



Hybrid Cloud For Enterprises: Why, What and How.

White Paper

There is a lot of talk about hybrid cloud these days, with all vendors enabling it in their own way. It has become one of the most hyped, talked-about topics in the enterprise IT space. But the most important thing to understand is why customers need it in the first place. A close second is to know the various challenges in building a hybrid cloud solution. So as Simon Sinek said in his inspiring talk: it "Always Start with Why".

Why Hybrid Cloud

A hybrid cloud is defined as a cloud environment where an enterprise makes use of both an on-premises cloud and public cloud to run its applications. There are three main use cases where such a configuration benefits a business:

1) Cost Optimization

Hybrid cloud enables customers to deploy applications on the right cloud based on the application characteristics. This approach has great promise, but the devil is in the details of how to identify such

applications. In fact, we make similar decisions in other spheres of life every day as we think about buying, renting or leasing a car, a house or making loan or investment choices. Let us look at some application characteristics and the best-suited cloud for that application type.

Now, some applications fit into more than one characteristic here and you will have to decide the dominant feature in that case. For example, if you have an elastic application that needs workload locality, you may have to sacrifice the locality and run it on a public cloud. Similarly, if you have a short-lived workload that needs to work with data secured behind your firewall, you will have to use a private cloud.

The choice of one cloud does not always mean that you cannot run the specific application on another cloud; it simply provides the best option in terms of cost and performance. For example, one can run their CI/CD

Application Characteristics	Examples	Relevant Cloud
Short-lived	One time large computation	Public
Highly elastic	Seasonal online business	Public
Small scale with geographic diversity	Web app for a startup	Public
Consistent, long lived	Test & dev workload	Private
Needs low latency, high performance	Enterprise Applications, VDI	Private
Needs workload locality	A tightly-coupled multi-tier app	Private
Need more control over users, budget	Self-service cloud for higher education	Private
Needs integration with on-premise devices or data	R&D on in-house data	Private
Needs security and data compliance	Finance, ERP, HR apps	Private

pipeline (test & dev) workload using tools like Jenkins[2], Bamboo[3] on a public cloud as well. But running such a consistent, large workload there can be very expensive and it may also run slower due to lack of quality of service in a public cloud. Pipeline (test & dev) workload using tools like Jenkins[2], Bamboo[3] on a public cloud as well. But running such a consistent, large workload there can be very expensive and it may also run slower due to lack of quality of service in a public cloud.

2) Bursty workloads

Many businesses are cyclical in nature. Some common examples include accounting, e-commerce, tax filing, travel and education enrollment systems. The application scale and infrastructure requirements can vary drastically from month to month. One can run the base workload on an on-premise infrastructure and scale on a public cloud based to meet peak demand. Having a hybrid cloud setup is very useful in this case.

3) Consume Cloud Services

Many applications make use of different services to deliver business value. For example, an application that provides data sharing across branch offices may use cloud-based object storage as a reliable store and for data sharing across offices. Here, one can deploy the application on-premises with some amount of caching, but the cloud can be used as the main data store. Another example of consuming such a service is to do data ingestion in a public cloud and eventually ship that data to an on-premises cloud for permanent storage and analysis. This can be very cost-effective since the inbound bandwidth in public clouds is free and one can only ship the relevant data to an on-premises cloud after doing data filtering, compression and WAN optimizations.

Overall, given the increasing maturity of solutions in this space, the availability of rich set of cloud services and the ease of migration between clouds, one can deploy a sophisticated hybrid cloud strategy and derive a lot of business value from it.

A hybrid cloud is a must-have tool for enterprise IT — helping to effectively deliver business value and compete in a digital world.

Let us look at some of the ways of setting up a hybrid cloud and the technical challenges in doing so.

Hybrid Cloud: Key Technical Challenges

There are several technical challenges in setting up a hybrid cloud solution, and one must pay close attention to these while thinking about building such a solution. Figure 1 shows the high-level architecture of a hybrid cloud environment. Note that one of the key components is the connectivity between the two clouds and the second key component is the orchestration layer to deploy workloads across multiple clouds.

1) Secure network connectivity

Since you will be running workloads on-premises and in a cloud, you have to decide what type of network connectivity you need between the two environments. The first option is to set up VPN connectivity to the network in the public cloud from your own site or sites. In that case, the same L2 network in public cloud is extended and reachable from your own local network. Most public cloud vendors provide options for hardware or software-based VPN gateways[4].

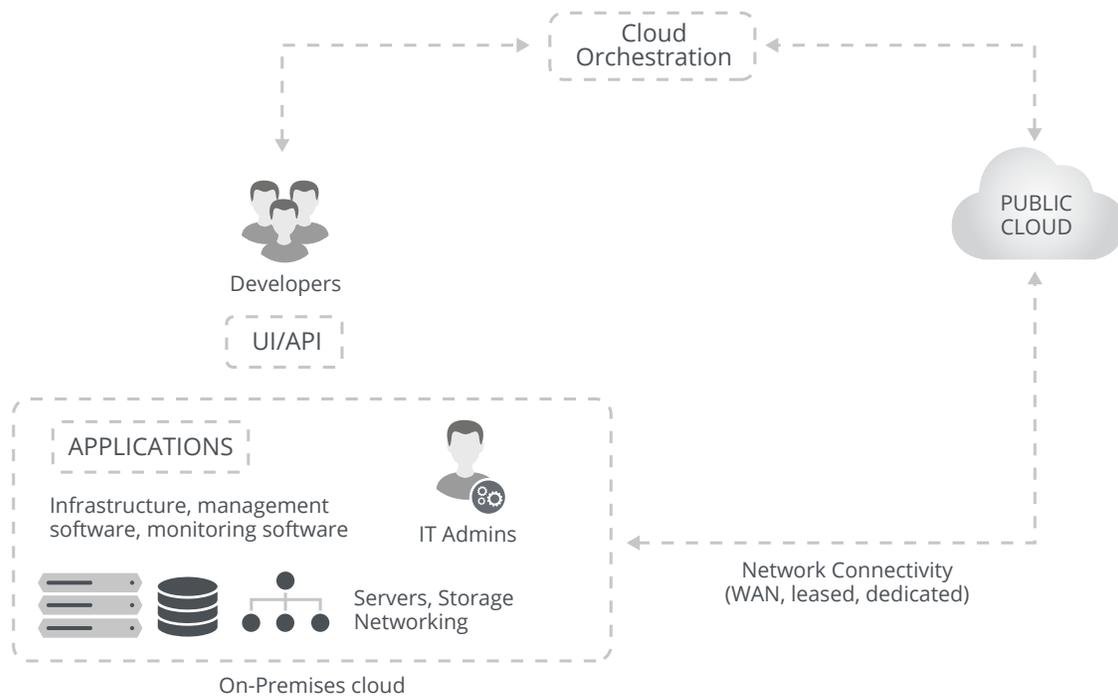


Figure 1: Hybrid Cloud Architecture

The second option is to just have L3 connectivity to the VMs in a public cloud with public IPs. This is more restrictive, since all VMs that need connectivity from your enterprise will need to have public IP addresses. Public IPs are a scarce commodity; they are quite expensive to get and they also expose your workloads to port scanners looking for security vulnerabilities.

2) Low latency network connectivity

Once you have dealt with security, which is the most critical issue, you need to look into the network latency between two environments. A higher latency will mean you have two islands: on-premises cloud and public cloud. With higher latency you can't run an application across two clouds in a seamless manner. Most

applications are sensitive to latency between various components, and having WAN latency and jitter can be very costly from a performance and reliability point of view. There are some examples where it works, however, such as consuming cloud storage as a backend from your on-premises application (for example using Amazon S3[5] as backend to store objects). Here, the latency requirements are loose by default. These days, multiple datacenter providers have set up low-latency, high-bandwidth links between their location and the public cloud. It is very useful to explore such options and choose a datacenter where such an option is available.

3) High bandwidth connectivity

In order to burst to cloud or migrate workloads to a public cloud, one of the key requirements is the time taken to migrate data from on-premises to public cloud and vice-versa. Doing migration over WAN can be very time consuming and cost prohibitive. We also know that the upload bandwidth in most cases is much more expensive than the download bandwidth. There are services like MegaPort[6] that provide lease-based, high-bandwidth connectivity between private datacenters and public clouds like AWS[7], Google Compute Engine[8] and Microsoft Azure[9]. This is a great model for customers who want migration between clouds. You should look for the option of renting space in a datacenter that provides such functionality.

4) Provisioning and migration functionality

You need a way to provision workloads across private and public cloud. In its simplest form, you can choose to do separate provisioning across the two platforms using their independent APIs and management UI, but this quickly becomes complicated. Ideally, you want a single way to deploy applications across more than one platform. Applications are often composed using multiple VMs running services and connected via one or more private networks, front-ended by a load balancer. Ideally, you should try to pick two platforms that support the same language to describe the application. Another alternative is to use a Cloud Management Platform that provides a higher-level abstraction for application deployment and translates that into the underlying cloud platforms. Some of the solutions in this space include ElasticBox[10], Cliqr[11] and Rightscale[12].

Once a workload is deployed, you may need to migrate it from one cloud to another. This is not a common operation and in fact most management platforms do not support it. You have to find third party solutions to take care of application migration from one cloud to another. This can also be time-consuming if lot of data needs to be moved. Try to structure your applications so you can burst onto a public cloud without too much data movement. There are a few options to handle that:

Network connectivity is the #1 challenge in building and adopting a hybrid cloud model.

1. Migrate stateless application components
2. Use a cloud-based data repository like object store in your application
3. Use a high-bandwidth connection for quick migration
4. Use a storage system that does asynchronous migration of data to cloud

With these, you can mitigate the issues related to application migration or bursting to the cloud. Next, let us look at various deployment options and available solutions from various vendors.

Hybrid Cloud Deployment Options

There are several ways in which one can deploy a hybrid cloud solution. Each one has pros and cons that one should know before making a decision.

There are three main components needed to build a hybrid cloud:

1. Private cloud infrastructure and software stack
2. Public cloud
3. A management layer across these clouds for provisioning, migration, visibility and planning

Stitching together various components

In this case, as a customer you can choose to deploy a private cloud from vendor A, a public cloud from vendor B and stitch them together using a cloud management platform from vendor C. In one case, A can be VMware[13] or an OpenStack-based cloud, B can be AWS[7] or Google Compute Engine[8] and C can be ElasticBox[10] or Rightscale[12]. Here you get a lot of flexibility, but it comes at a cost. You have to deal with multiple vendors, incompatible feature-sets, upgrades at different times, and finally training your team across all of these solutions. This solution should only be considered by large enterprises that have a lot of IT staff and specific requirements that cannot be serviced by more integrated solutions.

Same vendor public, private cloud & management layer

With some of the recent offerings, you can choose the same vendor for both private cloud deployment and use their public cloud. An example is using Microsoft Azure stack on-premises with Azure Cloud. Similarly, Oracle[14] recently announced cloud@customer integrated with Oracle public IaaS cloud.

VMware also used to have that with vCloud suite and vCloud Air[15]. VMware also announced partnerships with vendors like IBM and AWS where VMware infrastructure is installed in the datacenter of these partners. That provides the same VMware stack at both places, but if you want to consume public cloud services from AWS, that is a separate island and not managed by VMware's software. If you were looking to buy specific services from IBM or AWS, it would make sense to use that option instead of vCloud Air[15] or other smaller managed service providers.

In most of these cases, the problem is that it is quite limiting to use just one public cloud, and you should be able to use more than one of them. Public clouds are useful not only for infrastructure as a service but also to provide managed services. So pick this solution only if the public cloud solution fulfills all of your needs in terms of IaaS, PaaS and services.

Private/Hybrid cloud solution with public cloud choice

Finally, there are some vendors that provide an on-premises cloud that is integrated with AWS or other public clouds using their own management software. These vendors include VMware[13] and ZeroStack[16]. In this case, the private cloud and the management software is provided by the same vendor, so there are fewer moving pieces. In addition, you get your choice of public clouds. For most small to mid-sized enterprises, this is a great option due to greater flexibility and small operational overhead.

Summary

In the end, hybrid cloud is the best way to set up your IT if you want to optimize for cost, efficiency and control. Make sure that you make the right choices based on your applications. Aim to simplify the private cloud deployment and lower the operational costs associated with it. Ideally you want a private cloud that is as easy to use as a public cloud, and that gives you the flexibility to use either of them via a single platform that can also help you make the decision of which cloud is best suited for which application.

1. https://www.youtube.com/watch?v=u4ZoJKF_VuA
2. <https://jenkins.io/>
3. <https://www.atlassian.com/software/bamboo>
4. <http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpn-connections.html>
5. <https://aws.amazon.com/s3/>
6. <https://www.megaport.com/>
7. <https://aws.amazon.com/>
8. <https://cloud.google.com/compute/>
9. <https://azure.microsoft.com/en-us/>
10. <https://elasticbox.com/>
11. <http://www.cliqr.com/>
12. <http://www.rightscale.com/>
13. www.vmware.com
14. <https://www.oracle.com/index.html>
15. <http://vcloud.vmware.com/>
16. www.zerostack.com