination IP addresses, as well as TCP and UDP port translation.
- Dynamic Host Configuration Protocol (DHCP)—Using DHCP in vShield Edge provides the ability to support IP address pooling and one-to-one static IP address allocation based on the vCenter managed object ID (VMID) and interface ID of the requesting client. You can add a pool of IP ranges for assignment to DHCP clients on the vShield Edge. The following are all DHCP configuration parameters:
  - IP Pools
    - Range of IP addresses—A starting range and an ending range. For example:

      ```
      <IpRange>
          <StartAddress>10.147.58.101</StartAddress>
          <EndAddress>10.147.58.252</EndAddress>
      </IpRange>
      ```
    - Domain name
    - Primary and secondary name servers
    - Lease time—Specify a default lease time in seconds. A maximum lease time in seconds can also be set. The following example allows you to add your virtual machines as DHCP clients. This enables the vShield Edge to bind the MAC address of a virtual machine to a static IP address. The following are all DHCP binding configuration parameters. For example:

      ```
      <DefaultLeaseTime>3600</DefaultLeaseTime>
      <MaxLeaseTime>7200</MaxLeaseTime>
      ```
  - Static bindings (for binding the DHCP address to a VM statically):
    - VM-ID
    - Interface-ID—This specifies the identification of the network interface to be configured.
    - Hostname
    - IP address—This is the internal IP address to be assigned.
    - Domain name primary/secondary name servers
    - Lease time—This is the same as for IP pools.

vShield App

vShield App is a hypervisor-based, vNIC-level application firewall that controls and monitors all flows between virtual machines in a virtual datacenter. Firewall policies can be applied to vCenter containers or to Security Groups, which are custom containers created through the vShield Manager Web interface. Container policies enable the creation of mixed trust zone clusters without requiring an external physical firewall. Classic five-tuple firewall rules are also supported. VMsafe hypervisor introspection allows for complete visibility and control of inter-VM traffic.

Note: Currently, vCloud Director is only integrated with the vShield Edge product. vShield App can be used in conjunction with vCloud Director through careful design of the infrastructure.

vShield Endpoint

VMware vShield Endpoint can be used to offload antivirus functions to a hardened security virtual machine delivered by partners such as McAfee, Symantec, and Trend Micro. Endpoint uses VMsafe APIs to peer into the filesystem to scan and remediate viruses. This removes the need for agents in the guest and prevents antivirus storms from consuming precious CPU cycles during scanning or AV update activities. Offloading antivirus provides enhanced security, as often the first task of malware is to disable AV agents. vShield Endpoint's efficient AV architecture provides antivirus as a service for large-scale cloud environments.

vCloud Network Use Cases

This is how the various vCloud networking components relate to each other:

- A vApp network can be connected to an organization network either directly, fenced, or isolated.
- An organization network can be connected to an external network either directly, fenced, or isolated.
- The three types of network pools can back both vApp and organization networks.
- External networks can provide access to services such as the Internet, VPN/MPLS, and so on.
- vShield Edge devices get created automatically to provide fence/NAT, DHCP, or firewall services between various networks.
- Network pools and external networks are directly mapped to vSphere resources, as shown in Figure 38.

Use Case 1

Organization A would like to use vCloud Director for a software development and QA team. Members should be able to create their own copies of vApps from the vApp templates, but with the flexibility to isolate their own environment without adding any complexity in terms of customization or connectivity.

For this environment, the vCloud administrator has already created an Isolated_OrgA network, which is an organization network that is created as Private-Isolated for Organization A. This network can be backed by any of the network pool types.

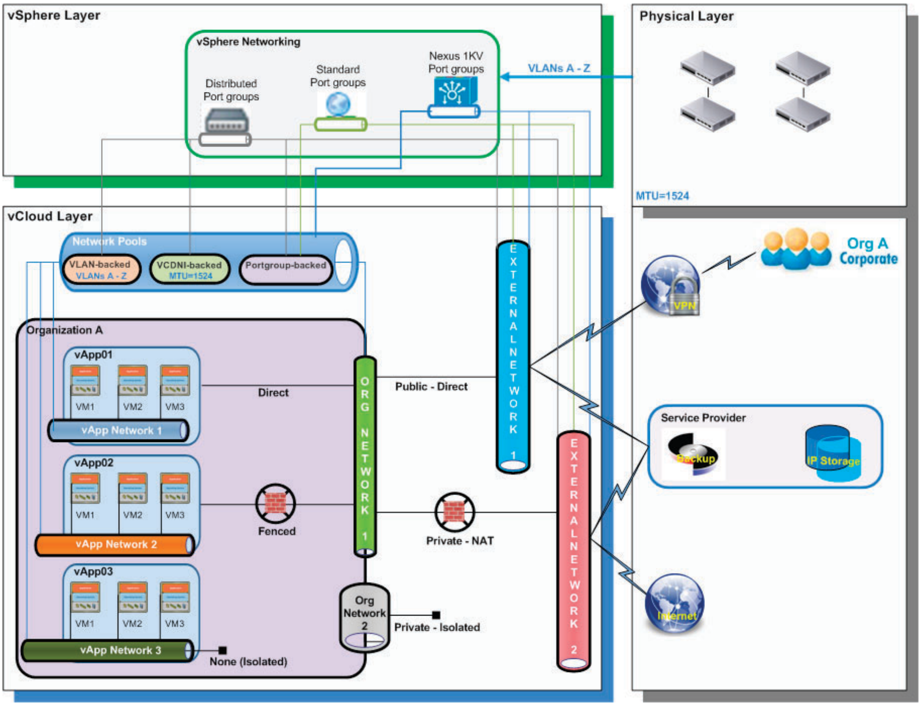

Figure 38. vCloud Network Relationship Mapping

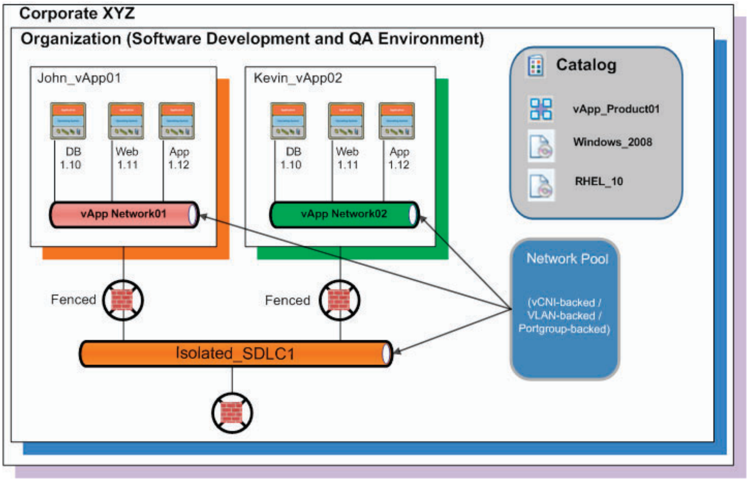

Figure 39. vCloud Network Use Case 1

One product will be developed by Organization A. The development team lead has created the vApp template and given access to the corresponding team members.

Developer John creates vApp01 from the vApp template and deploys it as a fenced connection, resulting in a vShield Edge created between the vApp network of vApp01 and the Isolated_OrgA network. Developer Kevin creates vApp02 from the same vApp template and deploys it as fenced, so another vShield Edge is created between the vApp network of vApp01 and the Isolated_OrgA network. In this case, both vApp01 and vApp02 are completely isolated by the vShield Edges.

vApp01 and vApp02 are deployed from the same vApp template, which has virtual machines with the same IP addresses, MAC addresses, and NetBIOS names. vShield Edge NAT and proxy ARP prevent conflicts by providing external NATed IP addresses. This enables developers and QA testers to quickly create multiple application environments without conflicts or the need for customization.

Use Case 2

Service Provider A wants to provide its Customer A's Tier 2 production environment with secured access to the Internet.

For this environment, the vCloud administrator has already created an external network, CustomerA_Internet01, and an organization network, CustomerA_OrgNet01, which is Private-NAT/routed to the external network. CustomerA_OrgNet01 can be backed by any of the network pool types, and CustomerA_Internet01 is provided with a public IPv4 IP addresses range that is exclusively assigned to this customer.

Customer A is hosting two different vApps, one for the Tier 2 production environment and the other for the DMZ environment, both used to connect to the Internet for both outbound and inbound traffic.

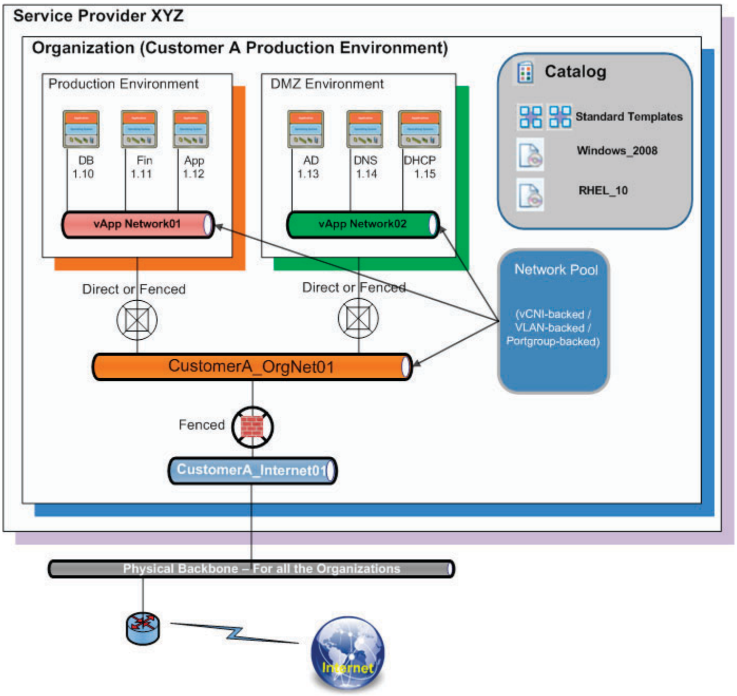

Figure 40. vCloud Network Use Case 2

Within the organization, both production and DMZ can either be fenced or directly connected (not recommended) to communicate with each other, whereas either's connectivity to the outside world via external network CustomerA_Internet01 is always fenced. So a vShield Edge is created that is used as a firewall to protect the traffic coming into the environment, and it can also control the outbound traffic, if needed.

Use Case 3

The third use case can be applied in both production disaster recovery and test lab environments where customers need access to their corporate network via an IPsec VPN or MPLS connection.

Multiprotocol Label Switching (MPLS) is a protocol primarily used in backbone networks to accelerate packet transmission. MPLS VPN and VPLS are protocols that leverage MPLS and enable point-to-multipoint, dynamic routing private networks. Quality of service can be used to shape network traffic, but there is no encryption of the data.

IPsec VPN is a point-to-point tunneling protocol that is usually configured over an Internet connection. All configuration is performed on the customer end. It provides encryption, but QoS is not guaranteed over the Internet.

In this use case, Service Provider XYZ wants to provide Customer B's pre-production environment with access to the VPN. Service Provider XYZ wants to keep the environment secure, but does not want to complicate it with several layers.

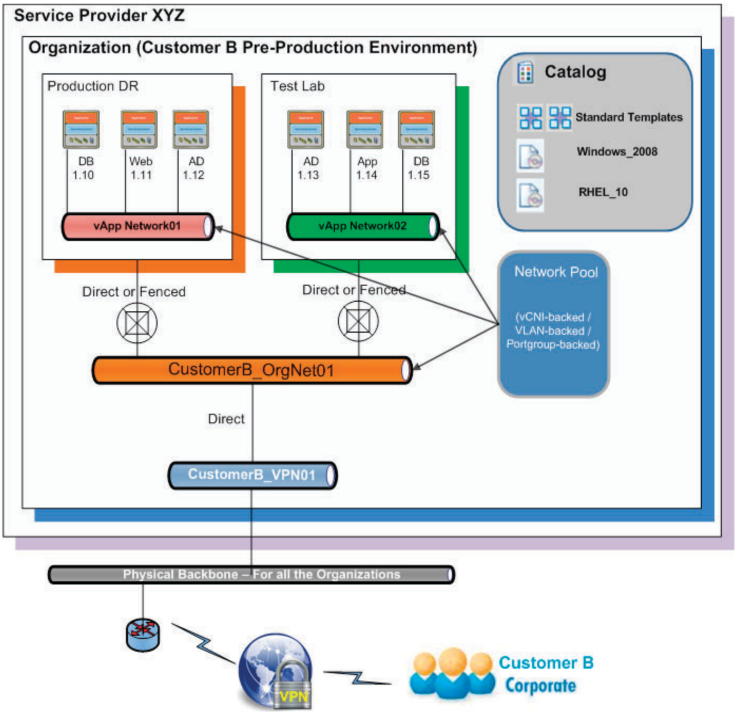

Figure 41. vCloud Network Use Case 3

In this environment, the vCloud administrator has already created an external network called CustomerB_VPN01 and an organization network called CustomerB_OrgNet01, and the organization network is public or directly connected to the external network. CustomerB_OrgNet01 can be backed by any of the previously discussed network types. All of the VPN connectivity has been established outside the cloud environment, and CustomerB_VPN01 is provided with a private IPv4 IP addresses range that is routed to Customer B's corporate environment.

Assume that Customer B is hosting two different vApps, one for the production DR environment and the other for their test lab environment. Both of the environments in this case are acting as Tier 2/3 applications that can be used in a disaster recovery scenario and for the test lab of their corporate production environment.

Within the organization, both production and DMZ can either be fenced or directly connected (not a recommended method) to communicate with each other, whereas either's connectivity to the outside world via external network CustomerA_VPN0101 is always directly connected. A vShield Edge is never created in this case: the traffic flows directly between the production DR and test lab environments to their corporate office.

Summary

VMware vCloud networking provides the flexibility and elasticity to support multi-tenancy. The concept of network isolation through overlay networks is driven by the need to scale beyond the limits of VLANs. Momentum for vCloud Director Network Isolation is building, with partners working to support vCD-NI across a range of network protocols and devices, irrespective of the final frame format. The next generation of cloud networking will likely resemble large-scale malleable L2 networks, with software innovations providing network isolation, simplified management, and automation between virtual and physical environments.

8. VMware vCloud Storage

Cloud computing intensifies the importance of data storage by transferring the responsibility of managing and maintaining the data from the consumer to the provider. End users expect that the output of computational tasks in the cloud will be stored in a secure and reliable manner. The lack of transparency about how data is stored and protected in public clouds is a key concern for most customers.

Storage design for cloud introduces unique challenges due to the heightened need for availability and security. Adding compliance and regulation to the mix may necessitate data encryption, replication, and faster data restoration. Conventional methods for architecting storage solutions cannot be strictly followed because of the self-provisioning nature of clouds. It is important that vApps maintain similar levels of performance and availability when migrated to or created in the cloud. Service level agreements (SLAs) between providers and consumers must be carefully crafted to reflect the capabilities that can be delivered.

Traditionally, storage layout is based on application workload profile and capacity requirements. In vCloud Director, the user self-provisions vApps into virtual datacenters, each with its own pool of storage resources. The physical storage details are abstracted from the user, who cannot specify the destination storage for vApps. Precise control of where vApps reside in the physical infrastructure is removed. This can lead to contention scenarios where intensive vApps access data that resides on the same physical spindles.

The goal is to build out a cloud storage layer that provides scalability, availability, performance, and cost efficiency. For vCloud environments, tiered pools of storage should be created for different workload profiles. This design offers flexibility and performance for the users while allowing for physical design optimization based on the expected workloads.

The following approach is recommended:

1. Perform a current state analysis for storage usage and trend.
2. Define the range of storage SLAs needed and appropriate pricing models.
3. Create multiple tiers of storage based on SLAs, workloads, and cost.
4. Map service tiers to provider VDCs in vCloud Director.
5. Design storage to ensure optimal availability.
6. Ensure that physical storage is modular and scalable.
7. Monitor storage usage and trends through capacity analysis tools.
8. Use storage performance tools to tune vApp storage workloads.

Storage Tier Design

To handle the elasticity and scalability required for a vCloud implementation, a modular tiered storage approach is recommended. The goal is to design for future growth while optimizing for performance. This requires determining the storage differentiation required for each service, then presenting pools of tiered storage to vCloud Director. Examples are shown in Tables 2 and 3, below, and Figure 42, on p. 66.

Item	Considerations
Availability	Disk mirrors: clones, snapshots, local, remote Mirrored fabrics, HBAs, storage array ports Isolated networks for networked storage Components: host, switches, arrays RAID, hot spares, FRUs, CRUs Multiple power feeds Backup generator Replication type: synchronous, asynchronous, journaling BC/DR solution: backup, restore, site failover Application criticality
Performance	RAID type: 1+0, 5, 6, X Disk group configuration Total mirrored cache Read/write cache tuning Storage processor speeds Disk type: SSD, FC, SATA Disk speeds: 15k, 10k, 7200rpm Port speeds (FC): 1Gb, 2Gb, 4Gb, 8Gb, 10Gb Port speeds (Ethernet): 1GbE, 10GbE FC, iSCSI, NFS, FCoE Application workload profiles

Scalability	Array processors Array disk capacity Array cache capacity FC switch port density Datacenter space LUN expansion Clone/snapshot space Array disk group, LUN limits Future storage requirements
Cost	Storage infrastructure Disk energy profiles Tiered storage Management tools People Support
Virtual Datacenter	Storage limit Thin provisioning Allocation Model

Table 2. Tier Considerations

SLA	Service	Cost	RTO	Storage	RAID	Applications
Tier 0	Premium	$$$$$	20 min	SSD, FC	1+0	Exchange, Oracle, SAP, SQL, OLTP
Tier 1	Enterprise	$$$$	1 hour	FC	1+0, 5	Web servers, SharePoint, Active Directory
Tier 2	Professional	$$$	2 hours	iSCSI, NFS	3, 5, X	Custom applications, QA
Tier 3	Basic	$	2 days	NFS	3, 5, X	Dev/Test

Table 3. Storage Tier Example

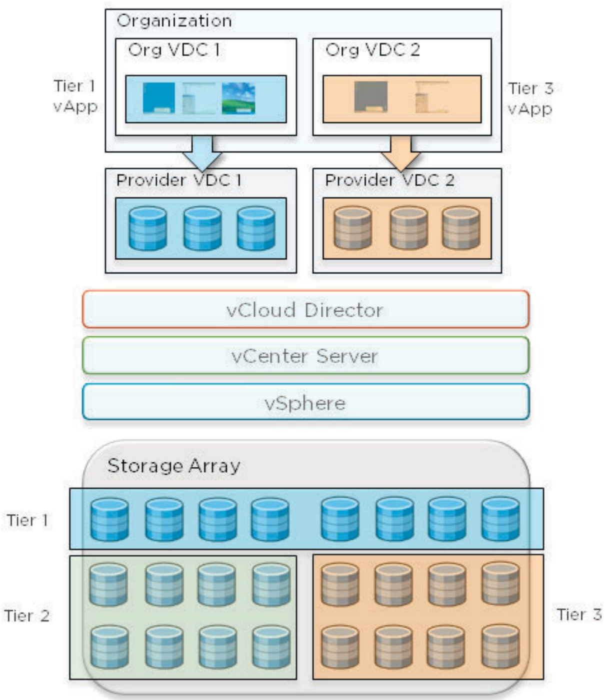

Figure 42. Provider VDC Tiering

Storage Configuration

Shared storage is required in order to deliver the full benefits of the vCloud platform. VMware vCloud implementations can use Fibre Channel (FC), iSCSI, or Network File System (NFS). Raw Device Mapping (RDM) is not supported.

Storage tiering is created by mapping provider VDCs to uniform datastore types. Organization VDCs are carved from each provider VDC and assigned to multiple organizations. Each organization in the vCloud can then draw from multiple tiers of storage resources through the associated organization VDCs. Figure 43 shows the screen used to add datastores.

In this example, uniform 500 GB iSCSI datastores have been added to a provider VDC. More predictable storage performance is provided by not mixing datastore types in the same provider VDC. All organization VDCs that consume resources from this provider VDC share the same set of datastores. When a vApp is deployed, it is placed on the datastore with the greatest available capacity.

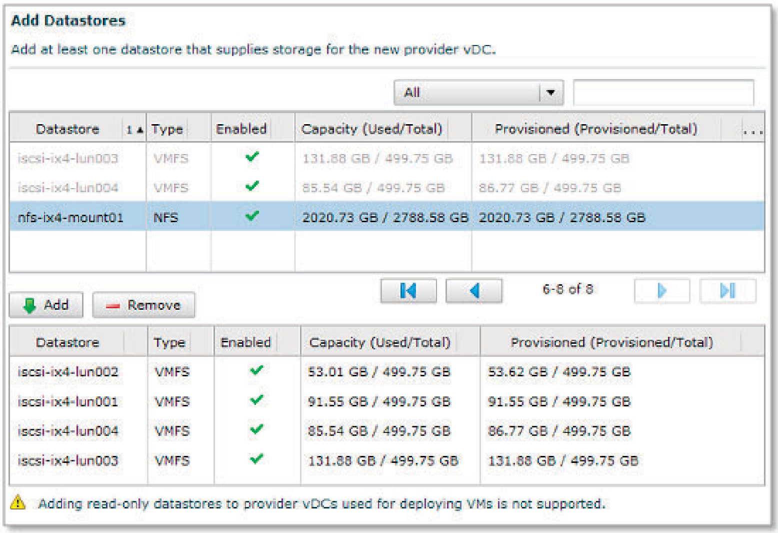

Add Datastores

Add at least one datastore that supplies storage for the new provider vDC.

Datastore 1 ▲	Type	Enabled	Capacity (Used/Total)	Provisioned (Provisioned/Total)
iscsi-ix4-lun003	VMFS	✓	131.88 GB / 499.75 GB	131.88 GB / 499.75 GB
iscsi-ix4-lun004	VMFS	✓	85.54 GB / 499.75 GB	86.77 GB / 499.75 GB
nfs-ix4-mount01	NFS	✓	2020.73 GB / 2788.58 GB	2020.73 GB / 2788.58 GB

Datastore	Type	Enabled	Capacity (Used/Total)	Provisioned (Provisioned/Total)
iscsi-ix4-lun002	VMFS	✓	53.01 GB / 499.75 GB	53.62 GB / 499.75 GB
iscsi-ix4-lun001	VMFS	✓	91.55 GB / 499.75 GB	91.55 GB / 499.75 GB
iscsi-ix4-lun004	VMFS	✓	85.54 GB / 499.75 GB	86.77 GB / 499.75 GB
iscsi-ix4-lun003	VMFS	✓	131.88 GB / 499.75 GB	131.88 GB / 499.75 GB

⚠ Adding read-only datastores to provider vDCs used for deploying VMs is not supported.

Figure 43. Adding vCloud Datastores

Availability

Continuous High Availability must be provided, as customers rely on cloud services for critical business applications. Redundancy is required for all vCloud storage infrastructure components:

- Host bus adapters
- SAN switches (fabrics)
- Storage subsystem (cache, processors, disks)

Multipathing

ESXi hosts can use multipathing for load balancing and path failover. Multipathing is performed at the SCSI mid-layer and not by the guest operating system. The vSphere Native Multipathing Plug-in (NMP) module has the following policies:

- Fixed
- Most Recently Used (MRU)
- Round robin (rr)

The multipathing option is determined by VMware and storage vendor best practices.

Storage vendors have leveraged the vStorage APIs to create MultiPath Plugs (MPP) that provide additional integration for increased performance and availability. Your storage vendor can inform you about MPP availability.

Check the *VMware Compatibility Guide* site to obtain vendor-specific configuration settings.

Performance

Storage performance is a critical requirement for cloud architectures. Though modern storage subsystems attempt to absorb the majority of I/O requests in cache, the back-end physical disk layout must be designed appropriately. Poor application performance in virtualized environments is typically the result of misconfigured storage. Application workload variables include:

- I/O size (KB)
- Sequential or random I/O
- Read/write mix

Response times for large block I/O sizes tend to be slower, but bandwidth is increased. Conversely, small block I/O sizes result in more I/Os per second (IOPS) but reduced bandwidth throughput. Most storage subsystems can handle more reads per second than writes. The read/write mix of applications impacts overall array performance. It is important to understand and predict the types of application workloads to be deployed.

VMware vSphere 4.1 introduced Storage I/O Control (SIOC) to provide QoS for VMFS datastores. Enabling SIOC on all datastores in a cluster prevents virtual machines from monopolizing storage I/O. SIOC does not support Raw Device Mappings (RDM), NFS datastores, or datastores with multiple extents.

VMware vStorage APIs for Array Integration (VAAI) was also introduced with vSphere 4.1 and enables storage-based hardware acceleration. VAAI allows vSphere to pass storage primitives to supported storage arrays, offloading functions such as full copy, block zeroing, and locking. VAAI improves storage task execution times, network traffic utilization, and CPU host utilization during heavy storage operations. VAAI can provide large performance boosts for cloud environments, especially with vApp deployments.

Storage Sizing

Some SLAs may dictate multiple copies of data (clones or snapshots), making it necessary to plan accordingly for the additional storage capacity.

The standard LUN size is based on expected average tenant customer requirements:

- Expected # of VMs per cluster
- Expected VM storage sizing
- Types of VM sizes (small/medium/large)
- Expected # of vApp templates
- Additional space for VM swap, snapshots, log, thin volume growth
- Expected size of catalog media

Catalog

The catalog accelerates deployment of applications by providing a repository of vApp templates and media to end-users. The vCloud Director transfer service is leveraged for moving files in and out of the cloud. Catalogs can be published to all organizations in the cloud or remain private to an organization. Organization administrators can further restrict catalog usage through access control lists.

The catalog provides an opportunity to standardize the various vApps used within the cloud environment. Cloud administrators should consult with organization application owners to determine which vApp templates should be publicly available. The vApp templates can then be built following company standard operating procedures.

Catalog files should be placed on an appropriate tier of storage based on deployment time requirements. The first step is to create a provider VDC that is associated with datastores. Then, for each organization, an organization VDC is associated with this provider VDC to hold catalog files. For multi-cell architectures, an NFS share is required to provide temporary storage for vApp uploads and downloads.

Thin Provisioning

Thin Provisioning is a technology, derived from vSphere, that is used to reduce the storage requirements for virtual machines.

For applications with predictable capacity growth, thin provisioning may provide a more efficient way of allocating capacity. Though datastore capacity utilization can be improved, additional management processes are required. vCenter alarms can be configured to alert when approaching an "out of space" condition so that there is sufficient time to source and provision additional disks.

Thin provisioning can be enabled for virtual datacenters, as shown in Figure 44. Once enabled on a VDC, all vApps deployed to that VDC are thin provisioned. Thin provisioning should not be enabled for VDCs that are designated for workloads with intensive I/O and unpredictable capacity growth.

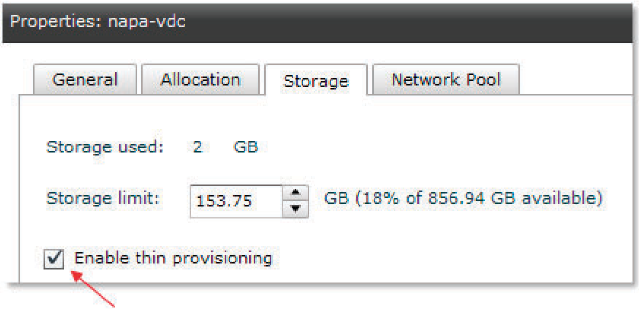

Storage used: 2 GB

Storage limit: 153.75 GB (18% of 856.94 GB available)

☑ Enable thin provisioning

Figure 44. VDC Storage Properties

Monitoring, Management, and Maintenance

Array performance monitoring tools identify application I/O profiles (bandwidth, read/write ratio, average I/O block size, random versus sequential I/O, throughput). Analyzing storage subsystem component usage can help uncover design flaws and imbalances.

Capacity planning tools are necessary for trending and forecasting future storage requirements. Processes are required to rapidly add storage infrastructure as needed.

Scalability

vCloud Director has the ability to scale up to multiple vCenter Servers and vSphere environments. The boundaries of the substrate layer are not represented in vCloud Director, but they still exist and must be accounted for when creating the storage layout. Achieving massive scalability exacerbates existing challenges such as backup and restore, disaster recovery, hardware refreshes, and component availability. Data protection and replication are essential for building reliable, large-scale cloud environments.

Design Considerations

In the absence of specific customer application requirements, Table 4 provides a guiding set of storage design principles for each storage pool tier.

Principle	Description
Availability	Storage subsystems should provide appropriate redundancy for the required service levels at the host, switch, and array. Mirrored fabrics are recommended. RAID configuration is dependent on the designated storage pool tier, which considers: ❖ Read/write ratios ❖ Average I/O size ❖ Sequential/random access ratios
Availability	Use Storage vMotion to enable the migration of storage from one storage platform to another. This allows for non-disruptive storage migrations and service tiering based on need.
Availability	Use multipathing within the virtual storage pool to enable a set of added availability to the current environment. The customer should have multiple paths from each ESXi host to the redundant storage fabric or NFS share.
Manageability	For the purpose of manageability, provide a standard deployment size, as well as a set of I/O profiles. Storage subsystem requirements for replication may impact this value.
Manageability	Through storage vMotion and abstraction of the storage layer, vStorage enables more flexibility in terms of storage lock-in. Storage may be chosen based on required service level, or to mitigate the effects of vendor lock-in by enabling transparent migration. Additional storage tiers can be added as needed.
Manageability	Use consistent LUN sizes. Create one datastore per LUN. Consider VM density, LUN restoration times, and vSphere maximums.
Performance	Storage subsystems should provide adequate performance to handle the aggregate workloads running on the virtual infrastructure. Periodic tuning may be necessary to maintain consistent levels of performance.
Performance	Use the appropriate multipathing policy (MRU, fixed, round robin) based on the array type. Check storage vendor best practices.

Performance	Configure the storage array cache properly. Perform LUN alignment based on storage vendor best practices. I/O load should be spread across multiple storage processors and spindles.
Performance	Configure appropriate queue depth per storage vendor best practices across the entire storage stack.
Recoverability	Make sure there is appropriate connectivity to backup/recovery infrastructure.
Recoverability	Verify that the storage hardware is on the compatibility list for vSphere.
Recoverability	Use the vStorage API with third-party backup software to enable off-host backups at both the file and the block level for storage.
Security	Shared storage should be secure for the virtual infrastructure and its supporting backup infrastructure. This may be accomplished through zoning, VSAN, VLAN, or ACLs.
Security	For Fibre Channel storage, use single initiator zoning to eliminate crosstalk and RSCN disruptions. Make sure that all hosts in a cluster have access to the same set of LUNs for vMotion compatibility. Minimize differences in number of storage paths (no more than four paths per LUN) to improve availability, performance, and manageability.
Security	Design vSphere roles to limit the access to storage by non-essential personnel.
Scalability	Employ monitoring tools to identify when additional storage capacity is required. Storage infrastructure should be capable of bursting when necessary. Thin provisioning on the storage array may be leveraged.
Scalability	vCloud Director maximums correlate with their vSphere equivalents. The maximum number of datastores in a provider VDC is bounded by the maximum number of datastores a host can see (assuming that all hosts in a cluster can access the same set of datastores).

Table 4. Storage Design Guidelines

Summary

Storage considerations change when designing for a cloud environment. Instead of focusing on the application workload requirements, more flexibility is needed, as vApp provisioning is now controlled by end users. Storage is also bound to the virtual datacenter construct, which represents a standard container for a pool of compute and storage resources.

Storage tiering enables the cloud provider to pool storage into separate virtual datacenters to provide differentiated services. Providing multiple service offerings to end users allows for better predictability of where workloads will land. After vApps are placed in the appropriate tier, a combination of vSphere and storage array optimizations can assist with balancing workloads. As cloud becomes more pervasive, VMware, along with storage vendor partners, continues to introduce increasing automation and orchestration within the storage stack.

9. VMware vCloud Director Logging and Monitoring

Logging and monitoring are essential operational components of any mission-critical application deployment. VMware vCloud Director offers both high-level and granular methods for viewing logs.

- The user interface is typically used for a quick overview of what types of events have occurred.
- Extensive logging capabilities provide a more in-depth and detailed view, and logging can be redirected to a central log repository such as a syslog server.

Logs are used to analyze the current state of the environment (monitoring), for auditing, and for troubleshooting purposes.

By default, each vCloud Director cell logs audit messages to the database, where they are retained for 90 days. However, it is recommended that an external syslog server be configured for log retention to avoid any dependencies during troubleshooting of the vCloud environment itself.

Log Files and Locations

The vCloud Director log files are stored in /opt/vmware/cloud-director/logs.

The vCloud Director log configuration file is located in /opt/vmware/cloud-director/etc/ and is called log4j.properties. If any changes are made to the configuration file, the vCloud Director daemon needs to be restarted.

Table 5 lists the log files that are created by vCloud Director, including a description of the content.

Log Name	Description
cell.log	Console output from the vCloud Director cell
diagnostics.log	Cell diagnostics log. This file is empty unless diagnostics logging is enabled in the local logging configuration.
vcloud-container-info.log	Informational log messages from the cell. This log also shows warnings or errors encountered by the cell.
vcloud-container-debug.log	Debug-level log messages from the cell
vcloud-vmware-watchdog.log	Informational log messages from the cell watchdog. It records when the cell crashes, is restarted, and so on.

Table 5. VMware vCloud Director Logs

Currently vCloud Director has six logging levels that can be defined. By default, the log files are set to DEBUG mode, which is sufficient for daily operations, general troubleshooting and auditing. Table 6 lists the various log levels, including a description of the level of granularity.

Logging Level	Description
FATAL	Logs only very severe error events that may cause the application to fail
ERROR	Logs error events that might still allow the application to continue running
WARN	Logs potentially harmful situations and warnings
INFO	Logs informational messages that highlight the progress of the application at a coarse-grained level
DEBUG	Logs informational events that are most useful for debugging an application at a fine-grained level
TRACE	Logs informational events at a level more fine-grained than DEBUG logging

Table 6. Log Levels

For daily use, WARN and INFO log levels may be sufficient. An increase of the log level increases the amount of diagnostic data, so it is recommended to also increase the log file size and the number of logs to keep. To control the logging verbosity for vCloud Director itself, edit $VCLOUD_HOME/etc/log4j.properties or use jconsole to connect and modify the loggers.

For support purposes, it is also possible to collect all logs associated with vCloud Director and package them in a single TGZ file. The vmware-vcsd-support script collects host log information as well as the vCloud Director logs. When the script is run it creates a file in the folder in the following format:

```
vmware-vcsd-support-2010-09-23.25214.tgz
```

The audit log file can be found in $VCLOUD_HOME/logs/vcloud-audit.log.

Log Rotation

Most logs get rotated when the maximum file size is reached (the exception is the Jetty request log, which is rotated daily). In vCloud Director, the default size is 25 MB for the HTTP requests and error logs. The maximum file size and the number of rotated files are defined in log4j.properties, and the settings are defined per cell.

Monitoring

There are multiple components attached to a vCloud offering that, when monitored correctly, can reduce administration overhead and result in a rich end-user experience. There are various aspects to monitoring, and several methods for pooling metrics and attributes into a single monitoring solution. Native vCloud components themselves provide some monitoring capabilities. Significant metrics can also be collected through the product APIs.

Each vCloud Director depends on the following to be operational:

- vCloud Director—dependent on vCloud Director database
- vCenter Chargeback Server—dependent on vCenter Chargeback database
- vCenter Server—dependent on vCenter database
- vShield Manager—to deploy vShield Edge virtual appliances
- VMware ESXi hosts—via vCenter Server
- Directory Service—Microsoft Active Directory
- DNS
- NTP
- Windows and Linux operating systems—hosting these services

VMware vCloud Director

The vCloud Director user interface provides views into the health and availability of the compute, network, and storage resources as well as the communication status of attached vCenter Server and vShield instances (Figure 45).

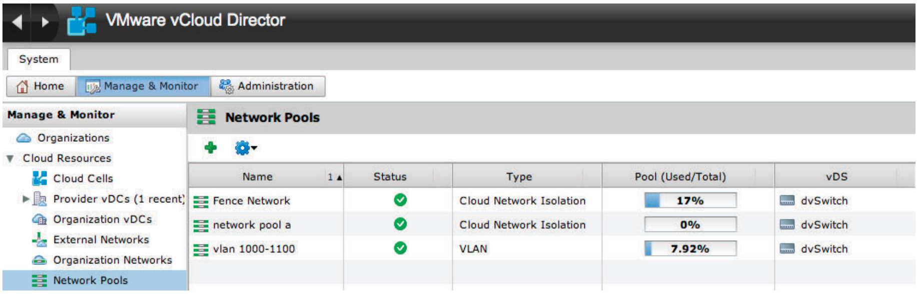

Name	Status	Type	Pool (Used/Total)	vDS
Fence Network	✓	Cloud Network Isolation	17%	dvSwitch
network pool a	✓	Cloud Network Isolation	0%	dvSwitch
vlan 1000-1100	✓	VLAN	7.92%	dvSwitch

Figure 45. Network Pool Availability

VMware vCloud Director primarily enables system administrators to monitor compute resources, but it also enables monitoring of network pools, storage, ESXi host status, and vCenter status (see, e.g., Figure 46).

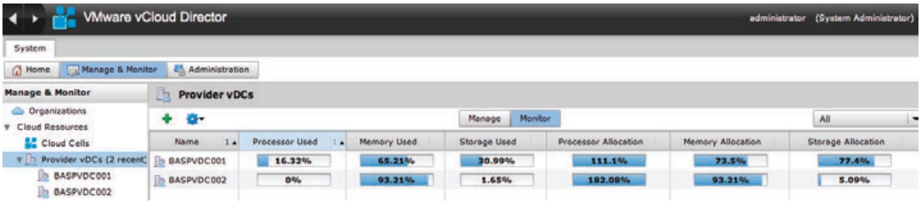

Name	Processor Used	Memory Used	Storage Used	Processor Allocation	Memory Allocation	Storage Allocation
BASPVDC001	16.32%	65.21%	30.99%	111.1%	73.5%	77.4%
BASPVDC002	0%	93.21%	1.65%	182.08%	93.21%	5.09%

Figure 46. Provider VDC Usage

A significant part of the vCloud Director monitoring framework is made available outside the user interface through use of JMX integration, guest level metrics, vCloud and vCenter APIs, and Syslog scrubbing.

Starting with the user experience, HTTP response times and codes for the user interface, API, and VMware Remote Console can all be captured and tracked against SLAs.

Exposing MBeans allows more visibility into the health of the vCloud Director instances and their relationship to the underlying vSphere layer through vCenter Server. Examples of metrics that can be utilized for monitoring include:

- Number of sessions over a specified time period
- REST API round-trip time; number of tasks, events, and operations for a specified vCenter instance
- vCenter connection state
- Frequency with which VPXD has been reset, either manually or through the user interface

Expanding on what is available from integrating MBeans, use of the vCloud API also allows the administrator to pull metrics displayed through the vCloud interface and correlate them over time. CPU, memory, network, and storage usage (allocated and actual) for a provider virtual datacenter, vCloud Director instance, or cell is easily extracted and charted. Number of organizations, provider virtual datacenters, organization virtual datacenters, and virtual machine counts can easily be categorized and integrated into a monitoring solution or custom portal.

VMware vCenter Chargeback

The vCenter Chargeback user interface provides the Chargeback administrator with a summary view into system health. vCenter Chargeback Servers, LDAP, Mail, vCloud Director and vShield data collectors, vCenter Servers, and associated databases are all monitored for availability. Figure 47 shows the System Health summary view.

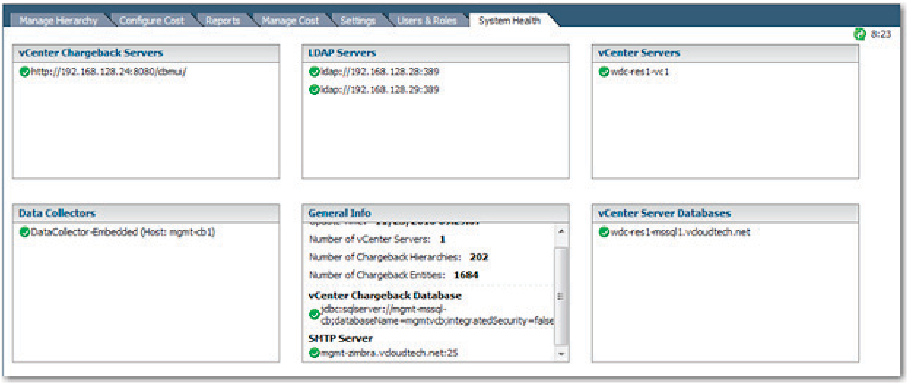

vCenter Chargeback Servers
✓ http://192.168.128.24:8080/cbmui/

LDAP Servers
✓ ldap://192.168.128.28:389
✓ ldap://192.168.128.29:389

vCenter Servers
✓ wdc-res1-vc1

Data Collectors
✓ DataCollector-Embedded (Host: mgmt-cb1)

General Info
Number of vCenter Servers: **1**
Number of Chargeback Hierarchies: **202**
Number of Chargeback Entities: **1684**
vCenter Chargeback Database
✓ jdbc:sqlserver://mgmt-mssql-cb;databaseName=mgmtvcb;integratedSecurity=false
SMTP Server
✓ mgmt-zimbra.vcloudtech.net:25

vCenter Server Databases
✓ wdc-res1-mssql1.vcloudtech.net

Figure 47. vCenter Chargeback System Health

Similar content can be extracted through use of the vCenter Chargeback API at http://[vCenter-Chargeback-FQDN]/vCenter-CB/api/systemHealth. This is useful when vCenter Chargeback components such as the vCloud Director, vShield Manager, vCenter Server data collectors, and load balancers are distributed remotely when scaling vCenter Chargeback deployments and when deploying a complete monitoring solution.

vCenter Chargeback components, monitored through the UI and API, are only part of what needs to be tracked. Tomcat and Apache processes also need to be monitored for health to make sure they are receiving enough compute resources based on the application load. The same analysis can also be applied to the data collectors and load balancers regardless of whether they are running local to the vCenter Chargeback Server or running remotely.

Windows, vCenter Chargeback, and Tomcat logs can also be linked to a syslog repository when scrubbing for events.

SNMP is not available as a monitoring solution from the vCenter Chargeback Server.

VMware vShield Manager

To establish a secure, multi-tenant environment, vShield Edge enforces user-defined firewall rules and provides certain basic services such as DHCP, NAT, VPN, and load balancing. The virtual appliance should be deployed at the edge of the datacenter, but the services and firewall rules apply throughout. Edge virtual appliances are managed by the vShield Manager, the centralized management interface for vShield products. The vShield Manager provides a RESTful API that can be used to configure an Edge appliance and retrieve traffic statistics. Because agents and VMware Tools cannot be installed on vShield Edge appliances and because SNMP is not available in version 1.0, monitoring of the vShield Edge appliances is only available through the vShield Manager user interface, the vCloud Director user interface, the vCenter Server, or through use of the vShield REST API.

Figure 48 shows a vShield deployment through vShield Manager.

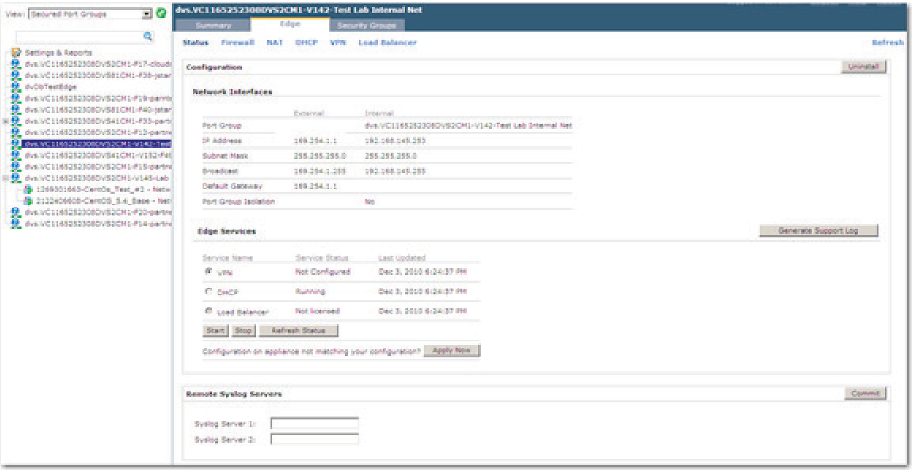

Network Interfaces

	External	Internal
Port Group		dvs.VC1165252308DVS2CM1-V142-Test Lab Internal Net
IP Address	169.254.1.1	192.168.145.253
Subnet Mask	255.255.255.0	255.255.255.0
Broadcast	169.254.1.255	192.168.145.255
Default Gateway	169.254.1.1	
Port Group Isolation		No

Edge Services

Service Name	Service Status	Last Updated
● VPN	Not Configured	Dec 3, 2010 6:24:37 PM
○ DHCP	Running	Dec 3, 2010 6:24:37 PM
○ Load Balancer	Not licensed	Dec 3, 2010 6:24:37 PM

Configuration on appliance not matching your configuration?

Remote Syslog Servers

Syslog Server 1:

Syslog Server 2:

Figure 48. vShield Edge Deployment through vShield Manager

To monitor a vShield Manager attached to a vCenter instance:

1. Send an ICMP echo to the vShield Manager management address.
2. Send an HTTP GET to port 80 or 443 and monitor response time as response code.
3. Check vSphere SDK tool status.
4. Determine whether there are REST calls for any functions.

To monitor the vShield Edge attached to an organization VDC:

1. Send an ICMP echo to external interfaces.
2. Send a UDP probe to the external interface or an alias on port 500 when the VPN is configured.
3. Send an HTTP GET if the load balancer service is running.
4. Review the external syslog (via internal or external interfaces) to scrub changes to the Edge virtual appliance (NAT and firewall rules, external IP address assignment).
5. Poll traffic statistics via the REST API and check the health of the Edge. The *vShield API Programming Guide* (http://www.vmware.com/pdf/vshield_41_api.pdf) provides details on the traffic schema.

Most of these objects can be monitored through the vCloud Director UI, but ultimately should be monitored through a centralized monitoring tool such as VMware Hyperic HQ Enterprise or Zenoss Cloud Monitoring.

These vCloud Director objects are important to monitor, and many are also exposed through MBeans. Many third-party monitoring solutions can read these metrics and statistics. vCloud Director is an important part of the end-to-end solution, but it is also part of a complete chain. Each object (service) in this chain should be monitored, preferably by the same tool in such a way that the overall SLA can be monitored.

10. VMware vCloud API

Cloud computing depends on the coordination of many complex technologies. For cloud computing to be truly ubiquitous, APIs to cloud interfaces must be standardized to allow for platform independence. Standardization avoids forcing developers to write applications that comply with multiple API specifications. Having a common set of APIs that provide access into the cloud is central to building a rich ecosystem of developers.

As the market is still in its infancy, numerous companies have submitted cloud APIs for consideration. API development must allow for extensibility to ensure that future cloud technologies can be accommodated. Advances in cloud computing will occur through standardization, accelerating adoption and usability.

What Is VMware vCloud API?

The VMware vCloud API supplies an interface for providing and consuming resources in the cloud. It delivers a different plane of abstraction from the vSphere/VIM API in that it is focused on use of virtual resources while hiding the underlying physical infrastructure. The vCloud API is designed for open standards, ensuring compatibility between cloud platforms that implement the API, such as vCloud Director.

The API is classified as follows based on the function and sphere of operation:

- User API—Contains operations to examine organizations and virtual datacenters and provides a way to manage and create vApps. The User API can create and manage vApp networks, vApp templates, and media. User APIs cannot be used to create an organization, VDC, organization network, or catalog.
- Admin API—Contains operations to manage and create organizations, VDCs, organization networks, and authorization entities such as roles, rights, users, and groups. It can be used to create, delete, or modify catalogs.
- Extension API—Is VMware-specific and is used to manage entities such as provider VDCs, network pools, vCenter Servers and their hosts.

The vCloud API is based on REST, XML Schema, and OVF. Representational State Transfer (REST) is a style of software architecture design that loosely couples services, is highly scalable, and uses the HTTP protocol. REST relies on the inherent properties of hypermedia to create and modify the state of an object that is accessible at a URL. VMware vCenter Chargeback and vShield Manager also provide REST-based APIs.

Language Bindings for vCloud API

Because the APIs are based on REST, they are not dependent on any programming language. Developers can consume the vCloud service using raw XML or by creating language bindings to the REST API calls.

To simplify the programming needs of the end user, VMware released the vCloud SDK for these APIs in Java, C#, and PHP programming languages. The vCloud SDK provides a set of libraries and numerous sample applications. The libraries hide the complexity of REST, HTTP communication, XML parsing, validation against XSDs, and other details.

Features and Benefits

- Availability of REST resources of vCloud API in Java, C#, and PHP
- High fidelity to vCloud API resource models
- Simple and clean design to help understanding and predictability
- No new client-side object model
- Resources associated with operations meaningful to clients
- REST verbs and URL semantics are hidden

vCloud API clients and servers communicate over HTTP, exchanging representations of vCloud objects. These representations take the form of XML elements in the body of the HTTP request and response. HTTP GET requests are used to retrieve the current representation of an object, HTTP POST and PUT requests are used to create or modify an object, and HTTP DELETE requests are typically used to delete an object.

Resources

The VMware vCloud API Community, http://communities.vmware.com/community/developer/forums/vcloudapi, contains the necessary resources to get started with the vCloud API. Sample code is available for each language-specific SDK.

Design Considerations

For design considerations, the following rules will maximize the benefits of vCloud API:

- Use the language SDK with which the customer has in-house expertise.
- Follow the chosen language's best practices.
- Use the vCloud SDK for backward compatibility.
- Bear in mind that because vCloud Director 1.0 does not support callbacks, the user has to poll the outcome periodically for success or failure before moving on to the next step in the workflow.
- A modular design helps in scaling and augmenting the functionality as more features become available in the vCloud API.
- Avoid hard-coding to minimize the amount of time required for QA.

Summary

The vCloud API provides control and flexibility to service providers and enterprise customers who want to provide differentiated services to their end customers. The vCloud API can be included as a part of an orchestration solution, which in turn can be a component of workload lifecycle management.

For example, an end user in an enterprise wants to deploy a workload to meet seasonal demand. The end user may log into a Workload Order Service and enter a few values such as the kind of virtual workloads to deploy, the duration of use, and the cost center of the business unit. The orchestration workflow can send this request to IT for approval and final deployment. At the end of the session, when the workload has served its purpose, it can be un-deployed, deleted, or archived to complete the lifecycle.

11. vCenter Chargeback

Elasticity of computing resources combined with Pay-As-You-Go pricing creates the illusion of unlimited resources on demand. Consumers can self-provision infrastructure resources and use them in a noncontiguous manner over variable periods of time. A virtual machine may be used for ten hours one day and be powered off the next day. For public and private clouds, this requires the ability to track consumption metrics and associate appropriate pricing schemes.

Public cloud providers leverage economies of scale, building out massive datacenters with commodity hardware to offer computing services to consumers at competitive prices. Pricing of services is a key aspect by which service providers can differentiate their cloud offerings. This requires a metering solution that can be easily implemented, offers flexibility when building reports, and can integrate with existing billing systems.

For enterprises, *chargeback* traditionally refers to the mechanism in which costs for hardware and software are charged back to the appropriate department or business unit. Virtualization requires additional granularity in order to track the numerous assets that could exist within hosts. Resource metering can be used to provide *showback* for enterprises hesitant to move to a full chargeback model. Showback highlights usage of the allocated infrastructure without attributing any costs. Ultimately, the ability to link resource consumption to costs of services helps influence user behavior, especially in environments where infrastructure resources can be rapidly self-provisioned.

VMware vCenter Chargeback provides the metering capability to measure, analyze, and report on resources used in private and public cloud environments. Cloud providers can configure and associate various cost models with vCloud Director entities. This cost transparency allows cloud providers to validate and adjust financial models based on the demand for resources.

Architecture

The main components for vCenter Chargeback include:

- vCenter Chargeback Server
- vCenter Chargeback database
- vCenter Chargeback data collectors
- vCenter Chargeback Web interface
- vCenter Chargeback API

Figure 49 shows the logical layout of components.

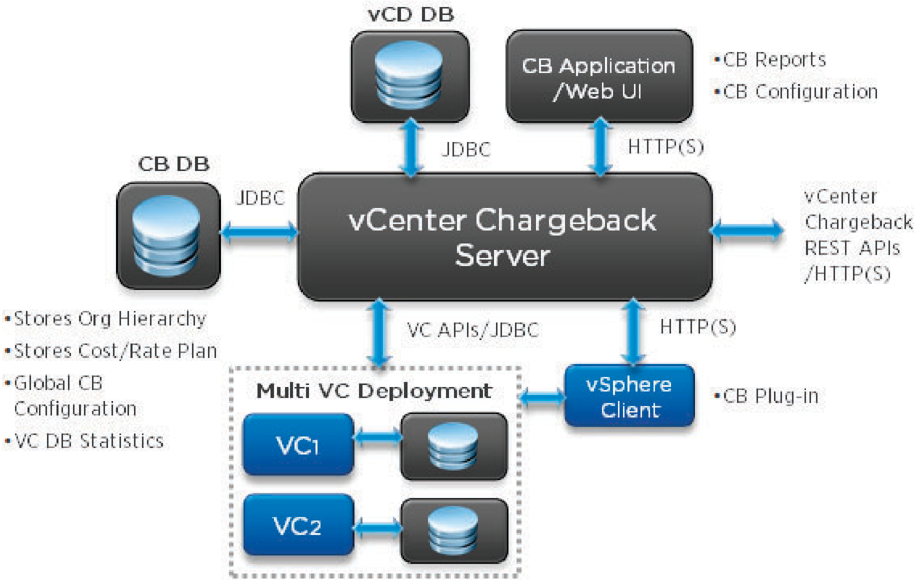

Figure 49. vCenter Chargeback Component Logical Layout

The vCenter Chargeback Server runs the Web interface, load balancer, and data collector services. Multiple instances can be clustered together to provide additional performance and availability. All information is kept in the Chargeback database, which stores organization hierarchies, cost/rate plans, and global configuration data. vCenter Chargeback runs on its own schedule independent of vCloud Director, samples information at intervals, and aggregates the data at other intervals.

vCenter Chargeback integration with vCloud Director is handled through two new data collectors (see Figure 50):

- ❖ The vCloud data collector connects to the vCloud Director database to retrieve resource allocation units for vCloud entities. Chargeback hierarchies for each organization are generated automatically.
- ❖ The VSM data collector connects to vShield Manager and gathers information on external network resources used by vCloud tenants.

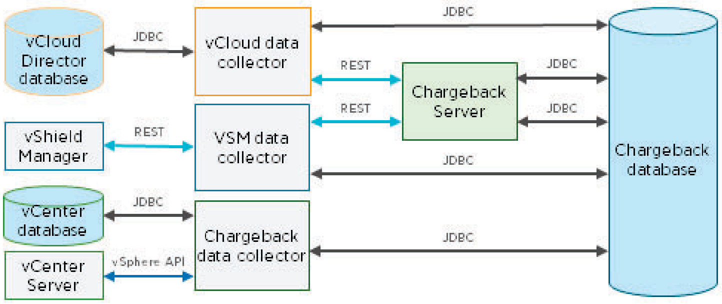

Figure 50. vCenter Chargeback Integration with vCloud Director

The Chargeback data collector connects to the vCenter Server database to gather virtual machine details. To reflect chargeback events for Pay-As-You-Go virtual datacenters, each vCenter Server that manages cloud resources must be added to vCenter Chargeback. It is possible to install multiple instances of data collectors, which will operate in an active/passive mode.

The vCenter Chargeback API is based on REST and provides a programming interface for Chargeback functionality. This includes hierarchy management, cost configuration, report generation, and report exporting.

vCloud Hierarchy

vCenter Chargeback organizes vCloud Director and vCenter Server entities through hierarchies. The vCloud data collector detects when new organizations are created in vCloud Director and automatically adds corresponding hierarchies to vCenter Chargeback. The mapping between organization and hierarchy is shown in Figure 51.

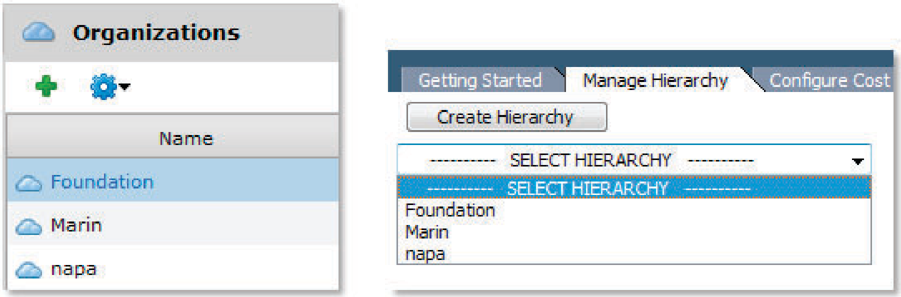

Figure 51. Organization to Hierarchy Mapping

Synchronization of entities between vCloud Director and Chargeback is automatic. If vApps, networks, or catalog files are deleted in vCloud Director, they are also deleted in the Chargeback hierarchy. If organizations are deleted from vCloud Director, the chargeback data remains, but with a "DELETED_" prefix and timestamp suffix added to the hierarchy name. If organization VDCs are deleted in vCloud Director, the chargeback data also remains. vCloud Director entities should never be deleted from Chargeback.

Configuration of the vCloud Director entities in Chargeback is determined by the allocation pool type of the virtual datacenters. Allocation units should not be modified from within vCenter Chargeback as they will be overwritten by changes in vCloud Director.

Cost Definition

Service providers create usage pricing models based on capital and operating costs factored with the profit margin desired for each cloud offering. A financial analysis is needed to derive the appropriate pricing model. For example, it may be required to inject the cost of power/cooling into the billing report or to charge a one-time set-up fee for new customers. Enterprises do the same if performing chargeback, but cost transparency, not profit, is usually the intended goal.

The VMware vCloud Datacenter and vCloud Express programs provide prescriptive guidance on how to implement prospective pricing models with vCloud Director and vCenter Chargeback.

The definition of standardized units of consumption provides a mechanism for scaling cloud infrastructure. This involves defining standard bundles (small/medium/large/custom) for different virtual datacenter allocation model types. Specific cost models for each consumption unit can be created and associated through vCenter Chargeback.

Cost Models

Cost models associate base rates, billing policies, and fixed costs to define how vCloud entities are metered and billed. Base rates define costs for resource metrics measured for a specific time interval. Metering and cost calculation are tightly coupled. vCenter Chargeback includes the following default cost models for vCloud Director allocation models (these cost models have no defined base rates):

- VMware vCloud Director Allocation Pool
- VMware vCloud Director Networks
- VMware vCloud Director Pay-As-You-Go—Fixed Charging
- VMware vCloud Director Pay-As-You-Go—Resource Based Charging
- VMware vCloud Director Reservation Pool

Choose the appropriate default cost model to apply to the vCloud Director entity instead of using the default Chargeback cost model. The default cost model has base rates defined for usage metrics that are not tracked for vCloud entities and should be avoided. Overage charges are not possible on vCloud hierarchies, since allocation units will be set by the vCloud Data Collector to match the VDC/vApp configuration.

Use the Pay-As-You-Go—Fixed Charging cost model if you intend to charge for configured resources regardless of power state. Use the Pay-As-You-Go—Resource Based Charging cost model if you intend to charge based on the amount of vCPU, memory, and storage allocated to the vApp while the vApp is powered on.

There are several options when determining how to bill for cloud resources:

- Create multiple cost models for defined units of consumption (small, medium, large) along with the appropriate monthly fixed costs. Assign fixed costs to the appropriate virtual datacenters.
- Use a single cost model per allocation pool VDC type and assign base rates. This assumes that the desired charge will scale linearly with different VDC and VM sizes.
- Configure VM instance fixed cost matrices that define the instance attributes and the associated fixed cost. This only applies to Pay-As-You-Go allocation type virtual datacenters.

Establishing policies for a finite set of cost models which are relatively static reduces complexity. An environment with a dynamic number of cost models would be difficult to manage.

Billing Policies

Billing policy expressions dictate how costs are associated with a particular computing resource metric. vCenter Chargeback comes with the following default VMware vCloud Director billing policies:

- Allocation Pool
- Networks
- Pay-As-You-Go Fixed-Base Charging
- Pay-As-You-Go Resource-Based Charging
- Reservation Pool

Other default billing policies, such as utilization-based, are specific to vCenter Server. Either the default vCloud Director billing policies can be used, or new billing policies can be defined.

Tables 7–11 show the expressions used to construct each of the default vCloud billing policies.

- "VM Power On/Off" refers to whether billing takes into account VM power states.
- If the resource attribute is set to Allocation, costs are calculated based on the allocation unit configured for that particular resource. If the attribute is set to Usage, vCenter Chargeback looks at the utilization for the resource, performs a summation for the values recorded during the billing interval, and then calculates the cost.

Resource	VM Power On/Off	MAX Operator	Attribute(s)
External network transfer	No	No	Usage
External network receive	No	No	Usage
All other resources	No	No	Allocation

Table 7. Allocation Pool Billing Policy

Resource	VM Power On/Off	MAX Operator	Attribute(s)
External network transfer	No	No	Usage
External network receive	No	No	Usage
All other resources	No	No	Allocation

Table 8. Networks Billing Policy

Resource	VM Power On/Off	MAX Operator	Attribute(s)
Fixed cost	No	No	Include
External network transfer	No	No	Usage
External network receive	No	No	Usage
All other resources	No	No	Allocation

Table 9. Pay-As-You-Go Fixed-Based Billing Policy

Resource	**VM Power On/Off**	**MAX Operator**	**Attribute(s)**
vCPU	Yes	No	Allocation
Memory	Yes	No	Allocation
External network transfer	No	No	Usage
External network receive	No	No	Usage
All other resources	No	No	Allocation

Table 10. Pay-As–You-Go Resource-Based Billing Policy

Resource	**VM Power On/Off**	**MAX Operator**	**Attribute(s)**
External network transfer	No	No	Usage
External network receive	No	No	Usage
All other resources	No	No	Allocation

Table 11. Reservation Pool Billing Policy

Cost Configuration

After cost models are defined, additional details concerning how vCloud entities are billed can be configured. Cost templates containing rate factors and fixed costs can be applied at any level in the hierarchy. Rate factors allow for a multiple to be applied to a computing resource during cost calculation. This helps reduce the number of cost models needed.

VM Instance

VM Instance costing allows the definition of a fixed cost matrix based on a hard bundle of vCPU (count) and memory (MB). VM Instance matrices are linked with a cost model and consist of the VDC selection criteria, a fixed cost table, and a default fixed cost. Selection criteria can be based on name pattern matching, custom attribute matching, or no criteria. VM Instance uses a stepping function**,** where the VM charge steps up to the next instance size. Costs are applied only for the duration when a VM is powered on and are not prorated. VM Instance costing can only be applied to Pay-As-You-Go virtual datacenters.

Fixed costs are applied based on the following algorithm:

1. Get the ordered list of all cost matrices for a cost model.
2. Get the list of all Pay-As-You-Go VDCs.
3. For each VDC, find the first cost matrix that matches the sort criteria.
4. Get the list of all VMs in the VDC.
5. Determine vCPU count and memory allocation of each VM.
6. Determine if the VM is new or if any allocations have changed. If no, process the next VM.
7. Look up the appropriate fixed cost based on vCPU count and memory allocation.
8. vCPU count is matched first, followed by memory size.

9. If vCPU is matched and no match is found for memory, the next higher match for memory is used.
10. If no appropriate fixed cost is found, the default fixed cost is applied.
11. Apply the new fixed cost on the VM.
12. Repeat steps 1 to 11 for the next cost model.

Multiple cost matrices can be created for a cost model and ordered by selection criteria priority. In this case, the criteria must be unique for each matrix. Figure 52 shows an example of a VM instance matrix.

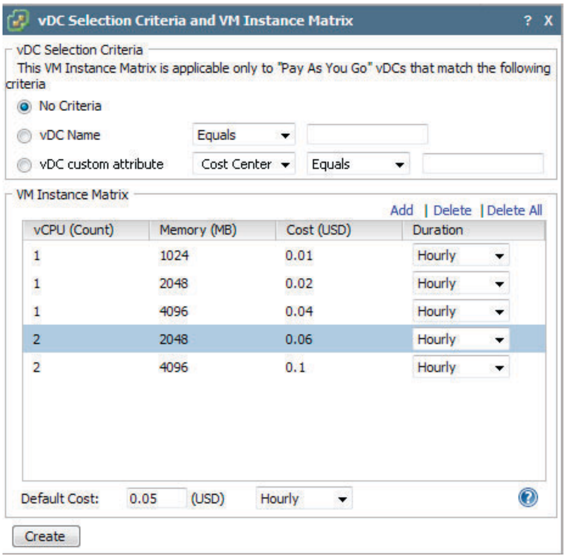

vDC Selection Criteria and VM Instance Matrix

vDC Selection Criteria

This VM Instance Matrix is applicable only to "Pay As You Go" vDCs that match the following criteria

● No Criteria

○ vDC Name — Equals

○ vDC custom attribute — Cost Center — Equals

VM Instance Matrix

vCPU (Count)	Memory (MB)	Cost (USD)	Duration
1	1024	0.01	Hourly
1	2048	0.02	Hourly
1	4096	0.04	Hourly
2	2048	0.06	Hourly
2	4096	0.1	Hourly

Default Cost: 0.05 (USD) Hourly

Figure 52. VM Instance Matrix

Reporting

Cost reports, usage reports, and cost comparison reports can be generated on-demand or on a scheduled basis. vCenter Chargeback reports provide visibility into resource usage, the cost applied to each entity, and overall billing total. To include multiple entities in one report, use Ctrl-Shift to multi-select the entities desired. All reports can be exported to Microsoft Word, PDF, or CSV format. The Chargeback REST API can be used to programmatically export XML reports.

User Management

Initially, a superuser administrator is created during the installation process. This superuser then adds additional users through the Web interface or Active Directory integration. Role-based access control is used to control the privileges assigned to each user or group. Synchronicity between vCloud Director and Chargeback is maintained by preventing users from modifying and deleting entities from vCloud hierarchies.

Integration

Because vCenter Chargeback does not perform any actual billing functions, integration with internal or external billing systems is required for public cloud implementations. The Chargeback APIs allow application developers to integrate vCenter Chargeback with existing billing systems.

Refer to the *vCloud Billing Integration Guide* for more details.

Design Considerations

- Define standard units of consumption and allocation models.
- Determine pricing through financial analysis.
- Map pricing scheme through Chargeback vCloud cost models.
- Use one Chargeback instance to manage one vCloud Director database.
- Cluster multiple Chargeback servers for higher Web interface availability.
- Install multiple data collectors for active/passive redundancy.
- The Chargeback Server is most heavily utilized during report generation.

Refer to the *Using vCenter Chargeback with vCloud Director Tech Note* for more details.

Summary

Resource metering restricts the behavior of the unbounded resource provisioning allowed through the Pay-As-You-Go model. If charging for resources is not the immediate goal, the collected information can be used for capacity management and future planning. VMware vCenter Chargeback delivers the metering capability to drive accountability in public and private cloud environments. The ability to define cost models, billing policies, and fixed costs allows for greater flexibility in defining how resources are measured and billed.

12. Applications in the Cloud

Applications are hosted within a VM in the cloud in the same way that they are hosted in a VM within vSphere. This combination of VM and application is packaged using the Open Virtualization Format (OVF) to allow portability between platforms.

Licensing requirements for compliance with software vendor guidelines must be considered and handled. Specific requirements need to be addressed with product vendors.

Application development for the cloud is supported by the VMware vFabric Cloud Application Platform. VMware ThinApp technology provides support for simplified application management, including consistency in application upgrades across an enterprise. Finally, migrations to and from the cloud platform must be addressed.

OVF

OVF is the industry standard Open Virtualization Format, which allows for portability between different virtualization platforms. This standard enables portability between different cloud platforms as well. It enables migration into the cloud by exporting a VM to OVF format and then importing the VM into the cloud layer as a vApp construct. The cloud layer cannot operate using OVF format files, but typically provides additional features and capabilities.

vApps, VMs, and Images

Virtual machines are the traditional way of representing virtualized systems, and they continue to be used within a vCloud.

Both existing and new applications can be packaged into vApps that can be deployed in private, public, or hybrid clouds. A vApp consists of multiple virtual machines that are packaged and maintained as a single entity within the vCloud. A vApp provides a standard way for describing operational policies for an application—enabling automatic execution by the cloud OS and facilitating migration between private and public clouds. vApps include metadata that enables all related VMs to be treated as one object.

Software Licensing Considerations

Cloud deployments from a service catalog require licensing for applications and operating systems. Most software application pricing is tied in some way to the physical realm, such as servers, processors, cores, or users. Because moving to cloud computing typically removes these constraints, an understanding of alternative licensing models is necessary.

Public Cloud Licensing

Public cloud providers offering Infrastructure as a Service tend to have licensing agreements in place with the major software manufacturers such as Microsoft and Red Hat. Providers frequently include the cost of the base OS as part of their VM cloud fees. An organization with existing copies of an OS may be able to get a reduction in cost as they transition their infrastructure to their public cloud provider.

Private Cloud Licensing

Application licensing in the private cloud, in particular, is a new and evolving area. Business units utilizing the vCloud Service Catalog to requisition new applications may incur additional OS software licensing costs as well. Consequently, it is imperative that IT monitor all application use and understand how that usage relates to the individual software manufacturer licensing requirements.

Reducing Licensing Costs

Some manufacturers offer significant cost reductions for using their products in a virtualized environment. These savings carry over to a private cloud if architected effectively, but IT must be familiar with both the specific software versions and the requirements in order to realize the savings. IT also must manage the private cloud to ensure that licensing benefits are optimized. For example, licensing Microsoft SQL Server Enterprise for the underlying CPUs of a vSphere host enables up to four instances of SQL Server (Standard or Enterprise) to run on that specific host. This licensing policy can result in a significant cost saving, but even more beneficial is Microsoft Server License Mobility, which allows running instances to be migrated as needed across hosts as long as the number of CPUs on the target host does not exceed the number of CPU licenses.

Preventing Increased Licensing Costs

Incorrectly architecting a private cloud can lead to increased software licensing costs. For example,Oracle Database software is licensed based on the number of physical CPUs and cores. Because processors tend to be under-utilized in a physical architecture, organizations can potentially reduce licensing fees by pooling workloads as part of a private cloud. However, care must be taken to restrict the Oracle VMs to a dedicated cluster (DRS/HA); otherwise, the roaming applications can result in very high licensing costs.

As another example, consider Microsoft Exchange. Because Exchange is licensed per server, increasing Exchange mailbox density per vSphere host by carving up Exchange into multiple virtual machines can result in substantial infrastructure savings, but IT administrators should be aware of potential increased licensing costs as they build out their private cloud environments.

VMware vFabric Cloud Application Platform

The VMware vFabric Cloud Application Platform is an application platform for both virtual and cloud deployments that consists of technologies such as vFabric tc Server, vFabric Hyperic, vFabric GemFire, vFabric Enterprise Ready Server (ERS), and vFabric RabbitMQ. The open source Spring Development Framework and a set of application services support vFabric. These services include application server, global data management, cloud-ready messaging, and application performance management.

vFabric tc Server

vFabric tc Server is the runtime server that supports the VMware vFabric Cloud Application Platform. tc Server is a hardened and optimized applicaton server for production workloads. It enables a quick build and deployment of applications for cloud deployment.

vFabric Hyperic

vFabric Hyperic provides application management within the vFabric Cloud Application Platform. It monitors infrastructure changes through automatic discovery and provides information about availability, performance, utilization, log events, and changes. Visibility includes physical, virtual, and cloud layers.

vFabric GemFire

vFabric GemFire provides data management for applications. It uses replication, partitioning, data-aware routing, and continuous querying to deliver real-time and scalable data access within a cloud environment.

vFabric Enterprise Ready Server

vFabric Enterprise Ready Server (ERS) is both the Web server and the load-balancing component of the VMware vFabric Cloud Application Platform. This is the most widely distributed Apache Web server package and provides the characteristics to support the platform.

vFabric RabbitMQ

RabbitMQ is part of the VMware vFabric Cloud Application Platform. It provides inter-system messaging to support cloud-based applications. The technology is open-sourced under the Mozilla public license and supports open, standard protocols rather than APIs.

End-User Computing

VMware uses the category “end-user computing” to cover desktop solutions and technologies. This category includes VMware View™, VMware ThinApp, and Zimbra. Detailed descriptions of the ThinApp and VMware View technologies can be found in

Foundation for Cloud Computing with VMware vSphere 4. Figure 53 shows how end-user computing is supported by VMware View and ThinApp.

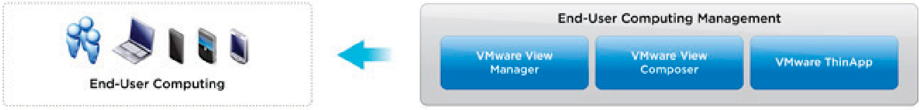

Figure 53. VMware End-User Computing

VMware View

VMware View provides desktop virtualization and can potentially be used as a cloud solution for desktop services. This technology has been in use for several years and is a proven Desktop-as-a-Service solution within a cloud.

ThinApp

VMware ThinApp enables provisioning of applications abstracted from the underlying operating system. The application is packaged as a standard EXE or MSI file, including registry settings or drivers. The application is executed in a sandbox environment. ThinApp applications can either be stored locally in the operating system or placed on a network share for simplified use and deployments. ThinApp packages reduce complexity and dependencies during development and update management.

Zimbra

Zimbra is a cloud SaaS collaborative technology. The Zimbra Collaboration Suite Appliance is deployed as a virtual machine and includes email, calendaring, contact management, and document-sharing features.

Migrations to and from the Cloud

Migrations to and from the cloud may include physical-to-cloud, virtual-to-cloud, and cloud-to-cloud migrations. Migrations can be performed manually or by using tools such as VMware Cloud Connector or Racemi DynaCenter.

Physical-to-Cloud Migration (P2C)

Migrating a physical machine to a cloud-based workload requires several steps:

1. Migrate to a virtual machine.
2. Save in OVF format.
3. Import the OVF format image into vCloud Director. An organization administrator or provider administrator can import the image for placement in the appropriate organization VDC in the vCloud deployment.

Virtual-to-Cloud Migration (V2C)

The steps for V2C are similar to P2C, but without the first step of migrating to a virtual machine.

Cloud-to-Cloud Migration (C2C)

With hybrid clouds becoming increasingly common, there is a need to simplify migrations between clouds. Manual processes have existed, but there are technologies available to assist in C2C migration. Two such technologies are VMware vCloud Connector and Racemi DynaCenter.

Manual Migration

Migrating from one cloud to another requires first exporting the cloud workload to OVF format. The OVF image can then be moved to another cloud platform and imported for deployment. Because OVF format does not include metadata of the vApp image, the final stage of the migration requires setup of the appropriate attributes such as SLA-related information and network/security configuration.

VMware vCloud Connector

VMware vCloud Connector (vCC) is part of the VMware vCloud product suite. It supports a hybrid cloud solution by using the vSphere Client interface as a single-pane-of-glass to view and manage hybrid clouds. The views include vSphere, as well as the private vCloud and public vCloud Datacenter service, both of which are based on vCloud Director technology. VMware vSphere customers can use their vSphere Client to view resources such as VMs, vApps, and templates across hybrid clouds, copy resources between clouds, perform basic power operations on resources, and access vApp consoles running in private or public vCloud instances. Support for administrators and end users is included.

Customers can use vCC for the following hybrid cloud use cases:

- Extend datacenter to public cloud—Customers can view resources across hybrid clouds, migrate workloads from datacenter to public vClouds to free up datacenter resources, deploy and operate elastic workloads in public vClouds, and bring workloads from public vClouds back to the datacenter when necessary.
- Populate private vCloud with workloads from existing vSphere environments.
- Migrate VMs and templates between vSphere instances (extends beyond Linked mode, which only provides resource visibility across multiple vSphere instances).

vCloud Connector is a client/server-based application. The server installs as a virtual appliance in vSphere, while its client is displayed via the vSphere Client connecting to that specific instance of vSphere. The user can use the vSphere Client to see and migrate VMs, vApps, and templates across hybrid clouds while the server handles all the crosscloud communication and interaction. The vCloud Connector architecture is shown in Figure 54, and the vCloud Connector interface is shown in Figure 55.

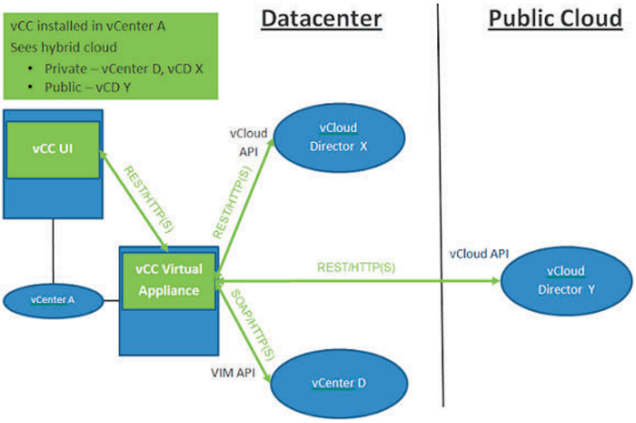

Figure 54. VMware vCloud Connector Architecture

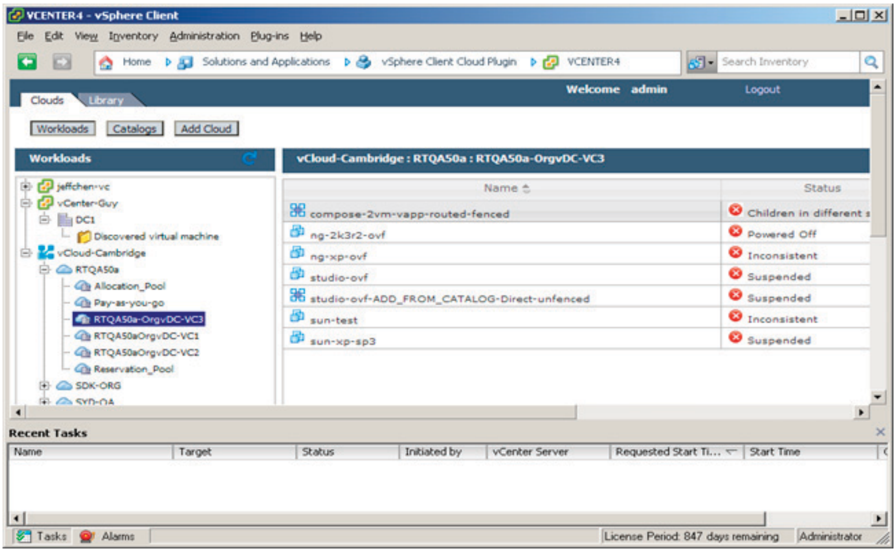

vCloud-Cambridge : RTQA50a : RTQA50a-OrgvDC-VC3

Name	Status
compose-2vm-vapp-routed-fenced	Children in different s
ng-2k3r2-ovf	Powered Off
ng-xp-ovf	Inconsistent
studio-ovf	Suspended
studio-ovf-ADD_FROM_CATALOG-Direct-unfenced	Suspended
sun-test	Inconsistent
sun-xp-sp3	Suspended

Name	Target	Status	Initiated by	vCenter Server	Requested Start Ti...	Start Time

Figure 55. VMware vCloud Connector Interface

Racemi Dynacenter

Racemi's server imaging and cloud migration technology enables enterprises to quickly capture, clone, and migrate their server images anywhere within their datacenter or to a private or public cloud. The Racemi automation software uses image-based provisioning to migrate server images regardless of the underlying operating system, application software, or configuration, or even the physical, virtual, or cloud infrastructure. And it can translate hypervisors, automatically moving server images among VMware, Xen, and Hyper-V in any combination or direction. The Racemi technology is available worldwide as part of CA Technologies Server Automation Suite and is available for licensed integration directly from Racemi.

Summary

There are a number of areas to consider when running applications within the cloud. Standardization, development tools, and migration tools support cloud deployments, cloud application development, and cloud operations.

13. Scalability

Scalability for a cloud can be defined as the ability to increase resources as demand grows. If 100 similar VMs can run on a 10 ESX host cluster, then 200 similar VMs should be able to run on a similar 20 ESX host cluster. This approach is straightforward when the requirements are the same for all customers. However, in an elastic cloud environment where the demand varies quickly and vastly, it becomes necessary to provide both a scale-up and a scale-out solution that can easily be adopted for various components without affecting existing customers.

Scalability of the vCloud environment depends on the vCloud components, the underlying virtual infrastructure, and supporting databases.

Scalability Considerations

Scalability solutions for vCloud Director must include components such as vSphere, Oracle Database, vCloud Director cells, load balancers, VMware vShield, and vCenter Chargeback.

It is advisable to separate management workloads from cloud resource workloads by creating a management cluster that hosts all services needed to stand up in the cloud environment. ESXi hosts designated to run cloud workloads can be grouped in a resource group, which is a group of hosts managed by a single vCenter Server.

VMware vSphere

Consider all of the datacenter, compute, storage and network maximums of vSphere for the management cluster and resource groups. The most critical variables include maximum number of hosts per cluster, datacenter, and vCenter; maximum number of registered and powered-on virtual machines; maximum number of LUNs per host; maximum volume size; and both vNetwork Standard and vNetwork Distributed Switch limitations.

Oracle

As of vCloud Director 1.0, Oracle is the only supported database. Oracle Real Application Cluster (RAC) can be used to provide scalability and availability for the vCloud Director database. Based on the load in the environment, calculate the number of nodes, the database size, connections, and so on. For example, Oracle supports up to 1,000 open connections to the database, but the default installation size is only 170. If the environment demands it, it is important to change this value to the required value or even to the maximum.

vCloud Director

We recommend starting the environment with at least two vCloud Director cells for redundancy, and at least one for every vCenter Server configured in the resource group. Generally, all of the vCloud Director cells are stateless and all of the information stays on the database, making scaling up and down within the cells a straightforward process. Variables such as the number of organizations and VDCs supported, the number of users and vApps or vApp templates supported per organization, and the number of supported concurrent and active consoles all need to be taken into consideration when sizing a vCloud Director environment.

Not all vSphere maximums are supported by vCloud Director. For example, the vSphere maximums for the number of powered-on virtual machines or datastores are not supported by vCloud Director. Table 12 lists the supported vCloud maximums. See the VMware Web site for updated information specific to your version of the vCloud technologies.

Category	Maximum Number
Virtual machines	10,000
ESXi hosts	1,000
vCenter Servers	10
Users	5,000

Table 12. VMware vCloud Director Maximums

Load Balancers

VMware vCloud Director does not provide a built-in load-balancing mechanism. Consider implementing a proxy-based load balancer solution that can be scaled on demand.

vShield Manager

There must be one vShield Manager per vCenter Server configured in the resource group.

vCenter Chargeback

For the vCenter Chargeback System, the scalability is a 1:1 ratio between the vCloud Director database and the vCenter Chargeback database, but there can be multiple vCloud cells, Chargeback data collectors, and Chargeback Servers. It is equally important to consider the maximum number of supported Chargeback entities, concurrent users supported across the system, maximum number of pages per report, concurrent reports, and so on.

Example Scenario

Even knowing all of the different components and their scalability ratios, there is no single approach that can be used to provide scalability for the vCloud. The following is an example scenario to show how scalability can be implemented. Though a service provider is used in this example, it also applies to enterprise customers.

Service Provider X created a vCloud Director solution, and for the initial phase only a Pay-As-You-Go model was offered to its customers. The Pay-As-You-Go model was used with a provider VDC named pVCD A-1 (Resource Group A) that consisted of a cluster with 8 ESX hosts (4 CPUs, 32GB RAM, 4 NICs). The load was unpredictable and memory constraints resulted as the load of the customers started to increase. The Service Provider X vCloud administrators could alleviate the memory constraints by increasing (scaling up) the memory on the existing 8 ESX hosts to 64GB or 128GB. And as the system demanded more resources, they could scale out by adding up to 32 ESX hosts in that cluster.

With the load of existing customers in mind, along with the growth of additional customers joining the cloud, the vCloud administrators decided to stop adding new customers to pVDC A-1 after reaching a 60% utilization rate. Instead, they created a second provider, VDC pVDC A-2 (Resource Group A), into which new customers were added. This scalable solution leaves up to 40% of growth capacity for existing customers in pVDC A-1 and completely avoids issues such as migrating existing customers off pVDC A-1 onto new pVDCs to give them more capacity.

The vCloud administrators of Service Provider Y reached the maximum number of VMs per vCenter and were at a point where they needed to add a second vCenter Server. They created pVDC B-1 (Resource Group B) and started adding existing customers to it, but soon realized that the customers were encountering network issues. This problem resulted because vCloud Director NI-backed network pools cannot be spanned across multiple vCenter Servers, due to Layer 2 connectivity limitations. Therefore customer organizations should be located behind the same vCenter Server.

Summary

There are many factors to keep in mind while designing scalability for the vCloud solution, but having both a scale-up and a scale-out approach for both the management cluster and resource groups is the key to a successful implementation. For future growth, capacity management should be performed and the environment scaled in standard increments of compute, storage, and network resources.

14. vCloud Security

The acceleration and adoption of Cloud Computing will dramatically increase the different types of security threats to the resources and information that organizations possess. We must not only understand and interpret these threats, but also provide appropriate compensating controls and procedures that can be employed to offset them, and do so in a way that is consistent with the nature and strategies of cloud computing that we are trying to safeguard.

Beyond identifying the possible threats and applying mitigation that is specific to each environment and design, it is important to continually update security policy documentation with the latest information. A comprehensive security architecture should be designed that includes all aspects of virtualization and physical security. VMware vCloud Director was designed with security and multi-tenancy in mind.

The infrastructure—including vCloud cells, the vCloud software stack, the vShield Manager, Chargeback, and databases—is faced with many possible threats. These threats can be mitigated through the *defense-in-depth strategy* of vCloud Director. This strategy leverages some of the best known and proven technologies to help provide the necessary protection for cloud infrastructure.

vCloud Director end-to-end security best practices, from architectural approach to implementation and operation, provide guidelines for securing the vCloud environment.

Developing a vCloud Security Strategy

VMware vCloud Director has been described as an abstraction from vSphere layer, and the same applies to vCloud security. All three layers (physical, virtual, and cloud) have various security concerns that need to be identified and mitigated. This requires identifying the existing security policies for both physical and virtual infrastructures, and then appending the vCloud security to include both internal and external threats.

Information security—including confidentiality, integrity, and availability (CIA)—defines the basic building blocks of any good security initiative that is designed to circumvent security threats, attacks, and vulnerabilities. When evaluating and mitigating risks, use the CIA information security framework (see Figure 56, next page).

Security attacks are generally aimed at confidentiality and integrity. They potentially give an attacker access to data and the ability to affect availability, and can result in denial-of-service (DoS) attacks. Web sites, including Yahoo and eBay, have been shut down due to *persistent DoS attacks*, which are similar to a DoS attack, except the attack is executed from multiple, distributed agent IP devices. Other cloud vulnerabilities include unsecured network interfaces and networks, excess privileges, misconfigurations or poor management, and unpatched vulnerabilities.

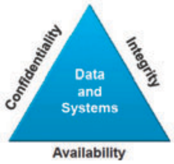

Figure 56. CIA Information Security Framework

Risk assessors must prioritize risks to determine the confidentiality level of the data as well as the likelihood and consequences of the data ending up in the wrong hands. The effectiveness of current controls must also be evaluated during this phase. A risk ratings mechanism is then used to recommend appropriate, cost-effective security controls as part of a plan to mitigate risks and comply with internal and external policies and requirements. The resulting security strategy includes methods for preventing, detecting, and responding to security events. The following are the basic high-level steps for developing a cloud security strategy:

1. Assess the security risks.
2. Develop a security strategy.
3. Implement your security controls.
4. Monitor your security environment.
5. Analyze and further update your security strategy.

VMware is committed to providing the best possible security for the cloud and offers online information about security and compliance. It is important to stay up-to-date with the latest information.

Figure 57 (next page) illustrates how the security layers function.

vSphere Security Functions

When implementing virtualization technology, organizations must ensure that they can continue to maintain a secure environment and meet their compliance obligations. This requires evaluating risks that might affect protected information and mitigating those risks through risk-appropriate standards, processes, and best practices. When evaluating and mitigating risks, use the CIA information security framework.

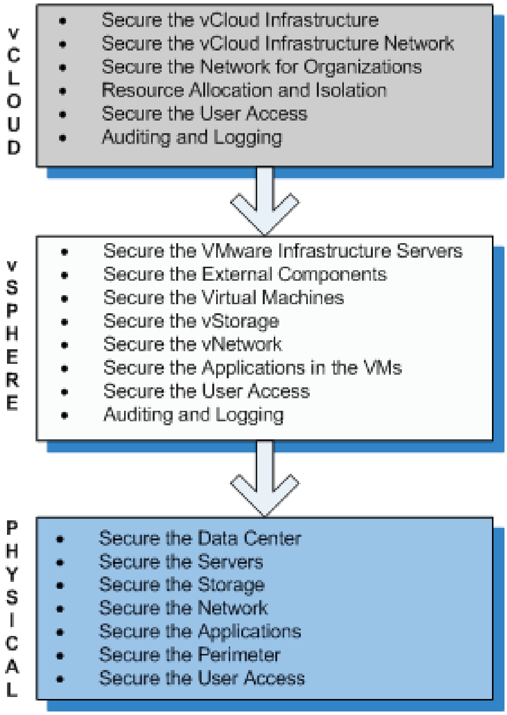

Figure 57. VMware vCloud Security Layers

vCloud Security Functions

Working with vCloud Director Users, Roles, Privileges, and Permissions

The vCloud Director security model enables assigning users, groups, roles, privileges, and permissions at different levels and to different objects within your cloud infrastructure. Properly configuring and assigning these roles, rights, and permissions enforces accountability. This defense-in-depth security architecture gives cloud tenants the privileges they need to be self-servicing and productive without compromising security.

Role-Based Access Control (RBAC) has been implemented in vCloud Director to address the problem of delegating system administration responsibility. Prior to RBAC, administrators were provided with the superuser (root) password, or setuid scripts were written with group permissions granted to specific administrators. Granting root permissions to a large population of administrators is dangerous because the root user can intentionally or unintentionally damage a server, as well as breach security. Custom setuid scripts are insecure and difficult to manage. The concept behind RBAC is to define subsets of administrative functions and assign them to particular users.

After a user is authenticated in vCloud Director, access roles govern the permitted actions. While various types of access control are available, vCloud Director is only concerned with role-based access control. RBAC is also known as *nondiscretionary access control.* It enables granular control over user permissions.

RBAC components include:

- *group*—A collection of accounts with rights to log in and perform other tasks within vCloud Director
- *role*—A collection of rights that a user or group is allowed to perform
- *right*—An allowed action or function within a role. A right allows a user or group to perform a certain and specific task.

Roles

Assigning vCloud Director roles is an essential task when building access control policies into the vCloud Director infrastructure. The vCloud Director roles are different from the ESXi host roles and vCenter server roles. VMware vCloud Director roles cannot be shared with other vCloud environments.

Auditing and Logging

VMware vCloud Director is made up of a combination of technologies that consist of several components that work together to deliver cloud services for tenants. As vCloud Director is not a single product, more than one auditing and logging mechanism is provided. A wealth of diverse information is logged by the vCloud Director components which can be leveraged to meet goals of transparency, compliance, and various types of analysis.

To set up auditing and logging:

1. Identity the relevant components that provide the auditing and logging functionality or capabilities, including:
   - VMware vCloud Director
   - vCenter Server
   - vSphere (ESXi)
   - vShield Manager (VSM)
   - vShield Edge (VSE)
2. Collect and store logs from each component and make sure that there is a clear retention policy.
3. Map events to compliance controls.
4. Utilize redundant collection. Use more than one collection node, either by sending all logs to two nodes or by having two nodes configured as a cluster.
5. Classify logs as provider infrastructure or customer organization-related so customer logs can be provided to the appropriate group.
6. Restrict access to logs to appropriate parties.

Tenants should be able to retrieve logs relevant to their organization based on:

- Date and time range
- Source component (for example, vCloud Director, VSM)
- Event type (for example, VM instantiated via an API)

Log Files

Log files from servers, networking devices, databases, and other subsystems can provide an indication of the cause behind a security event, but analyzing them can be tedious and time-consuming. Using an event correlation system such as CS-MARS, as well as other tools such as Swatch, enables better log entry analysis.

VMware vCloud Director and TLSv1/SSL

Using TLSv1/SSL capabilities with the vCloud Director significantly enhances the security of operations you perform.

The explosive growth of the Internet has created a need to securely protect sensitive communications sent over this open network. The TLSv1/SSL protocol has become a de facto standard for cryptographic protection of Web HTTP traffic. The IETF Transport Layer Security working group is also using TLSv1/SSL as a base for their standards efforts. In short, TLSv1/SSL aims to provide Internet client and server applications with a practical, widely applicable connection-oriented communications security mechanism.

Transport Layer Security (TLSv1)/Secure Sockets Layer Architecture

The TLSv1/SSL protocol exchange and connection is a complex process fully described in other publications. But in the context of how vCloud Director works, public key cryptography (based on certificates and RSA keys) is used in the early phases of the handshake to establish that the server is trusted and to securely exchange a symmetric key between client and server. After the TLSv1/SSL handshake is complete, traffic is encrypted using the symmetric key with an algorithm such as AES.

Note: Certificates are not used at this point. Nothing about this process is specific to vCloud Director; it is all part of the TLS standard.

VMware vCloud Director and HTTPS

vCloud Director uses digital certificates to enable secure communication based on TLSv1/SSL. Properly deployed, TLSv1/SSL provides privacy of communication (by using encryption) and also allows clients to verify the authenticity of the server with which they are communicating. Server authentication is necessary to prevent a *man-in-the-middle* attack, where the client is induced to connect to a server that is spoofing or proxying for the server with which it is supposed to be communicating.

For processing Web requests from a browser or a REST API client, vCloud Director supports version 1.0 of the TLS standard (TLSv1.0), as well as version 3 of the older SSL protocol (SSLv3.0). When vCloud Director acts as a client (e.g., communicating with a vCenter server), it uses TLSv1.0 only. VMware vCloud Director restricts the cipher suites used for encryption to those providing strong security (AES and DES3).

Verification of the server depends on having a certificate installed on the server that is signed by a recognized certificate authority (CA) and matches the host to which the client is connecting.

The vCloud Director keystore needs to be in the JCEKS format as VMware does not support other formats at this time. The keystore is a database of private keys and their associated certificates or certificate chains, which authenticate the corresponding public keys.

Note: Attacks against TLSv1/SSL can be launched if a targeted system supports weak ciphers. In such a situation, an attacker might be able to manipulate the system so that encrypted data is downgraded or even deciphered to achieve access to sensitive data.

VMware vCloud Director Security and vShield

The VMware vShield family of security solutions provides network security for vSphere and vCloud environments.

VMware vShield IPsec VPN

System administrators need to know the basics of cryptographic theories and IPsec protocols so that they can understand the key management and packet processing performed by the IPsec protocols. Those who want to dive deep into the mathematical core of cryptography can do so by referring to the appropriate RFCs, but from a practical point of view, it is not necessary to go into all the details and design decisions.

The IPsec protocols can be split into two main categories: packet handling and trust relationship management. The operating system kernel itself usually does packet handling because it requires speed, efficiency, and low latency, which are easier to offer at the low-level processing of the kernel. The trust relationship management is not as time-sensitive because it only happens at the start and at the refresh intervals of an IPsec connection (about once an hour). It also requires a lot of very complex code which is much better implemented outside the kernel as a regular program running on the computer.

vShield Edge IPsec VPN functionality is not available with the base Edge license that comes with vCloud Director. It is a separately licensed feature.

vShield IPsec VPN Configuration Example

Figure 58 shows a configuration example for a basic point-to-point IPsec VPN connection between a vShield Edge and a physical gateway VPN on the other end.

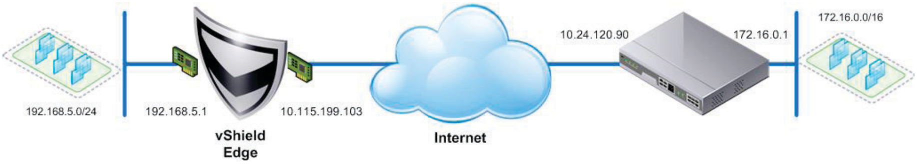

Figure 58. Point-to-Point IPsec VPN Connection between vShield Edge and Physical Gateway VPN

The vShield Edge connects the internal network, 192.168.5.0/24, to the Internet. The vShield Edge interfaces are configured as follows:

- External Interface = 10.115.199.103
- Internal Interface = 192.168.5.1

The gateway on the remote side connects the 172.16.0.0/16 internal network to the Internet. Its interfaces are configured as follows:

- External Interface = 10.24.120.90/24
- Internal Interface = 172.16.0.1/16

IKE Phase 1 and Phase 2 Parameters

The IKE Phase 1 parameters used by the vShield Edge are:

- Main mode
- TripleDES/AES (configurable)
- SHA-1
- MODP group 2 (1024 bits)
- Pre-shared secret (configurable)
- SA lifetime of 28800 seconds (eight hours) with no Kbytes rekeying
- ISAKMP aggressive mode disabled

The IKE Phase 2 parameters supported by vShield Edge are:

- TripleDES/AES (will match the Phase 1 setting)
- SHA-1
- ESP tunnel mode
- MODP group 2 (1024 bits)
- Perfect forward secrecy for rekeying
- SA lifetime of 3600 seconds (one hour) with no Kbytes rekeying
- Selectors for all IP protocols, all ports, between the two networks, using IPv4 subnets

VMware vShield Edge

To establish a secure multi-tenant environment, vShield Edge enforces user-defined firewall rules and provides certain basic services such as DHCP, NAT, VPN, and load balancing.

VMware vShield Edge is a virtual appliance that is deployed at the edge of the organization VDC so that the services and firewall rules apply throughout the organization. The vShield Edge virtual appliances are managed by the vShield Manager server, a vCenter Server-wide focal point for managing all vShield products. The vShield Manager provides a RESTful API with which you can configure an Edge appliance and retrieve traffic statistics. Because agents and VMware Tools cannot be installed on vShield Edge appliances and because SNMP is not available in version 1.0, monitoring of vShield Edge appliances is only available through the vShield Manager UI, the vCloud Director UI, the vCenter Server UI, or the vShield REST API.

- To monitor a vShield Manager attached to a vCenter instance, send an HTTP GET to the vShield Edge virtual machine and monitor response time as response code.
- To monitor the vShield Edge attached to an organization VDC:
  1. Use the external syslog (via internal or external interfaces) to scrub changes to the Edge virtual appliance (NAT and firewall rules, external IP address assignment).
  2. You can use the REST API to poll traffic statistics and check the health of the Edge. The *vShield API Programming Guide* provides details about the traffic schema.

Summary

Security and multi-tenancy are key factors in the design of VMware vCloud Director. Role-Based Access Control, auditing and logging, TLSv1/SSL, and the vShield family of security solutions comprise a defense-in-depth security architecture designed to help secure the vCloud environment. The recommendations in this chapter will help you deploy applications on the vCloud Director platform with confidence in their security.

15. Business Resiliency

Business resiliency is required in both the cloud and the virtualization layers. The separation of management and cloud resource workloads proposed in the Scalability chapter results in the management clusters and the resource groups being aligned with business requirements. Both business continuity and disaster recovery components must be considered when architecting for availability of the infrastructure and the workloads in the cloud, and for failover and recovery.

- *Business continuity* focuses on the prevention of failure and reduction in downtime.
- *Disaster recovery* (DR) focuses on the recovery of systems and infrastructure after a disaster. A disaster can be defined as failure of hardware or software. In a cloud environment, this includes the virtualization layer, the cloud layer, and the cloud workloads. There are two areas that require disaster recovery: the management cluster and the resource groups cluster. There are different approaches and technologies supported for each area.

Redundancy

As in the physical infrastructure and the virtual infrastructure, redundancy in components is critical. Consider implementing redundancy for compute, storage, and networking.

Using N+1 or N+2 server architectures provides greater availability. N+1 covers the loss of one server. N+2 covers planned downtime scenarios where one server has already failed.

Redundant communication links for storage and networking include storage multipathing and network link redundancy. For storage, a good option is multiple storage processors or storage heads. Multiple switches and other communications equipment, outside of the hosts, are needed within the infrastructure to which the vCloud environment is connected. For both storage and networking, multiple interface cards on the hosts can be useful.

Higher availability requires multiple vCloud cells and clustered Chargeback servers.

Management Cluster

From the underlying vSphere infrastructure level, the management cluster VMs can be protected with VMware DRS, HA, and FT. These technologies assume that management components are run as virtual machines. If they are not virtual machines, then other availability options should be used.

VMware DRS and HA are recommended for use. VMware HA should be enabled to support workloads in the cloud. VMware HA can protect against both ESXi failure and against OS failure as described in *Foundation for Cloud Computing with VMware vSphere 4*. VMware FT applies for single vCPU VMs in the management cluster.

VMware vCenter Heartbeat can be used to increase availability of vCenter Server. VMware vCenter Site Recovery Manager™ (SRM) can be utilized for handling of site-related failures of the management cluster components and VMs. Additional information about vCenter Heartbeat and vCenter SRM can be found in *Foundation for Cloud Computing with VMware vSphere 4*.

VMware Data Recovery (vDR) is based on the VMware vStorage APIs for Data Protection (VADP). This technology enables backup of virtual machines with centralized backup for improved manageability and a reduced dependency on a backup window while also eliminating the need for guest OS-based agents. These APIs replace VMware Consolidated Backup (VCB). They take advantage of the VMware storage VMFS snapshot capabilities. Figure 59 provides an overview of how VADP works. VMware Data Recovery utilizes VADP and can be used to back up the VMs in the management cluster. Other VADP-based backup technologies are available from third-party backup vendors.

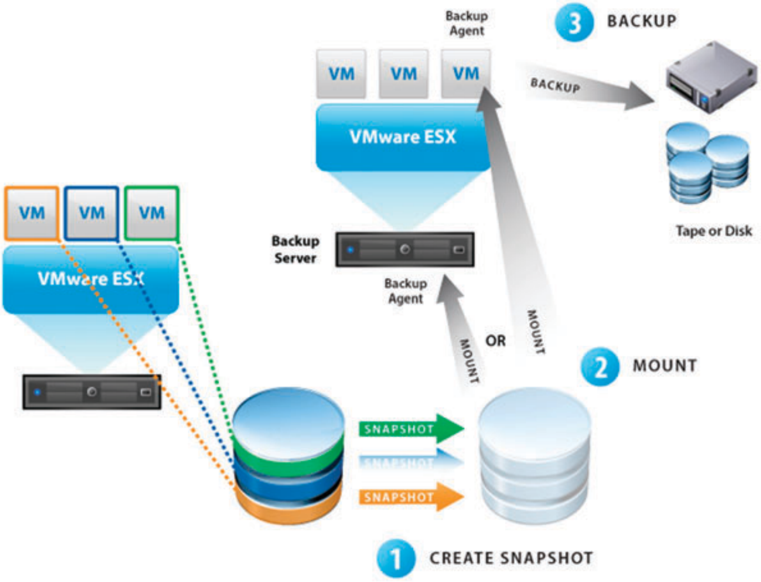

Figure 59. VMware vStorage APIs for Data Protection

Resource Group

The resource group clusters can be architected for business resiliency, but may not take advantage of all BC/DR features of the underlying vSphere level. The running workloads on the resource groups cluster can take advantage of VMware DRS and HA, both of which are recommended.

vApp Backup and Recovery

Backup and recovery of cloud-based applications requires some differences in approach from traditional backups for physical and virtual machines. This is because the cloud layer provides greater abstraction from the physical resources than does the virtual layer. There are also specific vCloud interface points that have specific integration considerations.

Backup vendors will integrate their products through VADP, which provides an interface for backup products to leverage the snapshot capabilities of VMware vStorage.

When choosing a third-party backup vendor, the following considerations align with vCloud backup/recovery requirements:

- VADP Integration
  - VADP provides change-block tracking capability to reduce backup time of vApps.
  - VADP supports backup of isolated VMs and vApps that use a vApp network or have no network connectivity.
  - Integration with VADP provides LAN-free and server-free backup. This produces a better consolidation ratio for vCloud and the underlying vSphere infrastructure.
  - When performing backups, use of VM UUID instead of VM name is recommended to support multi-tenancy and avoid potential namespace conflicts.
  - Integration with VADP, vCloud API, and vSphere API is important for providing a fully integrated and automated solution.
- vCloud Director Integration
  - Full integration enables use of vCloud multi-tenant security features such as integration with RBAC mechanisms used by vCloud Director.
  - Deployment of a vCloud architecture will likely require vendor support for both provider access (provider administrator) and consumer access (organization administrator and users).
  - The metadata of a vApp must be included to allow for recovery at alternate locations with the information tied to SLAs and security (such as network access). This metadata may be temporary or permanent. Details can be found on the VMware Web site.

VMware vCloud Director provides a new form of vApp that includes the ability to support multi-tiered applications. An example is a Web-based e-commerce solution that includes Web, application server, and database components. These are grouped together as a single e-commerce vApp. This e-commerce vApp is treated as one object with knowledge of the other VMs from which it is composed. Backup/recovery options, especially for full disaster recovery solutions, should support operations on the top level e-commerce vApp and enumerate the supporting objects to provide a backup of all tiers. This provides a consistency group for the entire application.

Some locations may have users who require the ability to do one-time-only backups. In addition, backups should support both full VM and file-level solutions.

High-Level Workflow

The following high-level workflow is based on work done by Michael White, BC/DR specialist at VMware.

The restoration of a vApp within vCloud differs from standard practices, due to the nature of cloud computing and vCloud technologies. The following provides details on how to perform backup and restore without guest agents.

Considerations for Success

When deploying vApps, be sure to specify descriptive VM names in the “Full Name” and “Computer Name” fields. The names may include tenant- or organization-specific information. If generic information is used it may be difficult to identify individual systems when performing backup/recovery configuration and operations. vApps and VMs that are provisioned by vCloud Director have a large GUID-template_name. This may result in similar VMs, making it hard for users and administrators to differentiate during restoration of an image.

Backup

Backup procedures are needed for the vApps that are deployed into the cloud. Traditional backup tools do not capture required metadata such as owner, network, and organization that are associated with a vApp. This results in recovery and restoration issues. Without this data, recovery must include manual steps, and configuration attributes have to be manually re-entered, slowing the recovery process and introducing additional risk.

Within a vCloud environment, vApps contain one or more VMs. Backup of vApps on isolated networks must be supported. Identifying inventories of individual organizations becomes challenging based on current methods that enumerate the backup items using vSphere, which uses a different UUID from that of vCloud Director.

Agent-based backups can continue to be used but will not work for all workloads in a cloud. This is due to the potential for isolated vApps or vApps connected to an isolated network. Another reason is that non-VADP solutions can impact consolidation ratios and network load. VADP-based solutions help alleviate these issues.

Several vendors are working on solutions with both VADP and vCloud integration. Check with your preferred vendor for status on their integration plans.

The backup solution needs to be aware of which VMs are part of the vApp to avoid any data or state corruption due to a restore from a different time frame.

Restore

vCloud Director cannot import a vApp back into the system if it is restored directly to its original location. Use the following procedure to restore the vApp:

1. Restore the complete VM/vApp to a new location that is accessible to the vCloud Director infrastructure.
2. Log into the vCloud Director as the organization administrator (provider administrator can also be used).
3. Import the restored VM/vApp, as shown below.

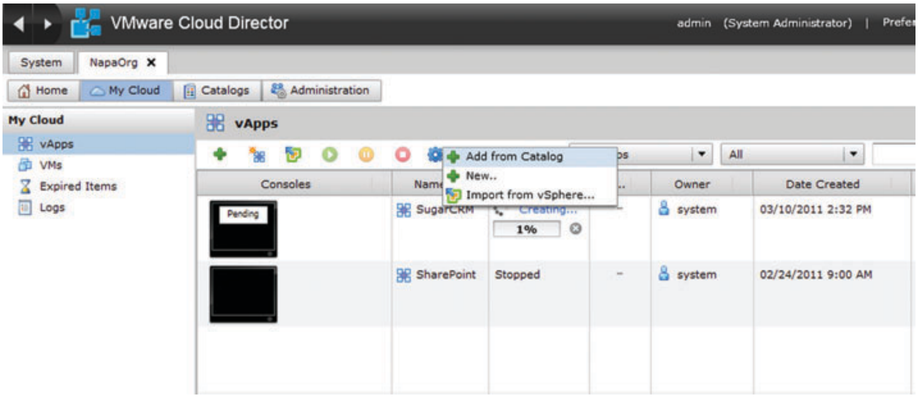

Consoles	Name		..	Owner	Date Created
Pending	SugarCRM	Creating... 1%		system	03/10/2011 2:32 PM
	SharePoint	Stopped	-	system	02/24/2011 9:00 AM

4. You may need to import this VM to an existing vApp.
5. After it is restored, change the owner of the VM/vApp to the correct owner.
6. Enable the network on this VM/vApp. This is done in the VM or vApp view, as shown below.

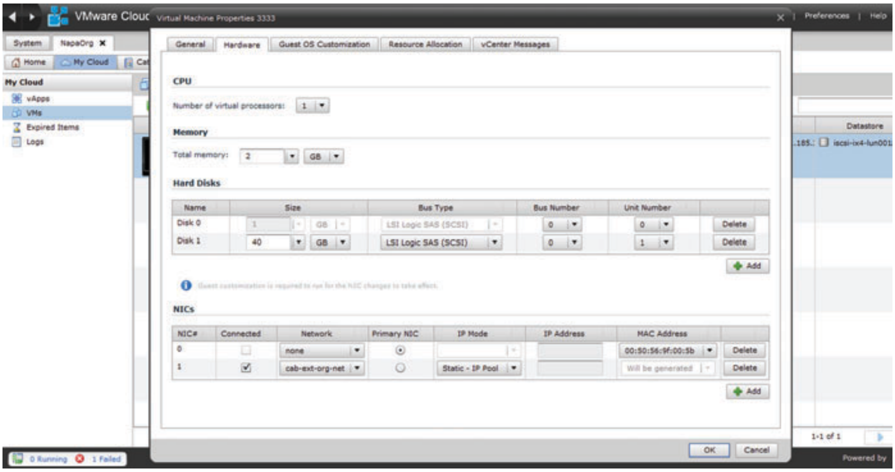

CPU

Number of virtual processors: 1

Memory

Total memory: 2 GB

Hard Disks

Name	Size		Bus Type	Bus Number	Unit Number	
Disk 0	1	GB	LSI Logic SAS (SCSI)	0	0	Delete
Disk 1	40	GB	LSI Logic SAS (SCSI)	0	1	Delete

Guest customization is required to run for the NIC changes to take effect.

NICs

NIC#	Connected	Network	Primary NIC	IP Mode	IP Address	MAC Address	
0	□	none	●			00:50:56:9f:00:5b	Delete
1	☑	cab-ext-org-net	○	Static - IP Pool		Will be generated	Delete

7. Have the owner of this VM/vApp power it on and use it to recover files. At this stage it can also be used as a replacement VM/vApp. Remember that the original VM may or may not be available to the user, depending on the disaster/recovery situation.

Summary

Business resiliency builds on top of the approach used for the virtualization layer. Due to further abstraction from the physical layer, the vCloud layer requires adjustments to the techniques used in business continuity and disaster recovery. Careful business continuity planning is essential for ensuring the reliability of cloud services.

Appendix. Third-Party Technology Integrations

VMware vCloud Director and Integrated Computing Stacks

Just as vCloud Director has standardized the architecture for building shared pools of virtualized resources for applications, many datacenter manufacturers are standardizing the infrastructure to deliver those virtualized services. Aligning standardization across application and infrastructure offers customers the opportunity to deploy cloud computing environments faster and at lower overall operating costs.

Integrated Computing Stacks is an evolving offering that brings together virtualization, server computing, networking, storage, security, and system management into shared pools of resources. They give enterprise, government, and service providers the foundational infrastructure to deliver private cloud or hybrid cloud services. The underlying premise behind integrated computing stacks is lowered cost of acquisition and operations through reducing the requirement to piece together all of the individual technology components. Additionally, the stack management systems increasingly are able to coordinate (operate, automate, manage, provision) the stacks as integrated resources, simplifying deployments and troubleshooting.

As the integrated stacks continue to increase in popularity, one of the most common questions is about the role of vCloud Director in building a cloud computing environment utilizing the stacks. This section briefly considers today's leading integrated stack solutions and how (and whether) they integrate with vCloud Director.

VCE—Virtual Computing Environment (VMware, Cisco, EMC)

Announced in the fall of 2009, VCE brings together technology from VMware, Cisco, and EMC in an offering called Vblock Infrastructure Packages (Vblock). Vblock delivers virtualization (VMware), computing and networking (Cisco), and storage (EMC) as pools for datacenter infrastructure optimized for virtualized workloads. VCE provides a management framework called Unified Infrastructure Manager (UIM) which provisions and monitors Vblocks as pools of compute and storage resources. VMware vCloud Director will be able to integrate with UIM to provision virtual machine resources (VM, vApp) to align with the underlying compute and storage resources. Operationally, customers will have the option of making vCloud Director the customer-facing portal or, alternatively, using both UIM and vCloud Director for internal operations.

HP—Converged Infrastructure

Unveiled at VMworld 2010, The HP Cloud Map for vCloud Director is an end-to-end cloud infrastructure solution integrating HP's BladeSystem Matrix and vCloud Director.

Converged Infrastructure brings together Virtual Connect, BladeSystem Matrix, FlexFabric, Matrix Operating Environment, and HP networking and storage. Through this integration, customers can accelerate the initial deployment of vCloud environments and optimize running vCloud environments by easily scaling vCloud resource pools to quickly adjust to changing business demands.

FlexPod for VMware—NetApp, Cisco, VMware

Building upon the Secure Multi-Tenancy (SMT) architecture announced in the spring of 2010, the FlexPod for VMware architecture was released in fall 2010. FlexPod delivers virtualization (VMware), compute, networking (Cisco), and storage (NetApp) in flexible pods of resources for virtualized datacenter environments. FlexPod components all provide plug-in management into vCenter and use an open API approach for overall system management and orchestration. VMware vCloud Director has been tested to work with FlexPod for customers looking to deploy private or hybrid cloud environments.

Dell—Virtual Integration System (VIS)

In mid-year 2010, Dell announced the Virtual Integration System (VIS) to deliver virtualized workloads across compute, network, and storage. Dell provides computing and storage components with networking components optionally coming from Dell or third-party vendors. To manage a VIS, Dell provides Advanced Infrastructure Manager (AIM). To provision services, Dell provides VIS Self-Service Creator (SSC). The combination of these tools allows IT operations to either offer customer-facing self-service portals or use all of the tools for internal operations. Although Dell is a large VMware OEM and vCloud Director can be deployed in VMware on Dell environments, customers may see AIM and SSC positioned as the preferred solution instead of vCloud Director, depending on customer requirements.

IBM—Cloudburst

IBM Cloudburst offers an integrated delivery service (on-premise). This solution provides integrated virtualization, computing, networking, and system-level management. Although IBM is a major partner of VMware and Cloudburst includes ESXi 4.1, Cloudburst ships with Tivoli Cloud Management software.

Oracle—Exalogic

Announced in the fall of 2010, Oracle Exalogic is targeted at “cloud in a box” services for Oracle applications. Exalogic uses Oracle VM, Oracle applications and middleware, Oracle Servers, and Oracle Storage interconnected with an Infiniband network. Because Exalogic does not support non-Oracle hypervisor environments, vCloud Director cannot be deployed in these environments.

VMware vCloud Director and Orchestration Tools

There are a number of companies that build management products which provide some level of planning, service-catalog, orchestration, and automation for virtualized infrastructure. Some of these companies include newScale, DynamicOps, FluidOps, Cisco/

Tidal Software, BMC, CA, and others. These companies build various levels of functionality to help IT operations simplify how pools of shared, virtualized resources are utilized.

In some cases, customers will have opportunities to integrate these products with vCloud Director. For example, a customer might choose to use the service-catalog capability from one of these products in conjunction with the orchestration within vCloud Director. In other cases, these products will have roadmaps to add broader feature sets that could overlap with vCloud Director, and customers may decide that the overlap causes too much complexity in their environment.

VMware vCloud Director Security with HyTrust

Cloud provider deployment can be accomplished with vCloud Director, HyTrust® Appliance, and HyTrust® Cloud Control™ Appliance.

The HyTrust Appliance provides consistent, granular access control across each and every available access method and protocol, highly flexible and enforceable virtual infrastructure policies, and detailed logs for every attempted change (whether allowed or denied) to the vSphere virtual environment.

HyTrust Appliance is a virtual security appliance that sits in the management network and provides unified administrative access control. Because all management traffic is passed through the HyTrust Appliance, it provides a perfect point for integration with stronger authentication solutions. HyTrust Appliance 2.1 supports the following authentication mechanisms:

- *Strong passwords*—These are supported through any LDAP version 3 compliant directories such as Microsoft Active Directory. Complex, multi-domain directory configurations are supported.
- *RSA SecurID 6- or 8-digit tokens*—HyTrust Appliance is RSA Secured certified: https://gallery.emc.com/docs/DOC-1882?viewTab=overview&version=6.
- *Single sign-on*—This is supported via Microsoft Windows pass-through authentication with smartcard integration.

In Figure 60, Mary is an enterprise administrator and manages their vSphere environment using the vSphere Client (connected directly to ESXi or to vCenter), Secure Shell (SSH), or any vSphere programmatic APIs and CLIs.

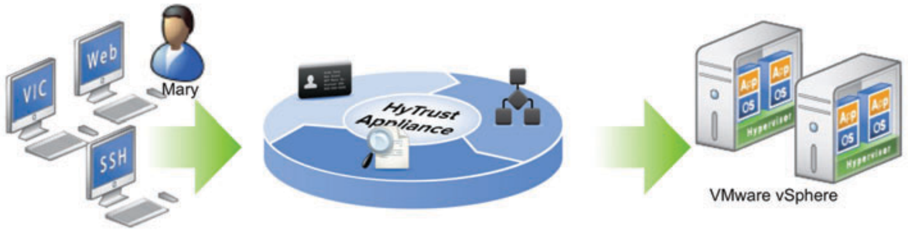

Figure 60. Enterprise Deployment of VMware vSphere and HyTrust Appliance

HyTrust Appliance is typically deployed in a transparent mode. Administrators perform their tasks no differently than if they were using vSphere directly. HyTrust Appliance also supports the Cisco NX OS CLI, enabling the same level of controls to be enforced even across a distributed virtual switch supplied by a third-party vendor such as the Cisco Nexus 1000V. This is a critical concern, as these third-party solutions may be providing essential virtual infrastructure services that operate within the hypervisor and/or across a cluster of hypervisors.

After administrators are authenticated, their group(s) membership is determined by the central user directory. Use of directory services enables centralization of user provisioning, deprovisioning, setting password strength policies, and management of functional and business group memberships. These groups are mapped to specific, custom roles that are defined in HyTrust Appliance. Roles are highly customizable and can be defined to limit what operations a specific administrator or a group of administrators are allowed to perform.

For example, Mary may be a network administrator for the Payment Card Industry (PCI) environment. All of the virtual resources associated with that environment can be labeled, and specific access and infrastructure policies can be configured. Figure 61 shows a HyTrust Appliance screen, where an access policy is configured to allow only administrators in the PCI_network_admin group to manage the PCI-related virtual switches on the hosts.

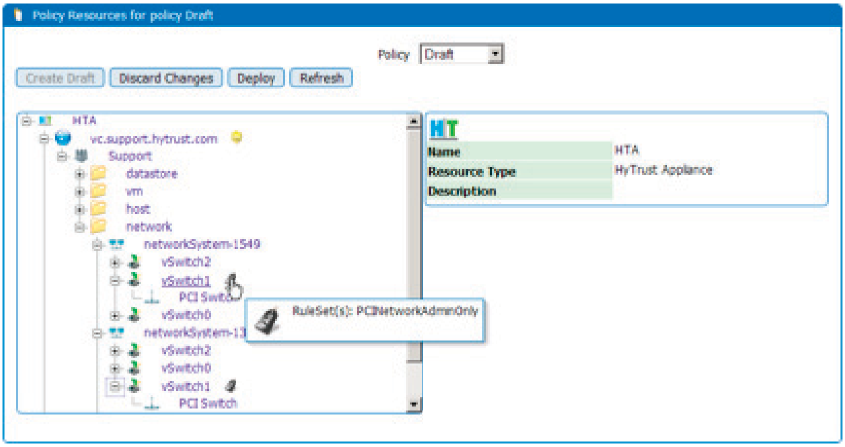

Name	HTA
Resource Type	HyTrust Appliance
Description	

RuleSet(s): PCINetworkAdminOnly

Figure 61. PCI Network Administrator Access Policy

Additionally, by configuring custom policies, administrative access can be further restricted to specific virtual resources, access method/protocol, and origination IP address/range, allowing management of the Cisco Nexus 1000V VSM, for example, only through SSH. If Mary attempted to access the Cisco Nexus 1000V VSM virtual machine using any other access method, she would be denied access. This is incredibly powerful ,since virtualized infrastructure resources such as virtual management appliances can be logically grouped together through HyTrust labels. Such a group could be configured with a special, much more restrictive access policy, preventing administrators from making mistakes such as powering off a virtual machine that provides critical services.

Label-based policy management provides tremendous flexibility, as access policies can be defined once and configured by label. Afterwards, every resource that is applied a particular label automatically inherits the label's configured policy. This frees administrators from hierarchical modeling and allows them to simply define arbitrary classification schemes and associate access policies. For example, all virtual machines, networks, and hosts dedicated to the Finance Department can be labeled "finance." Then access policies such as "only finance dept. administrators have the right to manage these resources" can be configured, as well as infrastructure segmentation policies. This is similar in concept to the VMware ESXi affinity and anti-affinity policies around High Availability (HA) and Dynamic Resource Scheduler (DRS). In HyTrust Appliance, policies can be defined to restrict a specific virtual machine to only be connected to a virtual network that is identified by a label, or to be instantiated on a host that is identified by a label. For example, a virtual machine with a "PCI" label may be restricted to a virtual network labeled PCI and be powered-on or cloned on a host labeled PCI. This prevents administrators from accidently or intentionally:

- Importing or migrating a sensitive virtual machine onto hosts that are not adequately hardened or protected, or
- Connecting a sensitive virtual machine to unprotected networks.

In another example, HyTrust Appliance can restrict Mary, who is a member of a PCI_security_admin_group that is mapped to a VIAdmin role, to only manage a vShield Manager virtual machine and other PCI-related management virtual machines (see Figure 62). Policies can also be configured to restrict these virtual machines to only run on a host that is labeled PCI, preventing accidental or intentional powering on or movement to hosts that are not labeled PCI.

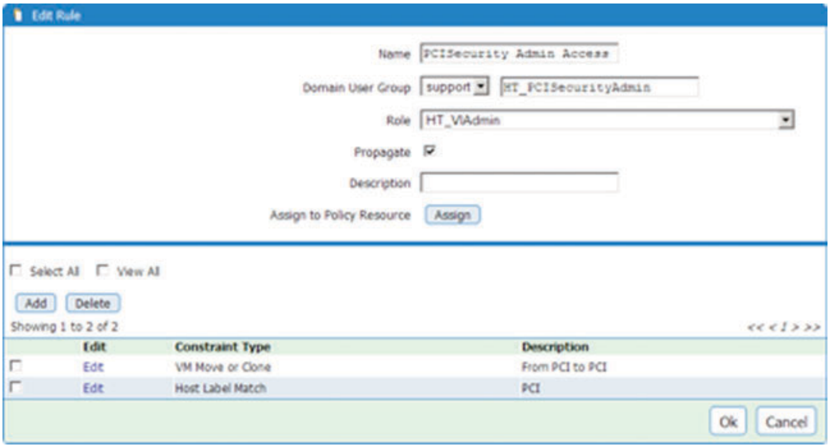

Name: PCISecurity Admin Access
Domain User Group: support ▾ HT_PCISecurityAdmin
Role: HT_VIAdmin ▾
Propagate: ☑
Description:
Assign to Policy Resource

□ Select All □ View All

Showing 1 to 2 of 2

	Edit	Constraint Type	Description
□	Edit	VM Move or Clone	From PCI to PCI
□	Edit	Host Label Match	PCI

Figure 62. PCI Security Admin Access and Infrastructure Segmentation

Because HyTrust Appliance intercepts all attempted management access and authenticates and authorizes individuals for every operation, it can generate detailed, audit-quality logs containing information about:

- When
- Who, in which role, attempted to perform which operation
- On what virtual infrastructure resource

- From which IP address ,and
- Whether it was allowed or denied

These logs are kept locally for temporary redundancy and troubleshooting. The logs can also be offloaded to multiple destination syslog receivers (for example, Splunk, RSA enVision, Arcsight, Loglogic, or any other SIEM/log aggregation solution). HyTrust Appliance has a built-in search engine that allows for quickly finding and/or filtering for specific log records. For example, in the HyTrust Appliance log viewer, as shown in Figure 63, all records at the WARN level can be found for an administrator. From that filtered set it is easy to discover attempted operations not permitted by role, or attempts to manage a virtual infrastructure resource without administrative access.

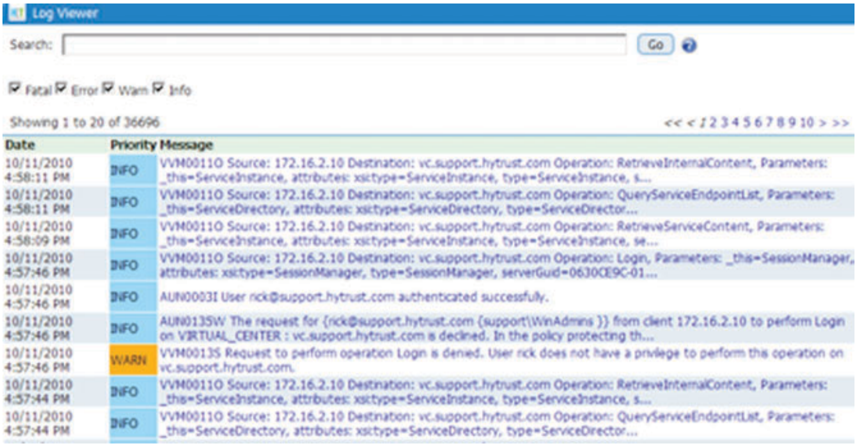

☑ Fatal ☑ Error ☑ Warn ☑ Info

Showing 1 to 20 of 36696

<< < 1 2 3 4 5 6 7 8 9 10 > >>

Date	Priority	Message
10/11/2010 4:58:11 PM	INFO	VVM00110 Source: 172.16.2.10 Destination: vc.support.hytrust.com Operation: RetrieveInternalContent, Parameters: _this=ServiceInstance, attributes: xsi:type=ServiceInstance, type=ServiceInstance, s...
10/11/2010 4:58:11 PM	INFO	VVM00110 Source: 172.16.2.10 Destination: vc.support.hytrust.com Operation: QueryServiceEndpointList, Parameters: _this=ServiceDirectory, attributes: xsi:type=ServiceDirectory, type=ServiceDirector...
10/11/2010 4:58:09 PM	INFO	VVM00110 Source: 172.16.2.10 Destination: vc.support.hytrust.com Operation: RetrieveServiceContent, Parameters: _this=ServiceInstance, attributes: xsi:type=ServiceInstance, type=ServiceInstance, se...
10/11/2010 4:57:46 PM	INFO	VVM00110 Source: 172.16.2.10 Destination: vc.support.hytrust.com Operation: Login, Parameters: _this=SessionManager, attributes: xsi:type=SessionManager, type=SessionManager, serverGuid=0630CE9C-01...
10/11/2010 4:57:46 PM	INFO	AUN0003I User rick@support.hytrust.com authenticated successfully.
10/11/2010 4:57:46 PM	INFO	AUN0135W The request for {rick@support.hytrust.com {support\WinAdmins }} from client 172.16.2.10 to perform Login on VIRTUAL_CENTER : vc.support.hytrust.com is declined. In the policy protecting th...
10/11/2010 4:57:46 PM	WARN	VVM0013S Request to perform operation Login is denied. User rick does not have a privilege to perform this operation on vc.support.hytrust.com.
10/11/2010 4:57:44 PM	INFO	VVM00110 Source: 172.16.2.10 Destination: vc.support.hytrust.com Operation: RetrieveInternalContent, Parameters: _this=ServiceInstance, attributes: xsi:type=ServiceInstance, type=ServiceInstance, s...
10/11/2010 4:57:44 PM	INFO	VVM00110 Source: 172.16.2.10 Destination: vc.support.hytrust.com Operation: QueryServiceEndpointList, Parameters: _this=ServiceDirectory, attributes: xsi:type=ServiceDirectory, type=ServiceDirector...

Figure 63. HyTrust Appliance Logs

To make sure that HyTrust Appliance is not somehow circumvented—for example, by direct physical access to the hosts—additional capabilities are supported:

- Root Password Vaulting—Locks down privileged host accounts and provides passwords for temporary use to enable time-limited privileged account access. Each host has a unique password, and only one temporary password can be active at a time. This allows highly auditable control over who requested and has access to specific hosts as "root." Because additional user or service accounts are not provisioned on the hosts, even with the ability to physically access hosts, administrators are prevented from abusing the systems without detection.
- Hardening of the hosts and VM containers—Hosts can be hardened per any one of the supported templates, which are customizable. Supported templates are: Center for Internet Security benchmark for ESX, VMware hardening guides, Payment Card Industry (PCI), and Sarbanes-Oxley (SOX). HyTrust Appliance allows automatic assessments, remediation performed manually, or remediation performed automatically per configured schedule. The results are viewable through the HyTrust Appliance management UI or downloadable in CSV format.

HyTrust Cloud Control

HyTrust Cloud Control is a solution that is currently under development. It will provide additional security and compliance capabilities to vCloud Director, similar to how HyTrust Appliance provides additional control and visibility for vSphere virtual infrastructures.

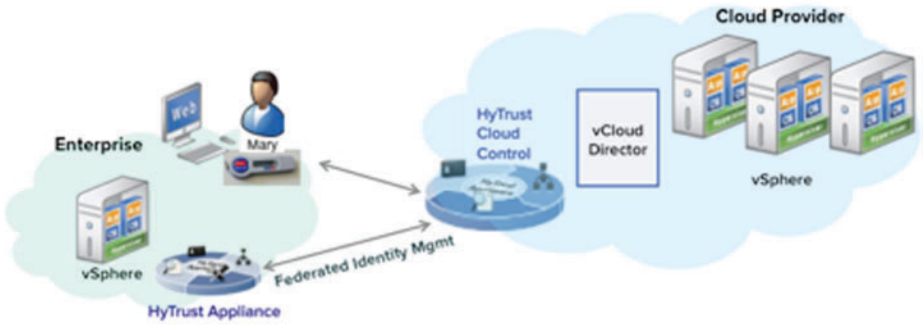

Figure 64. VMware vCloud Director and HyTrust Cloud Control

HyTrust Cloud Control was first presented as a proof-of-concept demonstration at VMworld 2010. At the event and subsequently, HyTrust Cloud Control has garnered much interest. The interest can be attributed to the following essential features that are currently not available in vCloud Director 1.0:

- The need for true transparency/visibility of the cloud provider infrastructure and change management activities per tenant to meet compliance requirements.
- Ease-of-use of administrative access provisioning and deprovisioning, stronger authentication, and more granular separation of duties. These apply equally to the cloud provider and the tenant administrators.
- Consistent security policy management within the private and the hybrid clouds; security policy migration with workloads automatically recognized and enforced in the cloud infrastructure.

HyTrust Cloud Control early access is planned for Q1, 2011. The main capabilities being considered are:

- Per-tenant/organization, audit-quality logs for change management operations performed within vCloud Director and vSphere
- Flexible policy management that spans the private and hybrid cloud infrastructures, and is persistent
- Support for additional authentication mechanisms such as RSA SecurID and smartcards
- Highly granular and customizable roles for both the cloud provider and tenant/organization administrators

Figure 65 shows a cloud provider deployment with vCloud Director and HyTrust Cloud Control located in the DMZ and the vSphere infrastructure in an internal network that is protected by a HyTrust Appliance.

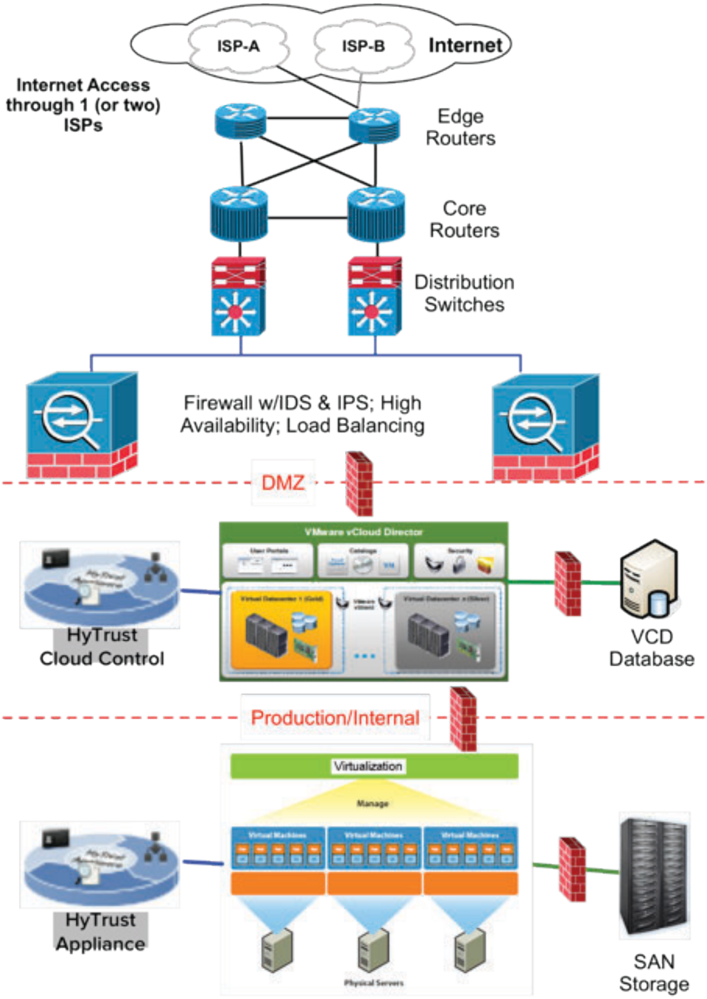

Figure 65. Deployment Architecture

Glossary

The following is a glossary of terms and acronyms used by VMware. For the sake of simplicity, multi-word terms that start with VMware, vSphere, vShield, vCloud, or vCenter are listed alphabetically using the following words in the term:

"VMware vCenter *term*" identifies technologies that are integrated with VMware vCenter, some of which are considered add-on products requiring additional licensing.

"VMware *term*" identifies technologies that are stand-alone products.

Where applicable, definitions are taken from the VMware glossary.

Allocation Pool—A pool of allocated resources in the virtual datacenter for which a certain percentage is guaranteed.

Cedar—A type of tree which is naturally resistant to decay and bugs.

vCenter **Chargeback**—A metering and costing engine which enables cost measurement, configuration, analysis, and reporting for virtualized environments. vCenter Chargeback provides the ability to map IT costs to internal business units or external customers.

vCenter **Chargeback API**—Provides an interface for vCenter Chargeback functionality. This includes management of the hierarchy, cost configurations, and reporting.

VMware **Converter**—Used for physical-to-virtual machine (P2V) migrations, as well as for importing virtual machines from other virtualization vendors (V2V). VMware Converter can import multiple machines concurrently and non-disruptively. Designed for large-scale consolidation, VMware Converter can be used with or without VMware vCenter Server.

Distributed Power Management (DPM)—Dynamically consolidates workloads during periods of low resource utilization to reduce power consumption. Virtual machines are migrated onto fewer hosts, and the unneeded ESX hosts are powered off.

Distributed Resource Scheduler (DRS)—Dynamically allocates and balances workloads across hosts in a cluster using vMotion to enhance performance, scalability, and availability.

VMware **ESX Server (VMware ESX Classic)**—Bare metal hypervisor that comes with a Linux-based service console to assist with management functions. VMware ESX was initially released in 2002.

VMware **ESXi Server**—VMware next-generation hypervisor, which provides the same feature set and performance as ESX classic with a significantly reduced footprint. ESXi removes the service console, improving security due to the reduced attack surface, and can be embedded on removable media or installed on a traditional storage device.

External Networks—External networks provide Internet connectivity to organizations. External networks are backed by port groups configured for Internet accessibility.

Fault Tolerance or **VMware FT**—Enables active clustering of virtual machines without requiring applications to be cluster aware. Provides clustering support of single vCPU VMs without requiring the embedded application to be cluster aware. FT utilizes VMware vLockstep technology. This technology uses an active secondary VM that runs in virtual lockstep with the primary VM. VMware vLockstep establishes and maintains this secondary VM. The secondary VM runs on a different host and executes the same set of instructions, in the same sequence, as the primary VM.

GemFire—An in-memory distributed data management platform and core component of VMware vFabric Cloud Application Platform. GemFire pools resources across multiple processes to manage application objects and behavior. Availability and scalability is improved through advanced techniques such as dynamic replication and data partitioning.

vCenter **Heartbeat**—Protects the vCenter Server, License Server, and database against hardware, OS, application, and network downtime. Failover and failback are provided for each. Protection is important especially when using VMware View, vCenter Lab Manager, or vCenter SRM as they require vCenter to be running at all times.

High Availability or **VMware HA**—Provides automated restart of failed virtual machines, regardless of the guest operating system technology. Provides fault tolerance in the event of an ESX host failure. VMware HA enables the automated restart of virtual machines on other hosts in a cluster upon host failure, minimizing downtime without the cost of application clustering.

Host—A compute platform supporting the execution of virtual machines. Includes standard physical servers as well as platforms specifically designed to support virtual infrastructure such as Cisco UCS.

Hyperic HQ—Provides complete discovery, monitoring, analysis, and control of all application, system, and network assets both inside and outside of the virtual machines. Hyperic HQ includes full VMware ESX and VMware Server support, analysis of utilization, and performance within a VM; correlation of events between hosts and guest operating systems; and control of VMs. This tool provides detailed analysis of how the virtual machine is performing and provides depth in latency analysis for an application.

Hypervisor—Hypervisor virtualization platforms have a partitioning layer that runs directly on top of the hardware and below higher-level virtualization services that provide a virtual machine abstraction. The hypervisor is installed on the computer just like an operating system. It provides the capability to create virtual machine partitions, with a virtual machine monitor running within each partition.

vSphere **Management Assistant (vMA)**—Linux appliance with pre-built management tools and the vCLI Interface. Allows scripting and agents to manage ESX, ESXi, and vCenter Server systems. vMA is a virtual appliance that includes the vSphere SDK and the vSphere CLI, logging capabilities, and authentication mechanisms.

Network I/O Control—Enables a pre-programmed response to occur when access to a network resource becomes contentious. Ensures performance of VM activities.

Network pools—A network pool is a collection of isolated Layer 2 virtual machine networks available to vCloud Director for the automated deployment of organization and vApp networks.

Network virtualization—VLANs (Virtual LANs) are used to segment networks on physical networks. This is different from the virtualized network devices available in VMware, although these two technologies can coexist.

Open Virtual Appliance (OVA)—A packaging format for virtual machines that allows virtual machine templates to be distributed, customized, and instantiated on any OVA supporting VMM. This format uses a single file (*.ova) using the TAR format.

Open Virtualization Format (OVF)—A specification that describes an open, secure, portable, efficient, and extensible format for the packaging and distribution of software to be run in virtual machines. OVF is a distribution format for virtual applications that uses existing packaging tools to combine one or more virtual machines with a standards-based XML wrapper. Stored as a package of files (*.vmdk, *.ovf).

vCenter **Orchestrator (vCO)**—Provides out-of-the-box workflows to help automate existing manual tasks. Workflows can be created, modified, and extended to meet custom needs.

VMware vCenter **Orchestrator API**—Allows for the programming of workflows for execution by VMware vCenter Orchestrator.

Organization—The unit of multi-tenancy representing a single logical security boundary. An organization contains users, virtual datacenters, and networks.

Organization administrator—Administrator for a vCloud Director organization responsible for managing provided resources, network services, users, and vApp policies.

Organization networks—Instantiated through network pools and bound to a single organization. Organization networks are represented by a vSphere port group and can be isolated, routed, or directly connected to an external network.

Organization Virtual Datacenter (organization VDC)—A subgrouping of compute and storage resources allocated from a provider virtual datacenter and assigned to a single organization. A virtual datacenter is a deployment environment where vApps can be instantiated, deployed, and powered on. Network resources are associated through a network pool mapping. Organization VDCs cannot span multiple organizations.

Pay-As-You-Go—Provides the illusion of an unlimited resource pool in a virtual datacenter. Resources are committed only when vApps are created in the organization VDC.

Port group(s)—Specifies port configuration options including VLAN tagging policies and bandwidth limitations for each member port. Network services connect to vSwitches through port groups. Port groups define how a connection is made through the vSwitch to the network. In typical use, one or more port groups are associated with a single vSwitch.

Provider Virtual Datacenter (provider VDC)—A grouping of compute and storage resources from a single vCenter Server. A provider VDC consists of a single resource pool and one or more datastores. Provider VDC resources can be shared with multiple organizations.

RabbitMQ—An open source message broker based on the Advanced Message Queuing Protocol (AMQP) standard. RabbitMQ is written in Erlang and uses the Open Telecom Platform (OTP) framework. Rabbit Technologies, Ltd. was acquired by SpringSource (a division of VMware) in April 2010.

Redwood—Home to the world's tallest trees, the Redwood National and State Parks are located along the coast of northern California. Also the code-name used during the development of vCloud Director 1.0.

Representational State Transfer (REST)—A style of software architecture that relies on the inherent properties of hypermedia and HTTP to create and modify the state of an object that is accessible at a URL.

Reservation Pool—Resources allocated to the organization VDC are completely dedicated. Identical to an allocation pool with all guarantees set to 100%.

vSphere Guest **SDK**—Enables development of applications that will run within a virtual machine using C or Java libraries. Enables customers to write smart applications that respond to changes at the virtualization environment layer. It is included with VMware Tools.

Security Assertion Markup Language (SAML)—An XML standard for management and exchange of security information between security domains. SAML can be used to federate identity management, allowing for a single sign-on user experience in cloud environments.

Single Sign On (SSO)—A property of systems whereby a single action of user authentication permits access to authorized systems across independent security domains. SSO is highly desirable from a usability and security perspective, as users will not be required to sign in multiple times to access cloud services.

vCenter **Site Recovery Manager (SRM)**—Provides disaster recovery workflow automation through a centralized management interface. SRM automates the setup, testing, failover, and failback of virtual infrastructures between protected and recovery sites.

Spring Framework—An open source Java application framework providing enhanced capabilities for creating enterprise Java, rich Web, and enterprise integration applications. Core features such as aspect-oriented programming, model-view-controller, and dependency injection help developers build applications faster. Spring provides a consistent programming model that is both comprehensive and modular.

Storage I/O Control (SIOC)—Provides a dynamic control mechanism for managing I/O resources across VMs in a cluster. SIOC can mitigate the performance loss of critical workloads during periods of congestion by prioritizing I/O to VMs through disk shares.

Storage vMotion—Storage vMotion enables live migration of virtual machine disk files across storage locations while maintaining service availability. Storage vMotion utilizes vMotion technology to optionally move the VM to an alternate ESX host which has access to both source and target storage locations. Storage vMotion can move the storage location of a virtual disk as long as the target is visible to the source and destination ESX host(s). The processes of the corresponding VM can stay on the same host, or the VM can be simultaneously migrated, using vMotion, to a new host.

VAAI—vStorage API for Array Integration

VADP—VMware API for Data Protection

vApp—A vApp is a container for a software solution, encapsulated in OVF format. vApps provide a mechanism for moving applications between internal and external clouds while preserving any associated policies. vApps contain one or more VMs, have power-on operations just like a virtual machine, and can contain additional attributes such as vApp networks and VM startup sequencing.

vApp Network—A network which connects VMs within a vApp, deployed by a consumer from a network pool.

vApp Template—A blueprint for a vApp, which can be customized when instantiated.

vCenter Server—The foundation for centralized virtualization management. vCenter Server aggregates virtual resources to provide a scalable and extensible platform. vCenter Server enables core features such as VMware vMotion, Distributed Resource Scheduler (DRS), High Availability (HA), and Fault Tolerance (FT).

vCLI (vSphere Command Line Interface)—Allows you to manage your virtual infrastructure using Windows PowerShell. This lets you script and automate actions you would normally perform in vCenter Server. There are approximately 200 cmdlets (PowerShell exposed procedures) to manage vSphere and ESXi functionality. There are many pre-built scripts available online that can provide functionality such as finding all VM snapshots, finding orphaned VMs, or even creating reports. vCLI was previously known as the VI ToolKit.

vCloud API—The vCloud API is a RESTful, pure virtual API for providing and consuming virtual resources from the cloud. It enables deployment and management of virtualized workloads in internal and external clouds. The two major components are the User API focused on vApp provisioning and the Admin API focused on platform/tenant administration.

vCloud Datacenter—A hybrid cloud solution enabling enterprises to extend their private cloud to the public cloud with flexibility, scalability, security, and operational efficiency. vCloud Datacenter services are co-branded by the service provider and VMware, providing enterprise customers with a choice of 100% compatible services that are based on VMware cloud architecture and certified by VMware.

VMware **vCloud Director**—A software solution providing the interface, automation, and management feature set to allow enterprise and service providers to supply vSphere resources as a Web-based service.

vCloud Director Catalog—A collection of vApp templates and media made available to users for deployment. Catalogs can be published and shared between organizations in the same vCloud environment.

vCloud Director Cell—A vCD cell is an instance of vCloud Director running on a physical or virtual machine. Multiple cells can be grouped together and connected to one vCD database for load balancing and availability.

vCloud Director Network Isolation (vCD-NI)—An alternative solution for network isolation. It provides tenant isolation on shared physical L2, using MAC-in-MAC Ethernet frame encapsulation to create thousands of overlay networks. In vSphere 4, vCD-NI is implemented using a kernel-level service in each ESX host which performs the encapsulation and decapsulation of vCD-NI packets.

vCloud Express—An Infrastructure as a Service (IaaS) offering delivered by leading VMware service provider partners. It is a co-branded service that provides reliable, on-demand, pay-as-you-go infrastructure.

vCloud Request Manager—Provides private cloud policy control through enforcement of business policies and procedures, maximizes efficiency and service delivery through process automation, supports approval workflows for provisioning requests, automatically tracks software license usage as vApps are deployed and decommissioned, and enforces standardized settings for organization using specific policies called "Blueprints."

VDC—See *Provider Virtual Datacenter* and *Organization Virtual Datacenter.*

VDR (VMware Data Recovery)—Provides data protection for virtual machines. VMware Data Recovery is fully integrated with vCenter Server and includes data deduplication to save on disk storage for full virtual machine backups. Includes file level restore or entire images as needed.

vFabric Cloud Application Platform—Combines the Spring application framework with a set of integrated application services (tc Server, GemFire, RabbitMQ, ERS, Hyperic) to form a comprehensive cloud application platform. VMware vFabric is designed to provide performance and portability across heterogeneous cloud environments.

Virtual switch—A software program emulating a physical switch to enable one virtual machine to communicate with another. See vNetwork Standard Switch and vNetwork Distributed Switch.

VIX—Allows development of programs and scripts to automate virtual machine and guest OS operations. VIX runs on Windows or Linux platforms. It manages vSphere, ESX, ESXi, VMware Server, and VMware Workstation through the use of C, Perl, and COM bindings. COM bindings include Visual Basic, VBscript, and C#.

vMotion—VMware vMotion enables the live migration of running virtual machines from one physical host to the other with zero downtime, continuous service availability, and complete transaction integrity. vMotion migration requires either the same processor family on both the source and target ESX hosts, or "Enhanced vMotion Compatibility" (EVC) on a cluster of hosts with technologies enabling vMotion-compatibility with older servers. Hosts need to be grouped within the same vCenter datacenter. The shared storage holding the VM virtual disk is presented to both the source and target hosts.

VMsafe—Provides an open approach to security through an application program interface (API). This enables selected partners to develop security products for VMware environments. VMsafe gives fine-grained visibility over virtual machine resources, making it possible to monitor every aspect of the execution of the system and stop previously undetectable viruses, root kits, and malware before they can infect a system. VMsafe provides inspection of virtual machine memory pages and CPU states, filtering of network packets inside hypervisors as well as within the virtual machine itself, and in-guest, in-process APIs that enable complete monitoring and control of process execution. Guest virtual machine disk files can be mounted, manipulated, and modified as they persist on storage devices.

vNetwork—Refers to the collection of networking technologies enabling optimal integration of networking and I/O functionality into vSphere. The vNetwork enhancements include the vNetwork Distributed Switch, VMXNET3 (third-generation paravirtualized NIC), IPv6 support extended to vmkernel and service console ports, bi-directional traffic shaping, network vMotion, and VMDirectPath.

vNetwork API—Provides integration with the virtual networking capabilities of vSphere to enable the development of advanced network tools.

vNetwork Distributed Switch (vDS)—Provides a switch which acts at a datacenter level across multiple ESX hosts offering centralized provisioning, administration, and monitoring. Simplifies network management by moving the virtual network configuration and management from the host level to the datacenter level.

vNetwork Standard Switch (vSS)—A software program emulating a physical switch to enable one virtual machine to communicate with another. It is a basic Layer 2 switch without routing.

vShield App—Is used to create security zones within a virtual datacenter. Firewall policies can be applied to vCenter containers or security groups, which are custom containers created through the vShield Manager user interface. Container policies enable the creation of mixed trust zone clusters without requiring an external physical firewall.

vShield Edge—Provides network edge security needed to support multi-tenancy. vShield Edge devices connect the isolated, private networks of cloud tenants to the public side of the provider network through services such as firewall, NAT, DHCP, site-to-site IPsec VPN, and load balancing.

vShield Endpoint—Leverages VMsafe hypervisor introspection capabilities to offload key antivirus and anti-malware functions to a hardened, tamper-proof security virtual machine, eliminating agent footprint. vShield Endpoint consists of a hardened security virtual machine (delivered by VMware partners), a driver for virtual machines to offload file events, and the VMware Endpoint security (EPsec) loadable kernel module (LKM) to link the first two components at the hypervisor layer.

vShield Manager—The centralized network management component of vShield. vShield Manager is a virtual appliance that is imported into the vCenter Server environment. vShield Edges are deployed automatically from VSM when specific network types are created from vCloud Director. Logging and reporting are built in to assist with compliance requirements.

vSphere—VMware vSphere is successor to Virtual Infrastructure 3. Known also as the first cloud operating system, vSphere 4 was one of the most complex software development projects of all time, consisting of over 3,000,000 engineering hours by over 1,000 engineers over a three-year period.

vSphere SDK—A software development kit that acts as an interface for ESXi, vCenter, and VMware Server to extend the management of the virtual datacenter. Programming languages supported include Perl, .NET, and Java.

vSphere Web Services SDK—Provides a Web service accessible through the vSphere API to develop client applications.

vStorage—Provides integration of advanced capabilities from storage vendors with the vSphere Cloud OS from VMware. This API enables customers to leverage array-based capabilities such as support for multipathing control, which enables advanced load balancing algorithms.

References

The following references provide detailed information about VMware vCloud and its components.

VMware Publications and Documentation

VMware vSphere 4

Product Documentation: http://www.vmware.com/support/pubs/vs_pubs.html

Product Overview: http://www.vmware.com/products/vsphere/

vSphere Evaluator's Guide: http://www.vmware.com/resources/techresources/10020

vSphere API: http://communities.vmware.com/community/developer/forums/managementapi

Technical Papers: http://www.vmware.com/resources/techresources/

VMware vCloud

Product Documentation: http://www.vmware.com/support/pubs/vcd_pubs.html

Product Overview: http://www.vmware.com/products/vcloud-director

vCloud Director Evaluator's Guide: http://www.vmware.com/go/tp-vcloud-director-eval-guide

vCloud Director Security Hardening Guide: http://www.vmware.com/resources/techresources/10138

vCloud Director KB Articles: http://blogs.vmware.com/kb/2011/09/vcloud-director-kb-articles.html

vCloud Request Manager: http://www.vmware.com/products/vcloud-request-manager/

vCloud API: http://communities.vmware.com/community/developer/forums/vcloudapi

VMware vCenter Chargeback

Product Documentation: http://www.vmware.com/support/pubs/vcbm_pubs.html

Product Overview: http://www.vmware.com/products/vcenter-chargeback/

Using vCenter Chargeback with vCloud Director: http://www.vmware.com/files/pdf/techpaper/vCenterChargeback_v_1_5_Tech_Note.pdf

Chargeback API: http://communities.vmware.com/community/developer/forums/chargeback

API Programming Guide: http://www.vmware.com/pdf/vCenterChargeback_v_1_5_API_Programming_Guide.pdf

VMware vShield

Product Documentation: http://www.vmware.com/support/pubs/vshield_pubs.html

Product Overview: http://www.vmware.com/products/vshield/

Online Communities

VMware vCloud Director: http://communities.vmware.com/community/vmtn/vcd

VMware vCenter Chargeback: http://communities.vmware.com/community/vmtn/mgmt/chargeback

VMware vShield: http://communities.vmware.com/community/vmtn/vshield

Books

Cloud Computing and Software Services: Theory and Techniques, Syed A. Ahson and Mohammad Ilyas, eds., ISBN 978-1-4398-0315-8, CRC Press, 2010.

Cloud Security: A Comprehensive Guide to Secure Cloud Computing, Ronald L. Krutz and Russell Dean Vines, ISBN 978-0-470-58987-8, Wiley, 2010.

Above the Clouds: Managing Risk in the World of Cloud Computing, Kevin T McDonald, ISBN 978-1-84928-031-0, IT Governance Publishing, 2010.

Foundation for Cloud Computing with VMware vSphere 4, John Arrasjid, Duncan Epping, Steve Kaplan, ISBN 978-1-931971-72-0, USENIX Association Short Topics Series, 2010.

About the Authors and Editor

THE AUTHORS

John Y. Arrasjid is a Principal Architect in the VMware Cloud Services Group. He specializes in Cloud Computing and BC/DR. John presents at VMworld, VMware Partner Exchange, and USENIX/LISA conferences. He serves on the USENIX Association Board of Directors. John is also a founding member of the VMware band Elastic Sky. He is a VMware Certified Professional and one of the first VMware Certified Design Experts (VCDX 001). John can be followed on Twitter at http://twitter.com/vcdx001.

Ben Lin is a Senior Consultant in the VMware Cloud Services Group. Ben works with Enterprise customers and Service Providers in the Americas, providing thought leadership and expertise in architecting cloud solutions. He focuses on IP development; creating services kits, reference architectures, and white papers for field, partner, and public consumption. He specializes in vCloud Director, vCenter Chargeback, and vShield products. Ben holds a Bachelor of Science degree in Electrical Engineering and Computer Science from UC Berkeley and is a VMware Certified Design Expert (VCDX 045). He can be followed on Twitter at http://twitter.com/blin23.

Raman Veeramraju was a Consulting Architect in the VMware Cloud Services Group. Raman primarily worked with service providers in the Americas and provided thought leadership in architecting public clouds. He was focused on IP creation and vCloud infrastructure design, and specialized in vCloud Director, vCenter Lab Manager, and VMware View. Currently he is a Practice Executive at Dell. He holds a Bachelor of Engineering degree in Electronics and Communication and has over 16 years of experience in several LAN/WAN, virtualization, and security technologies. He maintains and often writes on http://virtualyzation.com/ and can be followed on Twitter at http://twitter.com/ramantheman.

Steve Kaplan is Vice President, Data Center Virtualization Practice for INX. Steve has authored scores of articles, white papers, and books and is the author of the *VirtualMan* comic book series. He has spoken on virtualization at venues around the globe and maintains bythebell.com, a blog site emphasizing the economics of virtualization. Steve formerly ran an ROI consultancy and was the first person to approach a utility about providing energy rebates to organizations that virtualize. Steve holds a BS in business administration from UC Berkeley and an MBA from

Northwestern. He can be contacted at Steve.Kaplan@inxi.com or followed on Twitter at http://twitter.com/roidude.

Duncan Epping is a Principal Architect on the VMware Technical Marketing team. Duncan is focused on vStorage initiatives and ESXi adoption. Duncan is a VMware Certified Professional and among the first VMware Certified Design Experts (VCDX 007). Duncan is the owner of Yellow-Bricks.com, one of the leading VMware/virtualization blogs worldwide, and an active contributor and moderator on the VMTN Community Forums. He can be followed on Twitter at http://twitter.com/DuncanYB.

Michael Haines is a Senior Architect (Security) in the VMware Cloud Services Group. Michael principally works with Global Service Providers in Europe. He is focused on security and designing secure cloud infrastructures and specializes in vCloud Director and the vShield Security suite. Michael regularly presents at events and conferences around the globe, such as VMworld and VMware Partner Exchange. Michael is also well known as the author of two identity and security books and a host of papers whose subjects range from virtualization to naming services. He can be followed on Twitter at http://twitter.com/michaelahaines and emailed at mah@vmware.com.

THE EDITOR

Matthew Wood is a Senior Technical Writer and Editor for VMware Technical Services. Matthew works with architects and consultants to produce IP for services kits and solutions kits related to all aspects of VMware technology. He also writes original documentation for the VMware Services Software Solutions Group to support services automation tools such as VMware HealthAnalyzer and Migration Manager. Matthew has 38 years of experience working with technology companies, focusing especially on UNIX, virtualization, and management applications that support enterprise IT environments. He is an avid rock and mineral collector and photographer.

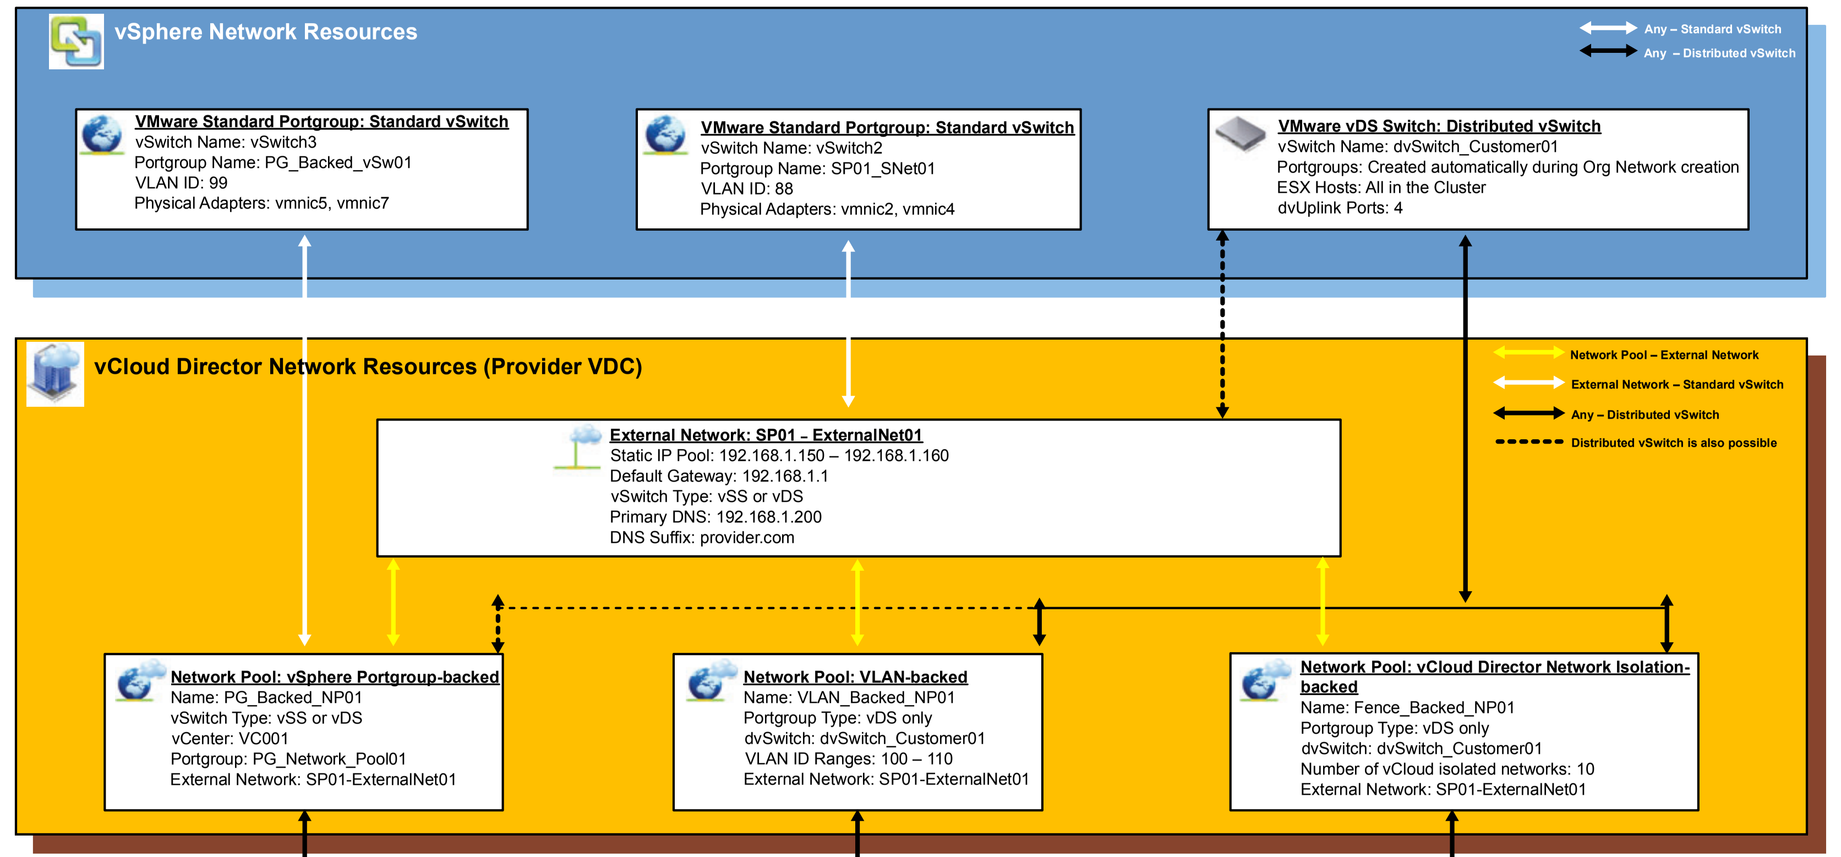
vSphere Network Resources
Any – Standard vSwitch
Any – Distributed vSwitch
VMware Standard Portgroup: Standard vSwitch
vSwitch Name: vSwitch3
Portgroup Name: PG_Backed_vSw01
VLAN ID: 99
Physical Adapters: vmnic5, vmnic7
VMware Standard Portgroup: Standard vSwitch
vSwitch Name: vSwitch2
Portgroup Name: SP01_SNet01
VLAN ID: 88
Physical Adapters: vmnic2, vmnic4
VMware vDS Switch: Distributed vSwitch
vSwitch Name: dvSwitch_Customer01
Portgroups: Created automatically during Org Network creation
ESX Hosts: All in the Cluster
dvUplink Ports: 4
vCloud Director Network Resources (Provider VDC)
Network Pool – External Network
External Network – Standard vSwitch
Any – Distributed vSwitch
Distributed vSwitch is also possible
External Network: SP01 – ExternalNet01
Static IP Pool: 192.168.1.150 – 192.168.1.160
Default Gateway: 192.168.1.1
vSwitch Type: vSS or vDS
Primary DNS: 192.168.1.200
DNS Suffix: provider.com
Network Pool: vSphere Portgroup-backed
Name: PG_Backed_NP01
vSwitch Type: vSS or vDS
vCenter: VC001
Portgroup: PG_Network_Pool01
External Network: SP01-ExternalNet01
Network Pool: VLAN-backed
Name: VLAN_Backed_NP01
Portgroup Type: vDS only
dvSwitch: dvSwitch_Customer01
VLAN ID Ranges: 100 – 110
External Network: SP01-ExternalNet01
Network Pool: vCloud Director Network Isolation-backed
Name: Fence_Backed_NP01
Portgroup Type: vDS only
dvSwitch: dvSwitch_Customer01
Number of vCloud isolated networks: 10
External Network: SP01-ExternalNet01

Raman Veeramraju
Virtualyzation.com

vCloud Director Network Resources (Use Cases)

Mapping
Organization Network Type ←→ External Network
Organization Network Type ←→ Network Pool
cloud01 → name of the instance given during the initial configuration

Public network – direct connection ←→ External Network

- Network Specifications including Static IP Pool, DNS Suffix, Primary DNS are directly obtained from External Network
- When VMs in a vApp are assigned to this network, they have a choice of either DHCP, Static – IP Pool or Static – Manual
- There is no Network Pool in this case

Actions on vCenter during the creation of this network

- No actions on vCenter

Actions on vCenter when the VMs in a vApp are deployed (Direct)

- Reconfigured and powered on the virtual machine
- Assigned an IP Address from External Network Static – IP Pool
- VM is placed under the folder structure → cloud01 / Org VDC / vApp /

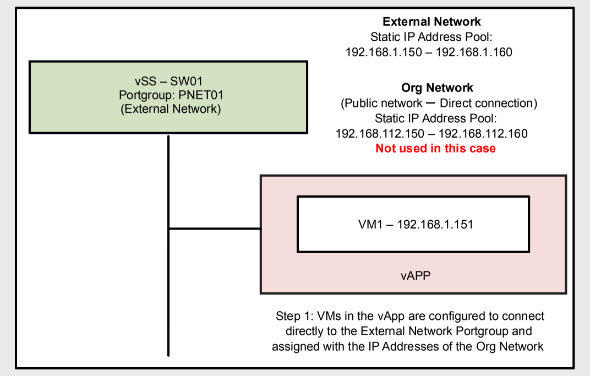

Actions on vCenter when the VMs in a vApp are deployed (Fenced)

- Added a Distributed virtual port group based on Org VDC Network Pool that this vApp is registered to
- Created a folder under Service VMs and placed the vShield Edge under this folder
- Connected the vShield Edge to the External Network on one end and to the vApp Network (new vDS here) on the other end and powered it on
- Connected the VMs in the vApp to the vApp Network and powered them on
- Both interface IP addresses of the vShield Edge are assigned from the Static IP Pool (Network Specifications) of the External Network
- Both Internal and External IP addresses for all the VMs in the vApp are assigned from the Static IP Pool (Network Specifications) of the External Network

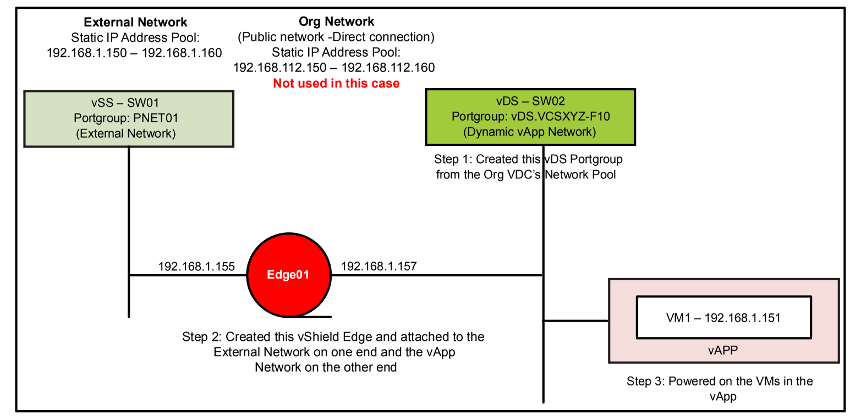

Private network – routed (NAT) connection ←→ vSphere Portgroup-backed

- Choose External Network and Network Pool (Portgroup-backed)
- Network Specifications for the Network Pool (Internal)
- Network Specifications for the External Network will be directly obtained from External Network
- External Management IP address for additional Services such as DHCP, Firewall, Port Forwarding (+ IP Masquerade)
- When VMs in a vApp are assigned to this network, they have a choice of either DHCP, Static – IP Pool or Static – Manual
- When VMs in vAPP are deployed, they will be assigned an Internal IP Address and an External IP Address from Org Network, if deployed in fenced mode

Actions on vCenter during the creation of this network

- Created a folder under Service VMs (00000029-vse-31)
- Registered the vShield Edge virtual machine in this folder (001629-vse_v29n31s2)
- Reconfigured the vShield Edge virtual machine under the Portgroup (vSS or vDS or Cisco Nexus 1KV) created by the admin
- Powered on the vShield Edge virtual machine

Actions on vCenter when the VMs in a vApp are deployed (Direct)

- Reconfigured and powered on the virtual machine
- Assigned an IP Address from Org Network Static – IP Pool
- VM is placed under the folder structure → cloud01 / Org VDC / vApp

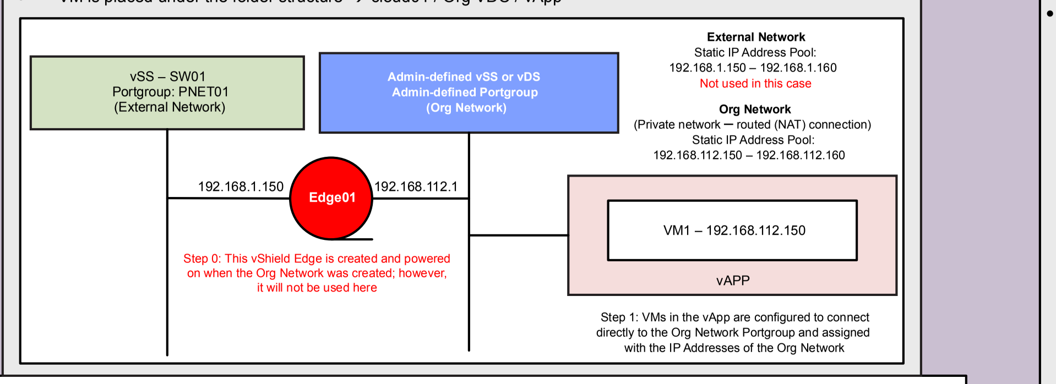

Actions on vCenter when the VMs in a vApp are deployed (Fenced)

- Added a vDS group (Dynamic vApp Network) based on Org VDC Network Pool that this vApp is registered to
- Created a folder under Service VMs and placed the vShield Edge under this folder
- Connected the vShield Edge to the Org Network (vDS portgroup created during the creation of this Org Network) on one end and to the vApp Network (new vDS here) on the other end and powered it on
- Connected the VMs in the vApp to the vApp Network and powered them on
- Both interface IP addresses of the vShield Edge are assigned from the Static IP Pool (Network Specifications) of the Org Network
- Both Internal and External IP addresses for all the VMs in the vApp are assigned from the Static IP Pool (Network Specifications) of the Org Network
- VMs can further be NATted to the External Network using vShield Edge (Edge01)

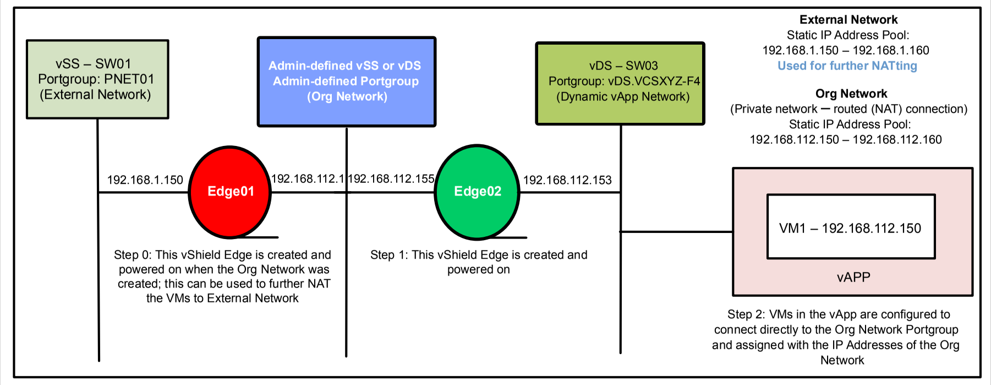

Private network – routed (NAT) connection ←→ VLAN-backed

- Choose External Network and Network Pool (VLAN-backed)
- Network Specifications for the Network Pool (Internal)
- A VLAN will be automatically assigned to this Network Pool
- Network Specifications for the External Network will be directly obtained from External Network
- External Management IP address for additional Services such as DHCP, Firewall, Port Forwarding (+ IP Masquerade)
- When VMs in a vApp are assigned to this network, they have a choice of either DHCP, Static – IP Pool or Static – Manual

Actions on vCenter during the creation of this network

- Added distributed virtual port group (vDS.VC100DVS1CM2-**V13**) to dvSwitch_Customer01 – with a VLAN ID 101
- Created a folder under Service VMs (00000030-vse-34)
- Registered the vShield Edge virtual machine in this folder (001630-vse_v30n34s2)
- Reconfigured the vShield Edge virtual machine
- Powered on the vShield Edge virtual machine
- Port XYZ link was connected in the dvSwitch_Customer01

Actions on vCenter when the VMs in a vApp are deployed (Direct)

- Reconfigured and powered on the virtual machine
- Assigned an IP Address from Org Network Static – IP Pool
- VM is placed under the folder structure → cloud01 / Org VDC / vApp

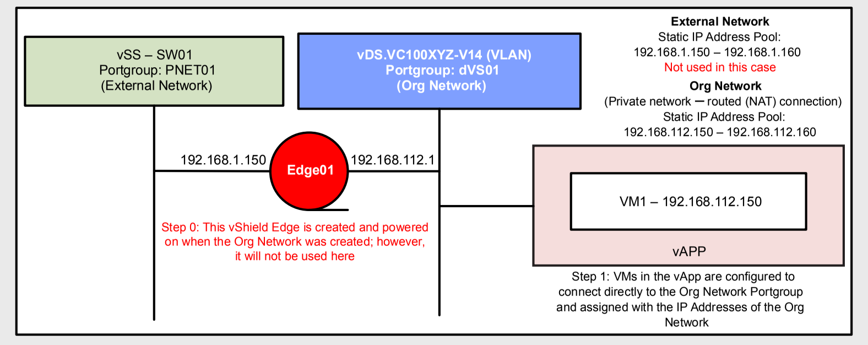

Actions on vCenter when the VMs in a vApp are deployed (Fenced)

- Added a vDS portgroup (Dynamic vApp Network) based on Org VDC Network Pool that this vApp is registered to
- Created a folder under Service VMs and placed the vShield Edge under this folder
- Connected the vShield Edge to the Org Network (vDS portgroup created during the creation of this Org Network) on one end and to the vApp Network (new vDS here) on the other end and powered it on
- Connected the VMs in the vApp to the vApp Network and powered them on
- Both interface IP addresses of the vShield Edge are assigned from the Static IP Pool (Network Specifications) of the Org Network
- Both Internal and External IP addresses for all the VMs in the vApp are assigned from the Static IP Pool (Network Specifications) of the Org Network
- VMs can further be NATted to the External Network using vShield Edge (Edge01)

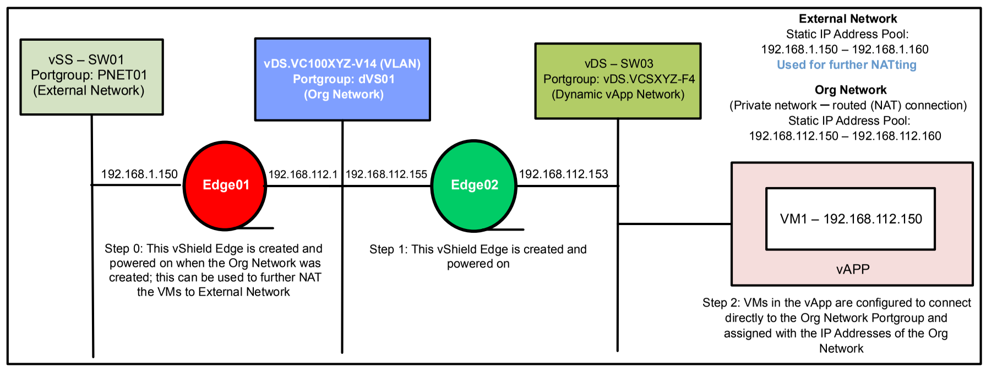

Private network – routed (NAT) connection ←→ vCloud Network Isolated-backed

- Choose External Network and Network Pool (vCloud Network Isolated-backed)
- Network Specifications for the Network Pool (Internal)
- A VLAN will be automatically assigned to this Network Pool
- Network Specifications for the External Network will be directly obtained from External Network
- External Management IP address for additional Services such as DHCP, Firewall, Port Forwarding (+ IP Masquerade)
- When VMs in a vApp are assigned to this network, they have a choice of either DHCP, Static – IP Pool or Static – Manual

Actions on vCenter during the creation of this network

- Added distributed virtual port group (vDS.VC100DVS1CM2-**F2**) to dvSwitch_Customer01
- Created a folder under Service VMs (00000030-vse-34)
- Registered the vShield Edge virtual machine in this folder (001630-vse_v30n34s2)
- Reconfigured the vShield Edge virtual machine
- Powered on the vShield Edge virtual machine
- Port XYZ link was connected in the dvSwitch_Customer01

Actions on vCenter when the VMs in a vApp are deployed (Direct)

- Reconfigured and powered on the virtual machine
- Assigned an IP Address from Org Network Static – IP Pool
- VM is placed under the folder structure → cloud01 / Org VDC / vApp

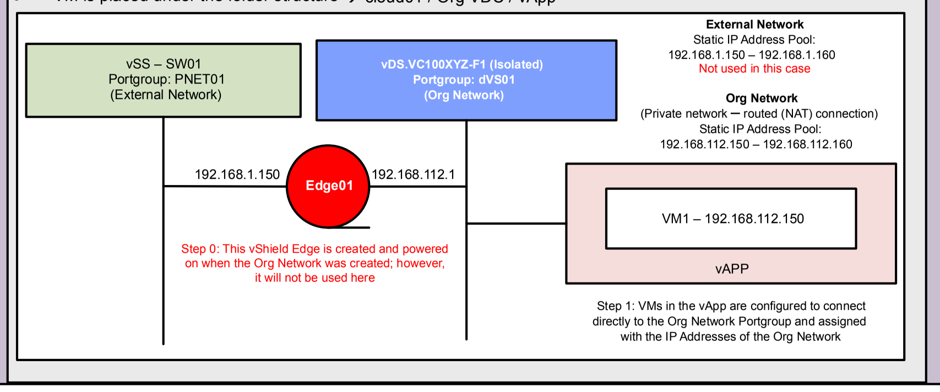
External Network
Static IP Address Pool:
192.168.1.150 – 192.168.1.160
Not used in this case
vSS – SW01
Portgroup: PNET01
(External Network)
vDS.VC100XYZ-F1 (Isolated)
Portgroup: dVS01
(Org Network)
Org Network
(Private network – routed (NAT) connection)
Static IP Address Pool:
192.168.112.150 – 192.168.112.160
192.168.1.150
Edge01
192.168.112.1
Step 0: This vShield Edge is created and powered on when the Org Network was created; however, it will not be used here
VM1 – 192.168.112.150
vAPP
Step 1: VMs in the vApp are configured to connect directly to the Org Network Portgroup and assigned with the IP Addresses of the Org Network

Actions on vCenter when the VMs in a vApp are deployed (Fenced)

- Added a vDS portgroup (Dynamic vApp Network) based on Org VDC Network Pool that this vApp is registered to
- Created a folder under Service VMs and placed the vShield Edge under this folder
- Connected the vShield Edge to the Org Network (vDS portgroup created during the creation of this Org Network) on one end and to the vApp Network (new vDS here) on the other end and powered it on
- Connected the VMs in the vApp to the vApp Network and powered them on
- Both interface IP addresses of the vShield Edge are assigned from the Static IP Pool (Network Specifications) of the Org Network
- Both Internal and External IP addresses for all the VMs in the vApp are assigned from the Static IP Pool (Network Specifications) of the Org Network
- VMs can further be NATted to the External Network using vShield Edge (Edge01)

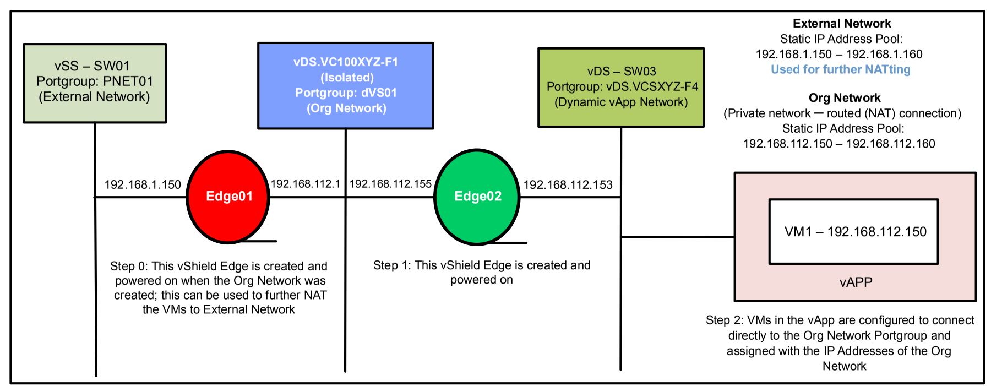

Private network – isolated ←→ vSphere Portgroup-backed

- Will be able to choose only Network Pool
- Network Specifications for the Network Pool
- Under Network Services Configuration, only DHCP is enabled; since this is an isolated network, Firewall, Port Forwarding, External IP Management will be disabled
- When VMs in a vApp are assigned to this network, they have a choice of either DHCP, Static – IP Pool or Static – Manual

Actions on vCenter during the creation of this network

- Created a folder under Service VMs (00000029-vse-31)
- Registered the vShield Edge virtual machine in this folder (001629-vse_v29n31s2)
- Reconfigured the vShield Edge virtual machine under the Portgroup (vSS or vDS or Cisco Nexus 1KV) created by the admin
- Powered on the vShield Edge virtual machine

Actions on vCenter when the VMs in a vApp are deployed (Direct)

- Reconfigured and powered on the virtual machine
- Assigned an IP Address from Org Network Static – IP Pool
- VM is placed under the folder structure→ cloud01 / Org VDC / vApp

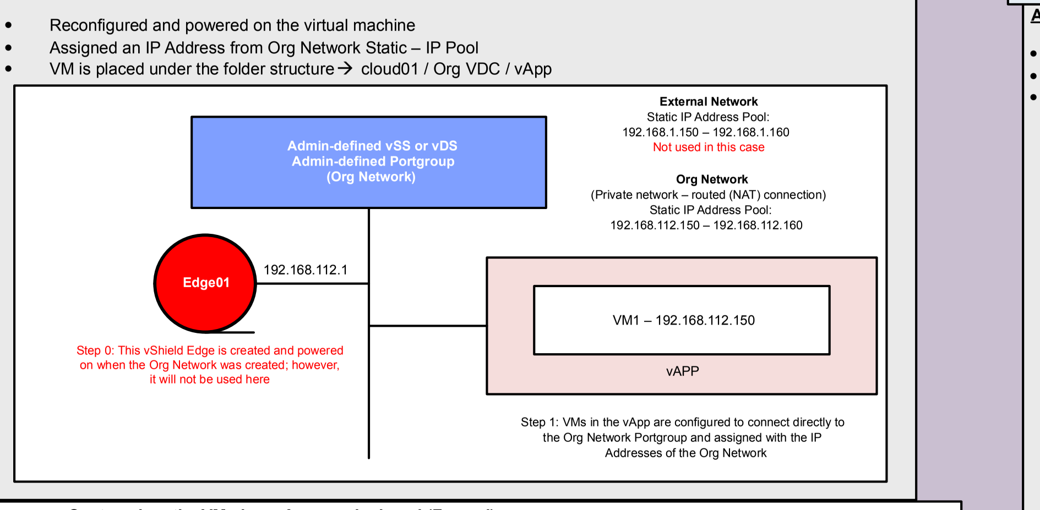
External Network
Static IP Address Pool:
192.168.1.150 – 192.168.1.160
Not used in this case
Admin-defined vSS or vDS
Admin-defined Portgroup
(Org Network)
Org Network
(Private network – routed (NAT) connection)
Static IP Address Pool:
192.168.112.150 – 192.168.112.160
Edge01
192.168.112.1
Step 0: This vShield Edge is created and powered on when the Org Network was created; however, it will not be used here
VM1 – 192.168.112.150
vAPP
Step 1: VMs in the vApp are configured to connect directly to the Org Network Portgroup and assigned with the IP Addresses of the Org Network

Actions on vCenter when the VMs in a vApp are deployed (Fenced)

- Added a vDS portgroup (Dynamic vApp Network) based on Org VDC Network Pool that this vApp is registered to
- Created a folder under Service VMs and placed the vShield Edge under this folder
- Connected the vShield Edge to the Org Network (vDS portgroup created during the creation of this Org Network) on one end and to the vApp Network (new vDS here) on the other end and powered it on
- Connected the VMs in the vApp to the vApp Network and powered them on
- Both interface IP addresses of the vShield Edge are assigned from the Static IP Pool (Network Specifications) of the Org Network
- Both Internal and External IP addresses for all the VMs in the vApp are assigned from the Static IP Pool (Network Specifications) of the Org Network
- VMs can be further NATted to the External Network via vShield Edge (Edge01)

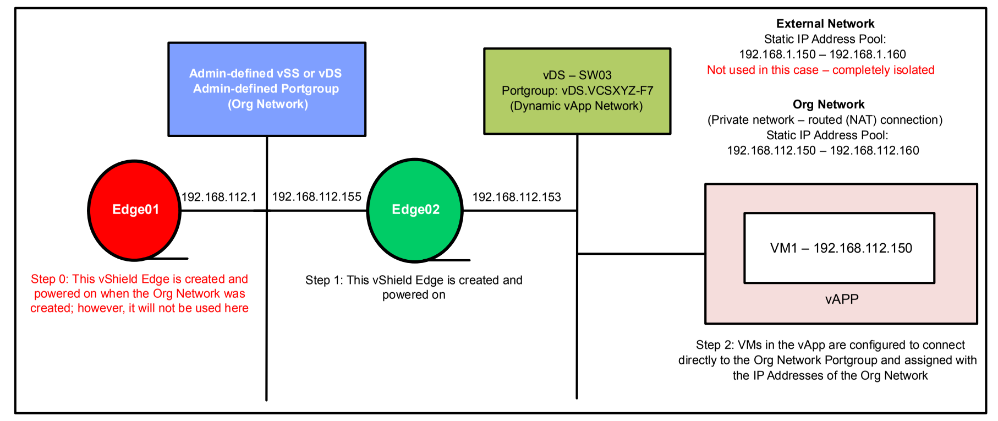
External Network
Static IP Address Pool:
192.168.1.150 – 192.168.1.160
Not used in this case – completely isolated
Admin-defined vSS or vDS
Admin-defined Portgroup
(Org Network)
vDS – SW03
Portgroup: vDS.VCSXYZ-F7
(Dynamic vApp Network)
Org Network
(Private network – routed (NAT) connection)
Static IP Address Pool:
192.168.112.150 – 192.168.112.160
Edge01
192.168.112.1
192.168.112.155
Edge02
192.168.112.153
Step 0: This vShield Edge is created and powered on when the Org Network was created; however, it will not be used here
Step 1: This vShield Edge is created and powered on
VM1 – 192.168.112.150
vAPP
Step 2: VMs in the vApp are configured to connect directly to the Org Network Portgroup and assigned with the IP Addresses of the Org Network

Private network – isolated ←→ VLAN-backed

- Will be able to choose only Network Pool
- Network Specifications for the Network Pool
- Under Network Services Configuration, only DHCP is enabled; since this is an isolated network, Firewall, Port Forwarding, External IP Management will be disabled
- When VMs in a vApp are assigned to this network, they have a choice of either DHCP, Static – IP Pool or Static – Manual

Actions on vCenter during the creation of this network

- Added distributed virtual port group (vDS.VC100DVS1CM2-**V12**) to dvSwitch_Customer01
- Created a folder under Service VMs
- Registered the vShield Edge virtual machine in this folder
- Reconfigured the vShield Edge virtual machine
- Powered on the vShield Edge virtual machine
- Port XYZ link was connected in the dvSwitch_Customer01

Actions on vCenter when the VMs in a vApp are deployed (Direct)

- Reconfigured and powered on the virtual machine
- Assigned an IP Address from Org Network Static – IP Pool
- VM is placed under the folder structure → cloud01 / Org VDC / vApp

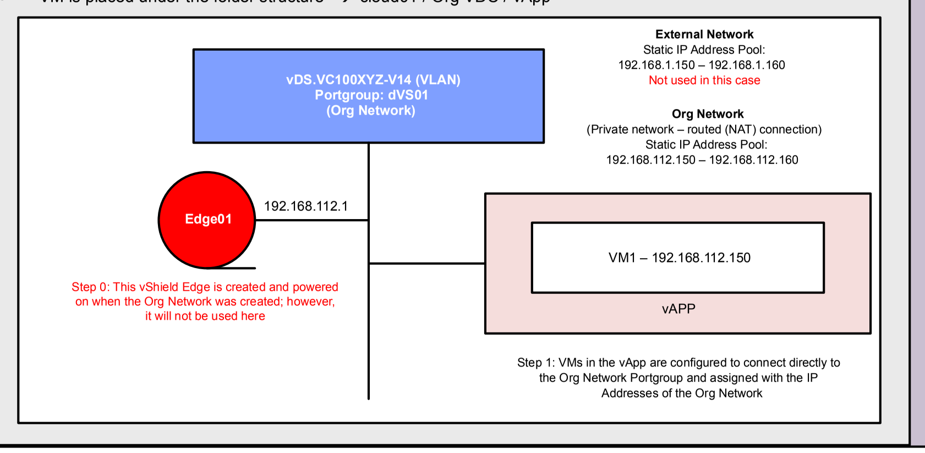
External Network
Static IP Address Pool:
192.168.1.150 – 192.168.1.160
Not used in this case
vDS.VC100XYZ-V14 (VLAN)
Portgroup: dVS01
(Org Network)
Org Network
(Private network – routed (NAT) connection)
Static IP Address Pool:
192.168.112.150 – 192.168.112.160
Edge01
192.168.112.1
Step 0: This vShield Edge is created and powered on when the Org Network was created; however, it will not be used here
VM1 – 192.168.112.150
vAPP
Step 1: VMs in the vApp are configured to connect directly to the Org Network Portgroup and assigned with the IP Addresses of the Org Network

Actions on vCenter when the VMs in a vApp are deployed (Fenced)

- Added a vDS portgroup (Dynamic vApp Network) based on Org VDC Network Pool that this vApp is registered to
- Created a folder under Service VMs and placed the vShield Edge under this folder
- Connected the vShield Edge to the Org Network (vDS portgroup created during the creation of this Org Network) on one end and to the vApp Network (new vDS here) on the other end and powered it on
- Connected the VMs in the vApp to the vApp Network and powered them on
- Both interface IP addresses of the vShield Edge are assigned from the Static IP Pool (Network Specifications) of the Org Network
- Both Internal and External IP addresses for all the VMs in the vApp are assigned from the Static IP Pool (Network Specifications) of the Org Network
- VMs can be further NATted to the External Network via vShield Edge (Edge01)

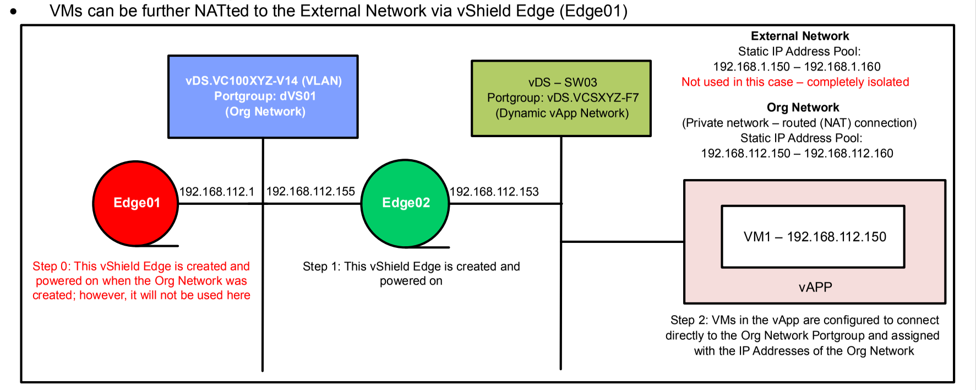
External Network
Static IP Address Pool:
192.168.1.150 – 192.168.1.160
Not used in this case – completely isolated
vDS.VC100XYZ-V14 (VLAN)
Portgroup: dVS01
(Org Network)
vDS – SW03
Portgroup: vDS.VCSXYZ-F7
(Dynamic vApp Network)
Org Network
(Private network – routed (NAT) connection)
Static IP Address Pool:
192.168.112.150 – 192.168.112.160
Edge01
192.168.112.1
192.168.112.155
Edge02
192.168.112.153
Step 0: This vShield Edge is created and powered on when the Org Network was created; however, it will not be used here
Step 1: This vShield Edge is created and powered on
VM1 – 192.168.112.150
vAPP
Step 2: VMs in the vApp are configured to connect directly to the Org Network Portgroup and assigned with the IP Addresses of the Org Network

Private network – isolated ←→ vCloud Network Isolated-backed

- Will be able to choose only Network Pool
- Network Specifications for the Network Pool
- Under Network Services Configuration, only DHCP is enabled; since this is an isolated network, Firewall, Port Forwarding, External IP Management will be disabled
- When VMs in a vApp are assigned to this network, they have a choice of either DHCP, Static – IP Pool or Static – Manual

Actions on vCenter during the creation of this network

- Added distributed virtual port group (vDS.VC100DVS1CM2-**F2**) to dvSwitch_Customer01
- Created a folder under Service VMs
- Registered the vShield Edge virtual machine in this folder
- Reconfigured the vShield Edge virtual machine
- Powered on the vShield Edge virtual machine
- Port XYZ link was connected in the dvSwitch_Customer01

Actions on vCenter when the VMs in a vApp are deployed (Direct)

- Reconfigured and powered on the virtual machine
- Assigned an IP Address from Org Network Static – IP Pool
- VM is placed under the folder structure → cloud01 / Org VDC / vApp

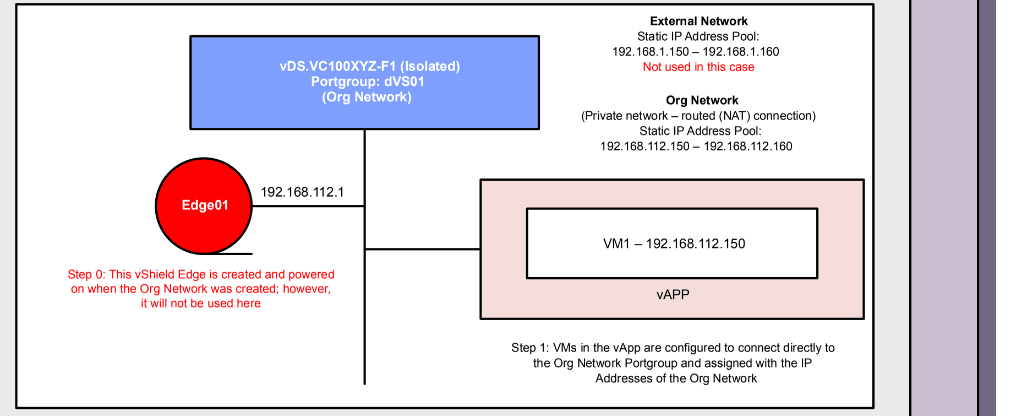
External Network
Static IP Address Pool:
192.168.1.150 – 192.168.1.160
Not used in this case
vDS.VC100XYZ-F1 (Isolated)
Portgroup: dVS01
(Org Network)
Org Network
(Private network – routed (NAT) connection)
Static IP Address Pool:
192.168.112.150 – 192.168.112.160
Edge01
192.168.112.1
Step 0: This vShield Edge is created and powered on when the Org Network was created; however, it will not be used here
VM1 – 192.168.112.150
vAPP
Step 1: VMs in the vApp are configured to connect directly to the Org Network Portgroup and assigned with the IP Addresses of the Org Network

Actions on vCenter when the VMs in a vApp are deployed (Fenced)

- Added a vDS portgroup (Dynamic vApp Network) based on Org VDC Network Pool that this vApp is registered to
- Created a folder under Service VMs and placed the vShield Edge under this folder
- Connected the vShield Edge to the Org Network (vDS portgroup created during the creation of this Org Network) on one end and to the vApp Network (new vDS here) on the other end and powered it on
- Connected the VMs in the vApp to the vApp Network and powered them on
- Both interface IP addresses of the vShield Edge are assigned from the Static IP Pool (Network Specifications) of the Org Network
- Both Internal and External IP addresses for all the VMs in the vApp are assigned from the Static IP Pool (Network Specifications) of the Org Network
- VMs can be further NATted to the External Network via vShield Edge (Edge01)

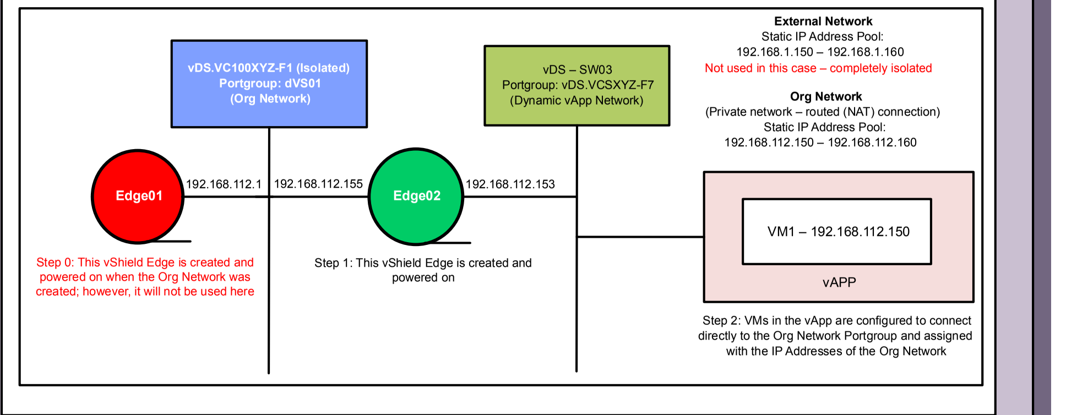
External Network
Static IP Address Pool:
192.168.1.150 – 192.168.1.160
Not used in this case – completely isolated
vDS.VC100XYZ-F1 (Isolated)
Portgroup: dVS01
(Org Network)
vDS – SW03
Portgroup: vDS.VCSXYZ-F7
(Dynamic vApp Network)
Org Network
(Private network – routed (NAT) connection)
Static IP Address Pool:
192.168.112.150 – 192.168.112.160
Edge01
192.168.112.1
192.168.112.155
Edge02
192.168.112.153
Step 0: This vShield Edge is created and powered on when the Org Network was created; however, it will not be used here
Step 1: This vShield Edge is created and powered on
VM1 – 192.168.112.150
vAPP
Step 2: VMs in the vApp are configured to connect directly to the Org Network Portgroup and assigned with the IP Addresses of the Org Network