

Opportunities for CSPs in Enterprise-Grade Public Cloud Computing

Differentiating the public cloud offer: from commodity to enterprise-grade services

Reference Code: OT00142-004

Publication Date: May 2012

Author: Peter Hall

SUMMARY

In a nutshell

Leading communications service providers (CSPs) have successfully found an important niche in public enterprise-grade cloud computing services, one that builds on their strengths in enterprise-grade networking and managed services provision. This niche consists of offering a middle ground between traditional public cloud computing delivered over the public Internet and customized and more expensive private cloud computing services. This report explains why the focus on enterprise-grade public cloud computing services is important to CSPs, and discusses a possible evolutionary path for managed enterprise-grade cloud services that continues to build on CSPs' strengths.

Ovum view

Cloud computing is becoming a new growth area for CSPs. Many have been monitoring the IT services market for many years, hoping to find new growth markets outside telecoms, but only the largest global and regional CSPs have succeeded in generating significant revenues from traditional IT services. Systems integrators (SIs) and specialist IT services players have not given up their markets easily, and for many CSPs the IT services market has proved difficult to penetrate. Public cloud computing has the potential to change this, as it provides highly productized solutions with automated provisioning and limited requirements for customization.

However, cloud computing is a highly competitive market with powerful players from both within and outside the traditional IT services market. Amazon Web Services (AWS) has grown to become the dominant player in public cloud computing, and presents formidable competition for new players coming to market without a background in IT hosting and data center services. However, many CSPs have already proved that they can find an important niche in the market by building upon their core strengths in enterprise-grade networking and delivering services to demanding SLAs.

Public cloud computing services have grown using the public Internet for delivery. In fact, many businesses consider the public Internet as an essential component of the cloud computing model. However, this has largely been alongside the growth of cloud computing for software-as-a-service (SaaS) applications and less demanding uses of infrastructure-as-a-service (IaaS) such as application testing and development. As enterprises become more comfortable with the cloud computing model, and as an increasing number of applications require high performance levels, end-to-end service assurance, and SLAs, the role of the network becomes much more important. Enterprise-grade networking providing enhanced security and resilience will likely replace the public Internet for an increasing number of enterprise cloud computing requirements. CSPs are in a strong position to offer such public cloud computing services and gain a competitive advantage over other players. CSPs also have the opportunity to lead the continued evolution of enterprise-grade public cloud computing by using technologies they currently employ in enterprise network services, such as applications performance management (APM), to maintain and build upon this competitive advantage.

Key messages

- The provision of enterprise cloud computing services is an important, high-growth market for CSPs faced with declining revenues from telecoms services. While most of the larger enterprise CSPs have ventured into some aspects of IT services provision over recent years, cloud computing broadens this opportunity to include most CSPs that address the business market. The provision of cloud computing services is a natural extension of the CSP's role in enterprise networking.
- An important opportunity for CSPs is the provision of enterprise-grade public cloud computing, which builds on their strengths in enterprise virtual private network (VPN) services and allows them to differentiate their offering from commodity cloud services.
- Enterprise-grade public cloud services are able to deliver considerably higher annual ARPUs (measured as average annualized revenue per customer) than commodity services, and with longer-term contracts, deeper customer relationships, and pull-through revenues from associated services such as consulting.

- A key requirement for enterprise-grade public cloud computing is the provision of SLAs that assure end-to-end service availability. This assurance cannot be provided with services that are delivered over the public Internet, which provides the connectivity for most public cloud services on offer today.
- While customers of most cloud computing providers can get end-to-end availability by renting a private circuit rather than using the public Internet, the responsibility for end-to-end availability is then usually split across two parties, with separate reporting points for problem management. A CSP can provide a single reporting point (and in some cases a single SLA) for the end-to end solution.
- Enterprises' greatest concerns around adopting cloud computing continue to be security, data privacy, and compliance issues. These concerns must be addressed by CSPs developing an enterprise-grade cloud computing service offer.
- We expect to see a continued evolution of enterprise-grade public cloud services that will encompass new features such as application performance reporting and management. Many CSPs have developed these capabilities to support their managed network services, so this evolution will continue to draw on their strengths, and we expect leading members of the CSP community to have an important role in driving this evolution.
- Intelligent network customer premises equipment (CPE) will also have a role in enabling new cloud computing capabilities such as enhanced security, application performance management (APM), and disaster recovery functions.



THE IMPORTANCE OF CLOUD COMPUTING SERVICES TO CSPS

Telecoms is a declining market

The traditional telecoms market is declining in most developed countries. Fixed telecoms is under pressure from declining voice revenues, while mobile telecoms, which has been an important growth market in most countries, is now seeing very low revenue growth in some developed markets, including Western Europe. This overall picture is in spite of some individual areas of strong growth, such as mobile data, consumer broadband, and enterprise Ethernet services. This means that CSPs need to look outside telecoms for growth in their addressable markets.

New revenue streams through IT-centric services

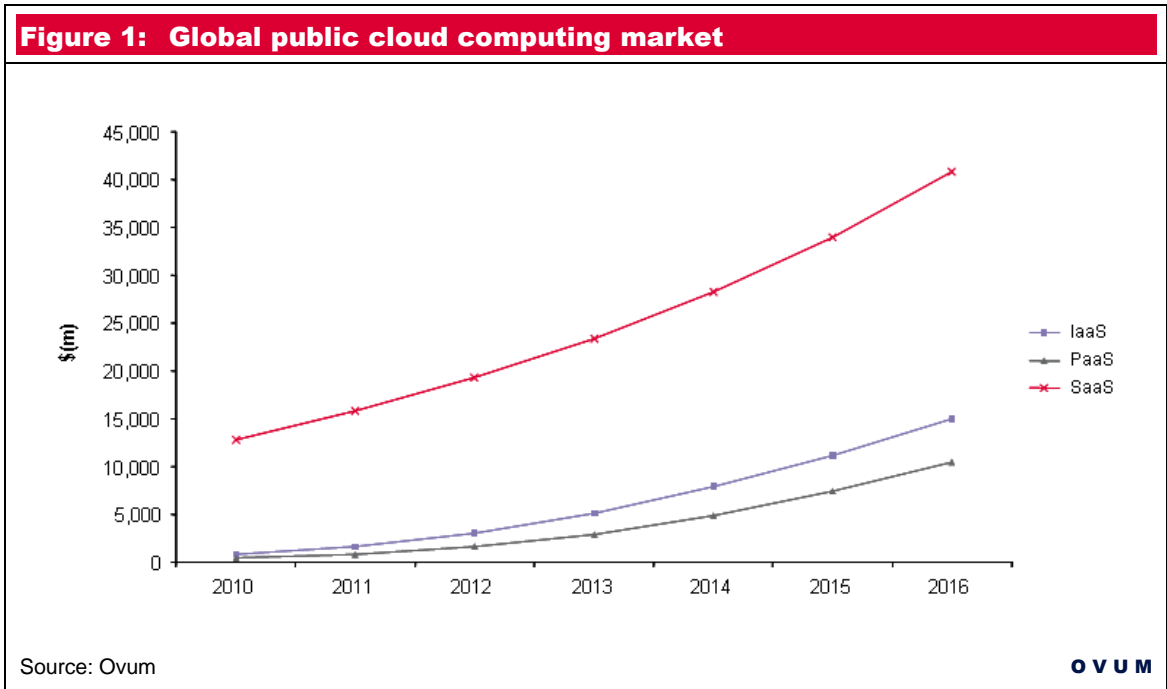
Leading global and regional CSPs realized the importance of IT services as a new growth market over a decade ago. Players such as Deutsche Telekom and BT developed an early position in the IT services markets through a combination of acquisition and organic growth. Many national CSPs have since adopted a similar strategy, though often on a more limited scale, with a range of managed IT services spanning security, hosting, and datacenter solutions. Cloud computing and related cloud services such as cloud-based unified communications and collaboration now broaden the opportunity for CSPs of all shapes and sizes to expand beyond communications markets in search of new growth opportunities.

Building upon core strengths in a high-growth market

Ovum forecasts that the global public enterprise cloud computing market will grow from \$18bn in 2011 to \$66bn in 2016, a CAGR of 29%. As Figure 1 shows, SaaS is the most mature component of the market with a forecast CAGR of 21%, followed by IaaS and PaaS at 57% and 66% respectively. This growth stands in contrast with the global market for fixed and mobile telecoms, which we forecast to be close to flat over the same period. While the cloud computing market is highly competitive, there is an opportunity for CSPs to succeed by focusing on their core strengths in networking, managed services, and delivering enterprise-grade services to high performance levels.

The cloud computing market is still in its infancy in terms of its overall evolution, with most enterprise interest and spend being in SaaS and much of the interest in IaaS being at the level of testing and development and proof of concept rather than in support of enterprise applications. However, this is changing, and we are seeing increasing adoption of cloud computing for mainstream enterprise applications. The opportunity exists for the CSP sector to become a key provider of enterprise-grade public cloud computing, drawing upon the skills developed over the

last decade in delivering enterprise-grade WAN services at a national, regional, and global level. Increasingly the network is being seen as a critical component in the next phase of evolution in public cloud computing, where high resilience and low latency will be critical factors.





THE CSP CASE FOR ENTERPRISE-GRADE CLOUD SERVICES

Differentiating from "commodity" cloud players and building on enterprise networking strengths

The majority of the public cloud computing services on offer today are provided over the public Internet. Regardless of the resilience of the service providers' cloud platforms, they therefore cannot be regarded as enterprise-grade, as service levels cannot be assured over the public Internet. Enterprise CSPs have the opportunity to leverage their enterprise-grade WAN services, such as MPLS-based VPNs and Ethernet VPNs, to provide cloud computing services with enhanced end-to-end security and resilience. Cloud computing services delivered over enterprise VPNs are still public cloud computing services, but are sometimes referred to as "virtual private cloud services" to distinguish them from standard public cloud services. In addition to enterprise-grade WAN connectivity they have other important features such as enhanced security through the use of VLANs and support for private IP addressing, which means that from a security perspective the cloud platform is inside the enterprise firewall.

The term virtual private cloud service has no standard industry definition, but is used to reflect the additional security and resilience features for all components of an end-to-end solution, including WAN, LAN, and computing/storage platforms. It differs from a private cloud solution in that these components are still virtualized, so it provides most of the cost benefits of public cloud computing. Additional capabilities may also be offered that are typically not found in public cloud computing, such as the automated provisioning of physical servers in addition to virtual servers. These physical servers support applications that cannot deliver optimum performance through virtualized compute platforms.

Enhanced wide area network security and resilience is proving attractive for many large enterprise deployments. Leading CSPs such as Verizon and Orange claim that over 60% of their IaaS solutions are delivered over their MPLS VPN networks, and that this proportion is increasing.

We are not trying to label individual cloud computing providers as "commodity" or "enterprise-grade"; most providers today have some capabilities that fall into both categories. Our emphasis is on describing those capabilities that could describe a service (or a particular configuration of a service) as enterprise-grade, and how this resonates with the core strengths of CSPs.

Building deeper relationships with enterprise customers

The role of the wide area network in delivering IT applications to the enterprise has grown with both the diversity of applications that need to be networked and the globalization of enterprises.

CSPs are increasingly seen by enterprise IT departments as more than just providers of the "pipes"; they have an essential role to play in delivering networked IT to high performance levels. The provision of cloud computing services deepens the relationship further, and provides an evolutionary path to a future where CSPs can take responsibility for the end-to-end performance of networked business applications. This path embraces the provision of the network; application performance optimization technologies; and computing, storage, and security platforms as part of an integrated service with end-to-end application performance SLAs.

Significantly increased ARPUs and long-term contracts

The experiences of global CSPs that are now in their second or third year of public IaaS offers indicate that a focus on customers with enterprise-grade requirements gives ARPUs as high as 10 or even 20 times those of typical commodity service requirements that are mainly around test and development. This is even before any pull-through from additional services is taken into account. Many of these enterprise-grade requirements also develop into long-term commitments where customers contract for service for one or more years rather than use a service on demand. Of course, some of these customers will start with a test or proof of concept, so the service model needs to be sufficiently flexible to also cater for this.

Generating pull-through for other services

The pull-through revenues from enterprise-grade cloud computing deployments can be significant, and the experience of CSPs with a mature public IaaS offer is that pull-through revenues of up to 50% of core cloud service revenues can be achieved. WAN services, security services, and consulting are examples of key areas of pull-through revenues. Consulting services typically include cloud readiness assessments (i.e. helping an enterprise identify applications that can be migrated to a cloud model) and security-related assessments such as evaluating the impact of data privacy and compliance requirements on cloud migration.

KEY REQUIREMENTS FOR A ENTERPRISE-GRADE CLOUD COMPUTING SERVICE

The provision of end-to-end service assurance

While there is no general definition for "enterprise-grade", a minimum requirement should be that the customer has a measure of the availability of the cloud computing service from an end-to-end perspective. Cloud computing availability SLAs usually refer to the cloud computing provider's platform alone, and exclude the impact of the WAN. Where the WAN is the public Internet this is no surprise, as no cloud computing provider can guarantee its performance. SLAs offered today by IaaS providers typically specify an uptime for the entire data center infrastructure up to and including the operating system running in the virtual machines.

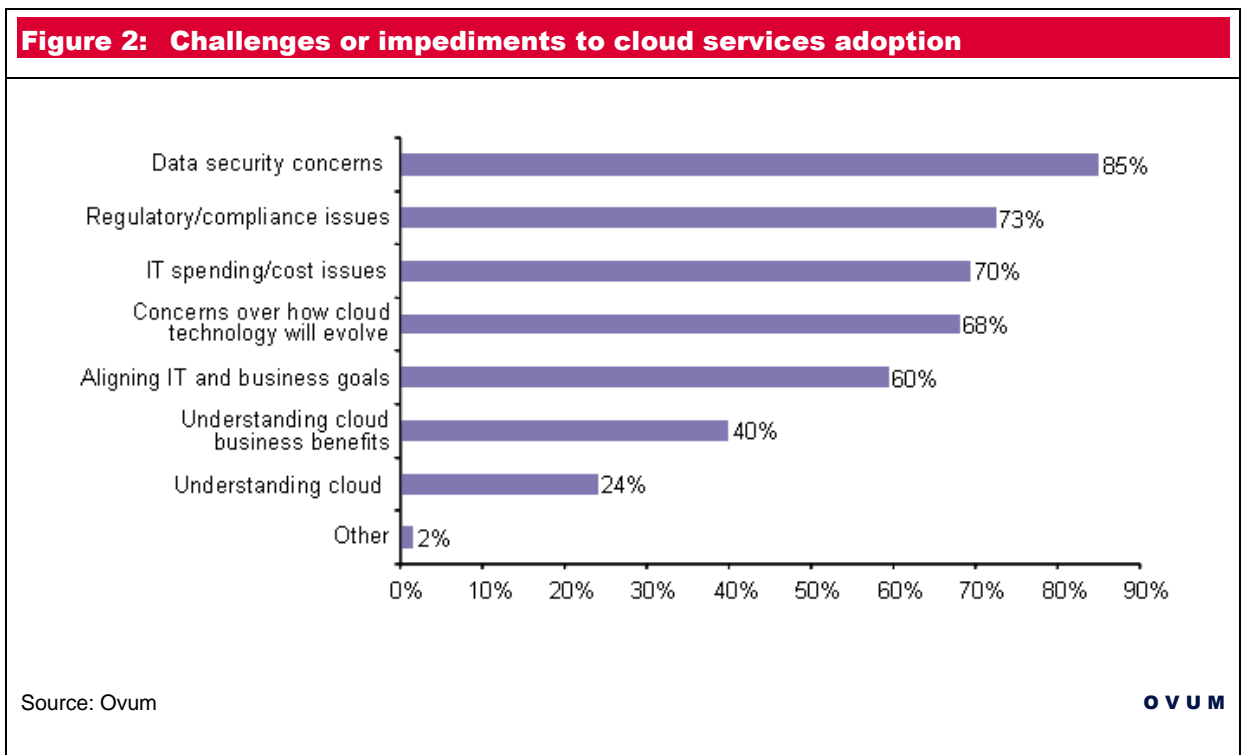
Where the network delivery of the cloud computing service is a network service that has its own SLAs, as with an Ethernet private line, VPN, or, as is more commonly the case today, an MPLS-based VPN, then end-to-end availability is determined by the performance of both the WAN and the cloud computing platform. Today it is rare to see a single end-to-end SLA, but CSP cloud computing offerings should be able to assure end-to-end performance through separate SLAs. For example, one will cover the cloud platform up to the point of peering with the WAN, and another the particular WAN solution used by the customer. Of course, a customer can achieve this by having separate providers for its WAN and its cloud platform, but having SLAs from the same service provider avoids a culture where one provider blames the other and the customer is left in the middle.

Another issue is the tendency of some cloud service providers to define their platform availability in terms of how service credits are paid rather than as a realistic measure of the availability that can be expected by the customer. This was once also a problem with network service providers, but two decades of providing SLAs for business-critical networking has instilled some discipline in terms of setting SLAs that reflect true historic network performance.

In addition to availability, cloud computing service providers commonly offer SLAs for other features such as online portal availability and time to initiate new requirements (e.g. adding virtual machines). However, another aspect of enterprise-grade services will relate to the performance of applications where the combination of computing/storage platform and WAN determines application performance. For example, a customer wanting to run a realtime unified communications application through a cloud service will be interested in additional performance parameters such as latency. The increasing requirements of global applications delivery will place greater demands on both the networking component and on the cloud platform, and CSP cloud providers will be well-placed to respond to this through SLAs.

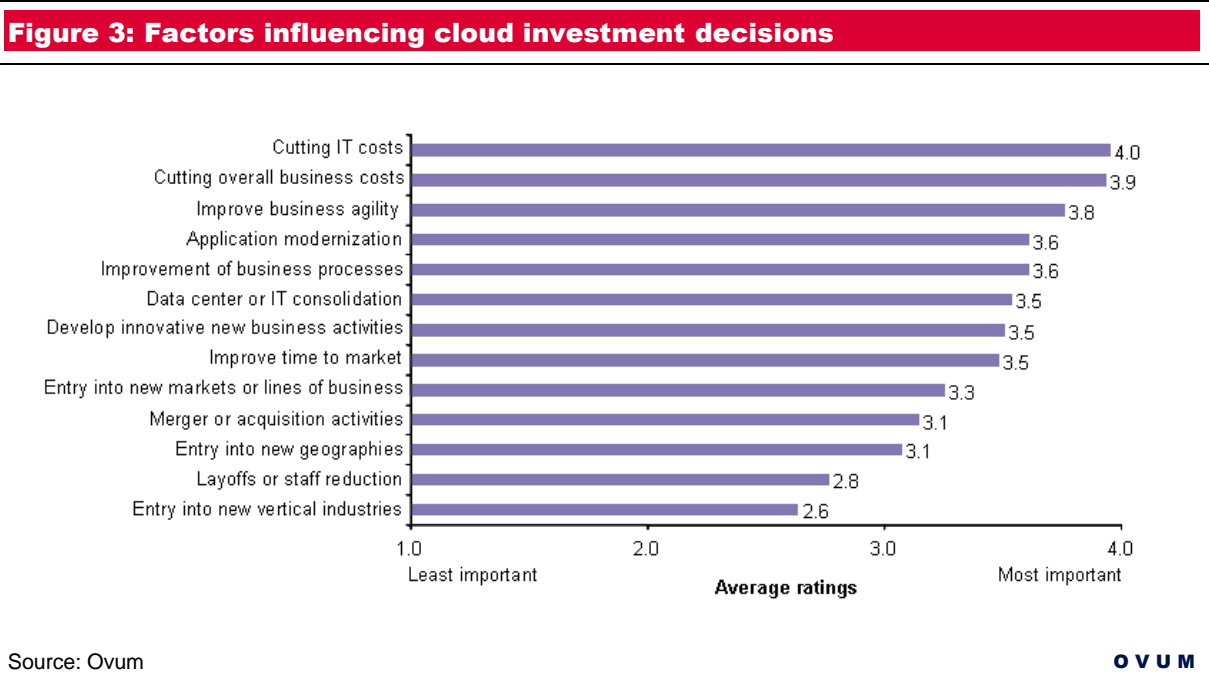
A strong focus and deep understanding of issues around security, data privacy, and compliance

To be a credible provider of enterprise-grade cloud computing service providers must also be able to demonstrate that they have proven and auditable processes and expertise regarding security, data privacy, and compliance. On occasion this will include the facility for customers to engage their own security audits on the service provider's facility. Figure 2 is taken from a survey of 200 IT decision-makers in Europe and the US conducted by Ovum in late 2011. It shows that data security and regulatory compliance are still the greatest barriers to cloud services adoption, despite considerable industry focus on these issues. For some enterprises this will either delay cloud computing adoption or focus adoption on private cloud.



However, Figure 3, which is taken from the same survey, shows that the greatest driver in cloud computing adoption is reduced IT costs, followed by improved business agility. Public cloud computing services offer the greatest cost savings, together with the benefits gained through improved business agility, but are seen as the least secure. Enterprise-grade public clouds implemented as virtual private clouds offer a middle ground, achieve most of the cost savings of public clouds, and go some way to reassure enterprises about security. A comparison can be made with enterprise VPNs, which today carry the majority of networked business applications traffic for enterprises throughout the world without the use of additional security measures such as

encryption. Security and compliance issues for some business verticals and applications will inevitably limit cloud computing adoption to private clouds, but there will be many situations where the adoption of enterprise-grade public cloud services is a satisfactory approach. This will depend on CSPs demonstrating that they can provide secure environments and robust security processes.



Flexible deployment models: public, private, and hybrid

Ovum is increasingly seeing that for many situations the choice between public and private cloud is not clear cut, and that for most enterprises there will be a role for both deployment models. Some enterprise applications will reside in a private cloud (either on-premise or hosted), and others in the public cloud. This hybrid approach is likely to become the norm for many enterprises, and a service provider that can offer this flexibility in its deployment models will be in a strong position. We are also seeing a need for greater flexibility in moving applications or workloads between the public and private cloud models. Some applications might initially reside in a private cloud and migrate over time to public (or vice versa), while elsewhere a private cloud application might burst out to the public cloud when additional unplanned capacity is required.

An increasing number of customers want a service provider that can offer both the public and private cloud components. For example, an e-commerce application might have a public component that is consumer-facing and a private component that provides back-office functions and interfaces to other internal business processes. Changing security or compliance needs may also require applications initially served by the public cloud to migrate to private. This flexibility is



going to be increasingly important, especially as enterprises will be under continued pressure to reduce IT costs and exploit the lowest-cost solutions that are fit for purpose. Requirements for flexibility are now being built into customer portals and automated, so that with the latest generation of cloud platform users can transfer workloads between private and public clouds on demand.

THE EVOLUTION OF ENTERPRISE-GRADE CLOUD COMPUTING SERVICES

The growing importance of the role of the network in public cloud computing

The last decade has seen a major transformation in enterprise networking, and a recent Ovum survey of 520 large enterprises globally indicated that 72% now use a modern carrier VPN service (such as MPLS VPN or Ethernet VPN) as their primary WAN technology. There has also been a transformation of the capabilities of VPNs in terms of SLAs and the provision of traffic prioritization through class of service. This has been driven by the growing demands on enterprise WANs to deliver business-critical applications to high performance levels throughout the enterprise. The globalization of enterprises and the growth of realtime applications such as voice and video have also meant that network latency has become an important parameter in WAN performance. These same trends are now happening with public cloud computing as it moves into a new phase of its evolution in which enterprises increasingly wish to deploy business-critical applications in the public cloud. Enterprise-grade VPN services are now becoming a key enabler of the transition to enterprise-grade cloud computing.

Another transformation in carrier enterprise networking that has taken place more recently is in the increased role of the WAN in inspecting, controlling, and optimizing the performance of individual networked business applications. This has led to some network service providers offering APM alongside network services such as MPLS VPN, and providing measurement and optimization of business applications performance through the deployment of network probes or application acceleration devices in customer data centers and key business sites. Major CSPs have introduced APM capabilities as managed value-added services so that customers and service providers can monitor (and in some cases optimize) the performance of individual networked business applications. We are now seeing this being applied to the cloud computing environment, for instance where applications optimization is combined with SaaS delivery through a content delivery network (CDN). It is also beginning to be applied in an IaaS environment, starting with customized solutions for private cloud deployments where APM is part of the service provider's overall solution.

CSPs have the opportunity to more closely integrate WAN services, APM capabilities, and cloud computing platforms to offer new managed services as part of a broader enterprise-grade cloud computing offer. Changes in the way APM is delivered are also helping to accelerate this integration. For example, there is a growing trend to incorporate APM capabilities in network CPE

as well as an increasing availability of APM applications that can be delivered from a virtualized platform.

The concept of managed public cloud computing

Some service providers already use the term "managed cloud computing", though it is not well defined. It tends to be used by CSPs that want to emphasize certain aspects of their service (usually IaaS), such as high availability or network security features. Today "managed" is more of a marketing term than an enhanced service capability. However, public cloud computing is beginning to evolve from a service that is largely self-managed from a customer's perspective (in that the customer is responsible for ensuring that they have requested sufficient resources in terms of virtual machines, storage, bandwidth, and so on) to one where the service provider can take some responsibility for ensuring that the business applications hosted by the service will achieve adequate performance levels. While service providers will today provide reporting on parameters such as virtual machine and storage utilization, it is still the responsibility of the customer to interpret this information in terms of its potential impact on application performance.

Over the next year or two we expect to see APM capabilities being applied to enterprise-grade public IaaS to provide optional enhanced management features. This would include application performance monitoring and optimization being built into the service provider's cloud platform, and linked to the self service portal. They could then be self-provisioned alongside computing and storage resources. Over time this could evolve into a fully managed solution where the service provider offers SLAs on the performance of the customer's individual hosted applications, and automatically adjusts computing, storage, and network resources as required. While many customers of public cloud computing will still want a simple solution that just hosts their applications at a low cost, there will also be many that want their cloud service provider to take more responsibility for an end-to-end solution with agreed application SLAs. This, of course, has many similarities with traditional managed application hosting. A key difference, though, is that, apart from using a lower-cost virtualized environment, CSP-managed services can closely integrate business-grade WAN services with the cloud platform to provide a highly flexible end-to-end managed solution.

The provision of application-based SLAs has been seen by many CSPs as the holy grail of enterprise networking. However, it is difficult to achieve this aim in the normal situation where enterprise applications are hosted out of the enterprise's own data centers, as the CSP has little or no control over the data center infrastructure. The approach of some CSPs has been to monitor the infrastructure in the customer's data center, report on end-to-end performance issues, and provide an SLA that excludes issues arising in the customer's data center. As enterprise



applications move to the cloud, and cloud services are provided by CSPs, end-to-end applications-based SLAs will become a reality.

Intelligent network CPE will have an increasing role in enterprise-grade cloud computing services

To date the role of network CPE has been largely ignored in the provision of cloud computing, as the CPE is normally transparent to the delivery of a cloud service. However, this is set to change as we see solutions emerging where intelligent network CPE, network transport services, and cloud service platforms work hand-in-hand to provide enhanced capabilities such as security, APM, and disaster recovery functions. A simple example might be additional network resilience, whereby a cloud service delivered over an enterprise VPN would failover to the public Internet in the case of link failure. This could potentially extend to a scenario where the CPE provides emergency storage backup, or even runs an image of the hosted application in disaster recovery situations. Again this is a great opportunity for CSPs to drive innovation in conjunction with networking equipment vendors.



APPENDIX

Disclaimer

All Rights Reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the publisher, Ovum (an Informa business).

The facts of this report are believed to be correct at the time of publication but cannot be guaranteed. Please note that the findings, conclusions, and recommendations that Ovum delivers will be based on information gathered in good faith from both primary and secondary sources, whose accuracy we are not always in a position to guarantee. As such Ovum can accept no liability whatever for actions taken based on any information that may subsequently prove to be incorrect.

Ovum operates under an Independence Charter and the report represents the views of Ovum. Further information on Ovum's Independence Charter is available on request.

Ovum may have been paid for the right to re-use the material.