Pipeline
Knowledge Is Power

Over-the-Top Services
by Wedge Green and Barbara Lancaster

The future, according to some scientists, will be exactly like the past, only far more expensive.
  - John Sladek

A Decade of "Writing on the Wall"

Perhaps the first person brave enough (or crazy enough) to voice the implications of the power of edge based services within the Internet, and what it might mean to Service Providers, was David Isenberg. He wrote *The Rise of the Stupid Network* in the summer of 1997 (10 years ago) and soon thereafter departed AT&T. This article was controversial and still is – in fact Isenberg has become something of the "Alan Ginsberg" of the Internet. The implications were very clear: the trends David identified would kill service provider business models and maybe the companies if service providers could not respond.

Wedge Greene debated David that fall at the IETF – arguing that a different, software-based, smart network was possible and could save Service Providers. In fact, both were right and wrong. Wedge's stated vision of many small interacting software services hosted on platforms is sweeping the Internet; however, service providers are not driving this deployment or even benefiting from it. David seems to be right that these trends are outpacing the ability of the service provider to alter its business model. What both engineers predicted correctly was that edge user-developers would become legion. Swamping in game theory is an effective strategy, but will the Stupid Network prevail in the long run? Will the good of the many outweigh the good of the few?

Can Spock... I mean the 'intelligent network'... be resurrected?

Today this debate rages. Partly this debate is about business models, partly about legality and regulation, partly about what technology to allow and to leverage, and lastly it is about consumer and business desire for services. But what does this mean for the management of networks, for OSS & BSS? What is changing? What is no longer realistic and what new opportunities are opening up? How can intelligent planning, good policies, proper technology and well crafted processes turn the tide of operator fear toward operator profits?

Users simply want to have access to their favorite services at any place from any

consumer equipment they own. It is increasingly hard to tell a consumer that it is economically different to connect to the network via a land phone, cell phone, or computer. In fact, younger users fail to separate out the three networks, inside their head, as separate networks. Thanks to cyber-punk science fiction and popular movies, they see not a layered and interconnected series of discreet networks, but one organic whole network. Service providers long have counted on the steady delivery of access and core services that everyone will want and use regularly, and that, most importantly, will have enough economic life-span to justify service development, management development, and network tailoring for best delivery. Unfortunately, today's users just are not waiting on this service delivery model. Traditional, solid NPI soon may be obsolete. As Japan saw with the introduction of DoCoMo, services today are very much like fads, racing like wild-fire through the user community. Many are short lived; most are displaced by marginally better or cooler services. A few become mainstays of user utility. What these have in common is offering no direct control by service providers, or generating any new revenue for network operators.

What is OTTS?

Over-the-top (OTT) services is the buzz-expression for services carried over the networks, delivering value to customers, but without any carrier service provider being involved in planning, selling, provisioning, or servicing them – and of course without any traditional telco booking revenue directly from them.

Currently, according to the Alexa Global Traffic Rankings the major destinations for traffic in the internet are, in order, Google, Yahoo, MySpace, YouTube, Facebook, Windows Live, eBay, Wikipedia.org, msn.com and Craigslist. All of these portals are not owned by networks and most host OTT services which provide value to end users, and for which the traditional carriers get no specific revenue. Portals are receiving revenue for advertisements but many critics maintain they are not paying enough for transport of the data that allows the users to click those ads. Social networks like Facebook, MySpace, Bebo, Friendster, Tagged, and Linked-In are capturing millions of user connectivity hours.

It is already clear that in the future the network operator will not "own the customer." This is the fall of a main tenet of service provider business models. It has profound implications on OSS/BSS. If the customer is not ours, why spend so much insuring customer satisfaction? Why deploy Contact Centers? If you cannot bill for a service, why should we care how a service performs? What will we do with these massively expensive revenue assurance platforms?

If the "customer ownership" is lost, what avenues are left for the service provider to leverage? Service providers probably can hold onto the "last mile" of access – at the home and at the mobile device, not as the *only* supplier of services, but as the provider of the best quality, most comprehensive access to information. Deep Packet Inspection is increasingly possible in real time for broadband networks. Likewise on cable plants, DOCSIS can be used to engineer QoS based on the type of data stream. These technologies might allow fair use traffic management, at least at the network edges, where connections are usually of lower speed, but will these be used to provide a quality experience for the user or as a toll bridge to charge all comers?

What else can networks offer besides QoS for bits? It is possible today to determine the specific applications that generate the traffic. QoS specifically tailored to services and applications (for example, shaping traffic and setting latency to maximize Oracle transaction rates) is a perhaps the strongest card the service provider has for the future.

For a while, Service Providers may be able to leverage specific customer information, particularly near-real-time information contained in services like presence and location, especially where presence, location, and the specific technology via which the customer is connected (right now) can be linked together. And the embedded internal platforms for OSS/BSS services, like monitoring, reporting, billing, and settlement transactions, might be turned outward as collaborative services (such as aggregated billing for distribution of a media companies content). And maybe, trusted company status could allow for service providers to continue supplying "trusted," secure, and always available services. But eventually, even these could be usurped by web companies.

OTT services actually are a mix of technologies, vendors, and platforms. It includes:

- Downloaded applications which use point-to-point technology and the compute-power of the users computer to deliver IP services

- Peer-to-Peer, mesh network delivered and supported applications, which leverage many user platforms, most of which do not directly participate in the service transaction, but which silently loan capacity

- Widgets hosted on Portals, meant to keep users inside the portal environment

- RSS feeds of all sorts of media

Each of these requires different policies for the Service Provider and each creates different management challenges in the network. Each may also require different models for revenue recapture.

Indeed, in the widget world, taking the time for software quality and user acceptance testing is a business recipe for suicide. The user community is the test bed. Developers get their application out in front of users first, and then fix it later,

and if it is successful, scale it later. J-Squared Media delivered its first product widget, a Sticky Note application, onto the Facebook environment only a week after Facebook opened up its API to third-party developers. Instead of quality control, these developers rely on social network discussion groups where users directly post praises, problems, and wish lists. This is NPI turned on its head! This is totally alien to the IT departments of traditional service providers. Ironically, though, it is familiar to service provider Operations Groups.

NOCs have long had an underground tradition of cobbling together scripts and tools to rapidly fill the gaps in management and customer service left by bloated, overdue, or ill-designed management applications. (Indeed, the first Frame Relay performance report was generated by an MCI operations team to answer a customer's skepticism on the flow of traffic over a recently turned up circuit.) If we enable the NOCs with the same platforms and tools the widget developer community has, can we accelerate management development? But who will be the test/feedback community? And can service providers ever afford to fail at such a "tossed over the fence" application?

Open Source platforms are also becoming a real thorn, or potential opportunity, for traditional service providers and network/service management. Google is plowing ground here with two recent announced platforms: *OpenSocial* for developing widgets and an open source mobile operating system, *AndroidPhone*. Wide acceptance occurred for OpenSocial (even Facebook supports it) and we can expect similar wide acceptance of AndroidPhone. Open source provides a distributed test environment where rapid widget (and agent) development can undergo quality review and subsequent tuning, but using this approach means service providers must "give up intellectual property" to competitors in order to get back tested, improved applications. Let's start this process by calling this "collaboratively sharing intellectual property" instead of "giving it away."

What these widgets and open source platforms currently lack is an enforceable, embedded management model. NGOSS (at least Finegrain NGOSS) championed the idea of every service/application/component having not just a known service interface, but also a standard management interface. Yes, NGOSS services are supposed to all be instrumented. If such deployment and management interfaces become a standard, predictable part of all OTT services, then the opportunity exists to manage not just them, but to tune networks based on their use and distribution. This means extending the notion of "Governance," now a normal part of SOA, ESB, and Grid platforms, to OTT services. Yet can this happen in this "wild west" environment?

Net Neutrality vs. the Walled Garden

Traditionally service providers have attempted to maintain a walled garden of services and devices. The US mobile operators are a notorious example today with their policy of "locking" phones to networks, but Apple has turned this model back against the mobile operators themselves in seeking to direct all services and management of the iPhone back to Apple – even to permanently disabling phones that break "locked" arrangements. Walled gardens allow for billing by service and by device. It also provides for a "captured" customer. Unfortunately, most users hate being limited by walled gardens.

It is easy to become alarmist in the face of the geometric increase in the rise of OTT services. 'Increasingly, the relationship between revenue, investment, and bandwidth utilization is eroding…. OTT services can engender either unprecedented network contention or massive infrastructure costs with no offsetting increase in revenue." [Barry Hardek, Vice President of Marketing & Sales, CableMatrix in TelephonyWorld] So, can and should operators act to throttle or suppress the growth of OTT services?

Vint Cerf has stated, "The Internet was designed with no gatekeepers over new content or services. A lightweight but enforceable neutrality rule is needed to ensure that the Internet continues to thrive." **Network neutrality** maintains that networks carrying broadband or IP traffic shall be free of restrictions on the kinds of equipment that may be attached, the types of communication allowed, and where communication is not unreasonably degraded by owners of the underlying infrastructure. As IP was developed by the government, there is a sound legal basis for such requests.

However, critics argue that 'net neutrality' potentially reduces incentives to build and upgrade networks and launch new network resident services. Basically, the network owners must receive compensation for building and maintaining the network - or market forces (not allowing for altruism) will punish these companies and perhaps drive them under. Then the government would be forced into either a bail out or a take over of the networks. (Remember what happened to passenger railroads.)

 "There is also a lack of incentive for access providers to invest in capacity on behalf of Google to deliver richer, heavier content (where Google collects the revenues)." [Simon Torrance, Telco 2.0] Google understands this and is determined to control its own destiny. Google's answer may be to build network itself.

Ironically, the use of P2P services to illegally copy and distribute copyrighted material may give the ammunition for network operators to throttle P2P traffic in general. The self-same users, who have come to depend on P2P mesh network delivered services, may be shooting themselves in the foot. On the other hand, lawful content distribution has become so complex to understand that building justifiable business policies that are defendable to irate customers is not currently possible.

Compromise may be found in the business model of differential charging to guarantee "Quality of Service." Unfortunately, this can lead to poor service if QoS is set below that needed for specific protocols and information streams; this is called throttling traffic. However, it is not very clear who pays the network owner for the QoS: sender, receiver, producer, user, network handing off the traffic, network receiving the traffic? As the government is pulled into this debate, will the government begin to regulate throttling or QoS such that it becomes the new tariffed service? The business model is still quite undetermined.

Telco 2.0

Telco 2.0 is a trademarked web portal of Simon Torrance. Simon has worked with operators and other clients to evolve a proposed business model that reacts to most of these issues driven by the advent of OTT services. This Telco 2.0 business model also assumes that a strong convergence will occur between the Media Content

companies and the Telco ecosystem. In this vision, the service provider becomes a "logistics solution provider" with a *dual-sided business model*. The SP receives payment from the content creators to ensure delivery of their specific service with appropriate QoS. It is also getting paid by the end customer for access and delivery of bits. It probably bills for the delivery of content and then aggregates the payments for the content rights distributor, skimming a transaction fee. And finally, SPs sell business intelligence information on customer use patterns back to the content creators and their competitors.

This optimistic model relies on an extended Ecosystem of partners. Specifically the Media, telecom service providers, consumer electronic telecom end-user-device manufactures, and platform vendors are essential members all collaborating in a bi-directional supply chain. The end user is an asset whose access is controlled by the service provider, which is the glue holding this construct together. This is a workable, but perhaps idealistic, business model. The TeleManagement Forum leadership and its Service Provider council apparently back this approach - This is what fueled the expansion of the TM Forum to include Media companies. But where are the OTT service developers and platforms in this model? We do not see them. Indeed, Telco 2.0 while plausible also does not seem to have a clear story of OSS/BSS involvement beyond revenue aggregation and billing for transaction services.

The recent TM Forum "*Content Encounter*" embraced Telco 2.0 in a super-Catalyst demonstration. It also included DRM and "media content envelopes" as a telco ecosystem service product, aimed at exciting media and content companies with control of, and usage revenue for, their intellectual property. It also demonstrated an operator providing everyone who pays for it with customer use information allowing classification of users and directed market action targeting these same end users. This seems the ideal "get paid at both ends" solution. However, the typical customer using OTT services hates DRM. So will they hate these enveloped media services and not buy these services? Is the service provider opting into the same "bad guy" image of the digital rights distribution companies? Or will the users just break DRM again, wasting everyone's development efforts?

Will users even need the envelope content information maintained in these media protocols? On the surface, this is great stuff. However, today it is easy enough to find music information via a web-service data lookup to known music library sites. Likely this kind of information service will only proliferate; the services themselves paid for via alternate revenue routes such as selling customer media use patterns back to media creators.

And this Telco 2.0 media logistics chain, and the missing OTT services, affects Cabelco's and Telco's in different ways. Cable delivery plants will be pressed to keep up with bandwidth capacity challenges as OTT video explodes; yet they are competitively accelerated by have existing relationships with media companies and an understanding of this market. Telcos with fiber and switched VDSL probably can support capacity demands, but may founder in understanding how to leverage legitimate media content. We are not used to seeing cabelcos and telcos collaborating, or even waiting for each other to catch up; instead they compete.

Lastly, every time the operators train up staff & engineers and then downsize, non-traditional operators (including OTT service vendors) get human resources: experienced human capital that can build networks and services not under the

economic control of traditional Operators. Google is not buying an Operator because they bought the talent instead and can now build a network tailored for their needs. For example, Henry Sinnreich, who led MCI's advancement of SIP, is now at Google helping them build services that do not need Operators.

SIP vs. IMS

SIP has recently been called the "third great protocol of the Internet." SIP enables a wide variety of services to be composed and delivered in near-real-time. 3GPP and IMS standards groups choose to embrace SIP instead of something like SS7. SIP now fuels IMS, but does SIP need IMS? Indeed, no. SIP is itself architecturally neutral and IMS is a specific (and rather elaborate) architecture. SIP can be deployed in P2P mesh models, in heavy-edge-client models, and in platform embedded widget models. Today, SIP is used somewhere in all of these. Also, development environments like Astrix are providing a bridge from software code to telco services – with development times in days rather than months. This is showing dramatically that users can get fancy services and receive converged data/media/voice without service providers. But will these be quality, managed services?

Some say this dual trend means the death of IMS, but IMS and 3GPP also include a Policy QoS model. This may serve to be *the* important feature which gets IMS deployed after all. And this is very closely related to the management of quality networks. Using identification of service type, a known registry of user needs and desires, and a flexible policy model of specifically how to tune the network to deliver consistent quality services is, in our view, the most solid business model for today's service provider. This should allow the provider to bill for consistent QoS for specific services delivered to specific users with specific product tariffs. Who the service provider bills then does not matter - as long as someone will pay… and if no one pays, then the Wild West, best-effort environment ensues. And the very teenagers driving OTT services may drive the best-effort network into effective gridlock. Then media rights owners and advertising driven portals, and even household parents, will step up to pay for the value of predictable quality.

So we see the advent of OTT services not as a threat to service providers, but potentially as an ironic opportunity. Success, however, will go only to the agile service providers who plan and act effectively, and not to those who continue to stick with their traditional tendency toward bloated network and management projects.

Key OSS/BSS needs

So, since we are not likely to see our industry end in ruin, let us return the implication of OTT services on OSS/BSS.

Current peak-usage metric models for managing and engineering network traffic probably do not reflect the assumptions and use patterns of current OTT services. Longer, more sustained-rate transmissions may become the norm instead of the exception. The old statistical models used to compute reasonable over-provisioning of access vs. core capacity will not hold in this new OTT service world. Billing based on inexpensive access capacity depends on over-provisioning for profit. But profits will be squeezed as access broadband becomes a near constant up and down

stream – a data fire hose of un-billable services, so over-provisioning will disappear and access charges will increase – or become service dependent.

Current firewall and security services are not designed to support the many widespread connections (to other user machines) that are set up by P2P mesh applications. These mesh networks are used by such dominate gaming products as World of Warcraft (WOW) to download updates while playing continues. These distributed updates generate equal upstream and downstream traffic. If WOW users will buy high-end machines to play competitively, someone will pay for tuned, consistent quality network connections to the WOW servers.

Information paths routed across many separate carriers with guaranteed service will be common place. There will need to be an increasing collaboration of information among these networks. Security gateways will become common, as will arbitrator gateways for transferring QoS contracts; so also the related SLA arbitration. Quad-play will become commonplace and lay a considerable burden on near-real-time management and the ramping to scale of that management. Settlement will become much more complex – often involving multi-party transactions with elaborate and obscure rules – making settlement a policy-driven application. [Just look up the distribution rules in DRM contracts in the media world to see just how elaborate this can get.]

Lastly, deep packet inspection leading to differential QoS treatment for differential charging will become the norm of all edge and gateway interconnection points. This, of course, will require great leaps in mediation services (aggregating billing transactions) and will increase the relative importance and value of very-high-rate mediation and aggregation billing components.

Lastly, customer service guarantees, given to end users, businesses, and upstream members of the logistics flow will require enhanced customer service. This will drive the advancement of new architectures and technologies in the Contact Center and require active, close integrations with OSS/BSS applications.

We think that Telco 2.0 is correct that the telecom ecosystem must move away from relying on the end user as the sole revenue source and towards a model of not just dual directional revenue, but indeed multidirectional revenue streams - and so also include the user as a resource and a partner as well as a customer. However, every business model needs to deliver value to all parties in order to be successful. Service providers have a long way to go to demonstrate that they deserve to be paid by both the supplier and the buyer for the same transaction.

As we have said before, the technical solution to these (and other) problems might be smart agent subsystems. You should begin planning for the deployment of smart agents on or near end-user devices and at access nodes (home gateways & aggregations points). These agents will access device-level inspection of services and then apply differential treatment based on policy. They will funnel upstream information on use. They will not just control traffic but also allow billing based on service policies. And they may seek and collaborate with other distributed agents to coordinate the delivery of converged services at scale.

It seems that P2P mesh networks may themselves become applicable in the management of networks. This technology can distribute smart software agents as needed, into the very points where P2P and other OTT services create traffic. P2P meshes will then distribute data to these agents for the QoS, security, and use-

rights polices used by the smart agents, enabling near-real-time service QoS control. Thus you gain service agility. The software service "intelligent" network can rise from the explosion of OTT ecosystem services. Just like Spock did.

If you have news you'd like to share with Pipeline, contact us at editor@pipelinepub.com.