



EQUINIX

Application Performance

A FRAMEWORK FOR CLOUD ENABLEMENT

EQUINIX WHITE PAPER

Introduction.....	3
The Ever-changing Nature of IT	4
Cloud “What” and “Why”.....	6
Cloud Computing Defined.....	6
Essential Characteristics.....	6
Cloud Service Models	7
Cloud Deployment Models.....	8
Enterprise Cloud Enablement.....	10
Identifying Users and Their Needs.....	10
Defining Performance Requirements.....	11
Three Steps to the Cloud-enabled Enterprise.....	12
Benefits of Cloud Enablement.....	16
Considerations for Cloud Implementation	17
Understanding Networks.....	17
Cloud Hubs	18
Deploying Globally	19
Platform Equinix: A Global Interconnection Platform.....	20
Summary and Recommendations	22

THE CLOUD-ENABLED ENTERPRISE

In today's economy, companies that can quickly and effectively respond to changing information technology needs have a distinct advantage over those that cannot, by being in a better position to take advantage of new opportunities. And if they can procure and consume only as much as they need with no added waste, they can also save a lot of money.

Cloud computing offers the promise of cost-effective capacity on demand— anytime, anywhere, on any device. A cloud-enabled enterprise is able to exploit cloud services to meet its information technology needs easily and effectively. However, taking advantage of cloud computing is not as simple as just signing up for an account. Companies wishing to move to a cloud-enabled IT model need to take an end-to-end systems approach. This begins with a thorough understanding of the applications and services their users require, since end-user quality of experience (QoE) is ultimately what matters. User experience is closely tied to application response time, which in turn depends on effective delivery of content. Once clear and objective targets have been established, companies must optimize their network infrastructure to satisfy performance requirements and enable consumption of network-based services such as cloud computing. Once these steps are taken, companies are ready to effectively exploit cloud services.

This white paper describes the motivation behind enterprises' cloud adoption, and outlines the considerations, best practices and necessary steps for effective implementation of cloud-based services.

THE EVER-CHANGING NATURE OF I.T.

The challenge of balancing cost pressures with business demands has led to an ongoing search for cheaper, faster and better ways of structuring and delivering information technology. New ways of satisfying IT needs—usually driven by technology advances—result in waves of change as depicted in Figure 1. Over the past 50+ years, there have been several such waves, typically spaced approximately 10-15 years apart.

The first wave of corporate IT—or Management Information Systems (MIS), as it was known at the time—began in the late 1950s. Many companies purchased and implemented central **mainframe computers** to automate specific corporate functions. These environments were carefully controlled, with end-users having little if any access.

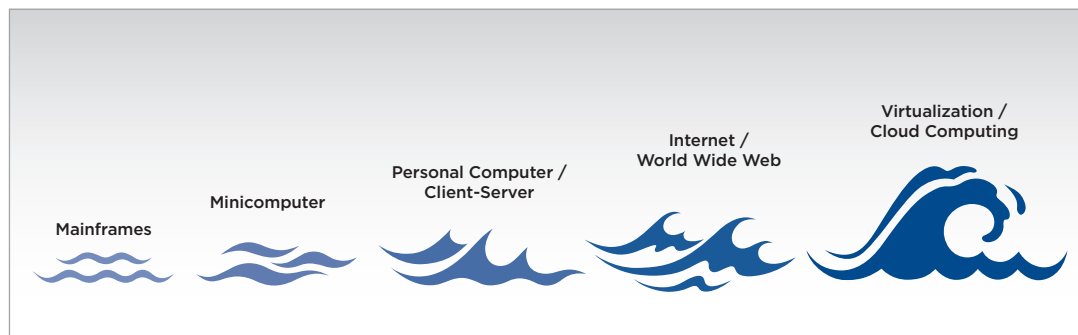


Figure 1: Waves of IT Change

The move from mainframe to **minicomputers** in the 1960s and '70s was the next wave. Minicomputers enabled smaller groups of users such as departments to have dedicated computing resources, thereby beginning the democratization of IT. And although end-users now had easier access to computing, there was still no question about the fundamental purpose of computing. It existed purely as a business tool.

The rise of the **personal computer** in the late 1970s and '80s led to client-server computing, the next wave of computing evolution. This transformation continued to push computing power closer to the end-user. Users now had access to something that had previously seemed unimaginable – their own dedicated computer. And while the advantages of this model were clear, for the first time, parts of the computing environment were decentralized. This led to new challenges in application design and delivery, and innovative approaches to software engineering were required to meet the demand.

The next wave of computing came in the 1990s and 2000s with the creation of the **World Wide Web**. Data sources and resources resided in many different places, with many different owners. In addition, end-users were now using computers more for personal in addition to work-related purposes. This transformation required new ways of architecting, securing and delivering applications. It also led to the widespread commercialization of computing as everyone from global corporations to single-site merchants established an online presence.

Today, we find ourselves in the midst of yet another computing transformation. The origins of this wave are tied to Moore's Law¹, the tendency of computing power to double approximately every two years. What once had been scarce and precious has become readily available and affordable.

Companies and users, however, could not easily take advantage of this abundance of capacity as long as the unit of consumption was a physical machine. As a result, companies saw the average utilization rate of their computing systems steadily decrease—often to less than 10 percent of capacity—despite the increase in available power. The initial solution has been **virtualization**—the ability to divide up physical systems into groups of virtual ones. This has allowed IT organizations to make far more efficient use of ever-faster systems.

The benefits of virtualization are clear, but they are limited to system owners. We are now entering the second half of the current wave—**cloud computing**.

1. "Moore's Law," http://en.wikipedia.org/wiki/Moore's_Law

CLOUD “WHAT” AND “WHY”

To understand the nature and appeal of cloud computing one need only look at common computing challenges. Businesses need a cost-effective, highly available, easily consumed source of computing services just like they need affordable, reliable electrical power. What they have today are groups of interconnected physical and virtual systems that require constant investment, management and support. End-users want reliable, easy-to-consume applications and services rather than complex, hard to maintain systems. Both groups would rather not need to know how the systems work. Simply put, they want utility computing services.

Cloud Computing Defined

Many definitions of cloud computing exist. One of the most widely accepted and most useful descriptions comes from the United States National Institute of Standards and Technology (NIST). The NIST Definition of cloud computing² states the following:

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models and (multiple) deployment models.

Essential Characteristics

As shown in Figure 2, the NIST cloud model defines the following essential characteristics and service models.

On-demand Self-service

A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

Broad Network Access

Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops and workstations).

Resource Pooling

The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state or data center). Examples of resources include storage, processing, memory and network bandwidth.

Rapid Elasticity

Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

Measured Service

Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth and active user accounts). Resource usage can be monitored, controlled and reported, providing transparency for both the provider and consumer of the utilized service.

2. "The NIST Definition of Cloud Computing," <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>

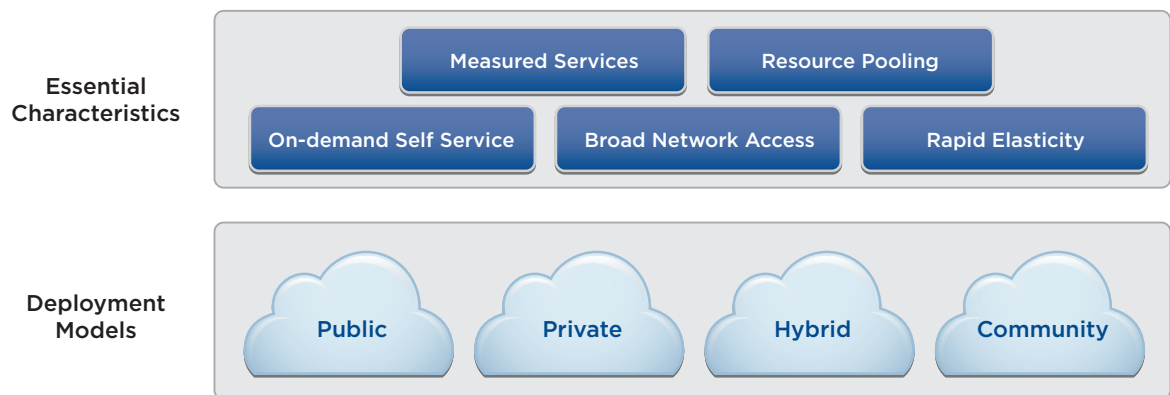


Figure 2: NIST Cloud Computing Service Models and Essential Characteristics

Cloud Service Models

The different cloud service models allow customers to decide which parts of their applications they want to own and manage and which parts they want someone else to be responsible for. In many ways this reflects an evolution that's been going on for years. As illustrated in the Figure 3, many aspects of cloud service models are analogous to existing, traditional service delivery models.

In terms of what the business owns versus what is provided by a vendor, **infrastructure-as-a-service (IaaS)** is similar to a managed hosting environment where a third party manages the physical systems and the application stack up through the operating system. **Platform-as-a-service (PaaS)** adds application components such as databases, runtime environments, middleware and developer tools. Finally, with **software-as-a-service (SaaS)** the entire application stack is owned, operated and managed by a vendor. The idea is basically the same as the application service provider (ASP) model from the dot-com era.

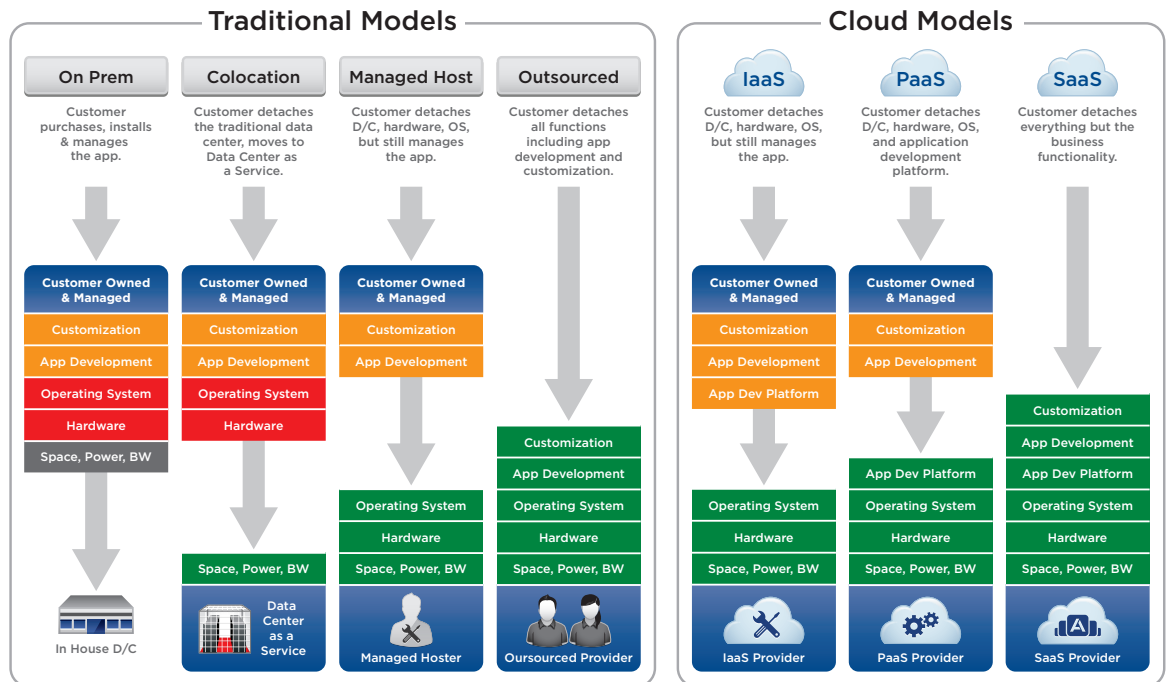


Figure 3: Traditional vs. Cloud Services Models

Cloud Deployment Models

The three main deployment models for cloud computing include public, private and hybrid. The basic definitions for each are:

- **Public** – IT capacity provided in a multi-tenant environment by a service provider and consumed as a service
- **Private** – IT capacity running on owned, dedicated systems and presented via a cloud consumption interface
- **Hybrid** – a combination of public and private clouds, possibly with the ability to move workloads from one to the other

Characteristics of Public Cloud

The appeal of on-demand capacity consumption is clear. Instead of capital-intensive build-outs with long lead times, application developers can move to a just-in-time consumption model. While this seems like a much easier way to operate, it comes with several challenges:

- **Long-term Cost** – Steady-state consumption of public cloud can be very expensive.
- **Limited Flexibility** – Public cloud providers offer pre-defined virtual machine sizes and service offerings. Your needs may not match.
- **Opaque Technology** – Most public cloud providers do not share the details of their underlying technology stack. This can be a concern to both developers and auditors.
- **Shared Security** – Simply put, multi-tenant environments present greater potential risk than private deployments.

So while public cloud can shorten development times and costs for new applications, it isn't always the best solution for ongoing operations.

A Look at Private Clouds

Like public cloud, private cloud computing offers simplified consumption and greater flexibility than a pure hardware-based deployment. Unlike public cloud, though, it runs on dedicated systems that are owned and managed by a single organization. These systems run a cloud software stack that abstracts the hardware into pools of virtual resources. Since most organizations have already virtualized at least a portion of their IT environment, many aspects of private cloud will be familiar. The advantages of private cloud include:

- **Cost Savings** – As previously mentioned, public cloud can be very expensive for steady-state workloads. When capacity requirements are well understood, it is more cost-effective to run at least a portion – the base load component – in a private cloud.
- **Added Flexibility** – Unlike public cloud, owners of a private cloud environment are free to configure virtual systems in any way they choose (within the limits of the underlying virtualization layer).
- **Technology Visibility** – A private cloud owner has clear knowledge of component hardware and software.
- **Enhanced Security** – Rather than a multi-tenant security model, private clouds are single-owner environments.

Hybrid—Mixing Public and Private

The combined use of both public and private clouds is often termed “**hybrid cloud**.” Because every environment is unique, there is no universal definition of hybrid cloud computing. For some companies, it may involve running certain workloads in a dedicated, private cloud and others in a public cloud. Others may take an “**Own the Base, Rent the Spike**”³ approach, using their private cloud to supply cost-effective base capabilities and bursting into the public cloud for high load or unusual demand conditions. Still others may employ a layered cloud design, running some components of a single application stack on physical hardware and others in the public cloud. Whatever the approach, hybrid cloud aims to balance control, cost and performance for effective service delivery.




Figure 4: Hybrid Cloud

3. “Own the Base, Rent the Spike,” <http://blog.equinix.com/2012/03/own-the-base-rent-the-spike/>

ENTERPRISE CLOUD ENABLEMENT

Clearly, the cloud has many characteristics that appeal to businesses. But this doesn't mean that existing ways of delivering services can be simply swapped for new ones. Change—especially when it involves systems required to run the business—takes time as well as careful planning and coordination.

Identifying Users and Their Needs

As previously stated, end-users are the ultimate arbiters of IT value. So the first thing an enterprise needs to do when developing a cloud strategy is to understand their users—all of them. The process is depicted in Figure 5. Users can come in many forms, including:

- Employees
- Customers
- Suppliers
- Partners

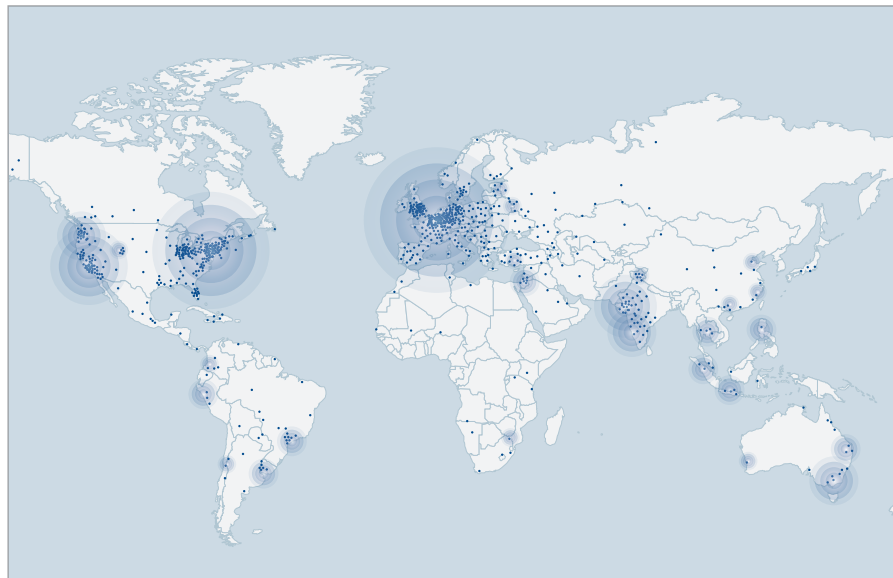


Figure 5: Example Plot of User Locations

Each of these may have different needs and expectations. For example, some employees may work at a company office location while others may travel frequently or work remotely. Suppliers may be subject to governance requirements from a different jurisdiction than that of the enterprise. Customers could potentially be anywhere.

As a prerequisite for moving to a services strategy, enterprises should understand the following characteristics of their user base:

- Location
- Type (employee, customer, partner, etc.)
- Applications/services used
- Access devices
- Usage patterns (time of day, length of interaction, etc.)

Defining Performance Requirements

Once user populations are understood, it's time to focus on the applications and services they need. The goal is to establish requirements for a good user experience. Perceived application performance may be closely tied to network latency. As shown in Figure 6, the impact can range from applications that are relatively unaffected by network delays to ones that are highly sensitive.

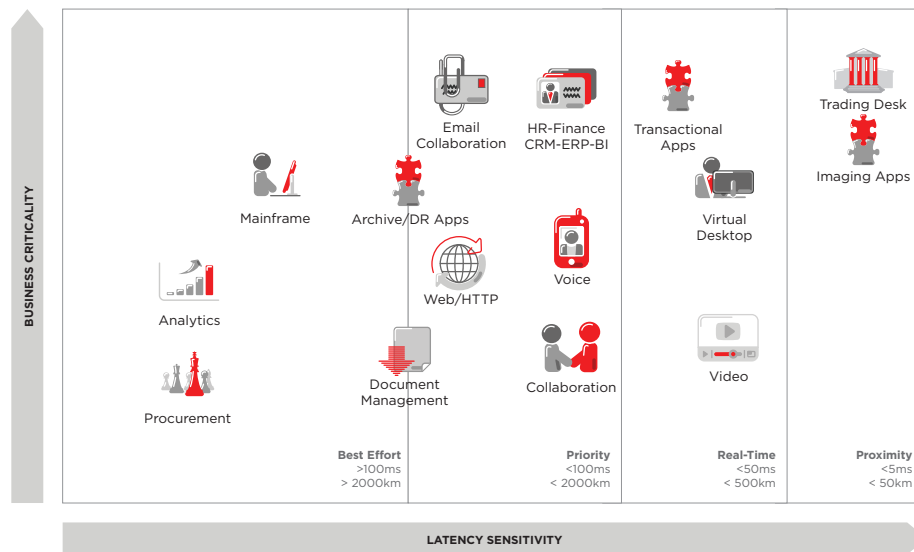


Figure 6: Relationship between Latency and Application Performance

The goal is to establish performance targets and latency limits for each major application. The perspective is that of the end-user, and the aim is to ensure a high quality of experience (QoE) for as many users as possible. In terms of the network, application targets should include:

- Target network latency
- Maximum acceptable latency
- Required average bandwidth
- Maximum bandwidth consumed

Once users and application requirements are well understood, enterprises are then in a position to implement cloud-based services.

Three Steps to the Cloud-enabled Enterprise

Step 1: Optimize the Network

Networks are the foundation of cloud computing. (Recall that “broad network access” is one of the essential characteristics from the NIST definition.) So the first step in moving to cloud enablement is to architect a network that is capable of delivering applications and services to their respective users—as defined in the previous section. After all, it doesn’t make sense to deploy network-based services until there is sufficient network capacity and quality to support them. Figure 7 shows an optimized application delivery network.

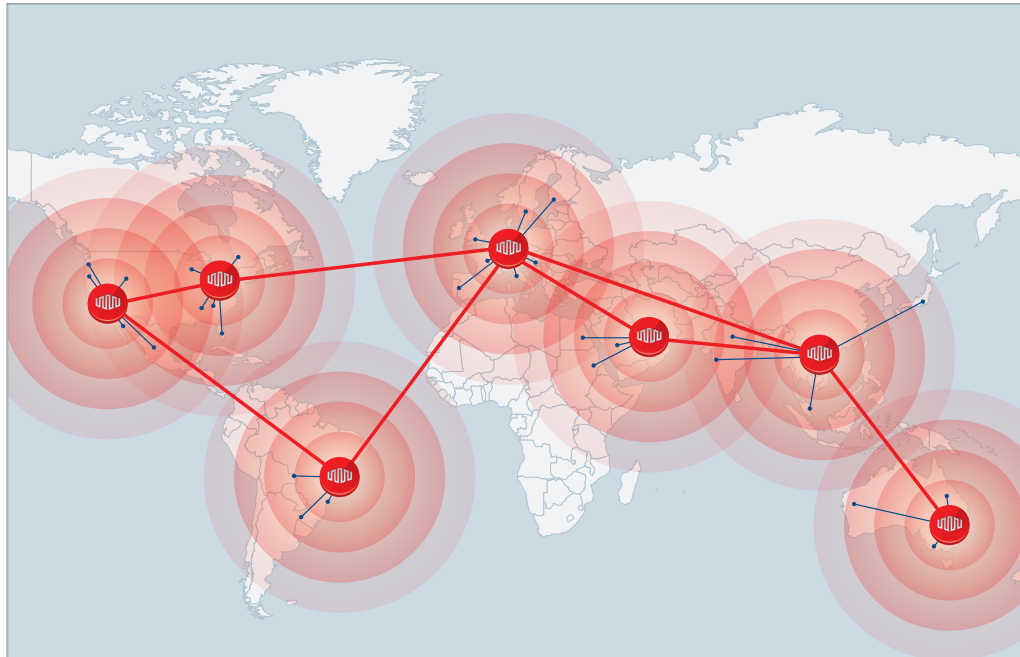


Figure 7: Network Optimization Using Network Performance Hubs

This is where the concept of network performance hubs (NPHs) makes sense. Simply put, a network performance hub (NPH) is a deployment of common off-the-shelf enterprise infrastructure (routers, aggregation switches, connectivity panels, equipment racks, etc.) to enable cost-efficient and high performance bundling, aggregation, and integration of various network and IT services.

NPHs are deployed in local proximity to specific end-user communities. Prime locations for NPH deployments include carrier hotels and commercial data centers (colocation providers) that host an expansive mix of service providers. These are often major interconnection and exchange points where the carriers’ core infrastructure nodes reside and where their primary network backbone routes are established.

For businesses, target services will likely include some or all of the following:

- Local access services (direct fiber, private line, DSL, Ethernet, etc.)
- Long haul and backbone network services (fiber, WDM, TDM, etc.)
- IP transport services (MPLS, IP-VPN, VPN, etc.)
- Internet services (IP transit, Internet access, peering, etc.)
- Voice and VoIP services
- Content distribution network services (CDN)
- Application delivery network services (ADN)
- Mobile platform gateways
- Voice, video and collaboration services
- Cloud and SaaS services
- Security perimeter/authentication services
- WAN optimization

By deploying a distributed grid of NPH nodes, enterprises establish a high-performance network backbone across which they can deploy and consume services. As illustrated above, each node will have a service radius that allows applications and services to be delivered within their QoE targets.

An optimized, distributed network has many advantages beyond simply performance. Many enterprises have existing wide area networks (WANs) that have grown over time in reaction to evolving needs, growth and acquisitions. They typically involve legacy technologies such as Frame Relay or MPLS. Newer technologies like carrier Ethernet give companies access to higher bandwidth, better performance and improved manageability—often at substantially reduced costs.

Network Optimization Strategy	Typical TCO Reduction
Elimination of local loop with NPH	15–30%
Using Tier 1 providers only in home markets	20–30%
Leveraging Tier 2 metro providers	35–60%

Table 1: Typical TCO Reductions Achieved Through Network Optimization

Step 2: Distribute Applications

With a high-performance network in place, enterprises are able to efficiently distribute their applications and data as required to meet performance goals. This may involve converging application infrastructure and deploying application performance nodes (APNs) that are connected to NPHs. Or they may choose to deploy enterprise private clouds in a number of key locations to serve data needs. Figure 8 illustrates applications distributed across NPHs.

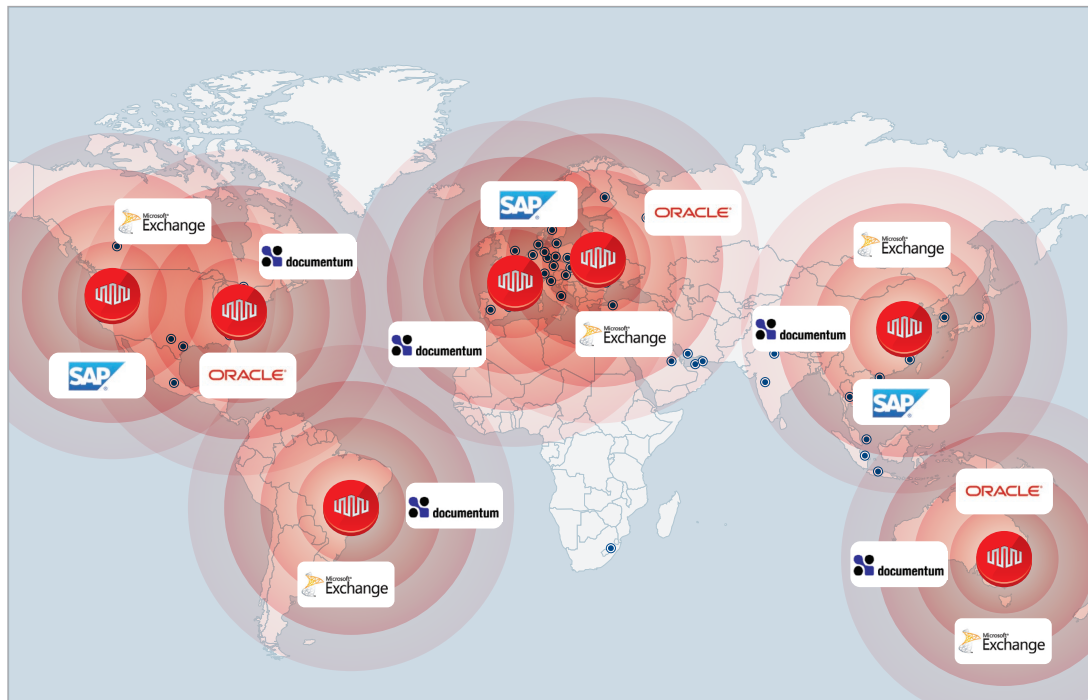


Figure 8: Using NPH Architecture to Distribute Applications

Typical benefits of moving to a distributed application environment include:

- 40 percent less variability in end-to-end transit time by making use of a NPH architecture⁴
- Up to 50 percent faster page load times on average, with substantially greater improvement for previously underserved user populations⁵
- 80 percent reduction in downtime by moving to higher quality networks and facilities⁴

In addition, by deploying network quality of service (QoS), enterprises can get maximum performance for business-critical and latency-sensitive applications.

4. "Optimizing Internet Application Performance," <http://info.equinix.com/rs/equinixinc/images/WP-Optimizing-App-Performance.pdf>

5. "More Bandwidth Doesn't Matter (much)," <https://docs.google.com/a/chromium.org/viewer?a=v&pid=sites&srcid=Y2hyb21pdW0ub3JnfGRldnxneDoxMzcyOWI1N2I4YzI3NzE2>

Step 3: Leverage Cloud Service Providers

Many companies wish to take immediate advantage of cloud service providers, but this should be the last step in a cloud enablement strategy. The move to cloud is premature until user needs are understood and foundational supporting infrastructure, like network performance hubs, is in place. Once the first two steps are completed, enterprises are fully cloud-enabled. They can then deploy private cloud and connect to public cloud services to deliver a high quality user experience, as shown in Figure 9.

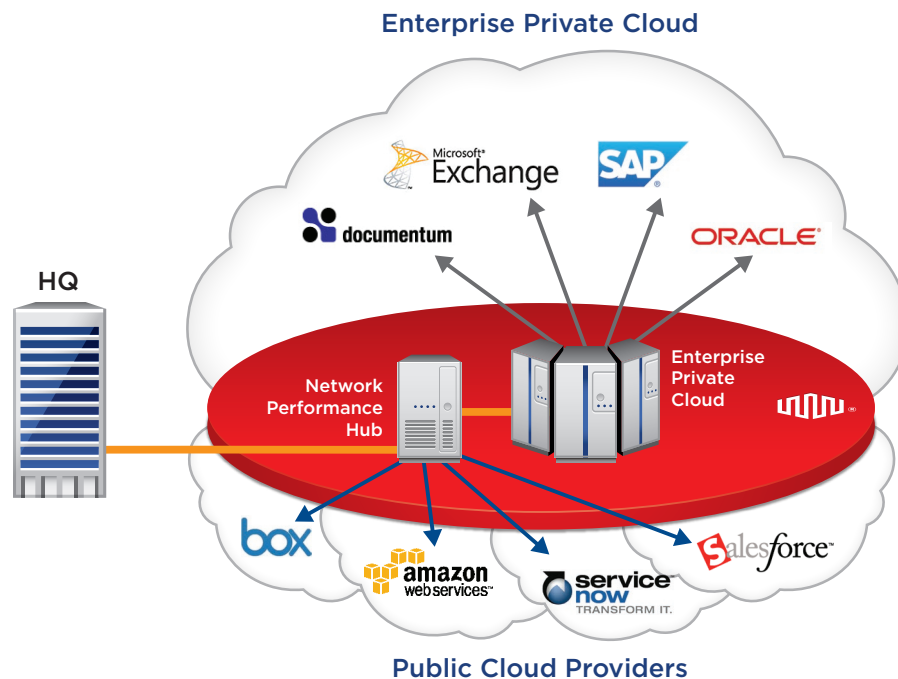


Figure 9: Augmenting IT Capacity Using Cloud Service Providers

The following are just a few examples of how enterprises can make effective use of cloud service providers:

- Deploying cloud-based backup to third-party storage services
- Augmenting compute and storage capacity for temporary or transient needs (own the base, rent the spike)
- Optimizing consumption of SaaS-based applications
- Implementing a distributed disaster recovery or business continuity architecture for critical applications and services
- Enabling rapid turn-up of development and test environments using public providers

In addition, depending on where the service providers are relative to the enterprise customer, it may be possible to take advantage of additional connectivity options such as a dedicated fiber connection to provide fast, secure, cost-effective interconnection.

Benefits of Cloud Enablement

Enterprises that move to a fully cloud-enabled service model benefit in many ways including:

- Improved user quality of experience globally
- Improved capability, manageability and performance of IT network and systems infrastructure
- Lower costs resulting from cost-effective, newer technologies as well as from capacity optimization
- New and enhanced service capabilities—without added IT infrastructure and complexity
- Reduced capital expenditures from switching to an operating expense (OpEx) consumption model
- Far greater choice of suppliers for a wide range of services and network options
- Increased business agility resulting from the ability to add new systems and services quickly

CONSIDERATIONS FOR CLOUD IMPLEMENTATION

Cloud computing clearly has many compelling features. At the same time, however, widespread use of cloud is still relatively new. This section provides guidance and best practices for companies interested in cloud enablement.

Understanding Networks

The first question companies must face when developing a cloud strategy is which services make the most sense for their particular needs. As pointed out in Step 1 of the cloud enablement section, having a well-designed network is foundational to the consumption of these services. So where should companies deploy their network? Some knowledge of telecommunications helps answer this question.

The Internet has become a ubiquitous part of most people's lives, but few people really understand how the Internet works⁶. A person might reasonably decompose the term into its components—"inter" meaning "between, among, jointly" and "net," short for network—and conclude the Internet is a network of networks. And they'd be correct. But exactly how do these networks network? That is, how does traffic from network A get to network B? The simple answer is "peering." In its most basic form, network peering⁷ occurs when two different networks agree to exchange traffic with each other. Without peering, users would need accounts with each and every service provider they wanted to use. So just where does this peering take place? For many reasons, network providers tend to peer their traffic at neutral Internet exchange points⁸ (IXPs). Figure 10 shows internet traffic patterns around the world.



Traffic Flows (Mbps)

5,000 2,500 1,000 100

Figure 10: Global Network Traffic Flows

6. Blum, Andrew, "Tubes: A Journey to the Center of the Internet," Ecco, 2012

7. "Peering," <http://en.wikipedia.org/wiki/Peering>

8. "Internet exchange point," <http://en.wikipedia.org/wiki/IXP>

Why does this matter? When the famous bank robber Willie Sutton was asked why he robbed banks, he is said to have replied, “Because that’s where the money is.” Similarly, for the best network performance one needs to go to where the networks are. That means connecting at IXPs, where network providers exchange traffic across their fastest and most reliable backbone circuits. And how does a non-telecommunications customer get access? By connecting within carrier-neutral data centers.

Carrier-neutral data centers sit at the crossroads of the Internet. They are the connection points for traffic producers and traffic consumers, where content, voice, video and data all meet. What better place could there be for a business to connect to employees, customers and partners? Rather than attempting to optimize their entire network in their own facilities using a single provider, carrier-neutral data center customers can choose the best provider for each region, location and type of traffic. And with the benefit of competition among network service providers, they avoid the lock-in that comes with single-network providers while usually getting substantially better rates.

In addition, many carrier-neutral data center providers offer a range of connectivity types. Customers may be able to connect directly with other customers and providers that reside in the same facility, bypassing network service providers and achieving faster, more secure and significantly cheaper transfer. Neutral hosts may also allow customers to access carrier Ethernet⁹, providing high-bandwidth, cost-effective layer 2 connectivity. They are also likely to offer a variety of layer 3 connectivity options along with a choice of capacity, quality and cost. Best of all, circuits can be provisioned quickly—often in days or hours rather than weeks or months.

So enterprises should seriously consider the benefits of coming to where the networks are. After all, traffic moves a lot faster over freeways than over back roads.

Cloud Hubs

Service providers have long understood the value of placing their services in locations of high network density. For cloud providers this has led to the formation of cloud hubs—places where many cloud services are deployed. This approach offers numerous advantages to service providers, including the ability to:

- Deliver services across many networks from fewer locations, thereby reducing service costs
- Connect with other, complementary services in the same location
- Connect with providers and consumers directly instead of through a network service provider
- Exchange (peer) traffic with other providers
- Deploy on the network providers’ highest quality backbone circuits

Enterprises can realize similar benefits by deploying at cloud hubs where they have access to a variety of cloud service providers—many only a cross-connect away. In such cases, the network performance hub also becomes a cloud access node through which enterprises can connect to and consume a wide variety of cloud-based services.

9. “Definition: Carrier Ethernet,” <http://searchtelecom.techtarget.com/definition/Carrier-Ethernet>

Deploying Globally

Enterprises that operate in many regions around the world have a difficult IT delivery challenge. They are usually expected to provide high-quality, highly reliable services while closely controlling costs. The more sites they support, the harder this becomes.

Moving to a cloud-enabled delivery model alleviates many of these challenges. By their very nature NPHs benefit from consistency. Global companies looking to deploy an NPH architecture are faced with a choice. They can either:

- A. Choose different hosting providers in different locations, or
- B. Find a single, global hosting provider for their entire deployment

From the standpoint of vendor and contracts management as well as operational consistency, option B offers clear advantages. For this reason, enterprises should consider working with a global partner when deploying network performance hubs and cloud access nodes.

PLATFORM EQUINIX™: A GLOBAL INTERCONNECTION PLATFORM

Platform Equinix™ refers to Equinix's global footprint of interconnected International Business Exchange™ (IBX) data centers comprising unmatched network density and the broadest cross-section of business ecosystems operating within our facilities. Many of the world's largest enterprises rely on Platform Equinix for mission-critical facilities and optimized application performance. Equinix provides some of the most reliable and secure data center facilities in the industry and has earned customer trust by delivering consistently excellent service.

Equinix's founding principle is to serve as a neutral location for network providers and their customers to meet and efficiently exchange traffic. Platform Equinix data centers are the most heavily networked interconnection points in the world and serve as the Internet's largest peering points. More than 1,000 network service providers leverage Platform Equinix to interconnect with each other and their customers around the globe.

Enterprise customers deploy their IT infrastructures on Platform Equinix to ensure operational performance, cost efficiency, reliability and security. Within Equinix data centers, they leverage the largest global ecosystem of service providers spanning networks and network services, the Internet, cloud and SaaS, mobile platforms and managed services. Through vast service provider choice, customers directly connect to their preferred providers to best leverage the most reliable, cost-efficient and best performing connectivity for delivering their critical applications and data to their end-users.

Equinix has built IBX data centers in 32 markets throughout the Americas, EMEA and Asia-Pacific to best serve its customers. Platform Equinix provides less than 10 milliseconds' latency to over 90 percent of the population of North America and Europe, as well as key population centers throughout Latin America and Asia-Pacific. This global platform allows enterprises to quickly and easily deploy their key applications closer to critical end-user communities, reducing fiber miles and ultimately improving overall application performance.

As the world's largest provider of carrier-neutral data center services, Equinix is in a unique position. It has the most data centers in the most locations and with the greatest concentration of network service providers of any carrier-neutral provider. Equinix also has the most Internet Exchange Point (IXP) members¹⁰. This combination makes the Platform Equinix the "**epicenter of connectivity**" for customers establishing network performance hubs to bundle, aggregate and integrate critical network services.

Today, Platform Equinix is home to more than 1,000 network service providers who rely on Equinix to connect to each other and to customers. CIOs can leverage Equinix to connect with multiple NSPs. Instead of being forced into a multi-year contract with a single NSP, they can establish a presence in an Equinix facility, interconnect like NSPs connect to each other and consume services from the most cost-effective and highest performing providers.

10. "List of Internet exchange points by size," http://en.wikipedia.org/wiki/List_of_Internet_exchange_points_by_size

Consuming services this way delivers:

- Rapid services connectivity, typically with same day/next day provisioning times
- Significant cost reduction through a competitive marketplace of service providers
- Consumption that is enabled through inexpensive cross-connects that are highly reliable, highly secure and ultra-low latency (essentially eliminating local loops)
- Improved reliability of connectivity to service providers
- Improved user experience by increasing connectivity throughput, reducing latency and packet loss and significantly improving the availability of services
- Greater local, regional and global performance through high-bandwidth backbone connectivity between NPHs
- Greater choice of service providers to enable diversity and to balance service cost, performance and quality
- More flexibility to evolve and scale the network over time

Each Equinix IBX data center is a vital and viral ecosystem where major networks, enterprises and business partners interconnect. Global enterprises, financial institutions, the largest networks and the Internet's foremost content companies all trust Equinix with their data, network and application assets. Equinix also offers industry-leading physical security, a global uptime SLA of 99.999% and proven expertise in configuring and supporting the most business-critical environments.



Equinix Global IBX Data Center Locations

SUMMARY AND RECOMMENDATIONS

The job of delivering IT services has long been a difficult one. The IT department is expected to provide secure, reliable and cost-effective services to users. Those same users, however, want easy access to the applications and information necessary to do their jobs. And they expect applications to perform well whether they're being accessed on a desktop computer at headquarters, on a tablet device at home, or on a smartphone in a cab halfway around the world. Traditional methods of IT delivery weren't designed to meet these demands.

Many enterprises are looking to cloud computing to enable access anytime, anywhere and on any device. The cloud-enabled enterprise can be more cost-effective, more agile and more capable. But successful migration to the cloud demands a thorough understanding of users and of the performance characteristics of their applications. For businesses with operations or users in multiple locations, it may also require migration to a more distributed application delivery architecture. With a strong network foundation in place, enterprises will be well-positioned to take advantage of cloud-based services to meet their IT challenges, both today and into the future.

Corporate HQ

Equinix, Inc.
One Lagoon Drive
4th Floor
Redwood City, CA 94065
USA

Main: +1.650.598.6000
Fax: +1.650.598.6900

Email: info@equinix.com

EMEA

Equinix (EMEA) BV
Luttenbergweg 4
1101 EC Amsterdam Zuidoost
Netherlands

Main: +31.20.753.7950
Fax: +31.20.753.7951

Email: info@eu.equinix.com

Asia-Pacific

Equinix Hong Kong Limited
Suite 6504-07,
65/F Central Plaza
18 Harbour Road
Wanchai, Hong Kong

Main: +852.2970.7788
Fax: +852.2511.3309

Email: info@ap.equinix.com

About Equinix

Equinix, Inc. (Nasdaq: EQIX), connects more than 4,500 companies directly to their customers and partners inside the world's most networked data centers. Today, enterprise, cloud, networking, digital media and financial services companies leverage the Equinix interconnection platform in 32 strategic markets across the Americas, EMEA and Asia-Pacific.

By connecting directly to their strategic partners and end users, customers are forming dynamic ecosystems inside Equinix. These interconnected ecosystems enable companies to optimize the performance of their content and applications and protect their vital digital assets.