CHAPTER 2

Hybrid Cloud Use Cases

As Enterprises are adopting both the Private and Provider Clouds (Public Clouds), they want the flexibility to place their workloads in either of these two clouds based on their needs, as well as company policy and/or compliance requirements. As the Enterprise business grows rapidly and requires additional compute resources, Enterprise IT wants to take advantage of resources in the Provider Cloud rather than building out additional Data Centers or adding additional compute resources in their Private Cloud. Also, in peak season, Enterprises require placing some of their workloads in the Provider Cloud to meet the demands but keep their sensitive data in the Private Cloud. However, if the enterprise is connecting to Provider Cloud via WAN, latency and bandwidth (BW) costs maybe a concern since most applications have strict latency requirements. It is very common to find Enterprises Data Center or Private Cloud co-located with Provider Cloud and therefore latency between application servers and tiers is not a concern.

This design guide emphasizes Capacity Augmentation use cases and sub-use cases that include Workload Offloading, Distributed Workload, and Planned Peak Capacity.

Workload Offloading

The Workload Offloading use case, with or without network and security services, focuses on the ability of Intercloud Fabric to help customers to use additional capacity of Provider Clouds to offload an existing application running in the Private Cloud, while extending network and security policies. The use case focuses on the offload of a complete 3-tier application (Web/App/DB services) from the Enterprise into the Provider Cloud. In some Service Provider environments, the Enterprise deploys firewall, load balancing, and routing services for data traffic extended into the cloud. Test cases for both, with and without services, were executed.

Note Intercloud Fabric is not positioned as a migration tool by itself. It includes an offload capability for the move of the VM and the seamless extension of the network and security to the Provider Cloud, while keeping the control point at the Enterprise or business customer. For one-time migration purposes where there is no need to extend the network and security or maintain the control from a portal in the Enterprise, Cisco recommends other tools from partners.

Distributed Workload

In a hybrid cloud scenario, applications are eventually deployed in a distributed fashion, in dispersed locations. Intercloud Fabric enables customers to take advantage of the ability to manage multiple Provider Clouds as a seamless extension of the Private Cloud, which makes it easier for distributed applications. This powerful ability creates the need for being mindful of requirements prior to distributing the application.

As part of the Distributed Workload use case, with or without network and security services, a web front end services of a 3-tier application is deployed and verified in the Provider Cloud, while the application and database services for the application reside in the Enterprise Data Center. In some Service Provider environments, the Enterprise deploys the firewall, load balancing, and routing services for the web traffic that extends into the cloud. Test cases for both, with and without services, were executed.

Planned Peak Capacity

In the Planned Peak Capacity use case, Enterprise customers use Service Provider Cloud resources to temporarily burst their workloads to meet any seasonal demands. The resources are released/decommissioned in the Provider Cloud when high-demand processing finishes.

Cisco Intercloud Fabric manages the creation and access to the VMs in the Provider Clouds, extending the network and Enterprise configured security policies, all while managing the life-cycle of the cloud positioned VM.

Cisco Intercloud Fabric exposes APIs on the business side that can be used by monitoring systems and/or cloud platforms to trigger instantiation of additional VMs to a certain application with configuration of the new servers and services as part of such application. This design guide does not demonstrate APIs or 3rd party tools.

Applications Used within Use Cases

Two 3-Tier applications were used throughout the testing and included a deployment of Microsoft SharePoint and a WAMP (Windows Apache MySQL PHP) placement. Each of these were deployed to the different provider environments, with some differentiation based on availability of services (Table 2-1).

Table 2-1 Service Providers, Services and Applications

Provider	Services	Application
Amazon EC2	ICF Firewall, ICF Router, HA Proxy[1]	3 Tier WAMP Stack / 3 Tier SharePoint
Microsoft Azure	HA Proxy	3 Tier WAMP Stack / 3 Tier SharePoint
Cisco Powered Provider (ICFPP)	HA Proxy	3 Tier WAMP Stack / 3 Tier SharePoint

1. HAProxy = Open Source Load Balancer

Further breakdown of these subcomponents is shown in Table 2-2 and Table 2-3, with the database resource varying due to provider OS support differences.

Table 2-2 SharePoint 3-Tier Application

Quantity	Resource	OS	Component
2-4	Web Front End (WFE)	Windows 2008 R2 SP1	MS IIS
1	Application (App)	Windows 2008 R2 SP1	MS SharePoint
2	Database (DB)	Windows 2008 R2 SP1	MS SQL Cluster

Table 2-3 WAMP 3-Tier Application

Quantity	Resource	OS	Component
2	Web Front End (WFE)	Windows 2008 R2 SP1	MS IIS
1	Application (App)	Red Hat Enterprise Linux 6.3	Tomcat/PHP
1	Database (DB)	CentOS 6.3/RHEL 6.3	MySQL

An open source load balancer application was deployed in the Enterprise and, depending on the use case, was offloaded to the service Provider Cloud to load balance and monitor traffic destined to each of the web front-end servers. The HAProxy application was installed on both a Red Hat Linux version 6.3 and CentOS version 6.3 virtual machine. It was deployed into the Enterprise's VMware environment. For more information regarding HAProxy and its functionality refer to the HAProxy web site.