

The Business Case for Deploying Cisco Metacloud



Executive Summary

Beginning with a brief overview of private cloud computing and the evolution of on-premises private clouds, we will compare Cisco Metacloud™ to options from VMware and the OpenStack community. We'll also describe Cisco's unique OpenStack as a Service (OaaS) operating model, and present Total Cost of Ownership (TCO) analyses that compare Metacloud to VMware vCloud Director and an OpenStack Do-It-Yourself (DIY) strategy.

“Over the past two years, private cloud has shifted from concept to reality.”

– Lauren Nelson, Private IaaS Cloud Lead Analyst
Forrester Research, November 2013

The Journey to Private Cloud

On-premises private cloud solutions allow IT organizations to retain control of company data and technology spending by empowering them to deliver new services as quickly and easily as external cloud providers.

It is a natural choice for companies concerned about data security, but it is increasingly of interest to large enterprises across all industries, particularly to those that have already made a substantial investment in their IT infrastructure and are looking to maximize the return on that investment.

Analyst data and client interviews suggest that organizations are only utilizing between 15 and 20% of their current server capacity – and only 30 to 40% of the capacity on physical servers that have been virtualized by implementing VMware, Hyper-V, and/or open source solutions such as KVM.

With that in mind, many proactive IT departments are currently implementing on-premises private cloud solutions that enable multiple business units to maximize physical compute, storage, and networking capacity within their own data centers. This strategy allows them to increase hardware utilization levels dramatically – often to as much as 70 or 80% – thereby securing the capacity they need for spikes and allowing them to defer or even eliminate the need for future hardware purchases.

On-premises private cloud computing has been evolving rapidly over the last decade and now involves a wide spectrum of technology solutions. Best practices include the following attributes:

- Self-service capabilities
- Automation of resource management functions
- Metering to support “charge back/show back” requirements
- Elasticity
- Ease of customization

Recently, Gartner reported that 32% of IT organizations are planning to implement a private cloud by 2015.¹

Key Takeaways

- On-premise private cloud solutions aren't just for companies concerned about data security.
- They're ideal for any firm looking to increase server utilization rates.
- They allow in-house IT organizations to deliver services as quickly as public cloud providers.

¹ Private Cloud Computing and the Future of Infrastructure, Gartner, Nov. 2012

“Private cloud has truly come into its own as a delivery mechanism that customers understand and are using to achieve the benefits of cloud where public options are either not available or viable.”

– **Allan Krans, Lead Cloud Practice Analyst**
Technology Business Research, November 2013

Private Cloud Deployment in the Enterprise

For businesses large and small, IT drives business results by unlocking vital revenue streams, rapidly capitalizing on new business opportunities, and recasting competitive landscapes. In highly successful businesses, IT has moved from being a cost center to a strategic partner, delivering tangible value and differentiation.

Cloud architectures are a critical component of this evolution, allowing businesses to reduce capital expenditures, scale capacity up and down quickly to meet usage spikes, and respond quickly to new business requests.

Users are no longer willing to wait for weeks or months for traditional IT departments to approve, procure, and image new hardware when they know they can place a call to a public cloud provider and have the compute capacity they need in a matter of minutes.

This increasing need for IT “response velocity” demands a solution that allows in-house IT departments to deliver services with the same speed and simplicity as external cloud providers, reducing or eliminating instances of solutions built and used inside organizations without organizational approval, known as “Shadow IT.”

Key Takeaways

- If IT departments are to become strategic internal business partners, they must be able to respond quickly to the evolving needs of their business.
- An on-premises private cloud deployment can greatly enhance the ability of an IT department to address new business initiatives.
- Being able to rapidly deploy a private cloud for an enterprise can alleviate problems associated with shadow IT.

“By any measure, private cloud computing has become a major trend since the concept was first publicly discussed in 2008. While it isn’t mainstream yet, many enterprises have deployed something, and have plans to expand. Market hype is being replaced with focused investments and realistic plans.”

– **Thomas J. Bittman, VP Distinguished Analyst**
Gartner Research, October 2013

Enterprise Options for On-Premises Private Clouds

There are two leading technology options in the market today for on-premises private cloud architectures: VMware’s vCloud Director and OpenStack.

VMware

vCloud Director provides a mature set of tools and management features, described in detail on [VMware's website](#).

However, vCloud Director is a proprietary and expensive solution that cannot be integrated with other public cloud offerings such as Amazon, or solutions from Rackspace and HP. A VMware hybrid cloud deployment is currently only possible between two VMware environments, limiting the flexibility that most IT departments require.

OpenStack

OpenStack is an open source cloud platform that controls large pools of compute, storage, and networking resources throughout a data center and is managed through a dashboard that gives administrators centralized control while empowering users to provision resources instantly. For more information, please read <http://www.openstack.org/software/>.

The current market momentum for OpenStack is considerable, with over 18,000 members, 300+ supporting organizations spanning 140 countries, and involvement from companies such as IBM, HP, EMC, Oracle, Brocade, Intel, Seagate, Juniper Networks, and Cisco.

The OpenStack Foundation is an independent non-profit entity providing shared resources to help achieve a mission of protecting, empowering, and promoting OpenStack software and the community around it. The OpenStack project currently has approximately 2,500,000 lines of code and is in the Juno (10th) code release. There have been hundreds of thousands of downloads of the software. A technical committee stewards the direction of OpenStack development and includes elected leads for each of the six core software projects. Large enterprise production OpenStack environments have been deployed at PayPal, Wells Fargo, Disney, CERN, Comcast, and Tapjoy [among others](#).

Technical operations teams currently evaluating or implementing OpenStack have a choice: They can elect a DIY strategy or partner with a company – such as Cisco – that specializes in implementing and supporting OpenStack.

The OpenStack project is complex and evolving quickly, with major releases occurring every six months.

As such, the DIY approach is impractical for most companies. Keeping current with the project, customizing the platform, and troubleshooting day-to-day issues requires significant internal resources (dedicated headcount), many of whom have other business-critical responsibilities.

Introducing Metacloud

Metacloud is a remotely engineered and operated private cloud – built for developers. If you like the public cloud, you'll love this solution. It includes full self-service access to compute, storage, and network resources – available through an easy to use web portal, OpenStack APIs, the command line, or integrated out of the box into your favorite DevOps and development platforms. Best of all? You don't need to spend a minute being distracted with how the platform is engineered, built, upgraded, and operated. Like public cloud, it just works. It's easy, fast, and predictable.

Your administrators will love Metacloud, too. With heavy experience in large-scale operations, Cisco knows it's important to understand what's going on behind the scenes. That's why we include tools to help your admins know what's happening at all levels of the stack. You get the ability to assess and optimize performance of applications by evaluating the health of your hypervisors, OpenStack services, network traffic, and capacity – all without having to be experts in how they work.

Metacloud includes all the features of OpenStack plus significant enhancements to the core OpenStack distribution, including:

1. Highly Available and Scalable Architecture

Metacloud includes a redundant and scalable architecture on top of industry standard open-source clustering tools. All OpenStack orchestration services can run any of 3+ control nodes, ensuring that our customers' clouds are always available and able to dynamically grow to meet their changing business needs.

2. Active Directory and LDAP Integration

Metacloud supports authentication against Active Directory, LDAP, and local accounts, allowing customers to leverage their existing user accounts. Authorizations are stored in Keystone, the OpenStack Identity Provider, allowing users to get started immediately without needing to make complicated schema modifications to their directory service.

3. Optimization for Data Center Ethernet

The Metacloud Network Manager is designed to conserve VLANs in a data center Ethernet. A tenant can initially be assigned a single VLAN and CIDR range. Once the initial CIDR is full, an additional CIDR range can dynamically be assigned to the same tenant and VLAN, allowing customers to dynamically grow their clients without consuming additional valuable VLAN resources.

4. Self-Service Token-Tenant System

The Metacloud Token-Tenant System allows non-admin users to create their own tenants and send tokens to their team members for access to the same tenant without an administrator needing to curate users. When combined with Active Directory or LDAP integration, it can provide a zero-touch, self-service environment for users.

5. Distributed, Replicated, Fault-Tolerant Storage

Metacloud integrates Ceph's RADOS Block Device (RBD) for reliable storage of VMs across commodity hardware. RBD-backed instances allow customers to be confident that their data is secure without having to invest in expensive high-end storage hardware.

Metacloud Architecture

A core component of the design and architecture of Metacloud is the Cisco Control Point (CCP). These control points provide the high availability, data resiliency, and rapid scale-out functionality of our solution for each availability zone.

All of the core OpenStack cloud services are highly available (Nova API, Rabbit MQ, etc.), and there is logic in Metacloud that dynamically balances cloud service requests across all controller nodes. If one controller node fails, the other controller nodes pick up the workload. This design and implementation is far ahead of others in this space and not available in the current "off the shelf" OpenStack distribution. It enables Cisco to support multiple hardware solutions and also allows us to provide our advanced support model (detailed in a following section of this document).

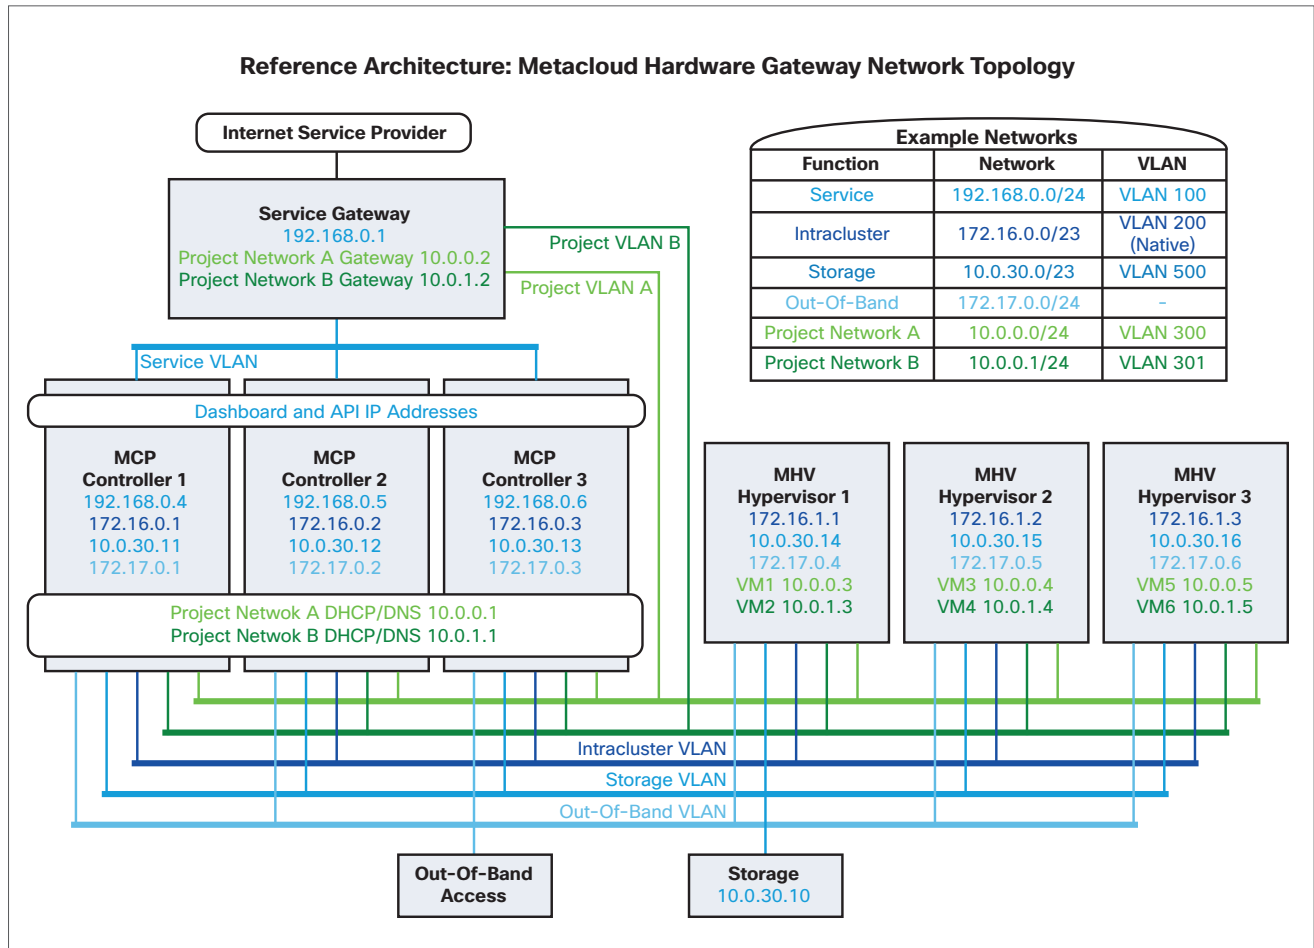
On the next page is a drawing of one of the two environments Cisco typically supports...a hardware gateway architecture. Cisco also enables enterprises to leverage existing network gateway solutions.

Key Takeaways

- A core component of Metacloud – the Cisco Control Point – enables data availability and resiliency as well as rapid scale-out capabilities.
- The dynamic load balancing and failover capabilities in Metacloud far exceed those available in the current "off the shelf" OpenStack distribution.
- Metacloud is capable of supporting both hardware gateway and network gateway architectures.

“Average length of downtime per disaster recovery event: 8 hours for non-cloud users; 2.1 hours for cloud users.”

– Aberdeen Group Cloud and IT Disaster Recovery Statistics, 2011



The Metacloud Difference

Companies doing business in the OpenStack space come to market with a wide variety of business models. Some are building OpenStack appliances. Some deliver OpenStack on a USB drive. Others use OpenStack as an opportunity to provide consulting services. Most employ very prescriptive reference architectures for server, storage, and network platforms. And almost all leave solution deployment, management, and support to the customer.

The Cisco OpenStack as a Service (OaaS) model is unique because we allow our clients to leverage their existing compute, server, and storage assets for their on-premises private clouds. Our solution - Metacloud - is hardware agnostic, so it works with products from HP, Dell, Cisco, IBM, Quanta, and even white box manufacturers. Metacloud also supports Virtual Data Center Pod environments from NetApp and Cisco and can support an EMC vBlock environment.

This allows our customers to continue to enjoy a return on the hardware investments they've made over the past several years and allows them to employ a multiple hardware vendor strategy for the future, which is desirable for negotiation leverage and risk reduction. This may seem simple. However, our competitors have challenges supporting heterogeneous environments and generally require new hardware purchases for production deployments.

Cisco is also different because we partner with customers to run the clouds we build. We don't drop a solution at the doorstep and expect our customers to figure out how to make it work. We take full responsibility for keeping our customers' clouds up and running - day-in and day-out.

We monitor their environments, troubleshoot hardware issues, deploy updates, implement feature enhancements, and partner with their IT organizations to help them deliver an outstanding level of service to the users they support. And we back it all in writing with service level agreements that ensure performance excellence.

Finally, Metacloud is unique in that this solution is production-ready in weeks rather than months. This is in sharp contrast to our competitors, whose solutions generally require a lengthy and tedious testing and QA process before they can be deployed for mission-critical functions.

Key Takeaways

- Metacloud is a hardware-agnostic solution that allows our clients to leverage their current hardware investments for compute capacity.
- Cisco partners with customers to operate the clouds that we build.
- Metacloud deployments are production-ready in just weeks.

Cisco's Proactive Support Model

Advanced Operational Support is an important differentiator between Cisco's solution and those of our competitors. Like many of the most powerful applications in widespread use today, OpenStack is a fundamentally complex technology. As such, running it effectively requires a level of expertise that simply isn't common among a majority of corporate IT teams.

With that in mind, Cisco includes subscription-based operational support with every solution we implement. From the day a customer's cloud is first deployed, we run it with them as an extension of their team, helping ensure the success of the project as well as their satisfaction with Cisco. Our support efforts can be broken down into four core functions: Operational Support, Upgrades, Bug Fixes, and Capacity Planning.

Operational Support

In a typical cloud deployment, customers would be handed software, given documentation on how to troubleshoot that software, and provided a hotline number for problems that they couldn't address on their own. It would be up to them to diagnose problems when they arose (often the most difficult part of the support process), describe them to the respective technicians on the other end of the support line, and then implement the fixes suggested by each technician.

Cisco approaches support from an entirely different point of view. We take full responsibility for the ongoing functionality of the clouds we deploy. So we're monitoring them 24x7x365 and we're notified of any problems in real-time.

Roles and Responsibilities	Cisco	Client
Provide enterprise version of OpenStack, fully supported by Cisco.	X	
Provide on-going preventative maintenance against platform, including providing updates.	X	
Provide 24 x 7 remote monitoring and support of Cisco's platform including remote triage. This includes all the key components of the Metacloud platform up to the hypervisor layer.	X	
Meet SLA requirements.	X	
Provide software enhancements including remote upgrades following the OpenStack release schedule.	X	
Provides all physical aspects of the environment including infrastructure and power.		X
Architect and design all dependencies between the application and the system software environment captured in the virtual machine ("VM").		X
Maintain and administer operating systems images for hypervisor VM's, which may also include items such as LAMP, MySQL, or other infrastructure components.		X
Operate and consume the cloud platform itself, including creating and managing projects, tenants, users, public IP's and resources within the cloud platform out to the user community.		X

Once a problem is detected, we immediately begin addressing it, and often have it completely resolved before our customers are even aware that anything has occurred. They are entirely free of the burden of figuring it out, calling support, and implementing solutions. The goal of our support model is to enable them to simply consume compute resources and focus their staff's valuable time on driving value for their core business.

Upgrades

OpenStack is continually releasing upgrades that increase its functionality. And while this is definitely a positive, the fact is that these releases come in varying levels of completeness. They may contain bugs, features that are incomplete, or drivers that are not compatible with the devices typically in use.

This creates a noteworthy burden for any IT team attempting to implement upgrades on their own. They have to evaluate every upgrade that comes in, determine whether it is ready for deployment, make fixes if necessary, and figure out how to apply it. Traditionally the amount of planning and preparation around such a cycle amounts to weeks or more.

As a Cisco customer, we handle it all for you. We assess every upgrade that comes in, whether from upstream OpenStack, kernel teams, or other resources, and ask, "What state is the functionality in?" The answer is generally one of the following:

1. **Ready to go.** In this case, we don't need to do any work on the upgrade. We simply install it as-is.
2. **Almost ready.** This upgrade needs some work. Cisco engineers will work to fix the deficiencies, then install it across our client base.
3. **Not ready.** In this case we do not install the upgrade.

When Cisco readies an upgrade for deployment, we don't only ensure it's ready for OpenStack - we ensure it's ready to integrate with the entire stack. That includes the hypervisor underneath OpenStack, the supporting tools, libraries...everything from the bare metal server up to and including OpenStack. Cisco takes a holistic view of the cloud so our customers can be confident that their upgrade will enhance the functionality of their platform without disrupting their business.

Once we've readied the updates for installation, all they have to do is notify us of the time that will work for them, and we take care of it - with virtually no downtime.

Capacity Planning

Cisco can also assist in capacity planning for a cloud's storage environment. Our advanced support model, which includes monitoring of CPU, memory, and storage utilization, provides our cloud operations team with valuable insights into the performance and capacity requirements of our customers' clouds. Based on trend analyses generated from this data, we can project growth rates and offer proactive advice on scaling their cloud infrastructure to ensure that they are well prepared to meet the needs of their growing businesses.

On-Premises Private Clouds - Economics and TCO

For on-premises private clouds, Cisco licenses its technology and support through a convenient and highly predictable monthly subscription model. Pricing is based on the number of physical compute sockets in a Metacloud environment. The monthly subscription fee is calculated by multiplying the per-socket fee by the number of sockets.

Clients also have an option to leverage Ceph, a highly available clustered storage solution. This solution is priced in a similar subscription model - a monthly subscription fee is calculated by multiplying the amount of raw storage by the per-GB capacity fee.

Many clients begin working with OpenStack in a phased roll out schedule. A typical Metacloud installation starts at 17-20 physical servers (just about a full rack). A full rack of current generation compute servers can generally support the equivalent of 1,600-1,800 small-sized VMs. This environment is referred to as an Availability Zone and the infrastructure can scale out to thousands of sockets in a data center.

Key Takeaways

- The Metacloud licensing model provides highly predictable monthly billing.
- Subscription fees are calculated by multiplying a per-socket fee by the number of sockets in the private cloud.
- Incremental features/technologies, such as Ceph, can be added to a Metacloud deployment using a similarly predictable billing mechanism.

TCO Analysis – Metacloud vs. VMware vCloud Director

Cisco and VMware have vastly different business models when it comes to private cloud solutions. Cisco offers its service through a subscription model and offers 24x7x365 support of the infrastructure while VMware has a traditional on-premises licensing model that requires significant upfront licensing and then requires annual maintenance fees for production support.

The VMware licensing agreement does not include proactive 24x7x365 support of their solution nor does VMware offer a service level agreement for uptime guarantees.

On-prem private cloud pricing assumptions

- 40 sockets – one rack – 20-node dual socket, Intel Xeon Gen8 server
- VMware vCloud Director Edition licensing
- 25% discount price per socket for vCloud Director
- 3-year maintenance included
- Estimated install fee included
- 1 administrator for VMware
- 24x7x365 support provided by Cisco

Results

Cisco licensing is over 60% less expensive.

Key Takeaways

- The traditional software licensing model used by VMware includes a significant upfront expense in addition to an ongoing annual fee.
- VMware DOES NOT provide a service level agreement for uptime guarantees or include 24x7x365 support.
- Metacloud offers disruptive economics compared to VMware vCloud Director.

TCO Analysis – Metacloud vs. OpenStack DIY

OpenStack source code is available for download at no charge and is a part of the Apache 2.0 license. However, free software does not equate to a free cloud. Keeping current with OpenStack requires one or more full-time internal resources dedicated to the rapid development pace of the project.

The OpenStack community is continually adding code and the technical committee is releasing new features every six months. Companies implementing OpenStack on their own need to evaluate each new feature and functionality upgrade to make sure it meets their needs and, in cases where it doesn't, add custom code to adapt OpenStack to their specific business workflows.

With that in mind, a key consideration when evaluating OpenStack DIY vs. Metacloud should be the cost of employing personnel dedicated to customizing, maintaining, and administering the project. Another is the time-to-market value of implementing a solution that is production-ready right away vs. building a solution that may be deployed in the next two to four quarters.

For the purposes of this analysis, we are comparing the cost of employing two full-time dedicated resources vs. the cost of deploying the Cisco solution for a 40-socket on-premises compute environment.

Metacloud Assumptions

- Installation included
- Monthly subscription price included
- Set-up fee included

DIY Assumptions

- One cloud operator at market rates
- Two OpenStack developers at market rates

Results

Metacloud is over 65% less expensive than deploying an OpenStack DIY solution.

The analysis above is extremely conservative and rationalizes the build vs. buy decision. A Metacloud solution is much less expensive than employing dedicated staff for an OpenStack DIY initiative.

Key Takeaways

- The Metacloud model avoids the cost of dedicated staffing for OpenStack projects.
- Cisco can significantly reduce time-to-market by deploying a solution that is production-ready on Day One.
- Software updates are managed by the Cisco engineering team.

Final Thoughts

Cisco employs a dedicated team of world-class developers and cloud operators that instantly become an extension of our customers' IT organizations. Our team is continuously innovating and adding new features to Metacloud, ensuring that our customers always have the latest cloud capabilities.

We can help mitigate time-to-market pressures by bringing a new customer's on-premises private cloud into production within weeks of executing an agreement.

To request a custom quote or schedule a demonstration, just visit the [Metacloud webpage](#), submit the appropriate form, and a Cisco representative will be in touch.

“Public and private cloud solutions each come with their own pros and cons. The public cloud becomes prohibitively expensive at scale and doesn't utilize on-premises hardware, while the private cloud demands intensive resources and knowledge just to maintain and manage. Cisco (formerly Metacloud) intends to deliver only the pros, offering solutions that provide the convenient and instantly-updated public cloud experience, along with the flexibility, security, and cost effectiveness of the private cloud – all while protecting legacy data center infrastructure investments.”

– [Jeff Vance, CIO.com](#)
September 2013