

NetEvents 2010: Video, LTE Networking, Cloud Computing & Virtualisation

Comm

WRITTEN BY CHARLES F. MOREIRA

SATURDAY, 03 JULY 2010 10:21

Video, networking, LTE (Long Term Evolution), cloud computing and virtualisation figured prominently at the NetEvents Asia-Pacific (APAC) Service Provider VIP Summit and the Press Summit at the Ritz-Carlton Millenia hotel in Singapore from 19 to 21 May.

Analysts speak

A connected future is emerging, according to Adeel Najam, industry analyst, Frost & Sullivan.

It will see increasing capacity and capability, with continued improvement in multi-function devices, increased sophistication of specialised devices, higher bandwidth networks planned for both fixed and mobile and higher compression and miniaturisation of consumption platforms.

It will also include changing consumption patterns, with more opportunities for consumption any time, anywhere, including multitasking, while consumers will have greater control over their consumption experience and be at the centre of it.

It will also see more ubiquitous behaviour. Just as flat rate plans have changed user behaviour in the fixed Internet PC world, cheaper mobile tariffs and better handsets will spur ubiquitous data usage in the mobile world.

As that happens, mobile broadband will have to manage its cost per bit to be able to compete when usage exceeds 4 to 5 GB per month and mobile is expected to reach fixed-type usage behaviours in two to three years, based on a study modeled by Cisco.

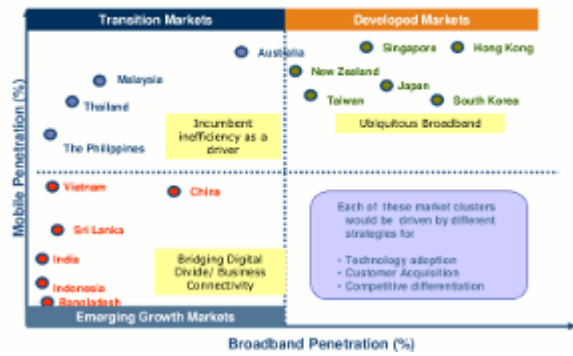
As it stands, fixed broadband consumption in the developed countries is 10GB per user per month, while in developing countries it's 4GB per user per month, while mobile broadband averages 1.5GB per user per month and the 500MB per user per month for the iPhone.

In two to three years time, the usage of pay-per-use users will average 200MB per month, while those with limited fair-usage data plans will average 1GB per month, and those with fully unlimited plans will average 3 to 5 GB per month.

This will lead to a change in the economics of mobile broadband, with backhaul capacity per provisioned user rising from 3 Kb/s to 40Kb/s per provisioned user at 4 to 5 GB per month, or a 12 to 15 times increase in backhaul capacity to cope with the increased traffic, and it will require more cells, hence base stations in the radio access network to provide the same amount of coverage, given the limited allocated spectrum.

Now the role of wireless broadband varies across the Asia-Pacific region, based upon the percentage of mobile penetration versus broadband penetration.

Role of wireless broadband would vary based on the type of market



Courtesy Frost & Sullivan

In the bottom left quadrant are the emerging growth markets, including China, Vietnam, Sri Lanka, India, Indonesia and Bangladesh, with relatively low mobile and broadband penetration; here mobile broadband will serve to bridge the digital divide and provide business connectivity.

In the upper left quadrant are the transition markets, including Australia, Malaysia, Thailand and the Philippines, with high mobile penetration, including over 100% but low broadband penetration and here, incumbent inefficiency drives mobile broadband uptake.

Lastly, in the developed markets such as Singapore, Japan, South Korea, Hong Kong, New Zealand and Taiwan, mobile broadband serves to enable ubiquitous broadband.

Each of these three market clusters will be driven by different strategies for technology adoption, customer acquisition and competitive differentiation.

Now an emerging technology which will enable better mobile broadband is LTE (Long Term Evolution), a highly evolved form of 3G, capable of over 100Mb/s speeds.

Deployment of LTE networks have already begun this year, the first commercial launches are expected in 2011, with wide-scale deployment in 2012.

Leading cellular mobile operators from both the GSM and CDMA camps across North America, Asia and Europe have already committed to LTE. They include France Telecom, Vodafone, T-Mobile, TIM, AT&T, NTT DoCoMo and China Mobile from the GSM camp and Verizon Wireless, Telus, Bell, KDDI, China Unicom and China Telecom from the CDMA camp.

A cloudy future for private cloud?

While the term “cloud computing” is the latest industry buzzword shouted from the rooftops and slavishly parroted by many IT media, especially those in need of some sensation to sell copy or attract eyeballs, Camille Mendler, vice president of Global Service Strategies with the Yankee Group is less than enthusiastic about its future, especially not that of what’s called “public cloud.”

Camille Mendler

Firstly. A slight digression on what the hell is cloud computing, to paraphrase Oracle Corporation chief executive officer Larry Ellison, when he was questioned on the subject a year or more back? In reply, Ellison



said. "The computer industry is more fashion-driven than women's fashion" and that if "orange is the new pink," then Oracle wouldn't fight the idiocy but would have to play along with it.

Well simply put, cloud computing's most well-known form involves the use of applications hosted centrally and accessed remotely over the Internet through a web browser. Simplest examples are free web mail services such as GMail, Yahoo! Mail, Google Apps, Google Calendar and so on, as well as paid customer relationship management services such as Salesforce.com.

The above class of cloud services are also known as software as a service (SaaS), the term which Oracle uses, instead of "cloud computing."

BTW. This writer remembers the term "application service provider model" (ASP model) used to denote the same thing over 10 years ago.

Anyway, other classes are platform as a service (PaaS), where a computing platform and/or software stack are provided centrally as a platform for deployment of the customer's applications instead of them owning their own, and the third class, infrastructure as a service (IaaS) provides a centralised computing infrastructure, typically a visualised platform environment to host a customer's applications.

Cloud computing is also divided into public cloud and private cloud categories. The former usually involves a one-size-fits all type of service provided to anyone who registers or subscribes, while the latter is a cloud service owned and controlled by a company, and accessible only to its staff or those which it does business with, such as its suppliers, distributors and perhaps its end-customers.

In her presentation, How Good is your Cloud Cover at NetEvents, Mendler acknowledged that while cloud computing certainly is the next Industrial Revolution, at the same time she pointed out the limitations of cloud computing,

She cited Victorian economist Stanley William Jevons, whose Jevons Paradox states that an invention which improves efficiency of a product's consumption also increases its volume of consumption. For example, the coal-fired steam locomotive, a product of the Industrial revolution, at the same time increased the consumption of coal.

Thus cloud computing has become a commodity item, much like coffee, pork bellies, orange juice and electricity, with the virtual machine instance being the unit of executing software in a server environment, and it's already being traded much like physical commodities in the traditional commodities market.

For example, Zimory – Deutsch Telekom's marketplace freely trades IaaS for buyers and sellers; With Amazon Spot Instances, buyers bid for spare EC2 capacity and run instances while bid exceeds spot price; while with Amazon Reserved Instances, buyers pay a pre-set reservation fee plus an hourly charge, billed on actual usage. These three respectively correspond to free markets, spot markets and forwards markets in traditional commodities trading.


However, at the same time, the traditional rules of governance, risk and regulatory compliance still apply to cloud computing, and vast differences exist in cloud service provision.

A 2010 report by the Yankee Group showed that only one in five of large enterprises is adopting public cloud, while 56% of enterprise usage is in SaaS, 25% in PaaS and 20% in IaaS, while 50% use private cloud services.

Another issue is that few cloud services have been road-tested for mission-critical reliability. Enterprises surveyed have no issue with the concept of cloud computing per-se but are concerned as to how service is

assured.

As such, their top five reasons for choosing to run their own private cloud over subscribing to public cloud service are security (73%), data privacy (55%), manageability (41%), cost effectiveness (31%) and compliance (30%).



Courtesy Yankee Group

Also in a 2010 audit of 41 public cloud service providers of SaaS, PaaS and IaaS services, the Yankee Group found that only 54% offered service level agreements (SLAs), 48% met European Union data protection adequacy, 43% give notice to change contract terms, only 9% give over 72 hours notice before conducting maintenance, 7% limit planned maintenance events per year, 6% warrant virus-free service, 2% offer proactive fault notification and none pay liquidated damages for service outages.

“Lastly, compliance is not only about security but also refers to the environmental footprint,” said Mendler. “As it stands, governments want to tax enterprises for fossil fuel usage, so how 'green' a virtual machine instance was to create will become more important to enterprises.”

We also had group interviews with representatives of participating companies and industry organisations.

Enabling video interoperability

“Online video is the distribution of video content across various networks to different users and screens,” said Daniel Bar-Lev, chief executive officer of the Video Convergence Forum (VCF), a global industry alliance for online video delivery solution providers - including content owners and hosters, broadcasters, equipment and software vendors, user device manufacturers and ad-insertion companies.

Daniel Bar-Lev

Established in February this year, the VCF's mission is to accelerate the interoperability of all video solutions in order to enable the delivery of any media over any network to any device. It also develops specifications and implementation agreements to promote all aspects of online video interoperability.

The VCF identifies three types of video content. The first being user generated video like what we find on sites such as You Tube, MySpace, Nokia Ovi and others. The



second are enterprise videos such as advertising clips, training videos, videos displaying properties, the links to which property developers and real estate agents can send to their clients to view and more. Thirdly are professionally created, revenue-generating videos for display on TV or in cinemas.

The multiple screens on which these videos are shown include PCs, TVs, mobile devices like the Apple iPhone, iPad, digital signage in supermarkets, cinemas and displays in elevators, and between the source and the viewer, these videos must be set up and adjusted for the type of delivery network, the viewing device, captions and advertisements may have to be inserted, and this can prove very costly, especially with a large number of viewers.

"We are seeing simultaneous and rapid growth in at least three aspects of the online video market: types of user device, network delivery infrastructure characteristics and localisation and this forces media content owners to prepare the same content in hundreds, even thousands of formats to serve all the potential markets," said Bar-Lev.

Add up all those transcoding and storage costs, monetisation capabilities and service assurance demands – these are critical issues for content owners and their delivery partners as they attempt to squeeze maximum value from their huge content archives, while serving their customers with high quality video" Bar-Lev added.

Well, online video is exploding, with ComScore reporting almost 20 billion views during January this year in China, Japan, Singapore, Malaysia, Hong Kong and Australia. This is matched in the US by 30 billion videos viewed online in March. In the UK alone, viewing of online video rose 37% over 12 months to reach 5.5 billion views in February – again according to ComScore.

According to Vidya Nath, Global Industry Manager, Digital Media Group, at Frost & Sullivan (F&S):- Based on current estimates of online video views in China alone, F&S expects a ten fold increase of viewing figures in that country over the next five years.

Also 'Over-the-top' or OTT video as a popular concept has taken firm roots, but it is unclear is how it is going to split with the ever growing number of device types and access methods, such as mobile devices and digital signage.

Besides that the regional landscape is strewn with so many cultures and languages that it requires localisation of advertising and subtitles, which poses a huge monetisation opportunity and an operational challenge to match.

The operational challenge begins with the question of the almost exponential proliferation of video profiles that describe how the video is edited for content, advertisements inserted, subtitles placed, as well as encoding for different target devices and delivery networks implemented.

According to Bruce Devlin, CTO of Amberfin, the basic question of how much 'localised' content should be prepared in advance and stored, and how much content should be prepared 'just in time' for video delivery is a complex one.

The better the model for standardised categorisation of online video profiles, the more effectively content owners, aggregation companies and other members of the online video ecosystem can match their business rules and workflows, as well as storage and monetisation capabilities to the actual market potential they are addressing.

In other words, the online video community can use standardisation of online video profiles to both control costs, and take advantage of fast appearing user markets.

Devlin leads the Video Convergence Forum's technical work on video profile definitions and specifications.

"If it costs US\$1 to set up and watch each video on a platform, it could well cost US\$30 billion per month just to set up and watch all those videos in the U.S. alone and we don't know how much it altogether costs exactly and who's paying for it," said Bar-Lev.

Also, setting up those videos consumes power and just consider that two Google searches alone consumes as much power as boiling a cup of water, which works out as a lot of electricity for say 60 to 70 Google searches.

So with the number of profiles defining different videos growing exponentially, combined with the different languages and types of content suitable to the culture and mores of different countries, so video companies could end up with 1,000 profiles, where the set up costs for each fragment is the same, so producers will have to decide whether they will realise and return on investment of each of the videos.

"The solution then is to standardise and automate much of the setup work to reduce costs, much like how Henry Ford automated the production of the Model-T, which dramatically reduced its cost and made it affordable for the mass market, so the market for cars grew," said Bar-Lev.

"Likewise, with video producers catering to Web 2.0, IPTV, Internet TV, mobile phones, Internet tablets and other platforms, content providers, mobile service operators, as well as consumers will benefit from the standardisation of profiles," Bar-Lev added.

The traditional way to fit a video to a viewing platform is to have a plug-in which fits a socket and this is where the VCF comes in to standardise ways to manipulate content to match the viewing device. The VCF basically seeks to break down the current walled gardens between content providers and viewers to create an open eco-system.

The VCF's key focus is on standardisation and marketing. It does not yet have a date of release of its standardised profiles, while on the marketing front it's creating a database and later this year will have an online video directory where video producers and related parties can register details of their products online.

Once the profiles and database are ready, during the setup, the system will detect the type and characteristics of the device and inform the transcoding company of the specifications of the device's screen and look up the file on the device, for example a mobile phone and encode the for it accordingly.

Since transcoding on the fly each time is a very expensive process, the solution will be to prepare and save the formats for each device in advance and store them but this creates other issues of storage and processing costs.

"So if we can reduce the number of formats through standardisation of plugs, sockets, processes and workflows, we can lower storage & processing costs, thus enabling higher availability," said Bar-Lev.

Video communications for all

The current video call and video conferencing options are either free with unpredictable quality at the low end, are of low quality, with high latency, no eye contact and use expensive equipment in the case of typical boardroom videoconferencing systems, and while telepresence systems at the high end provide a realistic, high quality experience, the equipment used is very expensive and in addition they also require very expensive dedicated networks connected to a MPLS (Multi-protocol Label Switching) backhaul, altogether costing hundreds



of US dollars per hour per endpoint, hence are unaffordable for most.

Ofer Shapiro

So Ofer Shapiro, Avery More and Alex Eleftheriadis founded Vidyo, Inc. in 2005 with the mission to make natural video communications universally available on a slew of affordable end points ranging from netbook, notebook and desktop PCs, iMACs and multi-screen boardroom telepresence platforms.

On 26 May Vidyo announced its VidyoTechnology software development kit (SDK) which enables developers to build multi-point videoconferencing applications on Android based smartphones and Moblin Internet tablets running Intel Atom Processor Z6xx series-based platforms (formerly Moorestown) and on ARM processor based platforms.

It will also support MeeGo based Internet tablets and mobile Internet devices when available. MeeGo is the result of the merger of Intel's Moblin and Nokia's Maemo Internet tablet OS, announced at the recent Mobile World Congress in Barcelona.

“Altogether, Vidyo's technology enables up to 200 users of all these PC, iMAC and mobile devices to participate in virtual meetings with each other over regular broadband connections such as DSL, over 3G, WiMAX and from WiFi hotspots,” Ofer Shapiro, Vidyo chief executive officer told NetEvents.

To participate, users would first have to download and install the appropriate Vidyo client software from their organisation's portal on their platform, as Vidyo itself does not operate the service but sells its solutions to organisations and service providers which do.

The VidyoPortal is a Web-based environment administrators can use to manage the VidyoConferencing system. Ordinary conference participants can use it to initiate meetings – via the Web – from anywhere. VidyoPortal's flexible interface features single-click-action buttons that take care of everything required to initiate a conference.

Also, regardless of whatever endpoint system users might have, VidyoPortal provides a consistent environment that's compatible with Microsoft Internet Explorer, Firefox and Safari.

End users can access portal through a range of endpoints that run VidyoDesktop or VidyoRoom. For first-time users, there's a simple sign-in or authentication procedure. Guest users can also be allowed access to the system if their organisation's policies permit that.


The core of Vidyo's technology is its patented VidyoRouter architecture which employs H.264/Scalable Video Coding (SVC) to let endpoints running it to provide a high quality user experience over regular broadband connections.

It thus offers users error resiliency and eliminates the latency found in existing multipoint conferencing unit (MCU) based implementations enabling natural, high quality video to work over the Internet, LTE, 3G and 4G networks for a fraction of the cost of dedicated telepresence systems by eliminating the need for the the MCU which codes and decodes information in traditional multi-point conferencing systems.

Videoconferencing on a phone with Vidyo

It also reduces transmission times, especially of bulky content such as presentations and videos by storing them on regional VidyoRouters in different parts of the world to reduce network traffic.

“So it basically reduces telepresence costs from dollars per minute to cents per minute and this will help drive videoconferencing use,” said Shapiro.



For example, Vidyo announced its Videoconferencing Touch Screen Executive Desktop at NetEvents and Shapiro claimed it brings the full 1080p high-definition videoconferencing experience to off-the-shelf desktop PCs for a total of under US\$1,000 per endpoint. The VidyoDesktop Executive client also runs on higher end PC workstations with dual high-definition screens and on iMACs. Vidyo demonstrated VidyoDesktop Executive at Interop in Las Vegas earlier in April, where it was selected at Best of Interop Finalist in the Collaboration category.


Meanwhile, its VidyoRoom HD–220 technology deliver video at 60 frames per second over general-purpose IP networks and also decode and display multiple HD participants as well – and all at a level of video quality which the company claims is “unequaled by systems that require dedicated bandwidth to perform at their best.”

Its entry level VidyoRoom HD–100 encodes 720p & 30 fps, yet decodes HD at 60 fps. The HD–100 can display a shared application on a second monitor. The VidyoRoom endpoints are all x86 based systems that use the same software base as our VidyoDesktop™ clients.

The VidyoRoom system is simple to use, easy to configure and displays video in multiple layouts — voice-activated with continuous presence, all participants equal, or any combination of these. And each participant may dynamically change their layout at any time during the meeting independently from any other participant. Flexible conference control options make it a snap to manage, using either the VidyoPortal™ or a remote control device.

VidyoRoom also interoperates seamlessly with VidyoDesktop clients as well as mobile clients.

Vidyo also lets customers continue to make full use of their earlier investments in videoconferencing equipment. Its VidyoGateway enables a customer's existing videoconferencing equipment such as their Polycom, Tandberg or other MCU-based systems to interact with their new Vidyo technology-based H.264/SVC systems. The gateway also supports legacy H.264 and H.263 video conferencing, as well as G.711, G.722, SIP and H.323 signalling.



Vidyo in healthcare

It's also used in healthcare applications in different parts of the world. For example, its VideoHealth telemedicine technology enables doctors to examine patients at home remotely, while its Videra videoconferencing service lets healthcare workers in the Nordic countries provide homecare to the elderly remotely. The system is designed to be easy to use, where the patient just needs to press a button to make a videocall.

In Latin America, cardiologist Dr. Javier uses his Virtual Web MD service developed to address the cardio expert shortage in that region via <http://virtualwebmd.com> which enables cardiologists to conduct patient examinations through the integration of an electronic stethoscope and to conduct physician training.

No video, well no problem as VidyoVoice lets voice-only participants with only single-line or mobile phones to join Vidyo conferences without requiring additional equipment in their network. Participants can simply dial into either a toll free or local access number to connect into a Vidyo user's room. VidyoVoice currently provides toll-free access in the US and local access numbers in 39 countries.

Its VidyoOne is a single U1 appliance at a single fixed price for all of the hardware and software components required for desktop conferencing, making it easy and affordable for any organisation to deploy their own VidyoConferencing solution.

VidyoOne intelligently routes video packets to endpoints where decoding takes place, delivering the highest quality video possible by using SVC technology.

Customers just need to connect VidyoOne to their for everyone in the organisation to communicate from anywhere they have an Internet connection.

"Gartner predicts that over 200 million workers worldwide will run corporate supplied videoconferencing solutions from their desktops by 2015," said Shapiro.

Mobile interactive services

"Mobile subscriptions are up and the increasing number of smartphones is driving up access to mobile videos, thus driving up bandwidth demands which in turn poses challenges to mobile operators which invest in upgrading the bandwidth of their networks at a loss," said Eamonn Kearns, Dialogic Corporation vice-president of Sales, Asia Pacific.

Established in 1984, Dialogic was a leader in computer telephony products and solutions, as well as voice, fax and video value-added services. It was acquired by Intel in 1999 and in 2006 was sold to private equity.

Today, the company provides solutions to optimise bandwidth and enables various mobile interactive services such as mobile-commerce, mobile-entertainment, information network services, high-definition (HD) voice in mobile networks, Web 2.0 interactivity such as social networking and more.

An example of HD voice is Skype which uses HD codecs and Orange in Europe uses HD voice as a competitive advantage and Dialogic is an enabler in the solution – ie. a video portal to look at Web content.

"Also, voice SMS where users record their voice and send it via SMS to the recipient is quite popular in the Philippines and Indonesia and it's suitable for places with low literacy rates," said Kearns.

"A third-generation SMS is video SMS where users can similarly record and send a video via SMS," Kearns added.

Both voice and video SMS only work over 3G networks and are handset-independent, unlike MMS introduced with GPRS handsets over 10 years ago but which due to its handset dependency, was not very practical.

Dialogic is currently sponsoring the 2nd Dialogic Innovator Challenge where developers compete to develop the ultimate mobile video sports application and post it on You Tube where viewers will vote on the winner who stands to win the US\$10,000 grand prize. For details visit www.dialogic.com/den/groups/innovation/default.aspx

Dialogic currently is in the process of merging with Veraz Networks, a VoIP, IP and packet telephony solutions provider pending regulatory approval. The merger is expected to be completed in the second half of this year and its name will be Dialogic on the NASDAQ Exchange.

MEF highlights Global Interconnect

Kevin Vachon, chief operating officer of the Metro Ethernet Forum (MEF) updated NetEvents participants on its Global Interconnect Programme and on the global growth of Carrier Ethernet services.

MEF's Kevin Vachon

The Global Interconnect Programme supports the MEF's mission of accelerating the worldwide adoption of Carrier-class Ethernet networks and services and seeks to deliver all elements relating to multi-operator implementation of MEF Carrier Ethernet services through a comprehensive informational programme.

The MEF is a global industry alliance comprising more than 165 organizations including telecommunications service providers, cable MSOs, network equipment/software manufacturers, semiconductor vendors and testing organisations.



Its mission is to accelerate the worldwide adoption of Carrier-class Ethernet networks and services. The MEF develops Carrier Ethernet technical specifications and implementation agreements to promote interoperability and deployment of Carrier Ethernet worldwide, and it also conducts certification programmes to enable standardisation and over 100 companies have so far been certified.

Standardised MEF Ethernet services are point-to-point E-Line Service used to create virtual private lines (EVPL), Ethernet private lines (EPL) and Ethernet Internet access; multi-point to multi-point E-LAN service used to create multipoint Layer 2 virtual private networks (VPN), transparent LAN service and multicast networks; and its point-to-multipoint E-Tree service use to create rooted multi-point Layer 2 VPNs, broadcast networks and telemetry networks.

Currently, 21 service providers across the Asia Pacific are MEF members, including Telekom Malaysia, China Telecom, Globe Telecom, KDDI, NTT, Philippines Long Distance Telephone (PLDT), PT Telekom, Singtel, Starhub, Symphony, telecom New Zealand, Telstra and others.

So far, 35 operators worldwide and over 400 equipment from 78 vendors are MEF certified.

Vachon highlighted the ratification of Phase 1 of MEF 26, the ENNI (External Network to Network Interface), Ethernet's growing role in mobile backhaul networks, and the recent launch of the first Ethernet Exchanges, offering a cost-effective alternative to the one-to-one connections between service providers.

ENNI basically is a standardised interconnection between different operators' Carrier Ethernet networks and MEF 26 ENNI specifications enables end-to-end delivery of Carrier Ethernet services across networks, while MEF 33 specifies how class of service are alignment between operators' networks.

“Among the many business opportunities created by the recently ratified MEF 26 standard, is the opportunity for

providers to now offer wholesale services, upgrade their Ethernet offerings to certified Carrier Ethernet networks and interconnect with multiple providers,” said Vachon.

“This increases providers’ options as well as delivers obvious savings from the use of standardised connections and processes as opposed to the current lengthy and uneconomical individual interconnect agreements,” Vachon added.

ENNI basically is a standardised interconnection between different operators' Carrier Ethernet networks and MEF 26 ENNI specifications enables end-to-end delivery of Carrier Ethernet services across networks, while MEF 33 specifies how class of service are alignment between operators' networks.

“Among the many business opportunities created by the recently ratified MEF 26 standard, is the opportunity for providers to now offer wholesale services, upgrade their Ethernet offerings to certified Carrier Ethernet networks and interconnect with multiple providers,” said Vachon.

“This increases providers’ options as well as delivers obvious savings from the use of standardised connections and processes as opposed to the current lengthy and uneconomical individual interconnect agreements,” Vachon added.

ENNI is absolutely key to the future growth of Carrier Ethernet, according to Michael Howard, Principal Analyst, Carrier and Data Centre Networks, Infonetics Research.

As operators increasing migrate to all-packet networks, they are investing heavily in delivering Ethernet services to the customers. The ability to span these services in a timely and economical manner over multiple operator networks is a multi-billion dollar opportunity in waiting.

When determining a potential Ethernet Service Provider (ESP) to partner with, operators must understand what Ethernet services the potential ESP offers, the technical details of the ESP’s offerings, business details, for example getting a quotation, how quickly can service be ready/interval, lit building status, and determine the geographic availability.

When ordering they must know what info is needed on the order for the specific service being ordered, including CoS, value-added services and so on, order status and the reasons for rejection of an order.

During provisioning, they know the network configuration and how to resolve technical issues.

Carrier Ethernet is now almost universally acknowledged by industry analysts as the only economically viable solution for mobile backhaul as operators struggle to cope with the changing traffic demands created by the migration to Smartphones by both business and consumers. As well as offering lower capital and operating expense, Carrier Ethernet’s scalability allows flexible response to changing traffic demands at mobile sites, as well as architectural scalability to meet future network needs.

According to a 2010 Infonetics Research Report, the worldwide average per new installed Ethernet wireline connection grew from around 8 Mb/s in 2005 to around 30 Mb/s now and is expected to grow to around 50 Mb/s in 2013. Over the same period, average bandwidth over legacy Switched Optical Network/Synchronous Digital Hierarchy (SONET/SDH) and Wavelength Division Multiplexing (WDM) connections would grow from around 5 Mb/s to around 10 Mb/s, while average traffic per Plesiochronous Digital Hierarchy (PDH) and Asynchronous Transfer Mode (ATM) over PDH connection would grow from around 5 Mb/s to around 15 Mb/s.

Also, a study by Nokia Siemens Networks of 19 operators across Europe, the Asia Pacific and the Americas

showed 100% year-on-year average growth in daily HSDPA data traffic from October 2006 to July 2009 as operators are rolling out increased capacities over EDGE, EV-DO, HSPA, WiMAX and LTE networks.

“With this big jump in mobile data traffic, the industry agrees on using Ethernet in the backhaul from base stations,” said Vachon.

Phase 1 of the MEF 22 mobile backhaul solution is an implementation agreement which provides generic specification for Ethernet backhaul architectures for mobile networks (2G, 3G, 4G), User-Network Interface (UNI) requirements, service requirements, clock synchronization, common terminology & standardised toolset.

Phase II comprises a granular focus on new mobile technologies such as LTE (Long Term Evolution), its scope includes performance recommendations, resiliency, sync, detailed SOAM recommendations, Multi-MEN, security and migration models leading to cost efficiencies.

Vachon also highlighted the recent announcement of Carrier Ethernet Exchanges. Implementing the MEF 26 ENNI standard, these exchanges provide Carrier Ethernet interconnection at strategic points in the network enabling multiple virtual connections over a single physical connection.

Ethernet exchanges are considered to be an essential prerequisite for service providers who looking to simply and lower Ethernet service implementation costs and deliver flexibility in scaling their services up or down as required.

The MEF has a website at www.MetroEthernetForum.org

Microwave cheaper than E1 and fibre

There's an ongoing dramatic expansion of network capacity and operators are installing very large pipes in anticipation of increased capacity demand, according to Amir Zoufonoun, president and chief executive officer of Exalt Communications.



Amir Zoufonoun

For example, mobile operators are spending US\$20 billion worldwide every year on backhaul using copper leased lines.

Several trends are driving dramatic month-over-month increases in the amount of traffic carried over cellular networks. The first is the widespread availability and promotion of mobile broadband services at prices that are allowing mass market penetration of WAN-connected laptops. The second is the explosion in the use of smartphones, exemplified by the Apple iPhone,

which are consuming data at near-laptop rates. The third is the growth of rich media and Web 2.0 content, both of which subscribers are demanding over the WWAN.

Adding to these trends is continued worldwide subscriber growth, and the migration of subscribers from 2G to 3G networks, and the result is that backhaul requirements of today's most advanced networks bear little resemblance to those of only a few years ago.

According to an Exalt white paper, typical 2G sites in the U.S. use 2 to 4 T1 lines providing between 3 to 6Mb/s backhaul capacity, which while sufficient for voice traffic, is way too little for EVDO and HSPA base stations which require more like 45Mb/s backhaul. In fact Deutch Bank estimates that by the end of 2011, typical business district base stations will require 50Mb/s backhaul

The air interface capacities of 4G access technologies such as LTE and WiMAX are higher still, with six peak subscriber download speeds measured in the tens of megabits per second. The rollout of 4G networks, in combination with the layering of multiple wireless access technologies, will equate to maximum backhaul requirements in excess of 300 Mbps at the busiest sites.

Peak capacities aside, Visant Strategies estimates that by 2012 between 24% and 32% of base station sites in Western Europe and North America will support backhaul capacity in excess of 45 Mbps, with as much as 9% of all sites requiring more than 155 Mbps.

However, continued use of copper leased lines for mobile backhaul in advanced cellular networks has well known limitations, one of them being their poor scalability, where the addition of more lines at a geometric rate is not only extraordinarily expensive but also time consuming. The wait time for new service provisioning can stretch for weeks and even months depending on service availability in the area.

Reliability is another issue, as leased lines frequently fail to meet the 99.999% availability figure required not just by service providers but by other large users of backhaul capacity, such as utilities and state and local governments.

With regard to network evolution, leased circuits are typically designed to deliver TDM connections such as T1/E1, DS3/E3 and OC3/STM1, but not Ethernet. However, UMTS and EVDO can use Ethernet as the primary network interface, while WiMAX uses only Ethernet. When using leased lines to backhaul Ethernet traffic, that traffic must be converted from Ethernet to TDM at each end of a connection; the required conversion equipment incurs additional cost and introduces excessive latency.

As this latency adds up, it poses a serious challenge for operators, as latency budgets are shrinking with each new generation of technology. The obvious solution is to use native connections for both TDM and Ethernet so that both types of service can meet the requirements of their respective applications and high user expectations.

Lastly, the cost of leased lines is dependent not only on capacity but often on distance, as in the case of point-to-point connections that pass through multiple central offices (COs) in the service provider's network. Furthermore, the monthly cost of leased lines can skyrocket when those lines cross Local Access and Transport (LATA) boundaries.

Besides leased lines, other ways to provide backhaul links are to use microwave or optical fibre, but fibre is costly to lay and these costs are not falling much anymore.

For example, fibre costs around US\$200,000 per mile (US\$ 124,223 per kilometre) to lay in the U.S. given labour costs there and roughly about US\$50,000 per mile (US\$31,056 per kilometre) in a developing country like Malaysia, and this cost has stopped dropping and has flattened off at a high level," said Zoufonoun.

However, a microwave link costs from US\$5,000 to US\$25,000 to lay and it has a range of up to 40 miles (64 km) in line-of-sight and it can be had up and running within a day, compared to weeks to lay fibre, not to mention the permits required from municipalities to lay it.

Also, capacity requirements at the network edge are generally lower than at the core, so it makes more economic sense to use microwave at the edge, than to use fibre.

As for speed, there already are 1 Gb/s radios and within two years, there'll be 10Gb/s radios, while there already is a 500Gb/s radio for communications in space, so capacity isn't an issue with microwave links. Instead the limitations currently lies in the device technology.

"Compared to leased lines, microwave systems are routinely engineered to deliver carrier-class availability at the 99.999% level and above," said Zoufonoun.

"Also as we've seen above, microwave is a one-time purchase typically equivalent to less than a year of lease costs and rarely dependent on the link distance or LATA boundaries. Once installed, microwave has the added advantage of being fully managed and able to accommodate remote upgrades, thereby significantly reducing future expansion costs," he added.

Exalt has over 50 products in which the electronics are enclosed, with operational temperatures ranging from -40°C to +65°C .

For example, Exalt's entry-level ExtendAir line of point-to-point outdoor microwave radios are pole mountable, hence have a zero footprint, operate at 4.9 GHz, 5 GHz and ANSI/FCC and ITU/ETSI 11 – 23 GHz licensed band systems offering flexible native Ethernet and optional native TDM capacity. All ExtendAir systems feature guaranteed throughput availability, low latency and high-security data networking.

Its low-end model, the rcXX000 has a base configuration of 25 Mbps full-duplex Ethernet throughput and is software upgradeable to 100 Mbps Ethernet. It has 1x10/100BaseT port with power over Ethernet (PoE). List prices of ExtendAir radios start at US\$8,000.

The EX-s Series GigE split-mount, point-to-point microwave radios are ANSI/FCC and ITU/ETSI 6 – 40 GHz licensed band systems offering flexible native Ethernet and native TDM capacity and designed to support migration from TDM to IP and 3G to HSPA/4G/LTE.

The split-mount, 1RU design combines indoor accessibility with the RF and cost efficiency of Exalt's outdoor models while delivering state of the art Ethernet and TDM performance. Standard features include adaptive modulation and capacity aggregation.

For example, its EX-s GigE 1000 is a licensed ANSI/FCC and ITU/ETSI band microwave radio system for service providers, enterprises and government organisations. It has a base configuration of 50 Mbps full-duplex Ethernet throughput and is software upgradeable to 100, 200, 300 and 360 Mbps.

"These provide mobile operators with low cost per bit with 1 ms latency, pay-as-you-go capacity, symmetrical or asymmetrical transmission, FIPS 197-compliance and 256-bit AES encryption of security," said Zoufonoun.

Exalt's customers include tier-1 carriers such as AT&T and T-Mobile and government bodies such as the Royal Thai Air Force.


Automation of IP address management

"There will be one trillion devices connected to the Internet by 2013," said Padma Warrior, CISCO chief technology officer on 25 March, 2010.

Source: Computerworld MarketVibe Oct 2008

According to an October 2008 Computerworld MarketVibe research report, the annual IP address management (IPAM) costs per IP address to enterprises and small-to-medium businesses increases as each new IP address is added.

The report showed prices growing from US\$7.46 per address per annum for under 1,000 addresses to US\$9.86 per address per annum for over 10,000 addresses, with an average of US\$8.00 per address.



According to Gartner, the exploding number of IP address and increased network complexity resulting from virtualisation and cloud computing will force network managers to embrace more automation of IP address management over manual methods.

For example, provisioning an IP address manually for a new server takes from 30 to 40 minutes, with up to two days delay, while provisioning a new virtual server hardware and operating system can be done automatically in less than a minute with no delay.

To that end, that Santa Clara, California-based Infoblox, Inc. acquired Nectordia on 4 May, so that it could extend its leadership in Domain Name Server (DNS), Dynamic Host Configuration Protocol (DHCP) and IPAM automation into change and configuration management.

The trend towards virtualisation presents new challenges for companies. The term virtualisation refers to the running of multiple logical computers, quite often running different operating systems, all within a single physical computer and these virtual computers can be created or removed as required, and they thus require rapid assignment of IP addresses for the duration of their existence.

On 24 March, 2010, Infoblox announced the Infoblox IPAM for Virtualisation solution, an advanced IPAM solution specifically designed to track IP addresses assigned to virtual machines. It lets organisations discover virtual endpoints and their relationship to the physical network; for example, which physical host currently runs a specific virtual machine – important information for effectively troubleshooting network issues.

It also enables virtualisation teams to manage their own IP space without relying on or affecting the other teams; enables networking and server teams to get visibility into IP space usage in virtualised environments.

Infoblox IPAM for Virtualisation uses an agent-less architecture to perform virtual endpoint discovery in a VMware environment. The discovery process finds and populates the IPAM database with information about all virtual endpoints present in the environment, including IP and MAC (Media Access Control) address, domain name, virtual entity type, entity name, virtual host, cluster, datacentre, virtual switch, and virtual host adapter.

At the same time, Infoblox also announced the availability of Infoblox vNIOS for VMware, its operating system optimized for VMware, which allows deployment of a virtualised instance of an Infoblox appliance. When used in conjunction with Infoblox physical appliances in a datacenter, and its ideal for enterprise branch offices, where footprint, power and other resources are often limited.


The vNIOS Virtual Appliance includes a full suite of core network services, including DNS, DHCP, IPAM, FTP,

TFTP, and HTTP.

In addition to providing an easy alternative to deploying hardware appliances at the branch office for providing highly available DNS and DHCP services, Infoblox vNIOS for VMware provides benefits including: Reduced power and footprint by minimising the number of physical servers and appliances, lower expenses by reducing hardware, cooling and real-estate costs, and increased resilience and availability of VMware environments.

Enabling better intercommunication

Today's enterprise networks comprises a host of interconnections between servers running core infrastructure and network management services, as well as applications servers and the interconnections start to look like spaghetti and harder to modify and maintain as the network gets bigger and more complex.



Spaghetti Junction

The sharing of data among infrastructure systems, and between the infrastructure and applications, typically requires custom, point-to-point integration using SNMP, syslog, proprietary APIs and custom scripts, resulting in systems which are complex, brittle and difficult to maintain.

As a result much of the data available in IT systems today remain locked in isolated silos, leaving many organisations without visibility into their infrastructure or assets, which in turn compromises security, increases costs and slows the ability to respond to changing conditions and business needs.


To address that, on 24 May, Infoblox announced the availability of its Infoblox Orchestration Server, a dedicated appliance that automatically aggregates, correlates, and distributes data to and from different IP connected systems, in real time.

It basically provides an IF-MAP service, to which each individual server or system is directly connected, with communication dynamically channeled between them by the Orchestration Server.

Infoblox Orchestration Server's IF-MAP solution

IF-MAP is a standard from the Trusted Computing Group (TCG); members of which include Hewlett-Packard, McAfee, Juniper, Microsoft and others.

In the area of network security, today's business and regulatory requirements demand that organisations provide appropriate levels of network and application access for a constantly changing mix of employees, contractors, partners and devices in data centres and remote locations.



Implementing effective policies requires information about a user's employment status, role, and privileges, and in some organisations, it's even necessary to restrict a PC's access to the network if their user leaves a secure location or if they exhibit anomalous network behaviour, such as worm traffic.

Providing that kind of dynamic, policy based security requires real-time information from multiple systems, including company directories, the human resource system, the building access control system, the network infrastructure, security sensors and others.

Other applications, such as automating virtualised data centres, tracking people and assets and enabling smart grids require the same kind of data sharing and coordination among disparate systems from multiple vendors. As such, a tremendous array of business problems can be addressed by simplifying and standardising the means by which systems can share data; which is what the Infoblox Orchestration Server does.

“Infoblox is committed to making IF-MAP technology available and affordable for organizations worldwide. With our unique expertise in providing highly available systems for distributing and managing mission critical data for IP-based systems, Infoblox is well positioned to help IF-MAP take its place as an important global standard,” said Richard Kagan, executive vice president and general manager of Infoblox’s Orchestration Systems Business Unit.

Prices of the Infoblox Orchestration Server starts at \$6,595 in Europe and Asia Pacific.

Converged Infrastructure the future?

Hewlett-Packard (HP) had announced the completion of its acquisition of 3Com just over a month ago in April and with this it's able to integrate its HP ProCurve LAN edge products with 3Com's routing, security, data centre and enterprise campus core switching solutions to be able to deliver efficient and secure business services.

Its portfolio is a key foundational item for a converged infrastructure, which enables clients to eliminate silos and integrate multiple vendors' technologies, so they can deliver flexibility to support changing business needs.

“The first wave of networking in the 1980s and 90s was driven by Novell, 3Com, Sun, IBM, DEC and Nortel, and focused on connectivity,” said Jay Mellman, HP Networking division senior director of product and solution marketing worldwide.

“The second wave during the first decade of the 2000s was driven mostly by Cisco and consolidated on Ethernet and its extension.

“The third wave which will dominate the second decade (of this century) will be based on orchestrated service delivery and a standards-based approach led by HP,” Mellman claimed.



HP's Converged Infrastructure concept

HP's vision of IT of the future is that it will be built based on a converged infrastructure, connecting servers, storage, power & cooling, management software and networks, and HP aims to provide an end-to-end portfolio of networking products, solutions and services to enable this.

Its horizontal product solutions include data centre networking, enterprise routing, end-to-end security, LAN switching, branch office, mobility and unified communications solutions. These are based on its A, E, V and S family of networking products and its end-to-end networking, security, unified communications & collaboration and integrated business solutions.