Cloud computing in telecommunications

Cloud services suffer as a result of concerns over connectivity control, trust, security, privacy and local presence. Fortunately, network service providers are well positioned to address these concerns.

JAN GABRIELSSON, OLA HUBERTSSON, IGNACIO MÁS AND ROBERT SKOG

The concept of cloud computing has been widely discussed and defined in a number of ways. Perhaps the most generally accepted definition has been put forth by the US Government's National Institute of Standards and Technology.[1] The exact definition is not really important, however; it's the new and changing business models enabled by the emergence of the different technologies associated with cloud computing that really count.

Introduction

Three things that provide consumer capabilities are often mentioned in the context of cloud computing. They are Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). In addition, network service providers may provide connectivity in a cloud landscape **(Figure 1)**.

In essence, SaaS enables consumers to use provider applications that run on a cloud infrastructure.

With PaaS, consumers can create and deploy applications onto the cloud infrastructure. Moreover, the applications can be created using provider-supported programming languages and tools.

IaaS enables consumers to access processing, storage, networks and other fundamental computing resources in order to deploy and run arbitrary software, including operating systems and applications.

Cloud computing has many characteristics, such as highly dynamic utility computing or elasticity, and the ability to access and pay for more or fewer resources, such as central processing units (CPUs) and storage, as demand increases or decreases. Virtualization technology and broadband connectivity are what makes this possible. In simple terms, virtualization allows traditionally course-grained physical resources to be split into fine-grained virtual resources that can be allocated on demand over a broadband connection.

Becoming a cloud provider

There are two main reasons why network service providers should consider becoming engaged in cloud computing. The first is to reap the benefits of cloud computing for IT optimization (lower costs, and more elasticity and speed). The second is to exploit new business opportunities **(Figure 2)**.

BOX B
Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources.

As virtualization and data-center technologies and characteristics mature, some telecom applications will move into the cloud environment. Others will be delivered as SaaS. This will provide the benefits of cloud technologies to operators. It will also improve vendors' managed services offerings, which actually provide benefits that are similar to clouds to operators that do not want to invest in or run their own cloud solutions.

Network service providers are well positioned to address many of the concerns that relate to cloud computing. Given the control they exercise over access transport networks, they can provide secure quality-of-service (QoS) connectivity to and from the data centers. Moreover, their customer-connectivity offerings can be very flexible (elastic) in terms of bandwidth.

Network service provider brands are more strongly connected with trust and security than most internet brands. As opposed to large parts of the global internet industry, the nature of the communication industry, with its locally regulated markets, has given rise to a strong local presence, providing a good foundation for addressing enterprises and consumers.

Cloud networking

Cloud computing gives network service providers new opportunities to move up in the information and communication technology (ICT) value chain. By virtue of the control they have over the communication infrastructure, network service providers have two important advantages compared with cloud providers from other industries. First, they offer domestic and business

BOX A Terms and abbreviations

API	application program interface
CPU	central processing unit
EC2	Amazon Elastic Compute Cloud
IaaS	Infrastructure as a Service
ICT	information and communication technology
NaaS	Networking as a Service
PaaS	Platform as a Service
QoS	quality of service
SaaS	Software as a Service
SAIL	scalable and adaptive internet solutions
SLA	service level agreement
TIC	transparent internet caching
VLAN	virtual local area network
VM	virtual machine

connectivity. In networking terms, no other player is closer to the consumer. As a consequence, network service providers have the ability to instantiate cloud functionality that users can access with a minimum of latency. Second, network service providers can tailor connectivity to the needs of their users; for example, by dynamically upgrading the connectivity between sites involved in high-definition video production. In other words, they can allocate – on demand and for an arbitrary period of time – the cloud server resources, such as data centers, and the high-bandwidth pipes that interconnect the data centers and users.

Better still, network service providers can take cloud computing one step further by adding dynamically scalable communication to the cloud service offering. Emphasizing the importance of adaptive connectivity, this next phase of evolved cloud computing could be called cloud Networking as a Service (NaaS). This concept is based on the vision of a unified management framework for computing and communication. Rather than controlling network and computation facilities as separate entities, with little or no power over the underlying resources, cloud networking can optimize resource allocations by considering network and computing resources as a unified whole.

A common management framework can operate across this distributed resource pool, going beyond current approaches to cloud computing and virtualized networks. The management framework would enable computing and communication services to be provisioned together to meet the needs of services with a broad range of requirements (expressed, for example, by service level agreements or SLAs) such as large-scale data-centric applications and information-centric networking. With cloud networking, resource allocation and service migration can dynamically take advantage of available network, computational and storage resources. In short, cloud networking will make it possible to manage the demand for network resources instantaneously, and to meet changing service needs more quickly and efficiently than today.

Further, cloud networking solutions will extend today's cloud computing infrastructures (which already include flexible networking, albeit only with the data centers in which they are harbored) into network service provider networks – for instance, into access networks. This space can then be filled with scalable and adaptive internet solutions (SAIL).[3]

FIGURE 1 **The roles of a cloud provider.**

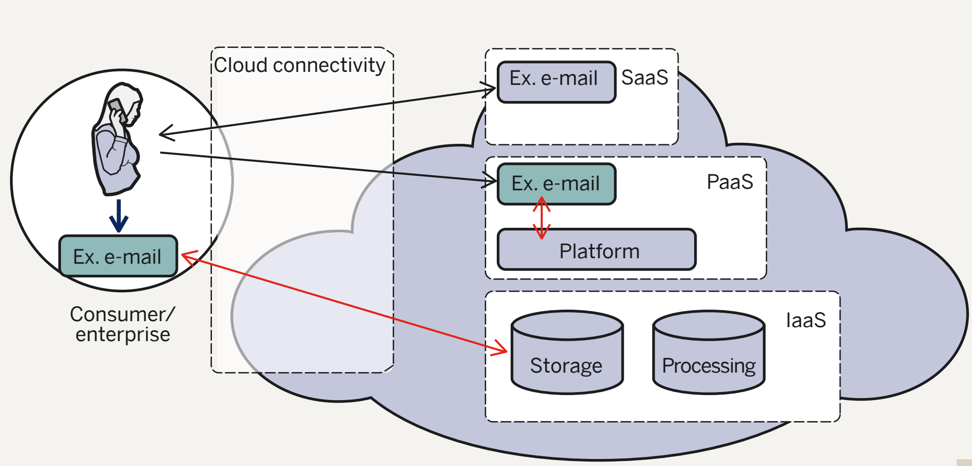

FIGURE 2 **Opportunities afforded by cloud computing.**

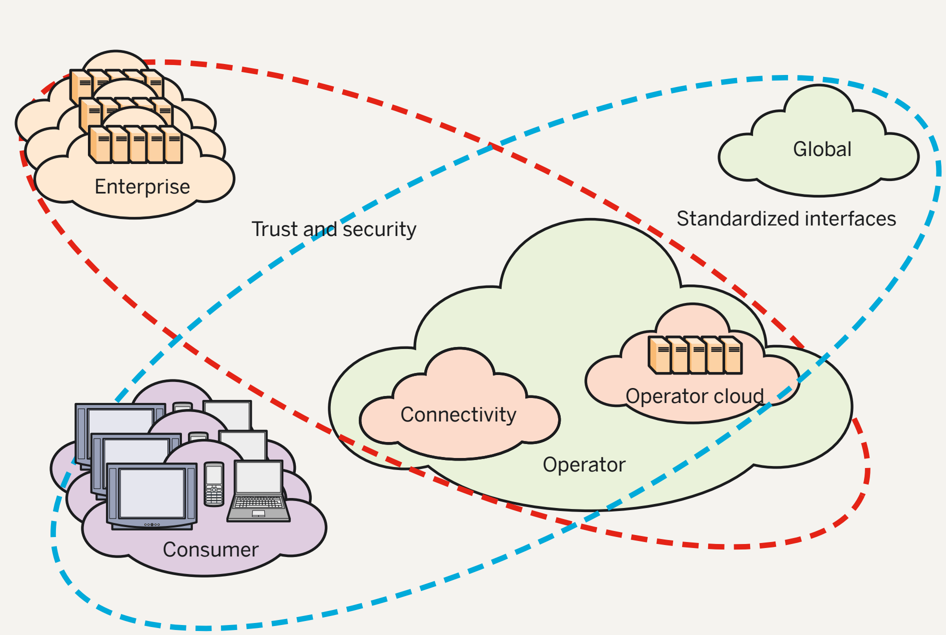

Virtualization over GEP2 blades

Ericsson's transparent internet caching (TIC) project has proven that the Ericsson's GEP2 server blade is a good platform for virtualization.

It facilitates virtualization without any locally connected system disks. This is an important factor for real-life implementations in the field because no moving parts means fewer potential causes of system failure.

Performance is also acceptable. The TIC project has been running about 50 highly active virtual machines per blade. The first resources to run out were CPU cycles and network capacity. Memory size was not an issue.

The complex network setup and the large number of nodes and node types pushed those working on the project to find a working solution for automatic configuration at boot time. Those involved in the project also had to find advanced scripting solutions to create the entire "world" from a few template virtual machines (VMs). The solutions have worked very well and could be used in similar projects where centralized control is necessary and where it must be possible to rapidly upgrade a large set of distributed nodes.

An entire, completely virtualized network is being run through the TIC project. Network virtualization techniques are used in the core network to slice links into smaller, completely separate ones. This is done using virtual local area networks (VLANs) in the virtual core routers to create virtual interfaces on the virtual Ethernet ports.

In many senses, the TIC/GEP2 virtualization cluster is the perfect platform on which to investigate and develop network virtualization management prototypes **(Figure 3)**.

Transcoding in the cloud

Important use cases for cloud computing include scenarios where load peaks are handled by offloading, by using

- internal resources: for example, in another data center, such as a private cloud; or
- external resources, such as public clouds.

In many practical scenarios, the public and private clouds used in a hybrid system will rely on different technologies. This, in turn, could imply different virtualization technologies, different application program interfaces (APIs) to request and control resources, and different network and connectivity environments. To illustrate the benefits of using cloud solutions to offload peak load and get a hands-on understanding of deployments in hybrid clouds, Ericsson has developed a prototype for scalable media-batch transcoding across multiple private and public clouds.

The prototype is based on a media-transcoding enabler at Ericsson Labs. The device consists of two parts:[2] a transcoding manager that dispatches the jobs, and a set of small, stateless and independent transcoding units that perform the actual media processing. It is possible to scale the capacity of the transcoding service up and down by adding and removing instances of the transcoding units.

FIGURE 3 **Topology of the TIC network.**

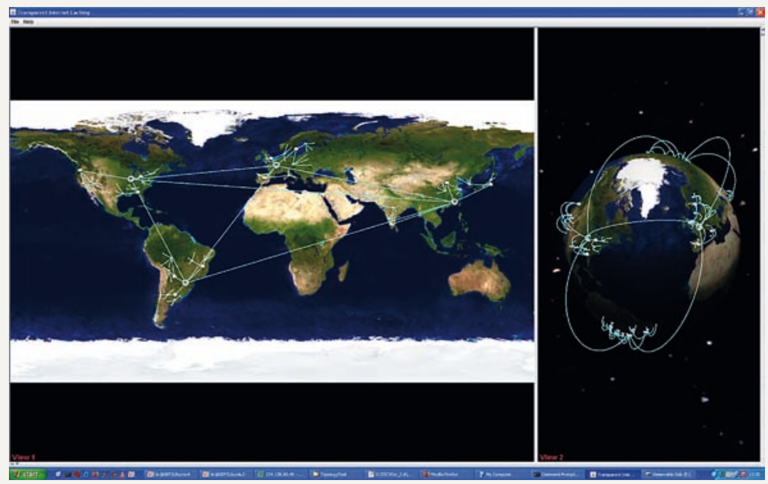

Current deployments require the new transcoder instances to be set up, configured and managed by hand. However, the prototype introduced a cloud resource manager to extend the transcoder setup. As seen in Figure 4, which depicts the overall architecture of the prototype, the cloud resource manager provides a single interface to applications, such as the transcoder manager, in order to request new resources and release those that are idle. The cloud resource manager allocates the resources by creating and configuring VM images in different clouds. Multiple APIs had to be implemented for this purpose, including those of VMware and Amazon Elastic Compute Cloud (EC2).

In order to deploy and configure new instances, the cloud resource manager contains a flexible deployment engine that matches the properties of clouds, available VM images, and the requirements of applications. As a result, so-called modifiers are applied to the preconfigured VM images in order to adapt and configure them for different applications. This enables deployments in hybrid cloud environments without having to generate specific VM images for each application and cloud.

The setup actively runs and uses resources across multiple sites and Amazon's EC2 public cloud. The prototype demonstrates the benefits of hybrid clouds and has shown ways of solving the associated technical issues. More advanced network service provider use cases will impose further requirements, such as latency, security, localization, SLA management and fault tolerance. The vision for ongoing research is to resolve these issues as well.

Security and trust management

The importance of security and trust, in particular where the private clouds of network service providers are concerned, motivates active research in several related areas. Ericsson is collaborating with academic partners to establish a cloud security taxonomy or classification system, and to derive security requirements for cloud computing. Special emphasis is being placed on the most challenging deployment scenario: hybrid clouds. Other activities focus on security for the important underlying technology enablers of clouds, namely virtualization and multicore platforms.

A main concern for this work is how isolation should be maintained in shared hardware environments, both for

FIGURE 4 **Architecture of the cloud-based transcoding demo.**

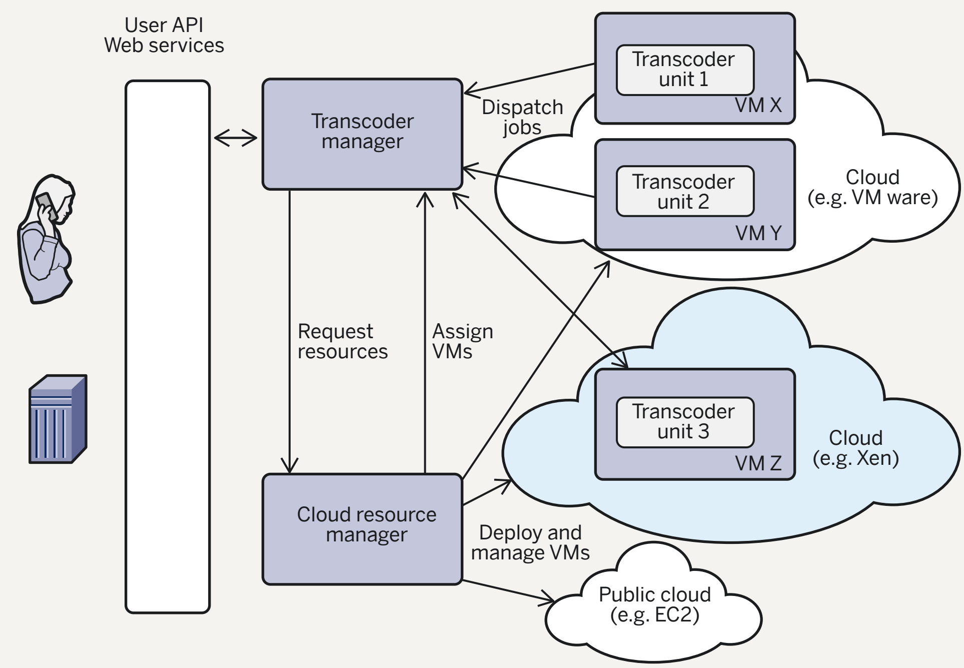

the virtualization layer (that is, the hypervisor) as well as for deployed VMs. One other track is intended for the development of models and methodology for analysis, and the design and the implementation of secure virtualization solutions.

A subordinate area particularly relevant to cloud computing focuses on efficient methods of securely managing trusted platforms that host virtualized systems. Trusted computing concepts, such as secure boot, trust anchors and run-time monitoring deserve special attention. The goal of these divergent avenues of research is to have trusted clouds, giving customers confidence in the underlying infrastructure, design and ability to dynamically enforce policies and respond to attacks.

Conclusion

As cloud computing gains traction, more and more applications and data are being moved from user premises and equipment to the internet. Without access networks, services and data will not be available for the users, and being connected is becoming a necessity. Connectivity increases the value of access networks and provides network service providers with new business opportunities.

Ericsson is actively researching different aspects of cloud computing. The goal is to provide the architectural components that will help network service providers realize these opportunities.

References

1. National Institute of Standards and Technology
http://csrc.nist.gov/groups/SNS/cloud-computing
2. Ericsson Labs, Media Transcoding
https://labs.ericsson.com/apis/converting-media
3. SAIL Project, 7th Framework EU program (FP7-257448)

Jan Gabrielsson,

who joined Ericsson in 1997, is an expert in software product strategies. He started out at Ericsson Research, working with agent-based telephony services, and later on with smart home projects. From 2001-2005, he worked with technology strategies at Service Layer R&D. At present, he works with portfolio and market strategies in the portfolio management group at Business Unit Multimedia. Jan holds an MSc in Computer Science from Uppsala University, Sweden.

Ola Hubertsson

joined Ericsson in 1992 to work with PDC system development and to head system management for the PDC business unit. Later on, he switched to network products for GSM and WCDMA before joining Ericsson Group Function Technology, where he works with network architectures and product performance. He is currently head of Architecture and Systems at Group Function Technology and Portfolio Management. Ola holds an MSc in Applied Physics and Electrical Engineering from the Institute of Technology at Linköping University, Sweden.

Ignacio Más,

who joined Ericsson in 2005, is a Senior Specialist in services and software at Ericsson Research. He has worked for the Internet Engineering Task Force (IETF) and in standardization, in internet protocol TV (IPTV) and messaging architectures, and is a member of the Ericsson System Architect Program (ESAP). He holds an MSc from both the Royal Institute of Technology (KTH), Stockholm, Sweden and the Telecommunications Engineering School at the Universidad Politécnica de Madrid, as well as a PhD in Telecommunications from KTH.

Robert Skog

is an expert in the service layer at Business Unit Multimedia. He joined Ericsson through a two-year trainee program for system engineers, and has since worked in the service layer area with everything from the first WAP solutions to today's connected home. Robert won the Ericsson Inventor of the Year Award in 2005. He holds an MSc in Electrical Engineering from the Royal Institute of Technology (KTH), Sweden.

Acknowledgments

The authors would like to thank the following people for their contributions: András Méhes, Johan Kölhi, Bob Melander, Manuel Couceiro, Alejandro Bascuñana, Miguel Ángel Pallarés, Román Ferrando, Daniel Catrein, Norbert Nieber, René Rembarz, Marc Vorwerk, Johannes Willig and Ioannis Fikouras.