



AT&T and the Future of Convergence

■ A CONVERSATION WITH JOE WEINMAN



In a wide-ranging interview presented by AT&T, the company's VP of Strategy and Emerging Services, Joe Weinman, offers his insights into telecom, IT and how they will converge.

■ *Convergence means different things to different people these days. What is your view of convergence?*

Many people focus on network convergence, where multiple legacy protocols are rapidly being replaced by IP / MPLS

as a single, multi-purpose network infrastructure. However, a broader set of themes has the power to be even more transformational. One of these is the convergence of IT and communications. Others are increased flexibility, agility,

and utility. Flexibility is increasing at all layers: bare-metal dynamic provisioning of servers, VLANs, firewalls, load balancers, storage, and even middleware and applications delivered via a software-as-a-service paradigm. In the WAN,

flexibility is increasing through virtualization, bandwidth on demand, QoS, route control, and even packet control. Agility, that is the provisioning speed to take advantage of that flexibility, is accelerating. And utility, or pay-for-use models as an adjunct to flat-rate payment models, is increasingly being applied throughout managed services.

■ *How is convergence influencing AT&T's market focus?*

Traditionally, our market focus was access, MAN and WAN voice and data services. However, this is being extended in two directions: WAN services are being extended forward into a greater customer premises equipment and mobile endpoint presence, and back into the cloud via network-resident services built on our hosting infrastructure. Closing the BellSouth (and Cingular) acquisitions cements our presence in terms of a broad range of enterprise mobility services. Additionally, we have ramped up customer premises solutions not only through managed routers, firewalls and other network elements such as IP PBXs, and also RFID (Radio Frequency Identification) edge servers, high definition immersive video conferencing, integrated branch solutions, and more. We now have over 30 hosting locations on four continents, which provide the foundation for an asset-based strategy focused on managed, hosted services, including electronic product code information services, utility storage

services, managed utility computing services, etc.

■ *In what other ways is the telecom environment evolving?*

Besides moving forward into a larger customer presence and into the cloud for hosted assets, the other key direction that we're moving in is up the stack beyond infrastructure. The next layer up could be termed a services layer, then an applications and content layer, and ultimately a management and security layer.

The services layer is the middleware that ties infrastructure components together. For mobile devices this would include presence and location services. For hosted services, it would include services oriented architecture components. For call and connection management, it includes IMS (IP Multimedia Subsystem). When we talk about IT and telecom converging, this is a key layer. Take, for example, a calendar-based variable messaging service – if your calendar shows that you're in a meeting, callers might be directed to your cell phone rather than your landline, or if your boss calls, you might hear a different type of ring.

Applications and content can reside on this infrastructure, whether it be search, CRM, ERP, content distribution, database access, or the like.

Finally, the management and security layer ensures that the other layers all work, whether due to changing external conditions or in the face of logical and physical threats.

■ *Where does the consumer fit into the changing telecom picture for AT&T?*

We're seeing growth into applications and content, and that's true whether we look at the consumer side or the business side of the company. With recent acquisitions and investments, we now have the opportunity to deliver broadcast quality video content to lots of consumer devices, be they immersive video conferencing rooms, large-screen televisions at home or smaller form-factor mobile devices like cell phones. In addition, many applications are starting to cross the boundary between consumer and business.

In addition to the consumer context, what we're seeing is that this new converged end-to-end infrastructure provides the ability for lots of new types of applications that leverage premises and mobile end points, network access and transport, and managed, hosted solutions. An example would be radio frequency identification and electronic product codes, where you can use a mobile reader to capture a transaction, such as a case of soup cans making its way through a distribution facility, then use further end-point devices such as RFID edge servers to capture that data off of the mobile readers, possibly operating in conjunction with a managed premises wireless LAN. Then you might feed those transactions across the WAN and ultimately into a managed hosted electronic-product-code information service that maintains those transactions and can support data mining and ERP systems integration.

■ *You've described a telecom environment with some very complex capabilities. How are these capabilities being enabled?*

Ultimately, making all of these layers work together requires a management layer, and we're seeing a lot of evolution around what the nature of that management is. First of all, it isn't just WAN management any more; it's managing to the mobile and/or customer premises devices as well as managing the hosted functionality in the cloud. Secondly, it isn't just about managing infrastructure any more, it's now about managing the services and applications and content. Finally, increased flexibility mandates not just network element management, but intelligent decision-making, policy-based orchestration, and optimization of the various components of the environment such as servers, storage, and services.

■ *What is AT&T doing to ensure it can provide this kind of functionality?*

We've already made a substantial investment of over US\$300 million in a system called iGEMS, which is the integrated Global Enterprise Management System. It has a substantial number of monitoring, management, and intelligent correlation tools built in for multiple layers, such as network elements, be they WAN or LAN; security elements such as firewalls and security defence appliances; storage such as enterprise arrays; servers; and even applications. It's going to continue to increase in

functionality to handle these more complex environments as next-generation dynamic pricing models and utility infrastructures take hold over the next 12 to 36 months.

■ *Can you describe some specific areas that AT&T is targeting?*

There are quite a few themes in terms of the types of things that we focus on. The first is clearly network-centric applications. As an example, we're now starting to move increasingly into video applications, such as video monitoring, IPTV delivery, or high-definition videoconferencing, which can use 15 megabits per second on a pretty steady basis. The result is that more than ever before we're seeing increasing bandwidth not just for data-centre applications, but also for business applications like voice and videoconferencing.

As such, reliability and quality of service are crucial, especially as we move to converged networks. It will be essential to make sure that VoIP doesn't interfere with video over IP, and ensure that lower class of service applications like file transfer or business data applications don't interfere with real-time applications. And it's not just network-centric applications that are pushing the envelope; it's also network-resident services like software as a service, video broadcasting, and e-commerce.

■ *What is AT&T's take on this global online gaming phenomenon?*

A lot of people think of gaming as a niche application – one

that is growing rapidly across multiple demographics, but still a niche application. We see it as more than that. Gaming represents the vanguard of what applications and interactive Web experiences will look like in the future. Today's text, audio, and graphics animation is going to evolve to seamlessly mixed video, graphics, and real-time interaction. For example, if you're on a customer support call with your washing machine company, rather than just seeing an image or even a small video chat window, you're going to get into an immersive, real-time, interactive experience where you're going to be able to interact with your customer service agent, and anyone else that he or she conferences into the discussion. While they might actually be sitting at home, in the next-generation distributed, IP-enabled contact centre, it's going to look as though they're all in the same room, or perhaps at the factory, showing you in real-time what button to press on your washing machine.

■ *What other impacts do you see as we go through these multiple types of convergence?*

One thing we're seeing a lot of interest in is next-generation high-definition videoconferencing, which provides a 'you-are-there' experience through high-definition 720p or 1080p monitors (depending on the vendor), multiple screens, life-size images with photo-realistic quality, and extremely low latency even at globally dispersed distances. We see that as an example of the type of application that requires functionality in the cloud – for

example, multi-point bridging, a rich set of flexible, highly performant network services, such as those provided by AT&T's global IP/MPLS network, and extremely capable end-points with a rich set of functionality. Tying it all together, there is integration through the services layer so that, for example, the videoconference is set up automatically when scheduled through the Outlook calendar.

Yet despite the recent interest in high-definition videoconferencing, we see it as one more component of 'video everywhere'. What we're going to have in the environments that people are deploying is standard-definition legacy videoconferencing; high-definition video, available with leading vendors; mobile video, which is now moving from low frame-rate, asynchronous video-clip transfer to real-time interactive mobile videoconferencing as well as mobile video broadcasting; desktop video solutions, both standard definition and now increasingly high definition with next-generation codecs; and on-demand (archived) and live enterprise broadcast video, as well as the digital asset management that goes along with those archived broadcasts.

Over time, there will also be the ability to do seamless video mobility. If you're taking a video call from your mobile device or cell phone as you come off a cellular network and then go to a wireless LAN, that call will be transferred seamlessly. In addition, it will ultimately be transferred from a low-resolution device, like a mobile video device, to possibly a high-definition screen with a high bandwidth connection. Other application functions

will go along with this, like video search and automatic speech-to-text translation for metadata tagging.

Another big trend is an increased focus on integrated offerings, where rather than having single piece parts that customers do themselves, there might be single boxes that provide increased functionality. In other words, when you deploy one physical element on the customer premises, it has various types of built-in functionality, like firewall, router, terminal adaptor, wireless access point, etc.

■ *How will advances in telecom help CIOs address some of their key needs in the future?*

CIOs are cognizant of the disconnect between fixed investment in static infrastructure and the vast majority of applications which have variable demand. Most CIOs would tell you that the average utilization of their environment is less than ten percent, which means that they spend, in effect, ten times as much as they need to for their physical environment. To deal with these issues, they traditionally have had only two choices: cut back a little on their environment, in which case they have insufficient capacity to meet peak demand, or overprovision in the range of peak plus ten percent, which means that they run at single-digit percentage utilization levels.

What the industry is starting to focus on is a variety of pay-per-use utility models where that variable demand is met by a varying utility infrastructure, rather than being met by fixed infrastructure. This means

What the industry is starting to focus on is a variety of pay-per-use utility models which enable variable demand to be met by a varying utility infrastructure, rather than a fixed infrastructure. This means there will always be sufficient capacity and CIOs won't have to over-provision to meet peaks



there will always be sufficient capacity and CIOs won't have to over-provision to meet those peaks. We anticipate that the next couple of years will bring customers deploying three different types of environments and economically arbitrating the benefits of each. The first one is the traditional dedicated environment with a flat-rate payment, which might be owned or leased equipment or a fixed-rate managed service with a monthly charge for patch management, configuration management, break/fix, HVAC power, etc. The second model is fixed resources with a variable payment model, which means you have fixed CPU resources, but you pay more or less each month for that environment, based on the average CPU utilization for that month. A number of our customers have found substantial economic benefit to this approach. The third model is dynamically allocated pay-per-use, where the size of the environment and payments vary – the traditional hotel or rental-car model. This means you have a resource, use it for some time period, stop using it and then pay for what you used.

Our studies have shown that a mix of all three environments reduces the cost and maximizes the flexibility and

scalability for customers. Probably the most critical model is the dynamically allocated pay-per-use resources, because any kind of variable demand really increases the economic attractiveness of these models.

■ *Where does virtualization fit into this scenario?*

Virtualization is an important technology. For example, multiple applications that normally would have a low CPU utilization on their own can be consolidated onto a single physical server using server virtualization software, which we can manage for customers. When you combine utility models that provide dynamically allocated resources together with virtualization, you end up with extremely fine grained and flexible capacity increments for computing resources.

In addition to the virtualized pay-per-use utility model just for servers, we're finding that that model is actually applicable to all areas of infrastructure, and even services built on top of that infrastructure. Storage services, for example, where customers can buy a number of gigabytes per month; network services, where we're moving to greater

flexibility through a variety of bandwidth-on-demand and class-of-service-on-demand offers; and middleware applications, such as BEA's WebLogic® and AquaLogic®, which users can access through AT&T hosting centre and can pay for either based on percentage of CPU utilization and/or minutes of CPU use.

Ultimately, what we see is a converged network and geographically dispersed computing infrastructure, which provides all kinds of benefits ranging from scalability to next-generation load-balancing.

■ *What is happening with convergence in the security area?*

Multi-layer security is another area in which we're seeing convergence. In the old days you had IT security approaches like server logs, and network security approaches like firewalls. AT&T has a service called Internet Protect that's got a lot of uptake in the market, which monitors traffic on the AT&T IP backbone – now approaching 10 terabytes per day. The service looks at any anomalies or malware signatures and then uses proprietary anti-DDoS (distributed denial of service)

capabilities to then protect that traffic at the edge, and either quarantine it or drop it at the edge to keep it from ever getting into customer networks. With convergence, what we're going to see is an evolution of increasing sources of information – not just signature/anomaly detection, but security defence appliances, honeypots, tar pits, and server logs. Additionally, we'll see increasing variety of components that are lines of protection, not just on the network perimeter, but also on the intranet and at the physical servers, firewalls, routers, etc, that are the various components of the distributed architecture.

■ *Sounds like there is a lot going on and a lot to keep an eye on.*

Absolutely. Today's highly interconnected environments continue to evolve at an accelerating pace. The good news is that there are compelling benefits for enterprises: revenue growth, cost savings, risk mitigation, and enhanced flexibility in proactively responding to shifting market conditions, and AT&T is front and centre in delivering these benefits to clients.