CHAPTER 1

Design Overview

This document presents design and implementation guidance for a private or public IaaS cloud data center. The intended deployment uses a multi-tenant, differentiated service tier model.

Cloud Tenants

A tenant is an entity that subscribes to cloud services. In the Enterprise private cloud deployment model, that entity is a department or sub-organization, such as development, test, research and development, or human resources. Figure 1-1 shows multiple users in the same department belong to the same tenancy. Within the tenancy, multiple workloads can be implemented by different users who belong to the same department.

Figure 1-1 Tenants and Workloads

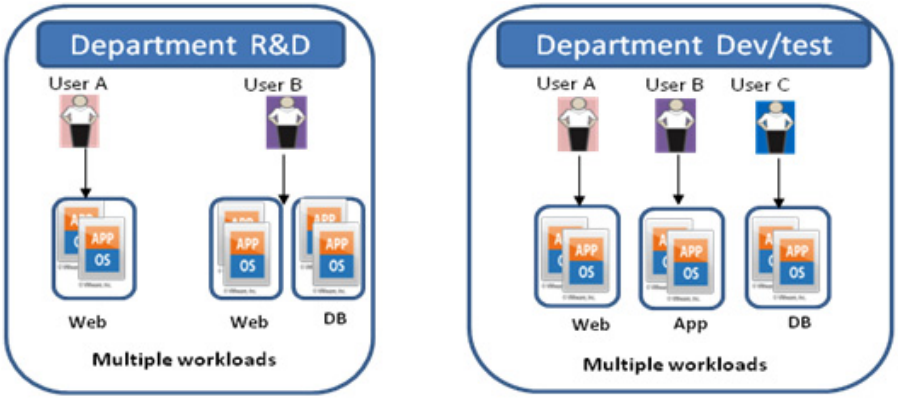

In the public cloud deployment model, a tenant is an individual consumer, an Enterprise, or a sub-organization within an Enterprise subscribing to the virtual private cloud services hosted by a Service Provider.

Each tenant must be securely separated from other tenants who share the common virtualized resource pool. However, workloads owned by one tenant are visible to others unless firewalls are configured to block communications among different tenant applications.

Differentiated Cloud Services

Cloud providers, whether Service Providers or Enterprises, want an IaaS offering with multiple feature tiers and pricing levels. The cloud is a source of highly scalable, efficient, and elastic services accessed on-demand over the Internet or intranet. In the cloud, compute, storage, and network hardware are abstracted and delivered as a service. End users consider the functionality and value provided by the service only; they do not need to understand or manage the underlying technology.

To tailor workload or application requirements to specific customer needs, the cloud provider can differentiate services with a multi-tiered service infrastructure and quality of service (QoS) settings. Such services can be used and purchased under a variable pricing model. Infrastructure and resource pools can be designed so that end users can add or expand services by requesting additional compute, storage, or network capacity. This elasticity allows the provider to maximize the user experience by offering a custom, private data center in virtual form.

Typically, cloud providers want to offer three, four, or five different service tiers and provide different service level agreements (SLAs). IaaS cloud services can be differentiated into pre-defined service tiers by varying support of the following features:

- **Virtual Machine Resources**—Service profiles can vary based on the size of specific virtual machine (VM) attributes, such as CPU, memory, and storage capacity. Service profiles can also be associated with VMware Distributed Resource Scheduling (DRS) profiles to prioritize specific classes of VMs. For example, a Gold service can consist of VMs with dual core 3-Ghz virtual CPU (vCPU), 8 GB of memory, and 500 GB of storage. A Bronze service can consist of VMs with a single core 1.5 Ghz vCPU, 2 GB of memory, and 100 GB of storage.
- **Storage Features**—To meet datastore protection, recovery point, or recovery time objectives, service tiers can vary based on provided storage features, such as RAID levels, disk types and speeds, and backup and snapshot capabilities. For example, a Gold service could offer three tiers of RAID-10 storage using 15K rpm Fibre Channel (FC), 10K rpm FC, and SATA drives. While a Bronze service might offer a single RAID-5 storage tier using SATA drives.
- **Application Tiers**—Service tiers can provide differentiated support for application hosting. In some instances, applications may require several application tiers of VMs. Often, each tier is placed on separate VLANs. For example, a Gold profile could have three application tiers on three separate VLANs to host web, application, and database (DB) services on different VMs. Each tier could provide five VMs each for redundancy and provide load balancing. A Silver profile could also have three tiers for web, application, and DB services, but each tier might have two VMs for redundancy and load balancing. In contrast, a Bronze profile could have three tiers but in a less differentiated manner, with the web, application, and DB services residing on the same VLAN or potentially on the same VM.
- **Stateful Services**—Customer or employee workloads can also be differentiated by the services applied to each tier. These services can be firewalls, encryption, load balancers, protocol optimization, application firewalls, WAN optimization, advanced routing, redundancy, disaster recovery, and so on. Within a service like firewalls, you can further differentiate among tiers as with inter-VLAN, intra-VLAN, or intra-host inspections. For example, a Gold tier might include firewall inspection, SSL off loading, IPSec encryption, server load balancing, and WAN optimization. A Silver tier might offer only firewall inspection and server load balancing.
- **Quality of Service Agreements**—Bandwidth control during periods of network congestion can be a key differentiator. QoS policies can prioritize bandwidth by service tier. Traffic classification, prioritization, and queuing and scheduling mechanisms can identify and offer minimum bandwidth guarantees to tenant traffic flows during periods of congestion. For example, a Gold service tier might be given the highest priority and a minimum network bandwidth guarantee of 50%. A Bronze service tier might receive best-effort treatment only and no minimum bandwidth guarantees.

The VMDC solution defines options for differentiating IT cloud services. In this reference architecture, these cloud services are called service tiers. Typically when we talk about service tiers, we look at the server CPU and storage options. But if a web application is hosted in the cloud model, load balancing and firewall inspection are also required. To achieve secure separation of tenant data, Layer 2 and Layer 3 features, such as virtual routing and forwarding (VRF) and VLANs, must be enabled. With this virtual network separation configured, service tiers contain virtual compute, storage, and network resources.

Deployment Models

The Cisco VMDC solution qualifies a three-tier model of Bronze, Silver, and Gold tiers comprising IaaS services. These tiers define service levels for compute, storage, and network performance (Table 1-1).

Table 1-1 Example Network and Data Differentiations by Service Tier

	Bronze	Silver	Gold
Services	No additional services	Firewall Services	Firewall and Load balancing Services
Bandwidth	20%	30%	40%
Segmentation	One VLAN per client, Single VRF	Multiple VLANs per client, Single VRF	Multiple VLANs per client, Single VRF
Data Protection	none	Snap - Virtual copy (local site)	Clone - Mirror copy (local site)
Disaster Recovery	none	Remote replication (With specific RPO/RTO)	Remote replication (any-point in-time recovery)

Using this tiered model, you can do the following:

- Offer service tiers with well-defined and distinct SLAs
- Support customer segmentation based on desired service levels and functionality
- Allow for differentiated application support based on service tiers

Cloud Service Tiers

In the Cisco VMDC solution, three service tiers are defined: Bronze, Silver, and Gold. Each service tier is a container that is assigned specific network, compute, and storage resources. In the following sections, we explain how to differentiate service tiers in your cloud and the resources that those tiers can contain.

The Cisco VMDC solution allocates network, compute, and storage resources according to the Bronze, Silver, and Gold service tiers. The following sections identify differences in resources among the three tiers:

- Network Resources, page 1-4
- Compute Resources, page 1-4
- Storage Resources, page 1-5

Network Resources

Figure 1-2 depicts the network components assigned to Bronze, Silver, and Gold service tiers in the VMDC solution. The Gold, Silver, and Bronze tiers present predefined baseline network configuration choices through a self-service portal from which a tenant can select the container.

Figure 1-2 Network Resources by Service Tier

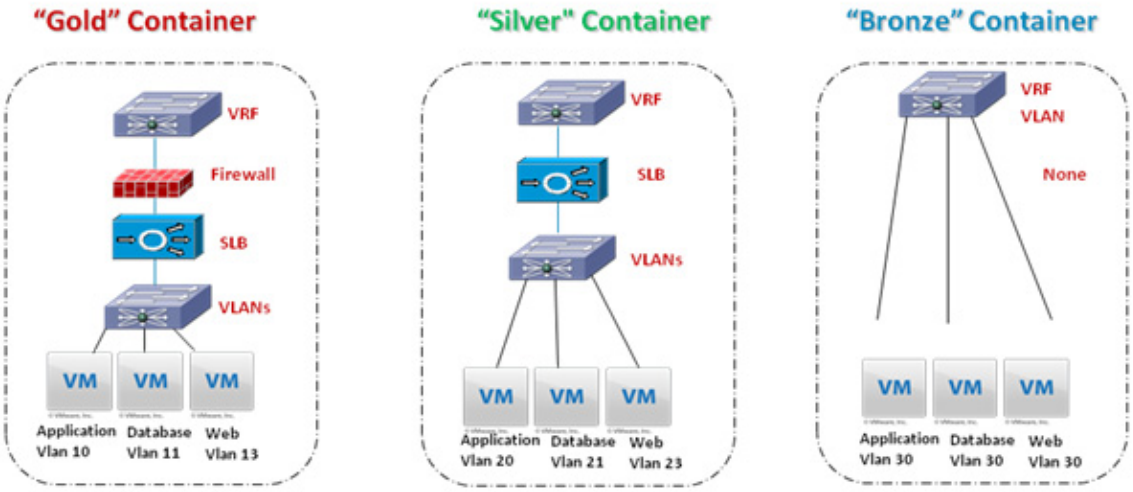

These network containers are enabled by end-to-end virtualization of the network infrastructure. The Cisco VMDC solution leverages end-to-end VRF-Lite, VLAN, and virtualized services, such as virtual firewall and load balancing contexts.

Each service tier uses a unique VRF-Lite instance per tenant to provide a dedicated virtual network (or virtual private data center). Depending on the application, multiple application tiers may exist, such as the hosted web, application, and database tiers of an e-commerce application. Each tier of this application resides in a separate VLAN within the VRF-Lite instance. For each Silver and Gold tenant, a unique VRF and three VLANs are provisioned. The Gold tenant is further differentiated by a dedicated virtual firewall and load balancing services, whereas Silver tenants receive only a dedicated virtual load balancing service. Being a best effort service, the Bronze tenant is assigned a unique VRF-Lite instance for each tenant, but all three tiers of an application must share the same VLAN, and no firewall or load balancing services are provided.

These service tier definitions form a baseline to which additional services may be added for enhanced security, PCI compliance, datastore protection, business continuity, or disaster recovery.

Compute Resources

In the VMDC 2.0 system, at the compute layer, service tier differentiation was modelled based on three compute workload sizes called Small, Medium, and Large. From an application perspective, key characteristics to consider are vCPU and RAM. Server virtualization runs multiple virtual servers on a single blade server.

The number of virtual machines (VMs) that can be enabled depends on the workload type being deployed and the CPU and memory capacity of the blade server. Cisco UCS B-series blade servers are two-socket servers based on the Intel Xeon series processor. Each socket has four cores with a total of eight cores, or 8 vCPUs, per blade. As Table 1-2 shows, 32 Small VMs per physical host were enabled

by allocating 0.25 vCPU for each virtual machine whereas Large has a dedicated vCPU for each VM, limiting the total Large workloads to 8 per blade server. Effectively, this comprises a compute oversubscription factor (OSF) of 4:1, 2:1 and 1:1.

Table 1-2 lists the workload options and compute resource sizes.

Table 1-2 Compute Resources by Size

	Small	Medium	Large
vCPUs per core	0.25 vCPU	0.5 vCPU	1 vCPU
VMs per blade	32 VMs	16 VMs	8 VMs
RAM (GB)	4	8	16

Storage Resources

The Cisco VMDC architecture defines three persistent storage workload sizes called Small, Medium, and Large. Datastore retention and availability is a major concern for customers of cloud-based offerings. Thus, a baseline premise of the system is that a tiered retention, protection, and recovery model will insure that storage availability and reliability may be tailored to meet tenant requirements.

Table 1-3 lists the workload options and storage resources sizes.

Table 1-3 Storage Services by Size

	Small	Medium	Large
Base storage (GB)	50	150	300
Storage growth increment (GB)	50	50	50
Backup (retention length options)	1 mo., 6 mo., or 1yr.	1 mo., 6 mo., or 1yr.	1 mo., 6 mo., or 1yr.
Data protection	None	Snap - Virtual copy (local site) SNAP copies every 8 hrs.; 36 hr. retention	Clone - Mirror copy (local site) - SNAP copies every 4 hrs.; 36 hr. retention
Disaster recovery	None	Remote replication Symmetrix Remote Data Facility (SRDF)	Remote replication SRDF

You can further refine the service tiers by differentiating the backup and recovery options. To ensure data protection and durability, Snap and Clone techniques can create point-in-time consistent copies of tenant volumes. To provide support for disaster recovery, snap volumes can be replicated to multiple locations. Table 1-4 presents example storage distinctions by service tier.

Table 1-4 Service Tier Distinctions for Storage

	Small	Medium	Large
Base storage (GB)	50	150	300
Storage growth increment (GB)	50	50	50

Table 1-4 Service Tier Distinctions for Storage (continued)

	Small	Medium	Large
Backup (retention length options)	1 mo., 6 mo., or 1yr.	1 mo., 6 mo., or 1yr.	1 mo., 6 mo., or 1yr.
Data protection	None	Snap - Virtual copy (local site) SNAP copies every 8 hrs.; 36 hr. retention	Clone - Mirror copy (local site) - SNAP copies every 4 hrs.; 36 hr. retention
Disaster recovery	None	Remote replication Symmetrix Remote Data Facility (SRDF)	Remote replication SRDF

Per Tenant Logical Flow

First, a tenant chooses the network container that provides him a virtual dedicated network within the shared infrastructure. A tenant has three selections to choose from for their network container - the provided network services and level of separation. Cisco VMDC defines the following network container selections: Gold, Silver and Bronze.

Second, a tenant chooses the compute and storage resources to run inside the container. Cisco VMDC defines three workload sizes for compute and storage resources: Small, Medium, and Large. A tenant's choice of a workload size depends heavily on the application being implemented in the container.

Therefore, a network container comprises the services, resources, and service path through the infrastructure. The concept of a network containers is introduced and embodied in the Cisco VMDC solution.

Solution Objectives

The Cisco VMDC solution targets the following objectives:

- Expedite ordering with a single bill of materials (BOM).
- Validate interoperability end to end.
- Support incremental investment via modular build out.
- Provide security at every layer.
- Provide an application aware network.
- Reduce power and space requirements per server.
- Increase performance for network-based backup and data replication.
- Support service tiers to quickly scale up without wasting resources.
- Enable workload portability.
- Support end-to-end service orchestration and automated provisioning.