rackspace

MANAGED VIRTUALIZATION
Powered By VMware®

The Incremental Advantage:

MIGRATE TRADITIONAL APPLICATIONS FROM YOUR ON-PREMISES VMWARE® ENVIRONMENT TO THE HYBRID CLOUD IN FIVE STEPS

CONTENTS

The Incremental Advantage:

MIGRATE TRADITIONAL APPLICATIONS FROM YOUR ON-PREMISES VMWARE® ENVIRONMENT TO THE HYBRID CLOUD IN FIVE STEPS

As an IT architect, you will likely be tasked with migrating some of your traditional (legacy) applications to a cloud environment that's hosted by a service provider. There are two basic options for accomplishing this migration.

One method is to completely re-architect your traditional applications for deployment in the public cloud. This may be a viable option, depending on your budget, time constraints, compliance concerns, number of available developers and other technical resources, and related issues. The key challenge is that re-architecting traditional applications can be a lengthy, complex process that eats up valuable development time and resources—and there's no upfront testing available to guarantee that the application will ultimately deliver the required performance. That's a lot of investment for an unclear payoff.

The alternative method is a lower-risk, incremental path to the hybrid cloud that involves splitting an application between a hosting provider's dedicated and public cloud environments. This incremental approach, described in detail in this paper, can greatly reduce your time commitment and the risk involved in migrating your traditional applications to the public cloud, while still giving you the cloud-bursting capabilities and other cost-saving benefits of the public cloud model.

Your precise path to the hybrid cloud will depend on the application. In this paper, we assume that you have an existing VMware environment and are migrating a three-tier application with separate web, application, and database components. We further assume that you're moving this application to the cloud due to one or more common requirements, such as elasticity, performance optimization, security, or cost optimization. To illustrate the incremental hybrid cloud approach, we'll use Rackspace® as the hosting service provider.

The end of the paper includes reference architectures to describe the eventual hybrid cloud solution. But first, we'll detail the following five steps to migrate your three-tier application to the hybrid cloud:

- **Step 1:** Profile your application

- **Step 2:** Examine key application characteristics

- **Step 3:** Build out the hybrid environment

- **Step 4:** Perform a proof of concept

- **Step 5:** Cut over to the hybrid environment

FIVE STEPS TO THE HYBRID CLOUD

The following five steps show how you can move a traditional application—in this case a three-tier application—to a hybrid cloud environment that connects a dedicated hosting environment to the public cloud.

Step 1: Profile your application

The first step is to profile your application to understand where it is at in its lifecycle and how to streamline it for more efficient cloud integration. You can also identify ways that you might be able to reduce costs through consolidation, modernization, or other alternatives. Application profiling typically involves creating a data flow model or diagram that highlights all the components and processes of your application. The diagram shows how data flows between the different processes, where data resides, and more.

The purpose of the application profiling process is to understand the components of your applications and the dependencies within those components. This information will drive architectural decisions related to which parts of the application will work in the cloud and which parts will need to be rebuilt or designed around. It will also help you identify levels of complexity and risk, and determine whether public cloud, private cloud, virtualization, or bare-metal infrastructure is the best fit for your application's components.

A variety of application profiling solutions are available, such as Rackspace's one- to three-day Application Profiling Workshop. Another popular option is Unified Modeling Language (UML), which creates industry standards around a variety of modeling exercises (learn more at www.uml.org). One of the most helpful types of modeling related to hybrid architecture decision-making is a deployment diagram, which can help you understand which processes lie on particular servers (learn more at www.agilemodeling.com/artifacts/deploymentDiagram.htm).

Step 2: Examine key application characteristics

Once you have completed the first step, you can carefully examine the web layer of your application for specific characteristics and requirements that will affect your deployment. Specifically, you will want to examine the nature and dependencies of each process, and how they interact with each other.

For example, in the web layer of your traditional application you may have a server that knows something unique about a given client or connection. That's a red flag for an application to be deployed in the public cloud because the web layer operates independently and may not be aware of the failure of another server or have access to the unique information contained on another server.

Another key consideration relates to the latencies and tolerances between processes. If the latency of a switching fabric is too heavy—in other words, if your application is designed such that all application-oriented processes

must live on the web head—then the processes may need to reside on the same servers or in close proximity. If not re-architected, this application characteristic may require a vertical rather than horizontal scaling strategy. You may also be forced to adopt a two-tier architecture where only your database traverses the dedicated fabric.

For a more detailed examination of the application characteristics you should consider before putting even a piece of your application in the public cloud, read cloud architect Wayne Walls' five-part blog on the "Pillars of Cloudiness," which begins with an examination of parallel computing: www.rackspace.com/blog/pillars-of-cloudiness-no-1-parallel-computing.

Step 3: Build out the hybrid environment

There are several good options for migrating your traditional application to a hosted environment, with the most common being image porting and data migration. As mentioned earlier, we'll describe this process with Rackspace as the hosting service provider.

Sizing the hybrid environment

The first step is to use resource monitoring tools to identify the baseline resource requirements of your application, including the CPU, memory, and disk requirements. Start by tallying the total resources you have available to each component of your application in its current state, and then compare that total resource allotment to trends in resource utilization. Take careful note of the peak requirements of busy utilization times as well as your most common baseline.

Having a good understanding of the resource needs and trends of your application will help you decide how much infrastructure you should deploy initially. For a large-scale application with notable spikes in utilization, focus on sizing your dedicated footprint to meet the baseline demand, and where possible leverage the public cloud for periods of higher utilization. The closer you tune the environment to match the demand pattern, the more cost savings you can achieve.

Image porting

The fastest, least-friction method for moving your web front end to a hosted dedicated environment is to simply move your image files over.

By using your existing VMware tools, along with Dedicated VMware® vCenter Server™ (Dedicated vCenter) from Rackspace, you can easily move VMDK files from your data center to Rackspace's data center. Image porting provides significant advantages with its ease and simplicity.

Data migration

You may be able to use image porting for all or some pieces of your application; for the remainder, data migration is the most common solution. Businesses and professional service partners have performed data migrations for a long time, and many options are available.

In general, you first stand up the new environment and then perform whatever configurations are necessary to get your application running on the new environment. You then copy all the unique or dynamic data from your existing environment to the new environment. While you are checking to make sure that the new environment is running as expected, the original environment is still in production. That means there's a delta—at some point you have to freeze the current production environment, copy the delta environment, and flip the switch to have all your data flowing as if the new environment is the production environment.

Data migrations can be a fairly cumbersome process. For qualifying customers with easily scriptable solutions such as Linux® web server front ends, Rackspace relieves much of the difficulty through free, automated script migrations from its Cloud Movers Team.

The Rackspace Cloud Movers Team migrates your source image data into a new image template in the Rackspace Public Cloud. This assumes that the images are already functioning in a way that will work well in the public cloud, with the image architecture uncoupled and the web server split from the application servers. If that's the case, you can simply drop your application and database layer into Dedicated vCenter through a regular migration, and the Cloud Movers Team will migrate images into the Rackspace Public Cloud.

The Cloud Movers Team approaches each migration on a case-by-case basis because not all images are a good fit for this service. The prerequisites for this service include:

• The source image and the new image in the Public Cloud must be like for like. For example, the source Microsoft® Windows Server® 2012 image would be migrated to a new Windows Server 2012 image template.

• The Cloud Movers Team can only migrate images that are based on guest OSes that are supported by the Rackspace Public Cloud. (See a full list of currently supported OSes at the bottom of this webpage: www.rackspace. com/cloud/servers/features.)

With your permission, the Cloud Movers Team logs into your on-premises servers and stands up the hosted infrastructure to match your own. They then run scripts on the servers that copy data structures, directory structures, and application configuration components to the destination on the cloud side. This data transfer is all done over the wire.

5 STEPS TO THE HYBRID CLOUD:

1 Profile your application

2 Examine key application characteristics

3 Build the hybrid environment

4 Perform a proof of concept

5 Cut over to the hybrid environment

For migrating data to the hosted dedicated environment, you have a few options. For instance, you might migrate VMDK files and related data over the wire via a secure VPN. Or, if you have a large data footprint, it may be more practical to ship a storage array to Rackspace, allowing the data to be imported directly into the environment without the latency of an over-the-wire copy.

Leveraging a third-party provider can relieve you of some of the most time-consuming, granular decision-making, so you can focus on larger architectural decisions regarding your new hybrid design.

Step 4: Perform a proof of concept

One of the main advantages of moving to the hybrid cloud instead of re-architecting for the public cloud is that you can perform a proof of concept to test your application in its new environment before going down the costly and time-consuming road of re-architecting.

A third-party provider is usually needed to help you perform load testing to capture the performance metrics you need. The specifics of the load test will depend on your application. If data doesn't impact your app performance, one time-saving possibility is to perform the load testing before copying all your production data and moving it over to the new environment.

If you've implemented scaling groups in your cloud environment, you will want to verify that the environment is scaling up and down as desired. If it's not configured to be automatic, make sure to thoroughly document your process for manually scaling the application up or down.

Step 5: Cut over to the hybrid environment

In this final step, you get traffic flowing to the new dedicated environment. This usually involves changing your DNS settings to point them from your existing IP addresses to the new environment.

If your application will support it, you'll want to take an incremental approach in this step to reduce your risk. You may, for instance, want to bleed off some users before you migrate all of them over to your new environment.

ARCHITECTURE OF THE HYBRID CLOUD

The hybrid cloud model gives you much greater flexibility in deploying your traditional applications. By splitting your application into separate components located in the dedicated hosted environment and the public cloud, you can achieve the cloud-bursting benefits of the public cloud while storing your most sensitive data and running your performance-hungry database and application components on dedicated gear.

In our three-tier application example, the application could be split into the following three tiers:

Tier 1: Front-end web layer

Typically you will want to place the front-end web layer of your application in the public cloud, which will help enable you to respond rapidly to unexpected spikes in demand and achieve associated cost-optimization benefits.

The hybrid cloud model enables you to further refine this traditional setup by splitting the web layer and placing at least some of it in your dedicated environment. For instance, a PCI-oriented customer could place a product brochure in the public cloud to take advantage of scalability; however, for security reasons, the payment processing (shopping cart) portion of the application would need to be in the dedicated environment (see Figure 1).
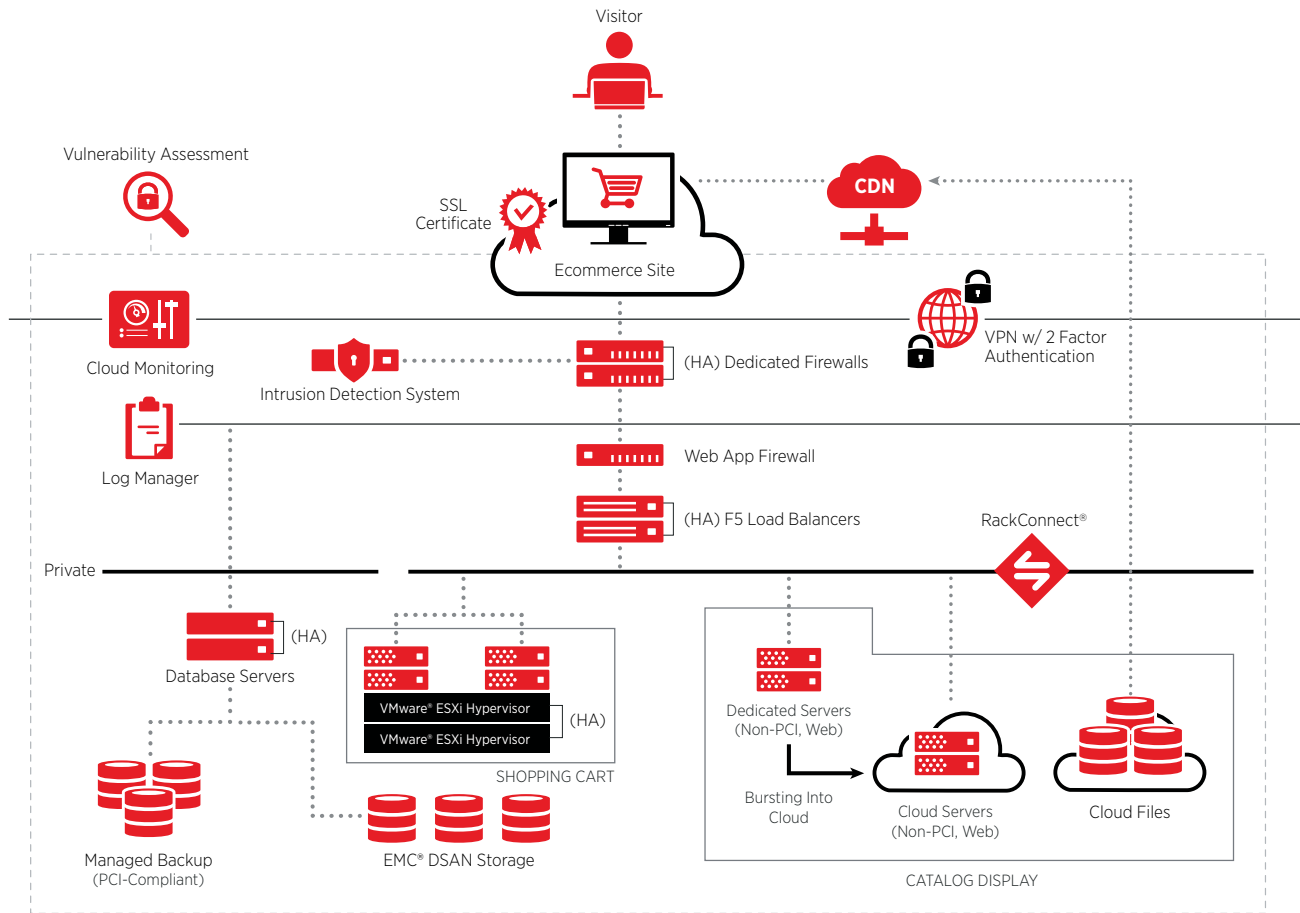
FIGURE 1: HYBRID CLOUD MODEL FOR AN ECOMMERCE SITE

RackConnect®-supported network devices

The RackConnect automation is compatible with any Cisco® ASA firewall (5505 Security+ or better) or F5® Big-IP load balancer. The optimal solution is the F5 Big-IP load balancer, which provides the ability to load balance between the dedicated Rackspace environment and the public cloud.

Your hosting provider will need to provide a way for your dedicated and cloud servers to communicate. In the Rackspace environment, that's done with a solution called RackConnect® that enables you to have both your dedicated and cloud web servers in the same load balancing pool. RackConnect makes it easy to eliminate your public cloud servers as needed to meet fluctuating demands and manage your dedicated servers—all from the same pool.

Other types of hybrid deployments that are spread over wide-area networks or physically distant facilities are limited by the speed of light, which introduces latency between the tiers of your application. RackConnect enables you to achieve extremely low latency by connecting your dedicated environment and the public cloud within the same data center.

The RackConnect automation system also configures the software firewall on each of the cloud servers, so that you don't have to set the software firewall rules (iptables or Windows Firewall) individually to ensure that the servers communicate in the way you desire.
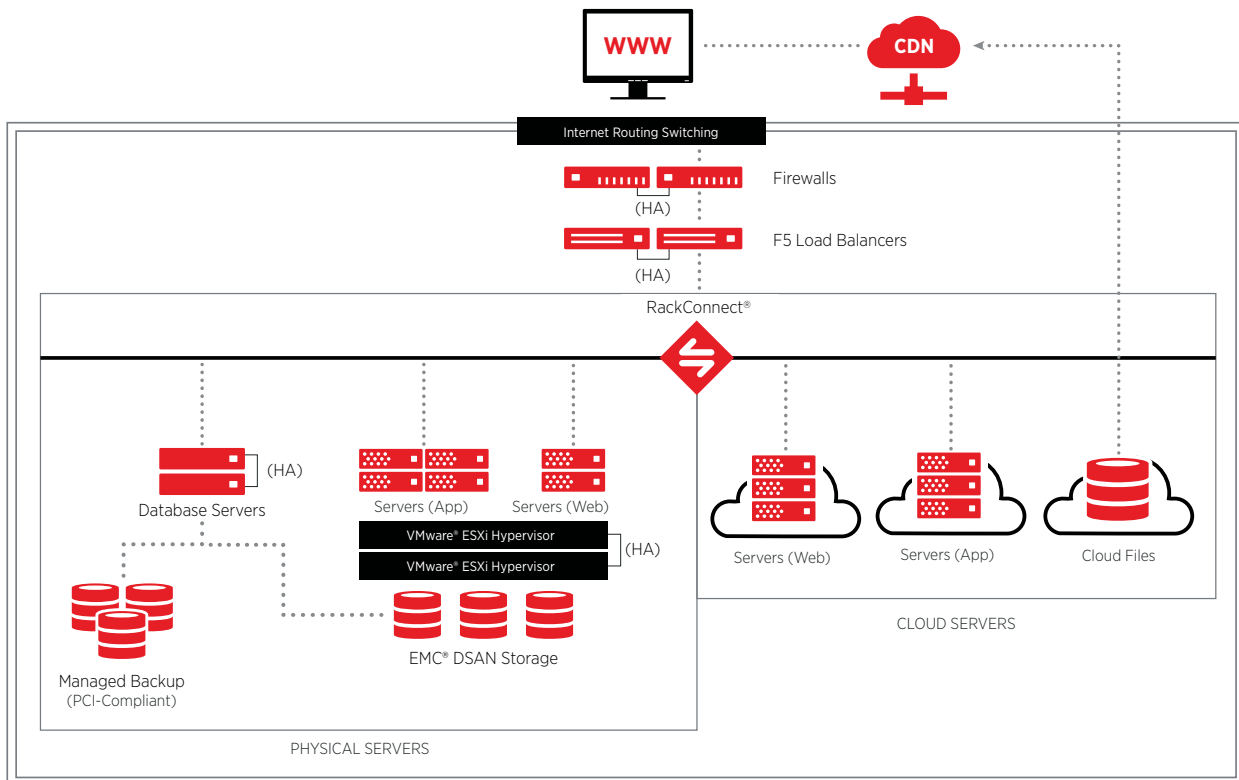
FIGURE 2: TYPICAL RACKCONNECT HYBRID DEPLOYMENT

Figure 2 shows a typical RackConnect hybrid deployment, in which new greenfield applications or web- and application-layer compute workloads reside in the public cloud, while high-performance connectivity is maintained to the backend traditional systems.

Tier 2: Dedicated components

You'll likely want to have at least some components of your web application reside in your dedicated environment. If you have a typical front-end Apache™ web server hosting your website, that's ideal for the public cloud. You could optimize for cost by splitting a portion of the front-end server traffic so that the dedicated footprint handles the constant base, and you leverage the cloud for unexpected traffic spikes. If you have security or compliance concerns, the single-tenant dedicated environment would be the best fit.

Figure 3 shows how Dedicated vCenter can be used to simplify management of components within the dedicated environment.
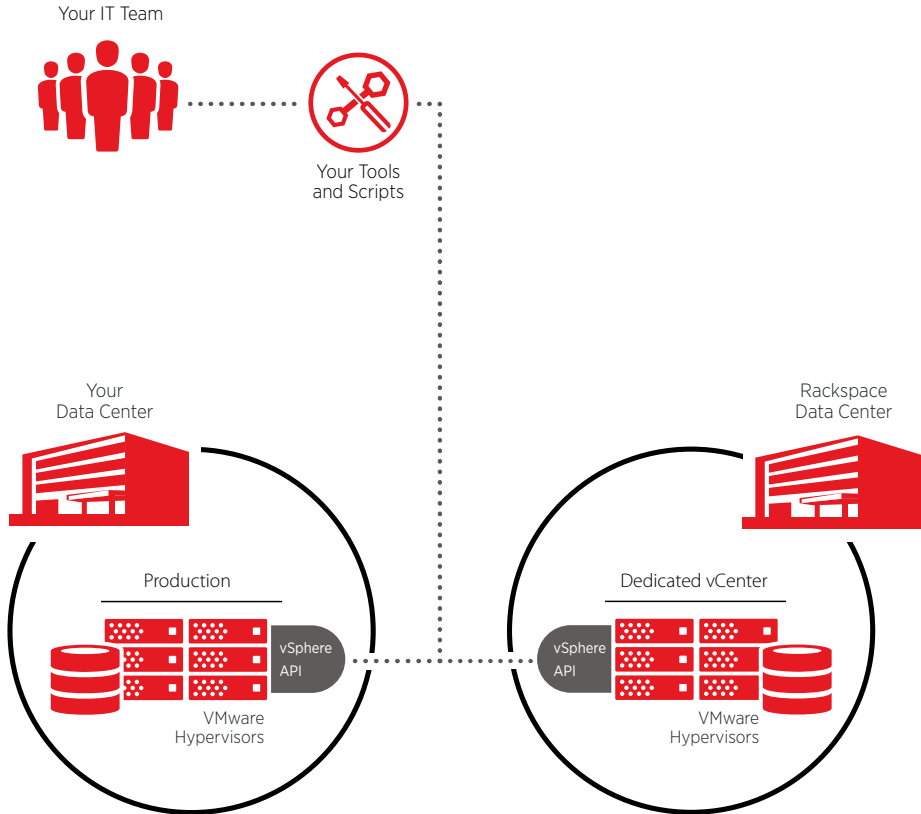
FIGURE 3: DEDICATED VCENTER SIMPLIFIES MANAGEMENT ACROSS ON-PREMISE AND HOSTED ENVIRONMENTS

Tier 3: Back-end database

There are a number of issues to consider when determining how to host the back-end database portion of your application to get the most out of your resources. Note that whichever method you choose, breaking out part of the workload to mix VMware with dedicated servers is a seamless process. It simply comes down to making a deployment decision.

If you find that you're scaling your database VM to the point where it's virtually the only VM running on a hypervisor—or if you're continually upgrading the hypervisor to manage one VM—it may make more sense to use a standalone server in your dedicated environment, which will reduce your overhead and latency.

Another consideration relates to performance issues, which might come up if you have multiple database VMs running on a hypervisor that all require significant disk I/O. In that case, it will likely make more sense to optimize performance using a bare-metal architecture in your dedicated environment. Such performance issues could arise at any layer of the application, although the database is the most common.

CONCLUSION

This white paper describes an incremental path for taking a traditional three-tier application to the hybrid cloud. This approach provides several advantages over completely re-architecting the application for the public cloud, including the potential to shave months off the re-architecting process and reduce risk by enabling load balancing and testing.

For IT architects charged with finding a cost-effective path to the cloud that helps protect existing VMware investments, a hybrid cloud may be the best solution. By combining a dedicated hosting environment with the public cloud, you can achieve all the benefits of the public cloud—including elasticity and cost optimization—while also retaining greater control, security, and flexibility over your deployment architecture.

As one VMware's largest VSPP partners, Rackspace has expert VMware Certified Professionals available and experience that comes with managing over 45,000 VMs. To find out more about how Rackspace can help you take the first step to the hybrid cloud, take a look at our Application Profiling Workshop.

ABOUT RACKSPACE

Rackspace® (NYSE: RAX) is the global leader in hybrid cloud and founder of OpenStack®, the open-source operating system for the cloud. Hundreds of thousands of customers look to Rackspace to deliver the best-fit infrastructure for their IT needs, leveraging a product portfolio that allows workloads to run where they perform best—whether on the public cloud, private cloud, dedicated servers, or a combination of platforms. The company's award-winning Fanatical Support® helps customers successfully architect, deploy and run their most critical applications. Headquartered in San Antonio, TX, Rackspace operates data centers on four continents. Rackspace is featured on Fortune's list of 100 Best Companies to Work For.

For more information, visit www.rackspace.com

Support: +44 (0)208 734 2700 | Sales: +44 (0)20 3553 3141 | Or Visit Us www.rackspace.co.uk
Rackspace Limited | 5 Millington Road | Hyde Park | Hayes | Middlesex | UB3 4AZ | UK