# The FUTURE of CLOUD is HYBRID OVER DISTANCE

## Multi-site Hybrid Clouds, the Next Stake in Ground in the Evolution of Cloud Computing

RACKSPACE®
HYBRID HOSTING

Written by Toby Owen, Hybrid Cloud Solution Leader,
Rackspace® Hosting



RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78218 U.S.A

rackspace® HOSTING

# Table of Contents

# 1. Executive Summary

In the evolution of cloud computing, hybrid cloud has made its mark; by changing the way businesses consume cloud resources, to fill gaps in the workload functionality of their dedicated technology. As the computing industry has advanced, the need to be able to connect multi-site hybrid clouds over distance has become more prominent. With distance comes lack of control over geography, increased latency, and concerns about connection reliability. Imagine the possibilities if an organization could use multi-site hybrid clouds seamlessly, without concern for these inefficiencies.

Last year, Brocade, EMC, and VMware collaborated to test and validate a solution that offers IT Operations the capability to dynamically migrate applications live across data centers, separated by a latency of up to 10 ms round trip (~550 miles), without business interruption. Once Rackspace caught wind of this initiative, we saw an opportunity to advance our vision to converge cutting-edge technology, to connect clouds, and to enable Enterprise-class customers to experience **Fanatical Support**® anywhere for their multi-site hybrid cloud solutions.

In this Proof of Concept test, Rackspace has explored the use of multi-site hybrid clouds over distance. We've continued the advancement of this technology in "the wild" by synchronously replicating data between data centers within a metro area, in order to move a live virtual machine between clouds.

# 2. Introduction

Over the past several years, cloud computing has moved from the early stages of exploration and discovery to become a major factor in nearly every IT project. Either it is the platform being used for new application development and deployment, or it is being evaluated as a potential solution for improvements to existing applications. Enterprises have been slower to adopt cloud than smaller businesses, since they have huge piles of "legacy" applications that were built around traditional IT architectures, which involve capital purchases, long development cycles, and many times are over-provisioned to ensure they stay available and functional during peak loads. Add to those challenges concerns about new technologies, an immature security model, lack of industry standards, and new development techniques and approaches with steep learning curves, and it's easy to understand the slow pace of cloud adoption in the enterprise.

Now that the cloud computing revolution is in full swing, some perspective is coming to light on what works best in the cloud, and how cloud can fit into existing IT deployment models. One overarching trend that has gained considerable momentum and focus is the use of hybrid clouds. Broadly, hybrid cloud can be defined as using cloud from an Infrastructure as a Service (IaaS) cloud provider WITH other platforms to deliver an application or workload to users. The promise of hybrid clouds has been a goal for many enterprises. The cost benefits of capacity on-demand of the cloud, the ease to spin up and down resources, and to pay for that capacity on an hourly or monthly basis is hard to ignore. Tie this with the fact that several virtualization technology vendors have built in support for moving live, virtual machines across a network, and this live motion represents a relatively easy way to transition applications and workloads between sites, promising even greater flexibility.

What many enterprises have found as they have explored hybrid cloud architectures more deeply, however, is that as distance between geographic sites grows the number of suitable use cases for a hybrid deployment shrink. So while many enterprises desire to move to a hybrid model, few have successfully done so.

The focus of this whitepaper is to explore the multi-site hybrid cloud model, seeking to uncover the use cases that will work, and what solutions are available to help meet the problems of data movement, latency, integration and management inherent in a distributed application architecture.

# 3. The Multi-Site Hybrid Cloud Model

A multi-site model for hybrid cloud involves attaching existing IT infrastructure, typically in a corporate owned data center, to a public cloud (IaaS) provider, either over a public internet connection (usually a VPN tunnel) or over a private leased line (usually a Metro Ethernet or Point-to-Point connection). The advantage of this approach is that core data used by the application can be maintained in place, as well as any legacy equipment that hosts parts of the application, such as an ERP system, a mainframe, or a large enterprise database, since many hosting providers don't support many of these platforms or applications (particularly the ones offering mature public clouds). The tradeoffs with this model include lack of control over geography, which can result in latency as distance between sites grows, time and expense to provision network connections, and reliability of the connection between sites (especially when using public transit).

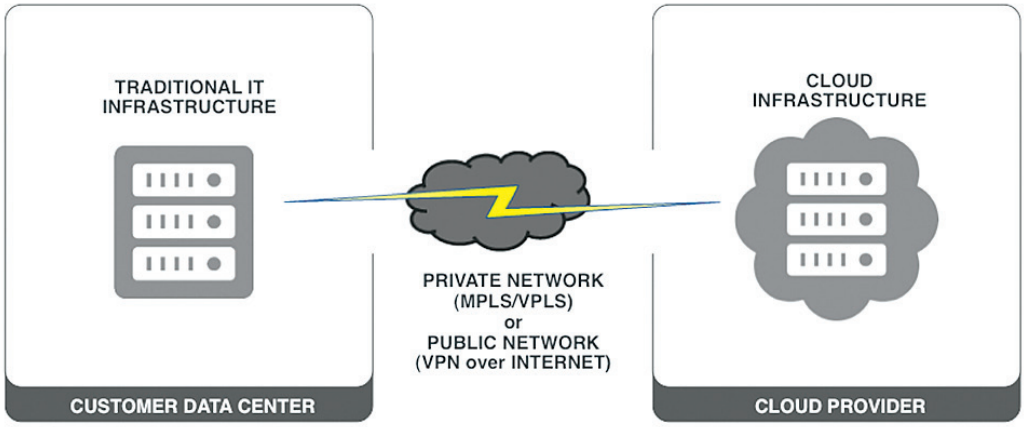


**Figure 1: A Multi-Site Hybrid Cloud**

There are several obstacles that exist for Enterprises transitioning to a multi-site hybrid cloud model. One of these barriers is that infrastructure architecture is highly dependent on application architecture. The models that work in a localized hybrid cloud may not work effectively across distance, due to the volume of data being passed, the frequency of updates, and the sensitivity of the application to delays in transactions. Live motion of a virtual machine is one such use case that works well in a Local Area Network (LAN), where bandwidth is plentiful and latency is negligible. But when attempting this over distance, latency becomes a limiting factor very quickly, since traditional means of motioning a virtual machine involve a copy of the machine configuration, followed by a movement of the data store. This data copy is what takes time, and if is the data is unavailable for too long a duration, the active virtual machine fails during motion.

Replicating the data in a synchronous fashion between the two sites allows a live motion to skip the data copy step – all that is transferred is the machine configuration and memory. The result is a very fast and reliable movement of an active virtual machine, since both sites are viewing the same data.

# 4. Hybrid for the Metro Area – A Leading Use Case for Multi-Site Hybrid Clouds

Live motioning of virtual machines can serve a variety of use cases, including:

- **Proactive disaster avoidance** becomes possible because machines can be moved quickly from the primary to secondary site.
- Maintenance without downtime can also be accommodated, either for a single application or even an entire data center, by migrating easily to a remote metro area site.

Other use cases also apply, even when not utilizing live motion of virtual machines, such as:

- **Zero downtime disaster recovery** (zero RPO) is possible when applications are actively deployed and load balanced between multiple sites, and because data is synchronized, fail back complexity can be reduced or eliminated.
- Synchronous data replication over a metro area connection by itself can unlock the ability to **outsource large development and test environments** by lowering the time and bandwidth required to copy test data back and forth.
- Low latency connections in a metro area can also facilitate **highly scalable applications**, where capacity can spike 10 times or more during peak times, by allowing additional capacity to be added at remote sites in the metro area, similar to a hybrid cloud within a single facility. This can be particularly beneficial from the financial standpoint, since the primary site can be sized to accommodate average loads, and burst load infrastructure can be consumed on a utility basis. Having data replicated at both sites as well as a contiguous layer 2 network between sites further reduces complexity of building burstable applications.

This architecture has been designed and tested previously by EMC, VMWare, and Brocade, using EMC VPLEX, VMWare vSphere servers using vMotion, and Brocade VDX fiber channel over IP (FCIP) to Ethernet switches in a lab environment. Given this solution stack, Metro vMotion can move a live virtual machine over a distance of up to 10 milliseconds, given a minimum bandwidth of 250 Mbps/migration, and a contiguous layer 2 network between the two sites.

Wanting to take this test farther, Rackspace has worked with EMC, Brocade, and Masergy to recreate this same solution "in the wild."

# 5. Hybrid Metro Proof of Concept

For the Hybrid Metro Proof of Concept, we created two environments for testing. The primary site (representing a customer data center environment) was in a colocation facility in Dallas, TX. The secondary site was located at the Rackspace Data Center, also in Dallas, TX. These sites are approximately 20 miles apart. Masergy Communications set up a Layer 2, 1 Gbps connection between the sites, using MPLS over Metro Ethernet. Network Latency for the Metro area connection averaged 1-2 ms, well within the threshold of 10 ms for synchronous data replication. In each site, we installed the following equipment: EMC VNX5300 storage array with 2TB storage

- EMC VPLEX VS2
- Brocade 7800 Extension Switch (FCIP switch)
- Two Dell R710 servers running VMWare ESXi 5.0.0
- Masergy Intelligent Bridge (MIB)

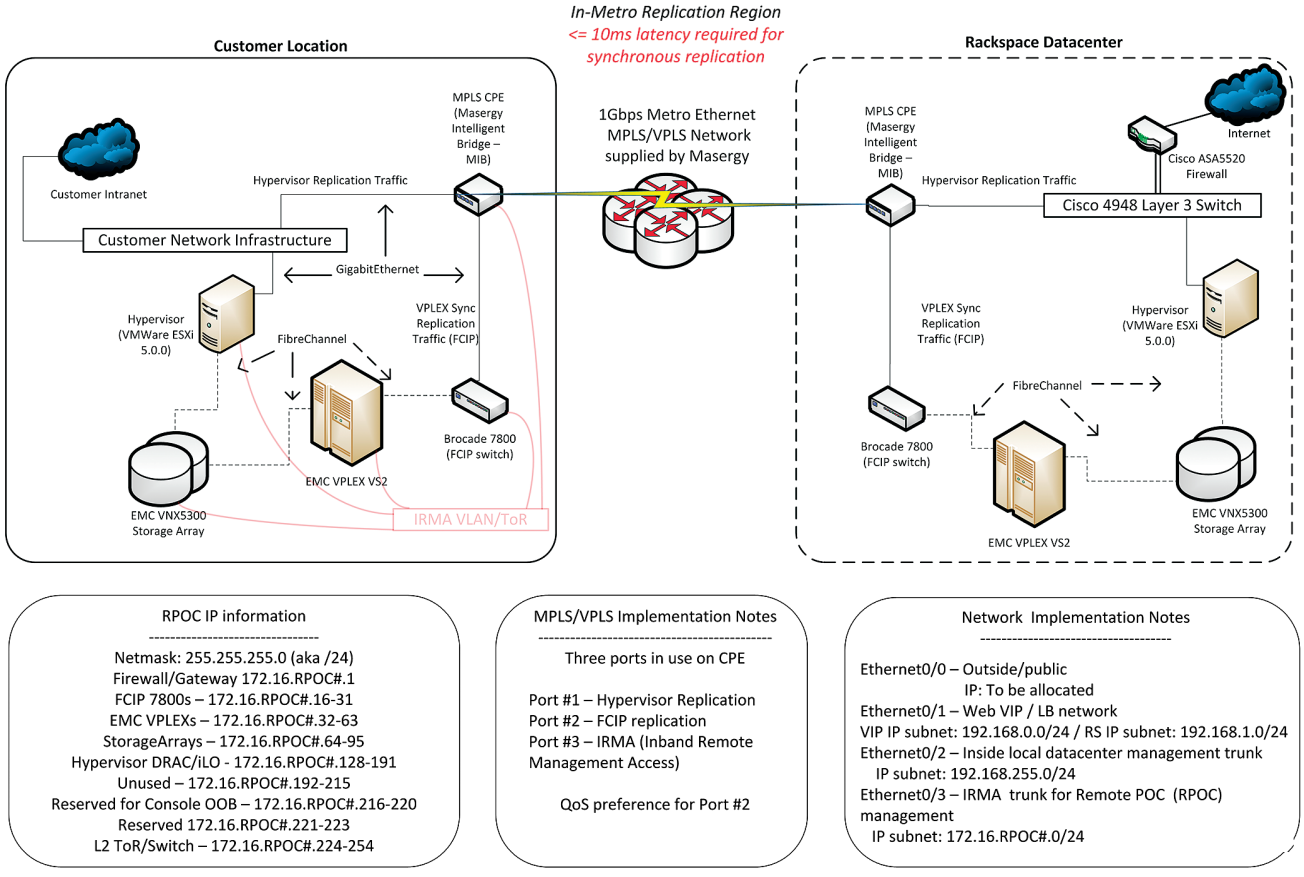


**Figure 2: Architecture for Hybrid Metro POC**

"The New Business Continuity: Moving Applications Across Data Centers Without Interruption" published by VMWare, 2011. http://www.brocade.com/downloads/documents/business-continuity-solution/new-business-continuity-moving-applications-across-data-centers-without-interruption.pdf

# 6. Hybrid Metro Testing

The Proof of Concept consisted of 2 phases – Setup and Testing.

### SETUP

The first step during setup was to configure synchronous data replication. Once the MPLS circuit was provisioned and connected, the VNX SAN arrays were initialized, and then set up to replicate. Replication over the WAN was handled via a Fiber Channel over IP (FCIP) connection.

Next, the Hypervisors and testing application were set up.

Two VMWare ESXi 5.0 hypervisors were installed on Dell R710 servers with 24gb of RAM and dual Quad Core Xeon processors at 2.27gHz. A single hypervisor was used to hold Windows® and Linux® Virtual Machines. The test application used was the DVD Store Version 2.1. This application is a complete e-commerce store simulation, including web server, application server, and database server layers as well as traffic/load generation and data capture.

### TESTING

Once the environment was configured, testing commenced, conducting the three tests described below.

### TEST 1: SINGLE SITE TEST WITHOUT REPLICATION

Testing was conducted using the DVD Store application with four threads running against a single CentOS server running MySQL and Apache. During Test 1, replication was disabled between the two locations to allow for a simulation of a single site. During Test 1 we averaged 1,830 orders per minute through the DVD Store with a max response time of 209ms and an average response time of 130ms.

After Test 1 was complete, the link was re-established and the two arrays were allowed to equalize.

| | Orders per minute | Maximum response time (ms) | Average response time (ms) |
|---|---|---|---|
| Average at primary site | 1,830 | 209 | 130 |

**Figure 3: Test 1 Results Summary**

http://en.community.dell.com/techcenter/extras/w/wiki/dvd-store.aspx

## TEST 2: SINGLE SITE TEST WITH REPLICATION

Testing was conducted using the DVD Store test application with four threads running against a single CentOS server running MySQL and Apache. This test was conducted while synchronous replication was running between the two sites. During Test 2 we averaged 1,350 orders per minute, with a maximum response time of 215ms and average response time of 171ms. SAN Replication accounted for an average of 5794.7 kB/s from primary site to secondary site.

| | Orders per minute | Maximum response time (ms) | Average response time (ms) |
|---|---|---|---|
| Average at primary site | 1,350 | 215 | 171 |

**Figure 4: Test 2 Results Summary**

## TEST 3: LIVE VMOTION OF THE DVD STORE APPLICATION

Test 3 consisted of a live motion of the active DVD Store application from the primary site to the secondary site. Testing was conducted using the DVD Store test application with four threads running against a single CentOS server running MySQL and Apache. This test was conducted while synchronous replication was running between the two sites. Before the vMotion was conducted, the site averaged 2,188 orders per minute with a maximum response time of 251 ms and an average response time of 109 ms. Migration of the site commenced, which took approximately 40 seconds. During vMotion, the connection to the store remained working and the site continued to process orders at a rate of 2,078 orders per minute, with a maximum response time of 1195 ms and an average response time of 114 ms. After vMotion completed to the secondary site, orders were being processed at at rate of 2019 orders per minute, with a maximum response time of 231 ms and and average response time of 118 ms. The DVD Store was then moved back from secondary to primary site. This vMotion also took approximately 40 seconds, again without interruption to the processing of orders.

| | Orders per minute | Maximum response time (ms) | Average response time (ms) |
|---|---|---|---|
| Average at primary site | 2,027.0 | 259.5 | 118.0 |
| Average during vMotion | 2,078.0 | 1,195.7 | 114.4 |
| Average at secondary site | 2,019.0 | 231.0 | 118.0 |

**Figure 5: Test 3 Results Summary**

# 7. Testing Analysis

The Proof of Concept yielded several important findings. First, synchronous data replication enables live motion of a virtual machine. Moving a live web application in less than a minute without interrupting active sessions shows the flexibility and resiliency of stacking data replication with high quality, high bandwidth private metro Ethernet and leading virtualization technologies. While maximum response times climbed during migration, average times stayed virtually unchanged, and total application throughput was not noticeably impacted. The implication of this finding for Enterprise IT organizations is massive. The flexibility and mobility of production applications no longer needs to be limited to a single data center or site. This opens the doors for major rethinking of data center expansion strategies, disaster avoidance and recovery plans, and globally distributed applications.

Second, while the testing was completed successfully, the engineering effort involved in setting up this application and environment was significant. The possibility in the future for customers to deploy this technology stack through partnering with a service provider could significantly lower the barrier to entry into these capabilities. And by building even more integration and automation around technology bundles such as those used in this Proof of Concept, the potential value of such a solution becomes even more compelling.

Finally, this test was conducted over a Metro Ethernet connection with an average of 1.7 ms latency. This can be stretched to greater distances. The tested solution stack has been proven to perform at up to 10 ms latency. This drastically expands the geographic reach of a Metro area Hybrid Cloud. And with asynchronous data replication, or other architectures, global distribution could certainly be possible, depending on each application's specific needs and traffic patterns.

# 8. Hybrid Cloud Network Design – An MPLS Perspective

As noted above, the Proof of Concept design and testing allowed local Ethernet switching at Layer 2 within a controlled environment. The practical application into an enterprise environment also requires Layer 2 connectivity, but across greater distances. Metro-E and Point to Point type services can potentially support the move into an Enterprise environment, but doesn't necessarily create a seamless addition to an existing network architecture. As a high percentage of Enterprise networks have adopted MPLS as the core technology, Rackspace approached Masergy with ideas on how to design a hybrid cloud solution to fit within a framework to support existing MPLS customers. The list of challenges included:

- Hybrid cloud capabilities over a single GigE MPLS Circuit – Virtualized Transport Layer to provide:
  - o Multiple and Logically Seperated Layer 2 VLANs for support of VPLEX (VMWare vSphere Servers using vMotion) and Brocade VDX FCIP connectivity
  - o Additional VLAN for Secure/Public Access into Environment
- Network Embedded (Cloud-Based) analytics for testing, traffic reporting and performance monitoring on an application by application basis
- Strict Priority QoS Model
- Optimized Circuit and LSP design to push the geographic limits of test environment providing low latency, but just as importantly, near zero jitter and 100% in-sequence packet delivery.

Within its Native MPLS Design, Masergy was able to provide connectivity between the Rackspace Data Center and an additional test location. As part of its standard offering, Masergy is able to provide multiple VLANs and services to every location supporting Private (Layer 2 and Layer 3) and Public VPNs with no limitations. Furthermore, through the use of their Intelligent Network Analyst tool, the team was able to verify the data streams under all conditions with granular reporting capabilities.

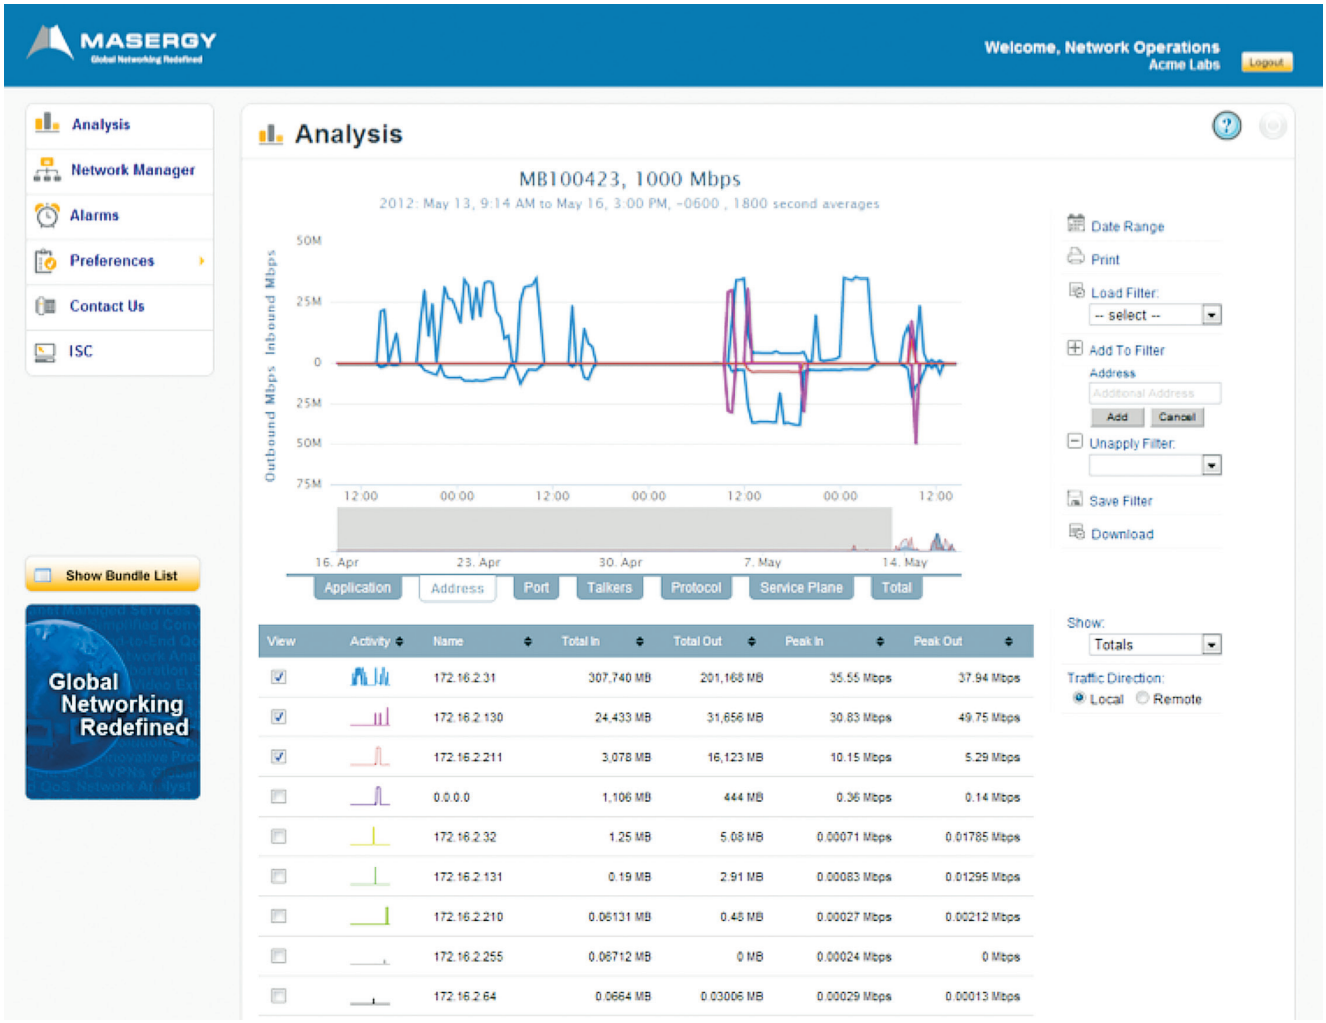


**Figure 6: Masergy Network Analyst Screenshot During POC Testing**

Hybrid clouds and the associated services are the future, but the network cannot be an afterthought. As cloud adoption proliferates, it is clear that no "one-size-fits-all" model delivers the dynamic applications that businesses crave. There are a variety of approaches to executing a cloud strategy based on the unique needs of organizations. However, even the most ambitious and innovative cloud strategies rely on one simple element – network quality. Without a sound and inherently flexible network design, business applications can't deliver the performance necessary to fulfill objectives. Installing a rigid or one-off network solution makes it difficult to adapt to the changing needs of cloud-based applications, limiting the very flexibility that drives cloud adoption.

# 9. The Future of Multi-Site Hybrid Cloud at Rackspace

Rackspace has developed a hybrid cloud solution called RackConnect™, which integrates dedicated hosting or private cloud infrastructure with the Rackspace Cloud. This solution has proven to be a significant benefit for customers who want to leverage public cloud resources as a portion of their hosted infrastructure. Customers have leveraged this solution to enable rapid scale up of applications for things like e-commerce solutions, marketing campaigns, Software as a Service (SaaS) offerings, rich media applications such as games, and social media apps. As a colocated hybrid cloud, RackConnect offers low latency, high bandwidth connections, allowing very creative use of infrastructure for workloads such as batch video transcoding applications, desktop as a service, and more. In addition, full integration and automation have simplified the provisioning and management of this hybrid infrastructure, and enhanced the security of a RackConnect customer's cloud deployment, by managing network access holistically and programmatically, with an easy to use Graphical User Interface. Hybrid clouds have become a highly demanded solution for our customers, from SMBs and Enterprises alike.

Rackspace in conjunction with NASA also helped create an open source project called OpenStack®, which seeks to define an open standard for cloud management. This project has grown to include more than 160 corporations who are actively advancing the codebase, which is now in its fifth release. Rackspace has also launched a next generation cloud based on OpenStack. In addition, Rackspace offers a private cloud solution based on OpenStack, called Rackspace Cloud: Private Edition. Customers can now have a private cloud deployment based on OpenStack, built and supported by Rackspace, which can be hosted anywhere in the world, either at Rackspace, a colocation facility, or in their own data center.

Rackspace also is the home of **Fanatical Support®** – world class service for IT and cloud hosting.

So what do RackConnect, OpenStack, and Fanatical Support have to do with one another? The answer is that these offerings will converge to create the future of hybrid cloud. Enterprises are demanding more cloud-based solutions, but are still saddled with mountains of legacy applications and infrastructure which they rely on at the heart of their businesses. These are difficult to forklift into the cloud all at once, but CIOs are demanding that their IT departments do more in the cloud. This drives the case for building hybrid clouds that can span distances, connecting and integrating on-premise applications with hosted and cloud-based infrastructure. Extending RackConnect beyond the data centers of Rackspace to enable global hybrid clouds becomes a new frontier that will unlock the cloud for many enterprises. The availability and increasing maturity of OpenStack Clouds, at Rackspace, other cloud service providers, and as private clouds for Enterprises will create a single management interface for managing remote and local cloud infrastructure. These future capabilities, including remote and local network and cloud management and integration, high quality, easy to provision Wide Area Networks, create the foundation for extending Fanatical Support to anywhere a customer needs world class support and service for their IT infrastructure and applications. Hybrid Metro Area Clouds represent one potential use case to accelerate the adoption of hybrid clouds by enterprises.

# 10. What Do You Need To Solve?

Does a Metro Area Hybrid Cloud represent a potential solution for your IT dilemma today? Does it trigger other use cases for hybrid clouds between sites that could meet your future IT needs?

If so, Rackspace wants to hear from you. Please follow the link below to complete a brief survey about what your needs are for hybrid clouds. As we seek to create the next generation of hybrid clouds, built on Fanatical Support, your feedback is crucial to our development process.

Survey link: **http://svy.mk/hybridmetro**

For more information, please e-mail us at HM.POC@rackspace.com.



RACKSPACE® HOSTING | 5000 WALZEM ROAD | SAN ANTONIO, TX 78218 U.S.A

# 11. Appendix

## ABOUT RACKSPACE

Rackspace® Hosting (NYSE:RAX) is the service leader in cloud computing and founder of **OpenStack**®, an open source cloud platform. The San Antonio-based company provides Fanatical Support® to its customers and partners across a portfolio of IT services, including Dedicated Cloud, Public Cloud and Hybrid Hosting. Rackspace has been recognized by Bloomberg BusinessWeek as a Top 100 Performing Technology Company and was featured on Fortune's list of 100 Best Companies to Work For. The company was also positioned in the **Leaders Quadrant by Gartner Inc**. in the "2011 Magic Quadrant for Managed Hosting." For more information, visit **www.rackspace.com**.

## ABOUT MASERGY

Masergy provides managed, secure global network services to enterprises that have complex needs across multiple locations. Masergy's integrated network and software solutions enable customers to seamlessly deploy and manage unified communications on a global basis.

Serving customers throughout the Americas, Europe, Asia, Africa and Australia, Masergy leverages advanced transport technologies to deliver global Ethernet services across a native MPLS network. The company pioneered customer-control network services and the delivery of multiple services over a single network connection. Masergy clients include global enterprises in a wide range of industries, including the financial/banking, energy, professional services, healthcare, entertainment, broadcasting, high technology and global manufacturing sectors. For more information about Masergy visit **www.masergy.com** or the company's blog **blog.masergy.com**.

## ABOUT BROCADE

Brocade (NASDAQ: BRCD) networking solutions help the world's leading organizations transition smoothly to a world where applications and information reside anywhere. (**www.brocade.com**)



## ABOUT EMC

EMC Corporation is a global leader in enabling businesses and service providers to transform their operations and deliver IT as a service. Fundamental to this transformation is cloud computing. Through innovative products and services, EMC accelerates the journey to cloud computing, helping IT departments to store, manage, protect and analyze their most valuable asset — information — in a more agile, trusted and cost-efficient way. Additional information about EMC can be found at **www.EMC.com**.

## DISCLAIMER

This Whitepaper is for informational purposes only and is provided "AS IS." The information set forth in this document is intended as a guide and not as a step-by-step process, and does not represent an assessment of any specific compliance with laws or regulations or constitute advice. We strongly recommend that you engage additional expertise in order to further evaluate applicable requirements for your specific environment.

RACKSPACE MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND, EXPRESS OR IMPLIED, AS TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS DOCUMENT AND RESERVES THE RIGHT TO MAKE CHANGES TO SPECIFICATIONS AND PRODUCT/SERVICES DESCRIPTION AT ANY TIME WITHOUT NOTICE. RACKSPACE RESERVES THE RIGHT TO DISCONTINUE OR MAKE CHANGES TO ITS SERVICES

OFFERINGS AT ANY TIME WITHOUT NOTICE. USERS MUST TAKE FULL RESPONSIBILITY FOR APPLICATION OF ANY SERVICES AND/OR PROCESSES MENTIONED HEREIN. EXCEPT AS SET FORTH IN RACKSPACE GENERAL TERMS AND CONDITIONS, CLOUD TERMS OF SERVICE AND/OR OTHER AGREEMENT YOU SIGN WITH RACKSPACE, RACKSPACE ASSUMES NO LIABILITY WHATSOEVER, AND DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO ITS SERVICES INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT.

Except as expressly provided in any written license agreement from Rackspace, the furnishing of this document does not give you any license to patents, trademarks, copyrights, or other intellectual property. Rackspace, Rackspace logo, Fanatical Support, RackConnect, and/or other Rackspace marks mentioned in this document are either registered service marks or

service marks of Rackspace US, Inc. in the United States and/or other countries. OpenStack® and OpenStack logo are either registered trademarks or trademarks of OpenStack™, LLC in the United States and/or other countries.

All other product names and trademarks used in this document are for identification purposes only to refer to either the entities claiming the marks and names or their products, and are property of their respective owners. We do not intend our use or display of other companies' tradenames, trademarks, or service marks to imply a relationship with, or endorsement or sponsorship of us by, these other companies.