Learn About Designing the Infrastructure for Hosting SaaS Applications

As an independent software vendor (ISV) providing software as a service (SaaS), you need secure, scalable, enterprise-grade infrastructure to host your services and to manage your *tenants*. The OCI solution provides a validated architecture that incorporates best practices to enable you to host your SaaS applications on Oracle Cloud.

About SaaS Applications

A SaaS application is an application that a vendor offers as a service in the cloud. Customers of the vendor subscribe to the service and use the application when they need it.

To a SaaS vendor, each subscriber (or customer) is a *tenant*.

A SaaS application can be deployed in the cloud by using the following architectural patterns:

- **A single, tenant-aware application instance**

  In this pattern, the SaaS vendor deploys a single application instance, which all the tenants use. The application handles the separation of the tenant-specific workloads and data.

  All the tenants get the same application version, built from a common code base. Because all the application deployments are based on the same code, the vendor can configure, patch, and upgrade the service efficiently. Scaling and operating the service is easy.
  However, building a tenant-aware application environment requires more effort initially. And this deployment pattern is not suitable for SaaS customers that require complete isolation.

- **Multiple tenant-specific application instances**

  This pattern is the focus of the solution.

  The SaaS vendor deploys and manages multiple isolated application instances. Each deployment is for a specific tenant. The SaaS vendor manages the individual tenant application instances through a common management layer.

  The vendor can choose to either build all the tenant application instances from a single code base or offer customized versions of the application to each

tenant. This pattern is ideal for SaaS customers that require complete isolation of the application environment.

Architecture

This architecture shows an Oracle Cloud Infrastructure tenancy that hosts multiple tenants of a SaaS vendor. All the resources in the architecture are in a single region.

The SaaS vendor's management infrastructure and the application resources of each tenant are isolated in separate compartments and virtual cloud networks (VCNs). Network isolation ensures that the applications and data are segregated from the other deployments in the tenancy. Compartments ensure logical isolation of the resources and enable granular access control.

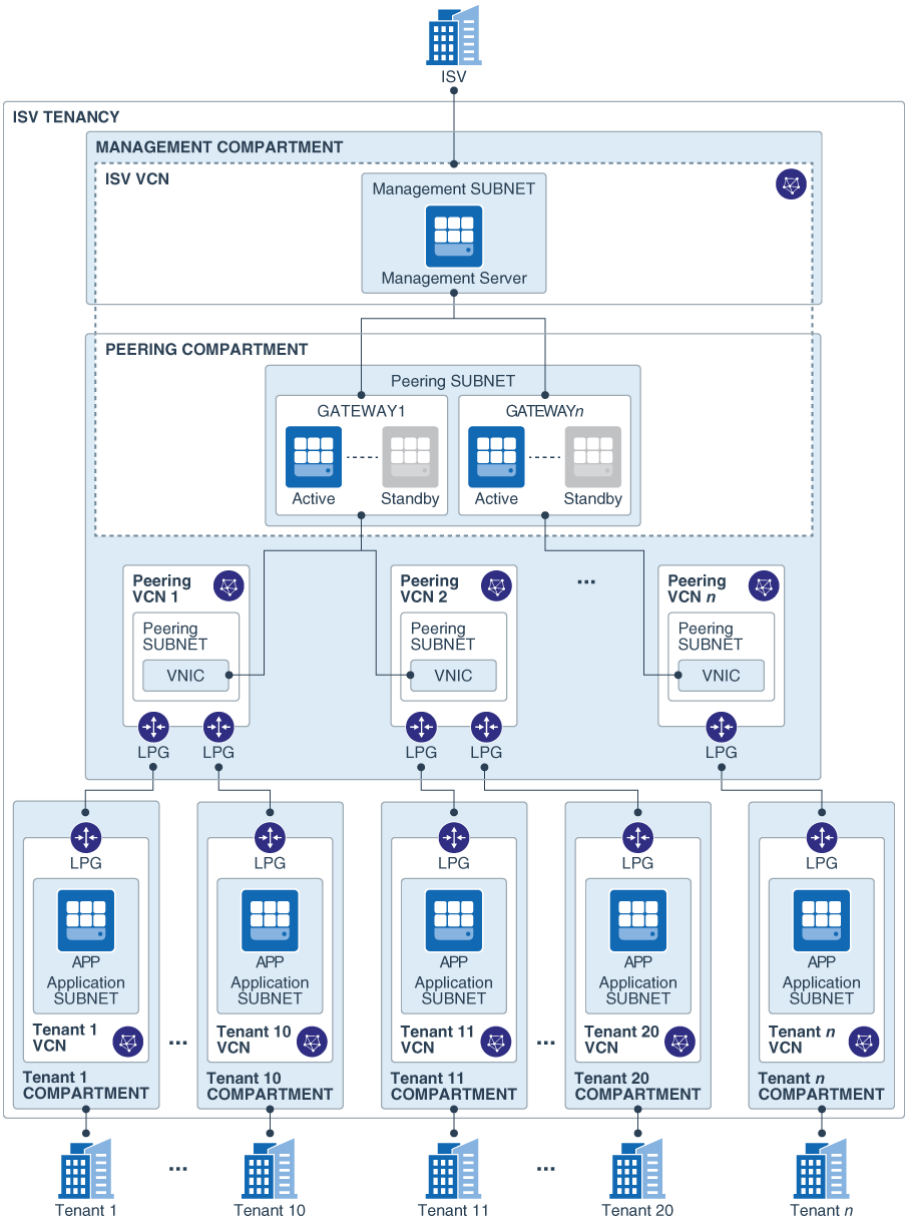

This architecture includes the following components:

- **Management compartment**

  The management compartment is a logical container for all the ISV-specific resources necessary for the common infrastructure used to manage the individual tenant application deployments. It contains the following resources:

  - **ISV VCN**

The resources required for the SaaS ISV to access and manage its tenants are attached to the ISV VCN.

- **Management server**

  The management server is a compute instance in a private subnet. You can run an infrastructure monitoring application on this server to monitor the tenant servers. The management server is attached to a private subnet in the ISV VCN. The management server can communicate with the servers in the tenant compartments through the routing gateways.

- **Peering compartment**

  For private communication between the resources in the ISV VCN and the tenant resources, a local peering relationship is necessary between the ISV VCN and each tenant VCN. But a VCN can have up to only 10 local peering relationships. To overcome this scaling limitation, the architecture uses routing gateways that can connect to multiple peering VCNs. Each peering VCN can have a local peering relationship with up to 10 tenant VCNs. So you can scale up the topology by adding routing gateways and peering VCNs in the peering compartment.

  - **Peering subnet**

    The peering subnet is a part of the ISV VCN. All the routing gateways are attached to this subnet.

  - **Routing gateways and peering VCNs**

    Each routing gateway is an Oracle Linux compute instance that routes traffic from the management server to the tenant VCNs, through a peering VCN.

    The primary VNIC of each routing gateway instance is attached to the peering subnet in the ISV VCN.

    The secondary VNICs of each routing gateway instance are attached to subnets of the peering VCNs. The maximum number of peering VCNs that a routing gateway can serve depends on the number of secondary VNICs that the shape of the underlying compute instance supports. For example, if a routing gateway runs on a compute instance that uses the `VM.Standard.2.4` shape, it can have a maximum of three secondary VNICs, and can serve up to three peering VCNs. Each peering VCN can be connected to up to 10 tenant VCNs. So a routing gateway that uses the `VM.Standard.2.4` shape can support up to 30 application tenants.

For high availability of each routing gateway, you can set up an active-passive pair of compute instances with a floating IP address, and use software such as Pacemaker and Corosync to ensure automatic failover.

- **Tenant compartments**

  The resources for each tenant are in a separate compartment. Each tenant compartment contains a VCN to which all the resources for that tenant are attached. So the resources for each tenant use a unique address space in a network that's isolated from all the other tenants in the topology.

  In each tenant compartment, you can provision a compute instance running an agent that can monitor the servers in the compartment and send metrics to the management server in the ISV VCN.

  When you add new application tenants, the necessary resources for the new tenant are provisioned in a new compartment.

  Tenants can access their application over the public internet or through a private connection (IPSec VPN or FastConnect). For access from the public internet, each tenant VCN requires an internet gateway. For access using VPN or FastConnect, a DRG is necessary. The architecture diagram doesn't show the internet gateways and the DRGs for the tenant VCNs.

About Required Services and Policies

The OCI solution requires the following services and access-management policies:

Service	Policies Required To...
Oracle Cloud Infrastructure Identity and Access Management	Create and manage compartments.
Oracle Cloud Infrastructure Networking	Create and manage VCNs, subnets, internet gateways, route tables, security lists, LPGs, and DRGs
Oracle Cloud Infrastructure Compute	Create and manage compute instances.