

# Decoding the Future of Network Centric ICT

Keeping up with change



## Issue 4

- 2 Keeping up with change
- 4 Flexible Infrastructure Planning
- 7 From the Gartner Files:  
Cloud IaaS: Networking Options

Featuring research from

**Gartner**

## Keeping up with change

Constant change has established itself as the new norm in business today - powered by our voracious appetite for technology, the need to communicate with one another, and our hunger for innovation. Looking back to past generations, it is staggering to see just how far we have pushed the boundaries and speed of innovation. For example, it took 52 years for electricity to be adopted by 50% of American homes. Surprisingly, it took 22 years for television to reach 50 million users. In contrast, it only took three years for the Internet to reach the same number of people, 1 million iPads to be sold in 28 days from launch<sup>1</sup> and 15 million by the time Apple was ready to release iPad2 in March this year<sup>2</sup>. The pace of technology diffusion and uptake, as well as our ability to adapt our lifestyles at work and play to new devices and applications, has truly been remarkable.

We are living in an age of change. As businesses make their recovery from the recession, they are no longer interested in just keeping operations running smoothly. They are now more interested to have their CIOs and senior IT managers put in place appropriate IT solutions to support and enable growth. CIOs with a core competency in driving change and innovation are in hot demand.

In recent Gartner research, findings concluded that "To succeed in 2011, enterprises must navigate the challenges and pursue the opportunities presented by new realities in business and IT. CIOs, especially, will be pressured to quickly grasp and lead this emerging vision." Gartner goes on to recommend that "CIOs, first and foremost, extract themselves from the "old realities of IT" (and their focus on operational issues), and redirect their time, energy and ability toward five key actions".

These actions are as follows:

- Action 1: Look inside and outside your enterprise and industry to determine the nature and scope of change
- Action 2: Build an innovation engine

- Action 3: Shift resources toward growth and innovation
- Action 4: Rethink risk management
- Action 5: Increase attention on governance<sup>3</sup>

It would seem CIOs are already aligning themselves with these actions. If you asked a CIO what his biggest headache was today the likely responses would include:

- "How do I use IT as a value creator?"
- "How can I accelerate growth and innovation with strategic IT investments?"
- "How do I reconfigure my IT environment to handle the cyclical demands characteristic of today's volatile times?"

All these concerns are focused on the need to drive and effect change and are also aligned with objectives such as delivering greater IT flexibility, IT governance, improving risk management and strategic fit with the business.

It is therefore of little surprise the more CIOs we speak with today, as compared to the previous year, the more Cloud Computing services are being evaluated to meet their concerns and challenges head on. Whether the ultimate flexibility of such a computing model relates to the IT infrastructure, applications or running Software as-a-Service (SaaS), many CIOs aspire for their enterprise to achieve the promise of an increasingly agile environment in the long term. After all, analysts including Gartner<sup>4</sup> have predicted that would be the case – that Cloud Computing would be amongst the top priorities for CIOs in 2011.

With the theme of 'flexibility' featured in this issue of our newsletter, we outline a discussion of what it means to have a flexible infrastructure, how it can address many CIO concerns and considerations when planning for an agile

<sup>1</sup>"1 Million iPads Sold in First Month", Philip Lam, March 2010, iPadInsider Blog

<sup>2</sup>"2010, year of the iPad is being played", Steve Job video at launch of iPad 2, 2 March 2011, www.quickpwn.com

<sup>3</sup>Gartner Inc., "The CIO's Role in Making the New Realities Real", Jorge Lopez and Kathy Harris, 24 February 2011.

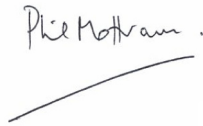
<sup>4</sup>Gartner press release, "Gartner Executive Programs Worldwide Survey of More Than 2,000 CIOs Identifies Cloud Computing as Top Technology Priority for CIOs in 2011", 21 January 2011. <http://www.gartner.com/it/page.jsp?id=1526414>

infrastructure. Often, there is nothing easier than to learn from someone who has done something before, so this newsletter issue also includes several Telstra customer examples. Furthermore, we are delighted to include some of the latest Gartner research providing insights into Cloud Computing, Infrastructure as a Service (IaaS) and Networking Options.

We hope you will gain some fresh perspectives and benefit from the insights shared in this newsletter. It is important to remember leading and managing change often requires flexibility to achieve and exceed desired outcomes. To do nothing or to move too slowly is to deny your business the opportunity to emerge as a market leader. And the pursuit of flexibility may present another set of challenges - but with no pain,

there is no gain. The winners in uncertain, volatile times will be those who have the ability to adapt and respond to change and who are able to move at the greatest possible speed.

Regards,

A handwritten signature in black ink that reads "Phil Mottram". The signature is written in a cursive style and is positioned above a long, thin, slightly curved horizontal line that serves as a separator.

Phil Mottram  
Executive Director  
Telstra International

## Flexible Infrastructure Planning

### Laying the foundation for a future-ready enterprise

Our ICT environment today is complex to say the least. The wired and wireless worlds are fast converging. New devices, applications and cloud-based services are coming into the market at a rapid pace. The complexity of ICT management is further compounded by the number of mobile devices being introduced into the enterprise environment, and the growth of the mobile workforce.

Driven by resource realities of limited or no increase in IT budget, many CIOs are working to transform and build flexible, scalable IT environments that can keep up with the pace of change. Flexibility is key for CIOs who wish to reap the benefits of an efficient enterprise, support new business processes, have better cost control, enhance productivity and business performance. A flexible infrastructure helps CIOs deliver on expectations to incorporate new technologies quicker and easier - from the use of social media to unified communications, mobile applications and cloud-enabled services – to providing pervasive support to users anywhere anytime and from any device, accelerate growth and enable business transformations.

A flexible enterprise infrastructure also allows CIOs to deliver on capabilities including:

- ability to run hybrid environments
- adjustable migration path
- parallel processing
- scalable availability to IT resources
- burstable network capacity
- ability to add on back-up resources when needed for business continuity and contingency
- simplified management of IT environment including mobile platforms.

### Flexibility begins at the infrastructure layer

In the journey to building a flexible infrastructure, CIOs need to be aware the infrastructure layers being architected and implemented need to inherently be flexible and tolerant of any failure. To ensure this is the case, CIOs need to begin

with the appropriate selection, configuration and management of ICT capabilities they need. Once this is done, CIOs then begin to evaluate how they wish to power their flexible enterprise infrastructure - whether by using emerging 'trendy' technologies such as cloud-based services or tried-and-tested managed services, virtualisation and consolidation services.

In order to ensure the flexibility meets enterprise-grade service-level expectations, CIOs should also not forget the role networks play, and their ability to drive up productivity and performance. Some of these network considerations include:

- Access to ultra-wide bandwidth and global links
- Real-time bandwidth allocation, billing and support
- Throughput and latency optimisation
- Availability and access to a mix of fixed, ubiquitous mobile networks and satellite services for a truly 'transparent' network to support today's global workforce
- Connectivity and networking options for optimal storage management and content distribution
- Data security given the number of 'smart' mobile devices plugged into the enterprise infrastructure.

### No pain, no gain

In our previous newsletters, we highlighted how cloud computing is network-centric in nature and it is an ICT deployment method for offering utility computing over the internet enabled by centralised, virtualised infrastructure in a private, public or hybrid environment. To recap, cloud computing is on a continuum of ICT deployment methodology, after out-tasked activities such as managed services, virtualisation initiatives and consolidation. As such, before businesses are ready to adopt cloud services, they would likely find themselves having to take evolutionary steps to enhance their infrastructural flexibility, rather than make a complete enterprise infrastructure overhaul.

While CIOs look to make a strategic impact on the business with cloud-based services, they are faced with the need to balance the risks involved with a generally 'safe' business culture. CIOs may be tasked to drive innovation and facilitate business transformations, but as they work toward setting up an infrastructure that is truly dynamic and responsive to volatile demands and changing business needs, any unpleasant surprises or business disruptions may well cost them their jobs. As such, when planning for a flexible infrastructure, CIOs need to consider not just the layers of the technology stack but also their IT processes, internal structure, how contracts are brokered with external service providers, as well as the business results they wish to affect.

As CIOs look at Infrastructure-as-a-service (IaaS) as an avenue for flexibility, they need to consider the requirements on their service providers for the delivery of a robust, and highly available IaaS layer.

Google offers a cloud-based email management service popular with consumers and some businesses. The idea of being able to plug into a highly available, reliable and secure email service with a predictable monthly fee is an attractive proposition. Imagine being able to dedicate limited internal IT headcount to more strategic projects, reduce your hardware and software maintenance issues, reduce user support issues, and ensure a predictable ICT cost, to name a few of the possibilities.

Unfortunately, in late February, Google's Gmail service suffered a glitch that temporarily wiped out email messages for some 150,000 users.<sup>1</sup> Though Google eventually managed to restore the messages, the outage created widespread angst among its users, many of whom depended on the service for both work and play.

Then in April, a separate glitch in Amazon Web Services' data centres brought down popular online services like Reddit, Quora, Hootsuite and Foursquare. Though less severe than Google's outage, the incident once again left thousands of users disconnected from cloud-based services.

Coming just months apart from each other, these two high-profile incidents raised questions about the reliability of cloud-based solutions, and just how bullet-proof they really are. As an outcome of building a flexible infrastructure, questions

have surfaced surrounding the wisdom of moving so fast and far from the traditional enterprise infrastructure model of owning and managing your own ICT capabilities.

### **Safeguard opportunity cost with service assurances and SLAs**

It is important to acknowledge the chances of a single data centre experiencing problems within the traditional enterprise infrastructure environment or a bank of servers malfunctioning is much higher than it is for a cloud IaaS provider. Similarly, it is easier to scale up or down with volatile market demands by working with a cloud IaaS provider than to try to do it with on-premise owned resources.

According to Gartner, "the traditional management of IT created a highly inefficient use of resources, for all the right reasons". Statistics by Gartner show that "on average, 40% of available rack space in data centres today is under-provisioned or unused. The same report also states that "On average, data centres provide 60% more power than they use, when considering effective compute capacity".<sup>2</sup>

The journey to building a flexible infrastructure is therefore one that requires careful planning and deployment in order to ensure minimum downtime for mission-critical services, and availability of contingency resources.

More importantly, CIOs need to ensure they have robust service assurances from their providers and Service Level Agreements (SLAs) in place as they piece together their path to a flexible enterprise infrastructure.

According to Gartner, "SLAs for cloud IaaS are structured quite similarly to SLA agreements for network services, hosting and data centre outsourcing. SLA metrics are typically calculated on a monthly basis. If an SLA is violated, you are entitled to a credit for part of your service fee. The amount of the credit is usually based on the length of time for which the SLA was violated, multiplied by the cost of the affected service. Longer single violations, longer cumulative violations or repeated violations may carry disproportionately higher penalties."

Gartner also recommends businesses consider two levels of SLAs - Cumulative Availability and Outage Length SLAs.

<sup>1</sup>"Google Gmail Outage Leaves 150,000 Users Without E-mail", Sharon Gaudin, PC World, <http://bit.ly/ig43F1>

<sup>2</sup>Gartner inc., "Shrinking Data Centers: Your Next Data Center Will Be Smaller Than You Think", David J. Cappuccio, 4 Mar 2011

1. **Cumulative availability.** This is the availability over the course of the entire month. This SLA is usually 99.5% or higher, and many cloud IaaS providers offer 99.99% or higher. We consider reasonable cumulative availability SLAs to be:

- 99.5% for single-data-centre, non-high-availability, non-production infrastructure.
- 99.75% for single-data-centre, high-availability, production infrastructure.
- 99.95% for multiple-data-centre, high-availability, production infrastructure.

2. **Outage length.** This SLA provides compensation for each outage exceeding a certain length. The typical minimum outage length for compensation is 10 minutes or more.<sup>3</sup>

### Network performance imperative

CIOs may be well-versed on the need to achieve greater strategic alignment, governance, enhanced performance and reduced operational risk. But it is easy to overlook the fact even the best-designed IT infrastructure is reliant on low network latency, resilience, and reliability in order to deliver fast, agile cloud services to their customers.

Consider how many universities today are tapping new customer base with ‘virtual classroom’ experiences or lessons over video-conferencing facilities. This has helped schools boost their revenues from overseas students who would otherwise find it too costly to physically relocate for their studies. Consider also healthcare institutions tapping cloud networks to collaborate on medical care and patient database. Medical specialists have gone even further to share their expertise on complicated surgical procedures with other doctors via video conferencing applications. In each of these scenarios, a high-performance network link, matched by network optimisation, have come to the fore.

Closer to home, Telstra’s network-centric IaaS is gaining traction with customers seeking to deploy private clouds in order to deliver seamless services to their customers in multiple locations around the world. Some of these customers include Visy (utility SAP on private cloud) and Komatsu in Australia (utility IT services on private cloud which also includes on-demand SAP services).

In Asia, Telstra is currently working with a number of customers to provide managed services via centralised IT infrastructure. We will be ready to help these customers transform from managed services to cloud services when they are ready. These customers include Cerebos (Supply Chain Management services on Telstra’s managed MPLS network), Savills (Managed MPLS network services) and a leading international bank (Managed global WAN services on Telstra’s MPLS cloud).

### Infrastructure planning at an inflection point

In summary, the journey to building a flexible infrastructure is a ‘work-in-progress’ project for CIOs. There has never been a more challenging time for CIOs making infrastructure planning relevant to today’s needs and ready for tomorrow’s priorities. Be it hardware availability, software applications or network capacity, there has never been a more difficult time to anticipate just how much technology is required to meet the enterprise’s demands.

As businesses prepare for the new reality of “doing more with less”, CIOs who focus their efforts on building a flexible infrastructure are headed in the right direction for success. A flexible infrastructure will allow CIOs to stay focused on making productivity gains and growth, while addressing the dichotomous nature of their current challenges. For instance, a flexible infrastructure supported by pay-per-use utility services helps CIOs maintain better cost control when faced with flat IT budget and the need for continuous software and hardware upgrades. Flexible infrastructure supported by out-tasked managed services will help CIOs enhance performance and support business transformation while reducing operational risk with service assurances and SLAs. A flexible infrastructure will also allow CIOs to support the explosion in bandwidth demand brought on by the influx of mobile and personal devices into the corporate environment as a result of having scalable network capacity arrangements with its service providers. And the list goes on.

<sup>3</sup>Gartner Inc., “Cloud IaaS: Service Level Agreements”, Lydia Leong, 7 March 2011

From the Gartner Files:

## Cloud IaaS: Networking Options

The network is a crucial part of cloud infrastructure as a service (IaaS). Not only does a network connect customers to the cloud, but it connects the different components of the cloud environment. IT managers must understand their networking options when selecting an IaaS solution.

### Key Findings

- Most IaaS providers offer private networking options, not just general Internet connectivity or an Internet VPN.
- Most IaaS providers include network security functions as part of their basic offering.
- Most IaaS providers can support application delivery controller and WAN optimization controller functionality, but these are usually add-on options, not part of the basic offering.

### Recommendations

- Choose an IaaS provider that offers a wide range of connectivity options for better flexibility and suitability for a broader range of use cases.
- Obtain SLAs not only for external network connectivity, but for the LAN between compute and storage elements.

### ANALYSIS

As described in “Evaluating Cloud Infrastructure as a Service,” all cloud IaaS offerings are not created equal, despite superficial similarities in the way the offerings are described. There is considerable variance in service provider design goals, the quality of technical implementations, and the cost-effectiveness and value for money of those implementations.

This is one in a series of reports detailing the differences in the technical architectures and business models of IaaS offerings. This document is focused on networking options associated with cloud compute IaaS.

### NETWORK FUNCTIONS

There is a variety of important network functions in IaaS solutions. These functions include:

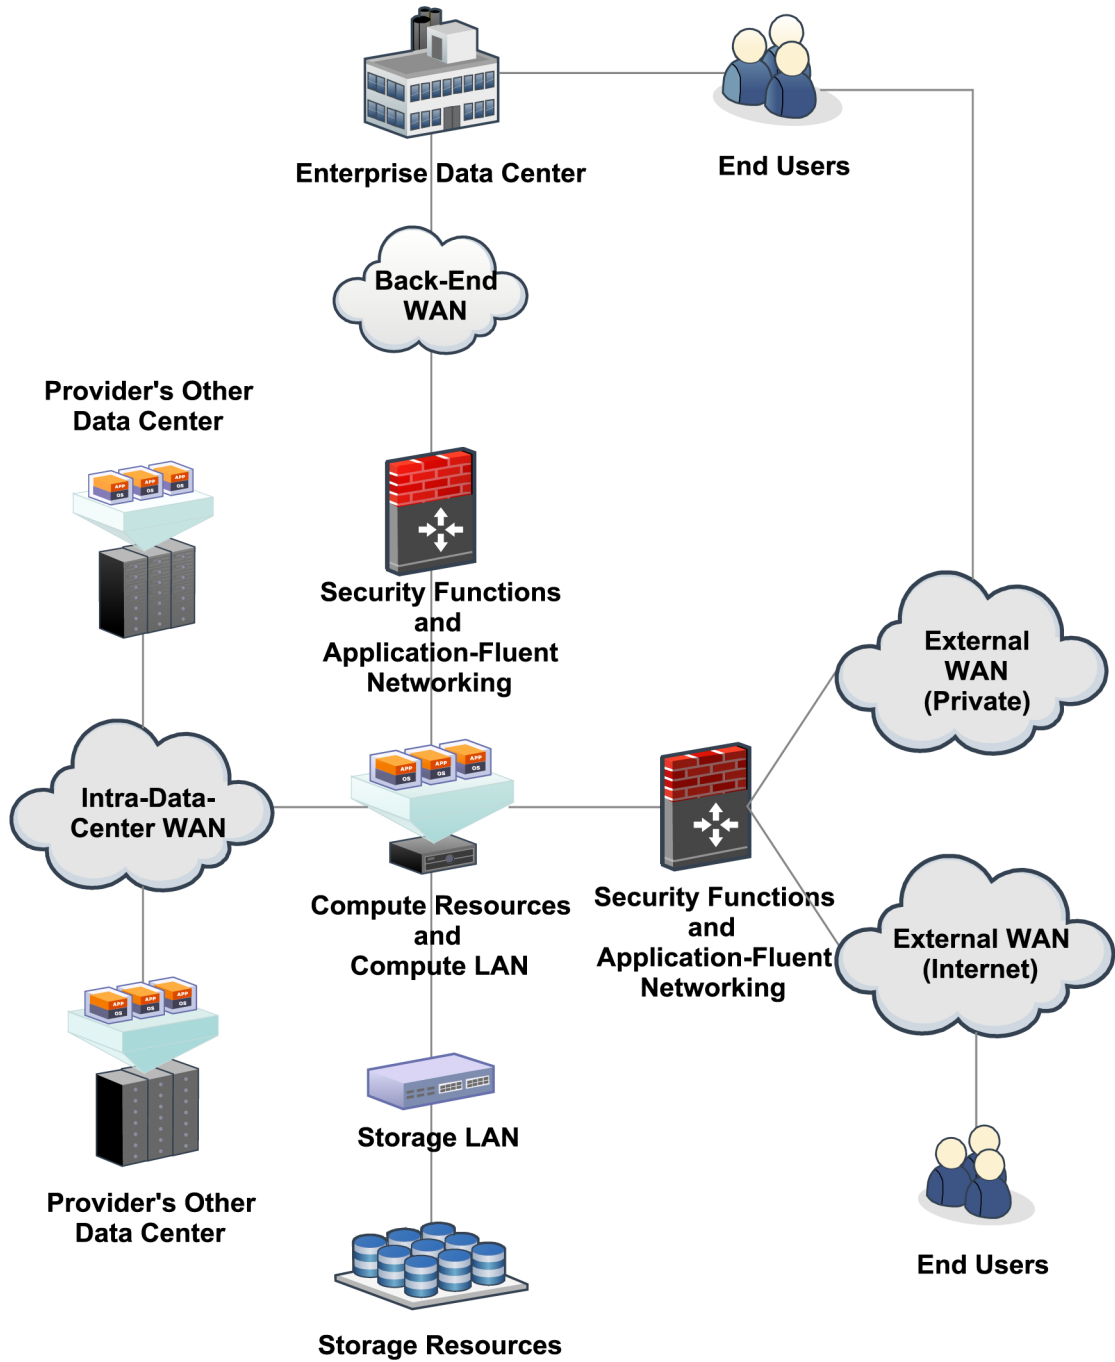
- **External WAN.** Connectivity between the cloud environment and the end user of content and applications.
- **Back-end WAN.** Connectivity between the cloud environment and the customer’s internal data center.
- **Inter-data-center WAN.** Connectivity between multiple data centers and cloud environments owned by the same service provider.
- **Intra-data-center LAN.** Connectivity between the cloud environment and other environments within the service provider’s data center.
- **Compute LAN.** Connectivity between compute resources within the same data center.
- **Storage LAN.** Connectivity between compute and storage resources within the same data center.
- **Application-fluent** networking functions, such as load balancing.
- **Security functions,** such as firewalling.

We will examine the common options for each of these functions, as delivered by IaaS providers. Figure 1 shows a conceptual networking diagram for cloud compute IaaS.

### External WAN and Back-End WAN

The Internet is typically used to connect customers to IaaS solutions. IaaS providers typically offer Internet bandwidth themselves; carriers offer their own networks, while most other IaaS providers offer some blend of multiple ISPs. In either case, the quality of the provider’s peering will matter for Internet performance; superior interconnectivity will result in higher overall availability and performance for end users.

**FIGURE 1** Conceptual Networking Diagram for Cloud Compute Infrastructure as a Service



Source: Gartner (March 2011)

Internet bandwidth is usually priced in one of two ways: like typical Internet connections, in megabits per second (Mbps) at the 95th percentile; or by the gigabytes (GB) transferred in and out of the cloud environment.

There is usually an SLA for Internet bandwidth that covers availability, packet loss and jitter. The SLA might also cover Internet performance to specific points of presence, usually measured in milliseconds of latency. These SLAs are comparable to the SLAs that carriers typically offer for business Internet connectivity.

Many customers want to be able to interconnect their cloud environment with their internal data centers using private connectivity of some sort. Some want to isolate their cloud environment from the general Internet, making their compute and storage resources accessible only via their private network; others simply want to add back-end connectivity for access to other corporate resources. Furthermore, customers may want to treat their cloud environment as an extension of their own data center, allowing cloud resources to receive Internet Protocol (IP) addresses from their own internal IP address pools, and potentially allowing them to be managed with the same tools.

Most IaaS providers offer Internet VPN services, such as IPsec tunneling or a Secure Sockets Layer (SSL) VPN. There is sometimes a surcharge for these services. There are purely software-based solutions as well, from providers such as CohesiveFT and CloudShare, which can be deployed whether or not the provider offers an Internet VPN as a service.

Most IaaS providers offer the ability to terminate private WAN connections in their data center and cloud environment, allowing customers to purchase traditional WAN services, such as a Multiprotocol Label Switching (MPLS) VPN and Ethernet, from the carrier of their choice. Even carriers that provide IaaS will usually allow customers to extend their private WAN services from other carriers into the cloud environment, so being a committed network customer of

one carrier does not preclude sourcing a cloud environment from another carrier. Customers will typically purchase this network service directly from their carrier and work with the IaaS provider to terminate the circuit.

Some IaaS providers can, when private connectivity or an IPsec tunnel is used, use a customer-provided IP address range for the virtual machines. This bridging functionality is necessary for such connectivity to be truly useful for spanning the enterprise's own data center and the service provider's cloud.

### **Inter-Data-Center WAN**

Many IaaS providers offer their cloud services in multiple data centers. Customers with global needs, or that need the highest levels of availability, are likely to deploy infrastructure in multiple locations. Therefore, the type and quality of connectivity between data centers matters to them.

Although some IaaS providers rely on the Internet to connect their data centers, many have their own private network between their data centers. They might own this network or obtain it from a network service provider. Customers should care about the service levels they obtain, but should not need to care whether the IaaS provider owns or merely operates the network.

Some providers will charge for bandwidth between their own data centers but others will not.

### **Intra-Data-Center LAN**

Many IaaS providers have multiple types of environments in their data centers. For instance, an IaaS provider might offer colocation, dedicated hosting and cloud IaaS within the same data center. It's common for customers to buy multiple offerings and to want to be able to combine them. For instance, a customer might have front-end Web and application servers on a shared IaaS fabric, database servers in the dedicated hosting environment and a credit card encryption "black box" in colocation; together, these components would form the entirety of an e-commerce site.

Most IaaS providers with multiple environments in the same data center have a way for customers to cross-connect between these environments and allow them to use a single IP address range and virtual LAN (VLAN) architecture across all these environments. This is a critical requirement for the many customers that have needs that cannot be placed on cloud virtual machines due to factors such as performance requirements, software licensing restrictions or functionality that comes only in the form of a physical appliance (as can be true for encryption devices used for e-commerce, for instance).

There may be a charge for the cross-connects, but there is normally no charge for the bandwidth between these environments.

### Compute and Storage LANs

Within the data center, a LAN connects compute resources within the cloud environment, and a LAN also connects compute resources to storage resources. Most IaaS providers do not run converged data center LANs, so these are usually two separate LANs.

The compute LAN is normally some form of Ethernet and usually has plenty of bandwidth. However, the network between storage and compute could be virtually anything; 10-Gigabit Ethernet, Fibre Channel and Internet SCSI (iSCSI) are common, and all have different levels of performance. As this back-end storage LAN is normally shared, there may be contention – in the case of some providers, sometimes severe contention – for storage bandwidth.

Providers typically do not charge for LAN bandwidth. They may offer an internal network SLA that covers availability, latency, packet loss and jitter between devices on their LAN.

### Application-Fluent Networking

IaaS providers usually provide customers with a way to gain access to WAN optimization controller (WOC) and application delivery controller (ADC) functionality.

ADC functionality is usually divided into two tiers. The base tier offers basic load-balancing functions. The higher tier provides the full suite of typical ADC functionality. This results in one or more of the following options being available:

- Basic load balancing as an on-demand, elastic service. This is managed by the service provider and integrated into the self-service portal; the customer does not need to know or care what the back-end implementation is, although it is most commonly shared appliances from a vendor such as F5. There may be a surcharge associated with this service.
- Full ADC functionality as a provider-managed, dedicated device in the service provider's hosting environment. There is usually a monthly fee, but the price may be usage-based.
- Full ADC functionality as a customer-managed, dedicated device in the service provider's colocation environment. The customer simply buys the device and pays for the colocation.
- Full ADC functionality as a virtual appliance in the cloud environment. It might be either provider-managed or customer-managed. There is usually a surcharge for this. Some virtual appliance vendors, such as Zeus, have arranged deals with IaaS providers that allow customers to purchase this functionality in an on-demand manner.

WOC functionality, which is important to many customers for whom reducing WAN latency is a crucial factor in application performance, is usually offered in one of the following ways:

- As a provider-managed, dedicated device, such as a Riverbed Steelhead appliance, in the service provider's hosting environment.
- As a customer-managed, dedicated device in the service provider's colocation environment.
- As a virtual appliance in the cloud environment. It might be either provider-managed or customer-managed.

As a result of this approach, WOC functionality is generally not available in an on-demand, elastic manner. Providers usually charge for the appliance itself and, if they are providing management of the appliance, for the managed service.

### Network Security

Most IaaS providers have automated mitigation of distributed denial-of-service attacks; this is normally included with the service, without a surcharge. Furthermore, nearly all IaaS providers have security measures within their internal LANs and compute environments to protect against, for example, network sniffing, spoofing and denial-of-service attacks. Finally, many IaaS providers try to offer a certain degree of virtual network isolation by using individual VLANs, virtual routers and virtual switches for each customer.

Most IaaS providers offer basic firewalling services to customers in the form of access control lists for IP addresses and ports. These are normally included with the service, without a surcharge. Most IaaS providers can offer additional intrusion detection and prevention functionality (intrusion detection system [IDS]/intrusion prevention system [IPS]); this may carry a surcharge that is nonelastic, and might require a dedicated device or virtual appliance.

### SOURCING CONSIDERATIONS

When considering your networking needs in a cloud IaaS environment, ask yourself:

- Do I want my application to reside on the Internet, on my private WAN, or both?
- Do I need private WAN connectivity, or will an Internet VPN suffice?
- Will all my infrastructure for a given application reside within the cloud environment, or will I need to connect it to other infrastructure in the service provider's hosting or colocation environment?
- What are my application-fluent networking needs?

Source: Gartner RAS Core Research Note G00210093,  
Lydia Leong, 7 March 2011

