



Public Cloud Service Definition

Public Version 1.5

TECHNICAL WHITE PAPER

Table Of Contents

Introduction	3
Enterprise Hybrid Cloud	3
Public Cloud Service Definition	4
VMware vCloud Datacenter Services	4
Target Markets and Use Cases	4
Challenges Solved	5
Service Definition	5
Service Offerings	6
Basic VDC	6
Committed VDC	6
Dedicated VDC	7
Compliance Definition	7
Compliance Controls	8
Compliance Visibility and Transparency	10
Compliant Architecture	11
Architecture Definition	11

Introduction


Cloud computing delivers convenient, on-demand access to shared pools of data, applications and hardware. The cloud computing paradigm—made possible by sophisticated automation, provisioning and virtualization technologies—differs dramatically from today's IT model because it decouples data and software from the servers and storage systems running them and enables IT resources to be dynamically allocated and delivered as a service, either in component parts—where users subscribe to specific applications or simply lease computing power—or as an integrated whole.

Cloud computing is changing the way IT resources are utilized. Users want the ability to access infrastructure resources how and when they choose. IT teams are asked to accommodate this shift in the consumption model but still deal with the security, compatibility and compliance issues associated with delivering that convenience to application business owners and developers.

Enterprise Hybrid Cloud

Private cloud is the cloud infrastructure operated solely for an organization. It can be managed by the organization or a third party and can exist on-premises or off-premises. *Public cloud* is the cloud infrastructure made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud is the cloud infrastructure composed of two or more clouds, private or public, that remain unique entities but are bound together by standardized technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).



The common misperception is that cloud computing implies an “external” cloud, based on public cloud services. The fact is that cloud computing is how you approach information technology. It is “a way of doing computing,” not a destination. Ultimately, most enterprises will benefit from adopting cloud computing within their own datacenters, building “private clouds,” and getting there in an evolutionary way through their existing virtualization journey.

Together with our leading service provider partners, we can also make hybrid clouds a reality, through a common platform built around VMware vSphere™ 4.1 (“vSphere”) and VMware vCloud™ Director, with common management and security models to give the enterprise the confidence they need, in an environment that provides on-demand application portability.

Public Cloud Service Definition

To run a public cloud, a service provider will first need to define the services that will be offered to enterprises that want to place their workloads in the cloud. VMware® works with service providers through several programs, most notably the VMware vCloud™ Datacenter Services program, to help ensure a minimum level of cloud service capabilities.

VMware vCloud Datacenter Services

VMware vCloud Datacenter Services is a VMware program designed to raise the bar and define a whole new type of enterprise-class cloud computing infrastructure services. In a market currently comprising commodity, low-performance, insecure public cloud offerings, vCloud Datacenter Services define a new enterprise-class cloud computing segment. It is the hybrid cloud solution for enabling enterprises to extend their private cloud to the public cloud with flexibility, scalability, security and operational efficiency. Through a common platform built around vSphere and VMware vCloud Director, with common management and security models, in an environment that provides on-demand application portability, enterprise customers and leading global service providers are delivering cloud-compatible, connected and integrated hybrid clouds.

VMware vCloud Datacenter Services are cobranded by the service provider and VMware; they initially were offered by a small number of service providers worldwide. This gives the enterprise customers of VMware a choice of 100-percent compatible services that are based on VMware vCloud architecture and are certified by VMware.

Target Markets and Use Cases

VMware vCloud Datacenter Services are designed to serve the corporate and departmental IT teams inside medium to large enterprises (1,000 to 50,000+ employees), and the government/federal sectors. The service enables these IT teams to augment their private cloud with public cloud capacity, in order to support preproduction workloads such as test and development workloads, as well as production workloads such as Web applications, marketing/brochure sites, and messaging/collaboration applications.

The service will support both new and existing workloads. The following are some—but not all — use case examples of actual cloud workloads:

- Preproduction environments
 - Develop/test/stage
 - Functional testing
 - Continuous integration
 - Training labs or demo environments
- Production environments
 - Web applications
 - Marketing/brochure sites
 - Multitiered Web applications
 - E-commerce Web sites
 - Corporate portals and intranet sites
 - Messaging and collaboration applications
 - SharePoint
 - Content/document management
 - Wikis/blogs

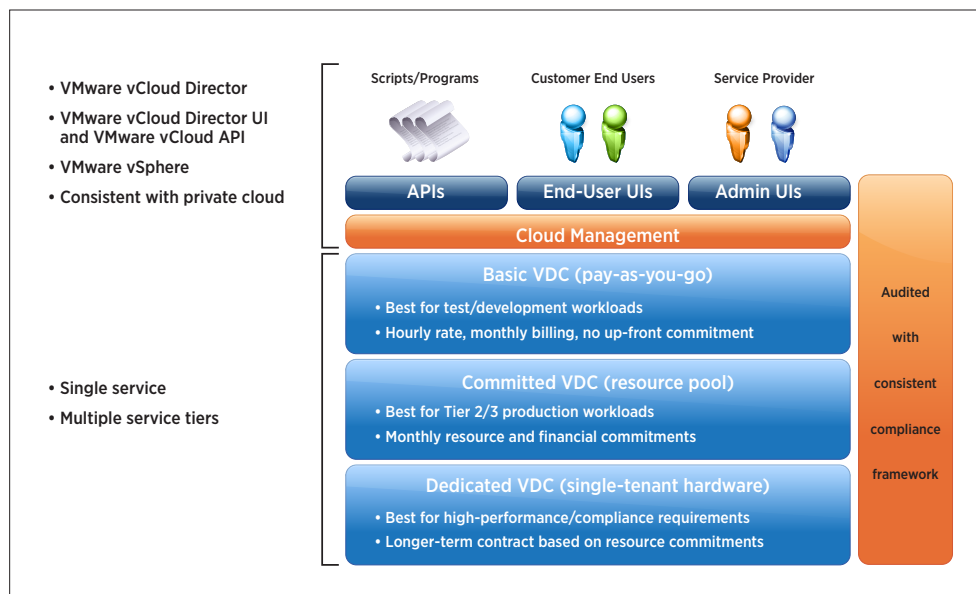
Challenges Solved

VMware solutions help resolve the following challenges to adopting the public cloud that exist for enterprises today:

- Trust/security
 - Alignment with existing processes and tools
 - Regulatory and standards compliance
 - Secure connectivity
 - Data location
- Management
 - Consistent identity-and-access management
 - Single-pane-of-glass resource management
 - Compatible platforms and API

Service Definition

The public cloud will need a service definition and service offerings. As an example reference, VMware vCloud Datacenter Services deliver three classes of on-demand, self-service virtual datacenters (VDCs):



The service is designed to make it as easy as possible for enterprises to move their workloads to vCloud Datacenter Services. Any existing VMware virtual machine or VMware virtual application (VMware vApp) can be run with little or no modification on vCloud Datacenter Services, and compatibility with existing enterprise VMware deployments is a key design objective. There is no requirement for an enterprise to deploy a private cloud—any VMware virtualized infrastructure is compatible.

All VMware vCloud Datacenter Services infrastructures and datacenters have been audited against a standard set of compliance controls for SAS 70 Type II or have received ISO 27000 certifications. In addition, all vCloud Datacenter Services will provide customers with relevant audit logs and compliance reports for their cloud environments, to ensure that enterprises can meet their own internal audit requirements.

Service Offerings

VMware vCloud Datacenter Services consist of three different service offerings. A single customer can have one or more of the three offerings:

- Basic VDC – provides unreserved “pay per use.” It is designed for quick-start pilot projects, or for workloads such as software testing that don’t need resource reservations/guarantee.
- Committed VDC – provides reserved compute resources (subscription model) with the ability to burst above committed levels if additional capacity is available. It offers predictable performance by reserving resources for workloads within a multitenant infrastructure, while also enabling on-demand self-service.
- Dedicated VDC – provides dedicated compute resources (using specific, dedicated hardware), sometimes known as *virtual private cloud*. It offers predictable performance by reserving dedicated resources, which is useful for situations where security or compliance requirements require physical separation.

	BASIC VDC	COMMITTED VDC	DEDICATED VDC
CONSUMPTION MODEL	Pay-as-you-go	Allocation pool	Reservation pool
CHARGEABLE UNIT	CPU, memory, storage	Resource pool	Resource pool
TARGETED USE CASES (BUT NOT LIMITED TO)	Preproduction (test/develop/stage)	Tier 2/3 production	Production, high-performance, security, compliance requirements
METERING	Hourly	Monthly	Monthly

Basic VDC

The Basic VDC service offering is an instance-based, pay-as-you-go resource consumption model. Each virtual machine provisioned in this VDC is charged separately, and separate billing records are produced for each virtual machine. Customers using the Basic VDC service will be charged for each hour or partial hour of consumption. For example, if a customer uses a machine for 5 minutes, they will be charged for one hour of usage. If a customer changes the virtual machine size after 5 minutes, this starts a new hour. Instance-based model refers to the bundling of vCPU and memory together into a single virtual machine instance and price.

Committed VDC

The Committed VDC service offering currently uses the allocation-pool consumption model (see VMware vCloud Director documentation for more details on allocation-pool model). A user is allocated a VDC that contains a certain amount of CPU (GHz), memory (GB) and storage (GB). The allocation-pool model is defined using two parameters: the reservation percentage and the total allocation (also called limit). The reservation percentage is how much in resources will be guaranteed or committed for the customer. The total allocation, or limit, is the maximum amount of resources the customer can consume.

For the Committed VDC, the reservation percentage is initially set to 75 percent of the total allocation/limit. This means the customer can burst up to an additional 25 percent of resources they originally requested. So, for example, if a customer buys a VDC with 10GHz of CPU resources, a VDC is created for the customer and 10GHz is allocated for the customer. This is the maximum amount of CPU the customer can ever consume. Of this 10GHz, the service provider should reserve 75 percent, which is 7.5GHz. This is the amount of CPU that’s guaranteed for the customer. The 25 percent, or 2.5GHz, will be available to the customer if the underlying cluster has available resources. This model enables service providers to charge a price that’s generally higher than for just 7.5GHz, given the additional burstable capacity, but it gives customers the benefit of potentially paying less than if the resources were fully guaranteed.

In the Committed VDC model, we have the following VDC sizes defined.

	SMALL	MEDIUM	LARGE	X-LARGE
CPU RESERVED	7.5GHz	18.75GHz	37.5GHz	75GHz
CPU LIMIT	10GHz	25GHz	50GHz	100GHz
MEMORY RESERVED	15GB	37.5GB	75GB	150GB
MEMORY LIMIT	20GB	50GB	100GB	200GB
STORAGE	400GB	1TB	3TB	6TB
VIRTUAL MACHINE LIMIT (CAN BE CHANGED)	20	50	100	200
APPROXIMATE VIRTUAL MACHINES (NOT LIMIT)	10-20	25-50	50-100	100+

The approximate virtual machine count is calculated based on the distribution ratio assumption listed at the end of this document. In most virtualized environments, memory is the gating factor in terms of resource utilization. Based on the distribution ratio that follows, we calculate that on average a virtual machine consumes 2.15GB of memory. For a 20GB VDC, that's approximately 10 virtual machines. If a customer runs more small virtual machines, that count can increase. So setting the virtual machine limit to be around 20 will give customers some flexibility to run more small virtual machines but still effectively manage the capacity.

Dedicated VDC

The Dedicated VDC service offering uses the reservation pool-based model. A customer works with a service provider to provision a cluster of servers dedicated to this customer. The hardware (network, storage, servers) is not shared with other customers. The customer gets full control over the reservation and limit of this set of resources. This service offering will be a fixed-price monthly subscription.

Compliance Definition

The security and compliance issue continues to be one of the biggest barriers for enterprise customers in adopting the public cloud. Most regulations and mandates in the industry, including SOX, PCI, HIPAA, COBIT and ISO, have two areas of requirements: transparency/visibility and control. Transparency is essential; cloud consumers must know “who” has accessed what data, as well as “when,” “where” and potentially “why,” based on documented evidence. PCI requirement #10 is a good example of the need for visibility and transparency. Control is also a necessary component of compliance for cloud consumers. For example, cloud consumers must be able to control who can access, configure and modify the cloud environment; what firewall ports are open; when to apply patches; and where the data resides. Cloud consumers, and especially enterprise customers, believe that “you can outsource responsibility, but you can’t outsource accountability.” At the end of the day, cloud consumers still are accountable for being compliant.

VMware vCloud Datacenter Services are designed to tackle this very problem. Along with service provider partners, they will do so in three areas:

- Ensure compliance through ISO 27001 certification or SAS 70 Type II audit, based on a standard set of controls
- Provide compliance logging and reports to customers so they have full visibility into their public cloud environments
- Architect the service so customers can have control of access to their cloud environments

The following sections provide a set of high-level requirements for VMware vCloud Datacenter Services. For detailed security and compliance implementation guidance, refer to the following documents:

- *VMware vCloud Control Matrix and Threat Model*
- *VMware vCloud Secure Deployment Guide*
- *VMware vCloud Compliance Logging Guide*
- *VMware vCloud Compliance Logging Reference Implementation*

Compliance Controls

In order to ensure that enterprise customers feel secure and safe in the public cloud, and that they have the necessary information and visibility into the service to meet their own internal audit requirements, VMware vCloud Datacenter Services must have one of the following:

- ISO 27001 certification, which proves that security management processes are in place, as well as a relevant subset of the ISO 27002 controls in place
- SAS 70 Type II audits based on the same standard, relevant set of compliance controls

VMware will supply the standard set of compliance controls (see following table), and the service provider is responsible for the actual ISO or SAS 70 audits, using third-party auditors. The compliance controls will be published to the enterprise customers so they understand that not only are vCloud Datacenter Services compliant but also that customers have full visibility into which controls the services were audited against.

ISO 27002	CONTROL TITLE
A.06.2.2	Addressing security when dealing with customers
A.07.1.1	Inventory of assets
A.07.1.2	Ownership of assets
A.07.2.1	Classification guidelines
A.07.2.2	Information labeling and handling
A.08.1.1	Roles and responsibilities
A.08.3.3	Removal of access rights
A.09.2.4	Equipment maintenance
A.09.2.6	Secure disposal or reuse of equipment
A.10.01.1	Documented operating procedures
A.10.01.2	Change management
A.10.01.3	Segregation of duties
A.10.01.4	Separation of develop, test and ops
A.10.03.1	Capacity management
A.10.03.2	System acceptance
A.10.04.1	Controls against malicious code
A.10.04.2	Controls against mobile code
A.10.05.1	Information backup

ISO 27002	CONTROL TITLE
A.10.06.1	Network controls
A.10.06.2	Security of network services
A.10.07.1	Management of removable media
A.10.07.2	Disposal of media
A.10.07.3	Information handling procedures
A.10.07.4	Security of system documentation
A.10.08.4	Electronic messaging
A.10.08.5	Business information systems
A.10.10.1	Audit logging
A.10.10.2	Monitoring system use
A.10.10.3	Protection of log information
A.10.10.4	Administrator and operator logs
A.10.10.5	Fault logging
A.10.10.6	Clock synchronization
A.11.1.1	Access control policy
A.11.2.1	User registration
A.11.2.2	Privilege management
A.11.2.3	User password management
A.11.2.4	Review of user access rights
A.11.3.1	Password use
A.11.4.2	User authentication for external connections
A.11.4.3	Equipment identification in networks
A.11.4.4	Remote diagnostic and configuration port protection
A.11.4.5	Segregation in networks
A.11.4.6	Network connection control
A.11.4.7	Network routing control
A.11.5.1	Secure log-on procedures
A.11.5.2	User identification and authentication
A.11.5.3	Password management system
A.11.5.4	Use of system utilities
A.11.5.5	Session time-out
A.11.5.6	Limitation of connection time
A.11.6.1	Information access restriction
A.11.6.2	Sensitive system isolation
A.12.1.1	Security requirements analysis and specification

ISO 27002	CONTROL TITLE
A.12.2.1	Input data validation
A.12.2.2	Control of internal processing
A.12.2.3	Message integrity
A.12.2.4	Output data validation
A.12.3.2	Key management
A.12.4.1	Control of operational software
A.12.4.2	Protection of system test data
A.12.5.1	Change control procedures
A.12.5.2	Technical review of apps after operating system changes
A.12.5.3	Restrictions on changes to software packages
A.12.5.4	Information leakage
A.12.6.1	Control of technical vulnerabilities
A.13.1.1	Reporting information security events
A.13.1.2	Reporting security weaknesses
A.13.2.3	Collection of evidence
A.14.1.2	Business continuity and risk assessment
A.14.1.3	Developing and implementing continuity plans
A.14.1.5	Testing, maintaining and reassessing plans
A.15.2.2	Technical compliance checking

Compliance Visibility and Transparency

Log management is built into many of the compliance frameworks such as ISO, HIPAA, PCI and COBIT. It is required to meet the requirements of these audit standards. Enterprise customers not only need visibility into their private clouds but also demand that service providers give them visibility into their public cloud environments. For example, enterprise customers are looking for all necessary logs and reports regarding user activities, access control, firewall connections and so on.

To meet the requirements of compliance with the controls listed previously, service providers must give cloud consumers visibility and transparency and include them in the service definition for their public cloud. To accomplish this, service providers must be able to supply relevant logs to cloud consumers. In general, services should have logs covering the following areas:

- Identity and access management
- User activities monitoring
- Change and configuration management
- Security and threat management
- Business continuity and availability management

- For reference, VMware vCloud Datacenter Services are based on a set of products that have been battle tested in many secure environments. Products such as VMware vCloud Director and VMware vShield Edge generate a set of logs that give customers visibility into all user activities and firewall connections. VMware will provide the necessary blueprints and best practices so that service providers can capture this set of logs and give customers the ability to download them. The logs must be available to customers for a minimum of six months.

In addition to the logs, service providers should supply basic compliance reports to enterprise customers so they understand all the activities inside their cloud environment. VMware will provide a set of best practices in this area to ensure that VMware vCloud Datacenter Services meet enterprise customer requirements.

See *VMware vCloud Compliance Logging Guide* and *VMware vCloud Compliance Logging Reference Implementation* for additional information.

Compliant Architecture

All VMware vCloud Datacenter Services offer unparalleled security. They are built on vSphere, the most secure virtualization platform with ELA4+ (vSphere 4.1 in progress) and FISMA certifications, and VMware vCloud Director, a cloud delivery platform offering secure multitenancy and organization isolation. With vCloud Datacenter Services, enterprises can exercise the defense-in-depth security best practice because the platform offers both per-organization firewalls and per-vApp firewalls. And all organizations are isolated with their own Layer 2 networks. Access and authentication can be performed against the enterprise's own LDAP/AD directory, which means that the enterprise can manage its own user base and provide role-based access according to its own policies.

VMware will provide a full set of architecture blueprints so service providers can implement a secure enterprise-class cloud infrastructure. See the following "Architecture Definition" section for more details.

Architecture Definition

In order to take full advantage of the hybrid cloud, VMware vCloud reference architecture is designed to be compatible with the private cloud stack VMware is advocating for enterprises. To support service providers in implementing this service, VMware supplies a full set of reference blueprints to them. This set of blueprints includes all documentation and best practices on understanding, architecting, sizing and implementing an enterprise-class cloud infrastructure. It represents not only hundreds of person-years of product knowledge but also many person-years of knowledge in building out scalable virtual and cloud infrastructures.

The VMware vCloud reference architecture is designed with the following products:

- VMware vCloud Director 1.0
- VMware vSphere 4.1
- VMware vCenter™ Chargeback 1.5
- VMware vShield Edge, which is embedded in VMware vCloud Director

VMware vCloud Datacenter Services also display the VMware vCloud Director user interface as the main provisioning and management interface for end users, as well as VMware vCloud API for automation.

The VMware vCloud reference architecture will include key documents for enabling partner and customer implementation of hybrid clouds. The key components are:

DOCUMENT	DESCRIPTION
Architecting a VMware vCloud	Defines the architecture and design for implementing VMware vCloud Director, built on vSphere, as well as the components and relationships to run a cloud
Reference Implementation	Actual instance of a VMware vCloud implementation at a detailed component level
Secure Deployment Guide	A detailed document on how to deploy VMware vCloud Director in a secure manner

See the previously referenced documents for detailed descriptions.

